
Ilvmlite Documentation

Release 0.13.0

Continuum Analytics

Apr 09, 2017

Contents

1	A lightweight LLVM python binding for writing JIT compilers	1
1.1	Introduction	1
1.2	Installing	2
1.3	llvmlite.ir – The IR layer	3
1.4	llvmlite.binding – The LLVM binding layer	19
1.5	Contributing	30
1.6	Glossary	31
1.7	Release Notes	32
2	Indices	37
	Python Module Index	39

A lightweight LLVM python binding for writing JIT compilers

Introduction

Overview

llvmlite is born of the desire to have a new Python binding to LLVM for use in [Numba](#). Numba used to rely on [llvmpy](#), but llvmpy became harder and harder to maintain, because of its cumbersome architecture and also because the C++11 requirement of recent LLVM versions don't go well with the compiler and runtime ABI requirements of some Python versions (especially under Windows).

Moreover, llvmpy proved to be responsible for a sizable chunk of Numba's compilation times, because of its inefficient layering and object encapsulation. Fixing this issue inside the llvmpy codebase seemed a time-consuming and uncertain task.

Therefore, the Numba developers decided to start a new binding from scratch, with an entire different architecture, centered around the specific requirements of a JIT compiler.

Philosophy

While [llvmpy](#) exposed large parts of the LLVM C++ API for direct calls into the LLVM library, llvmlite takes an entirely different approach. llvmlite starts from the needs of a JIT compiler and splits them into two decoupled tasks:

1. Construction of a [module](#), function by function, [instruction](#) by instruction
2. Compilation and optimization of the module into machine code

The construction of a LLVM module doesn't call the LLVM C++ API; rather, it constructs the LLVM *Intermediate Representation* in pure Python. This is the part of the *IR layer*.

The compilation of a LLVM module takes the IR in textual form and feeds it into LLVM's parsing API. It then returns a thin wrapper around LLVM's C++ module object. This is the part of the *binding layer*.

Once parsed, the module's source code cannot be modified anymore; this is where llvmlite loses in flexibility compared to the direct mapping of C++ APIs into Python that was provided by llvmpy. The saving in maintenance effort, however, is large.

LLVM compatibility

Despite minimizing the API surface with LLVM, llvmlite is impacted by changes to LLVM's C++ API (which can occur at every feature release). Therefore, each llvmlite version is targetted to a specific LLVM feature version. It should work accross all given bugfix releases of that version (for example, llvmlite 0.12.0 would work with LLVM 3.8.0 and 3.8.1, but with neither LLVM 3.7.0 nor 3.9.0).

Which LLVM version is supported is driven by [Numba's](#) requirements.

API stability

At this time, we reserve ourselves the possibility to slightly break the llvmlite API at each release. This is necessary because of changes in LLVM behaviour (for example differences in the IR accross versions), and because llvmlite as a young library still has room for improvement or fixes to the existing APIs.

Installing

Pre-built binaries

For your own convenience, we recommend you install the pre-built binaries provided for the [Conda](#) package manager. Official binaries are available in the [Anaconda](#) distribution; to install them, simply type:

```
$ conda install llvmlite
```

If you want a more recent version, you can also fetch the automatic builds available on [Numba's binstar channel](#):

```
$ conda install --channel numba llvmlite
```

If don't want to use Conda, or modified llvmlite yourself, you will need to build it.

Building manually

Prerequisites (UNIX)

You must have a LLVM 4.0 build (libraries and header files) available somewhere.

Under a recent Ubuntu or Debian system, you may install the `llvm-4.0-dev` package if available.

If building LLVM on Ubuntu, the linker may report an error if the development version of `libedit` is not installed. Install `libedit-dev` if you run into this problem.

Prerequisites (Windows)

You must have Visual Studio 2015 or later (the free "Express" edition is ok) in order to compile LLVM and llvmlite. In addition, you must have [cmake](#) installed, and LLVM should have been built using cmake, in Release mode. Be careful to use the right bitness (32- or 64-bit) for your Python installation.

Compiling

Run `python setup.py build`. This builds the llvmlite C wrapper, which embeds a statically-linked copy of the required subset of LLVM.

If your LLVM is installed in a non-standard location, first set the `LLVM_CONFIG` environment variable to the location of the corresponding `llvm-config` (or `llvm-config.exe`) executable. For example if LLVM is installed in `/opt/llvm/` with the `llvm-config` binary located at `/opt/llvm/bin/llvm-config` then set `LLVM_CONFIG=/opt/llvm/bin/llvm-config`. This variable must be persisted through into the installation of llvmlite e.g. into a python environment.

Installing

Validate your build by running the test suite: `run python runtests.py` or `python -m llvmlite.tests`. If everything works fine, install using `python setup.py install`.

llvmlite.ir – The IR layer

The `llvmlite.ir` module contains classes and utilities to build the LLVM *Intermediate Representation* of native functions. The provided APIs may sometimes look like LLVM's C++ APIs, but they never call into LLVM (unless otherwise noted): they construct a pure Python representation of the *IR*.

See also:

To make use of this module, you should be familiar with the concepts presented in the [LLVM Language Reference](#).

Types

All *values* used in a LLVM module are explicitly typed. All types derive from a common base class `Type`. Most of them can be instantiated directly. Once instantiated, a type should be considered immutable.

class `llvmlite.ir.Type`

The base class for all types. You should never instantiate it directly. Types have the following methods in common:

as_pointer (*addrspace*=0)

Return a `PointerType` pointing to this type. The optional *addrspace* integer allows you to choose a non-default address space (the meaning is platform-dependent).

get_abi_size (*target_data*)

Get the ABI size of this type, in bytes, according to the *target_data* (a `llvmlite.binding.TargetData` instance).

get_abi_alignment (*target_data*)

Get the ABI alignment of this type, in bytes, according to the *target_data* (a `llvmlite.binding.TargetData` instance).

Note: `get_abi_size()` and `get_abi_alignment()` call into the LLVM C++ API to get the requested information.

__call__ (*value*)

Convenience method to create a `Constant` of this type with the given *value*:

```
>>> int32 = ir.IntType(32)
>>> c = int32(42)
>>> c
<ir.Constant type='i32' value=42>
>>> print(c)
i32 42
```

Atomic types

class `llvmlite.ir.PointerType` (*pointee*, *addrspace*=0)

The type of pointers to another type. *pointee* is the type pointed to. The optional *addrspace* integer allows you to choose a non-default address space (the meaning is platform-dependent).

Pointer types exposes the following attributes:

addrspace

The pointer's address space number.

pointee

The type pointed to.

class `llvmlite.ir.IntType` (*bits*)

The type of integers. *bits*, a Python integer, specifies the bitwidth of the integers having this type.

width

The width in bits.

class `llvmlite.ir.FloatType`

The type of single-precision floating-point real numbers.

class `llvmlite.ir.DoubleType`

The type of double-precision floating-point real numbers.

class `llvmlite.ir.VoidType`

The class for void types; only used as the return type of a function without a return value.

Aggregate types

class `llvmlite.ir.Aggregate`

The base class for aggregate types. You should never instantiate it directly. Aggregate types have the following attribute in common:

elements

A tuple-like immutable sequence of element types for this aggregate type.

class `llvmlite.ir.ArrayType` (*element*, *count*)

The class for array types. *element* is the type of every element, *count* the number of elements (a Python integer).

class `llvmlite.ir.LiteralStructType` (*elements*)

The class for literal struct types. *elements* is a sequence of element types for each member of the structure.

Other types

class `llvmlite.ir.FunctionType` (*return_type*, *args*, *var_arg*=False)

The type of a function. *return_type* is the return type of the function. *args* is a sequence describing the types of argument to the function. If *var_arg* is true, the function takes a variable number of additional arguments (of unknown types) after the explicit *args*.

Example:

```
int32 = ir.IntType(32)
fnty = ir.FunctionType(int32, (ir.DoubleType(), ir.PointerType(int32)))
```

An equivalent C declaration would be:

```
typedef int32_t (*fnty)(double, int32_t *);
```

class `llvmlite.ir.LabelType`

The type for *labels*. You don't need to instantiate this type.

class `llvmlite.ir.MetaDataType`

The type for *metadata*. You don't need to instantiate this type.

Note: This used to be called `MetaData` but was renamed for clarity.

Values

Values are what a *module* mostly consists of.

llvmlite.ir.Undefined

An undefined value (mapping to LLVM's "undef").

class `llvmlite.ir.Value`

The base class for all IR values.

class `llvmlite.ir.Constant` (*typ, constant*)

A literal value. *typ* is the type of the represented value (a *Type* instance). *constant* is the Python value to be represented. Which Python types are allowed for *constant* depends on *typ*:

- All types accept *Undefined*, and turn it into LLVM's "undef".
- All types accept `None`, and turn it into LLVM's "zeroinitializer".
- *IntType* accepts any Python integer or boolean.
- *FloatType* and *DoubleType* accept any Python real number.
- Aggregate types (i.e. array and structure types) accept a sequence of Python values corresponding to the type's element types.
- In addition, *ArrayType* also accepts a single `bytearray` instance to initialize the array from a string of bytes. This is useful for character constants.

classmethod `literal_array` (*elements*)

An alternate constructor for constant arrays. *elements* is a sequence of values (*Constant* or otherwise). All *elements* must have the same type. A constant array containing the *elements* in order is returned

classmethod `literal_struct` (*elements*)

An alternate constructor for constant structs. *elements* is a sequence of values (*Constant* or otherwise). A constant struct containing the *elements* in order is returned

bitcast (*typ*)

Convert this pointer constant to a constant of the given pointer type.

gep (*indices*)

Compute the address of the inner element given by the sequence of *indices*. The constant must have a pointer type.

inttoptr (*typ*)

Convert this integer constant to a constant of the given pointer type.

Note: You cannot define constant functions. Use a *function declaration* instead.

class llvmlite.ir.**Argument**

One of a function’s arguments. Arguments have the following method:

add_attribute (*attr*)

Add an argument attribute to this argument. *attr* is a Python string.

class llvmlite.ir.**Block**

A *basic block*. You shouldn’t instantiate or mutate this type directly; instead, call the helper methods on *Function* and *IRBuilder*.

Basic blocks have the following methods and attributes:

replace (*old*, *new*)

Replace the instruction *old* with the other instruction *new* in this block’s list of instructions. All uses of *old* in the whole function are also patched. *old* and *new* are *Instruction* objects.

function

The function this block is defined in.

is_terminated

Whether this block ends with a *terminator instruction*.

terminator

The block’s *terminator instruction*, if any. Otherwise None.

class llvmlite.ir.**BlockAddress**

A constant representing an address of a basic block.

Block address constants have the following attributes:

function

The function the basic block is defined in.

basic_block

The basic block. Must be a part of *function*.

Metadata

There are several kinds of *metadata* values.

class llvmlite.ir.**MetaString** (*module*, *value*)

A string literal for use in metadata. *module* is the module the metadata belongs to. *value* is a Python string.

class llvmlite.ir.**MDValue**

A metadata node. To create an instance, call *Module.add_metadata()*.

class llvmlite.ir.**DIValue**

A debug information descriptor, containing key-value pairs. To create an instance, call *Module.add_debug_info()*.

class llvmlite.ir.**DIToken** (*value*)

A debug information “token”, representing a well-known enumeration value. *value* should be the enumeration name, e.g. 'DW_LANG_Python'.

class `llvmlite.ir.NamedMetaData`

A named metadata node. To create an instance, call `Module.add_named_metadata()`. Named metadata has the following method:

add (*md*)

Append the given piece of metadata to the collection of operands referred to by the NamedMetaData. *md* can be either a *MetaDataString* or a *MDValue*.

Global values

Global values are values accessible using a module-wide name.

class `llvmlite.ir.GlobalValue`

The base class for global values. Global values have the following writable attribute:

linkage

A Python string describing the linkage behaviour of the global value (e.g. whether it is visible from other modules). Default is the empty string, meaning “external”.

storage_class

A Python string describing the storage class of the global value. Default is the empty string, meaning “default”. Other possible values include “dllimport” and “dllexport”.

class `llvmlite.ir.GlobalVariable` (*module*, *typ*, *name*, *addrspace*=0)

A global variable. *module* is where the variable will be defined. *typ* is the variable’s type. *name* is the variable’s name (a Python string). *addrspace* is an optional address space to store the variable in.

typ cannot be a function type. To declare a global function, use *Function*.

The returned Value’s actual type is a pointer to *typ*. To read (respectively write) the variable’s contents, you need to *load()* from (respectively *store()* to) the returned Value.

Global variables have the following writable attributes:

global_constant

If true, the variable is declared a constant, i.e. its contents cannot be ever modified. Default is False.

unnamed_addr

If true, the address of the variable is deemed insignificant, i.e. it will be merged with other variables which have the same initializer. Default is False.

initializer

The variable’s initialization value (probably a *Constant* of type *typ*). Default is None, meaning the variable is uninitialized.

align

An explicit alignment in bytes. Default is None, meaning the default alignment for the variable’s type is used.

class `llvmlite.ir.Function` (*module*, *typ*, *name*)

A global function. *module* is where the function will be defined. *typ* is the function’s type (a *FunctionType* instance). *name* is the function’s name (a Python string).

If a global function has any basic blocks, it is a *function definition*. Otherwise, it is a *function declaration*.

Functions have the following methods and attributes:

append_basic_block (*name*='')

Append a *basic block* to this function’s body. If *name* is non empty, it names the block’s entry *label*.

A new *Block* is returned.

insert_basic_block (*before*, *name*='')

Similar to `append_basic_block()`, but inserts it before the basic block *before* in the function's list of basic blocks.

set_metadata (*name*, *node*)

Add a function-specific metadata named *name* pointing to the given metadata *node* (a *MDValue*).

args

The function's arguments as a tuple of *Argument* instances.

attributes

A set of function attributes. Each optional attribute is a Python string. By default this is empty. Use the `add()` method to add an attribute:

```
fnty = ir.FunctionType(ir.DoubleType(), (ir.DoubleType(),))
fn = Function(module, fnty, "sqrt")
fn.attributes.add("alwaysinline")
```

calling_convention

The function's calling convention (a Python string). Default is the empty string.

is_declaration

Whether the global function is a declaration (True) or a definition (False).

Instructions

Every *instruction* is also a value: it has a name (the recipient's name), a well-defined type, and can be used as an operand in other instructions or in literals.

Instruction types should usually not be instantiated directly. Instead, use the helper methods on the *IRBuilder* class.

class llvmlite.ir.Instruction

The base class for all instructions. Instructions have the following method and attributes:

set_metadata (*name*, *node*)

Add an instruction-specific metadata named *name* pointing to the given metadata *node* (a *MDValue*).

function

The function this instruction is part of.

module

The module this instruction's function is defined in.

class llvmlite.ir.PredictableInstr

The class of instructions for which we can specify the probabilities of different outcomes (e.g. a switch or a conditional branch). Predictable instructions have the following method:

set_weights (*weights*)

Set the weights of the instruction's possible outcomes. *weights* is a sequence of positive integers, each corresponding to a different outcome and specifying its relative probability compared to other outcomes (the greater, the likelier).

class llvmlite.ir.SwitchInstr

A switch instruction. Switch instructions have the following method:

add_case (*val*, *block*)

Add a case to the switch instruction. *val* should be a *Constant* or a Python value compatible with the switch instruction's operand type. *block* is a *Block* to jump to if, and only if, *val* and the switch operand compare equal.

class `llvmlite.ir.IndirectBranch`

An indirect branch instruction. Indirect branch instructions have the following method:

add_destination (*value*, *block*)

Add an outgoing edge. The indirect branch instruction must refer to every basic block it can transfer control to.

class `llvmlite.ir.PhiInstr`

A phi instruction. Phi instructions have the following method:

add_incoming (*value*, *block*)

Add an incoming edge. Whenever transfer is controlled from *block* (a *Block*), the phi instruction takes the given *value*.

class `llvmlite.ir.LandingPad`

A landing pad. Landing pads have the following method:

add_clause (*value*, *block*)

Add a catch or filter clause. Create catch clauses using *CatchClause*, and filter clauses using *FilterClause*.

Landing pad clauses

Landing pads have the following classes associated with them:

class `llvmlite.ir.CatchClause` (*value*)

A catch clause. Instructs the personality function to compare the in-flight exception typeinfo with *value*, which should have type *IntType(8).as_pointer().as_pointer()*.

class `llvmlite.ir.FilterClause` (*value*)

A filter clause. Instructs the personality function to check inclusion of the the in-flight exception typeinfo in *value*, which should have type *ArrayType(IntType(8).as_pointer().as_pointer(), ...)*.

Modules

A module is a compilation unit. It defines a set of related functions, global variables and metadata. In the IR layer, a module is represented by the *Module* class.

class `llvmlite.ir.Module` (*name*='')

Create a module. The optional *name*, a Python string, can be specified for informational purposes.

Modules have the following methods and attributes:

add_debug_info (*kind*, *operands*, *is_distinct*=False)

Add debug information metadata to the module with the given *operands* (a mapping of string keys to values) or return a previous equivalent metadata. *kind* is the name of the debug information kind (e.g. 'DICompileUnit').

A *DIValue* instance is returned, it can then be associated to e.g. an instruction.

Example:

```
di_file = module.add_debug_info("DIFile", {
    "filename": "factorial.py",
    "directory": "bar",
})
di_compile_unit = module.add_debug_info("DICompileUnit", {
    "language": ir.DIToken("DW_LANG_Python"),
    "file": di_file,
```

```
"producer": "llvmlite x.y",
"runtimeVersion": 2,
"isOptimized": False,
}, is_distinct=True)
```

add_global (*globalvalue*)

Add the given *globalvalue* (a *GlobalValue*) to this module. It should have a unique name in the whole module.

add_metadata (*operands*)

Add an unnamed *metadata* node to the module with the given *operands* (a list of metadata-compatible values). If another metadata node with equal operands already exists in the module, it is reused instead. A *MDValue* is returned.

add_named_metadata (*name*, *element=None*)

Return the metadata node with the given *name*. If it doesn't already exist, the named metadata node is created first. If *element* is not None, it can be a metadata value or a sequence of values to append to the metadata node's elements. A *NamedMetaDatum* is returned.

Example:

```
module.add_named_metadata("llvm.ident", ["llvmlite/1.0"])
```

get_global (*name*)

Get the *global value* (a *GlobalValue*) with the given *name*. *KeyError* is raised if it doesn't exist.

get_named_metadata (*name*)

Return the metadata node with the given *name*. *KeyError* is raised if it doesn't exist.

get_unique_name (*name*)

Return a unique name across the whole module. *name* is the desired name, but a variation can be returned if it is already in use.

data_layout

A string representing the data layout in LLVM format.

functions

The list of functions (as *Function* instances) declared or defined in the module.

global_values

An iterable of global values in this module.

triple

A string representing the target architecture in LLVM "triple" form.

IR builders

Contents

- *Instantiation*
- *Properties*
- *Utilities*
- *Positioning*
- *Flow control helpers*
- *Instruction building*
 - *Arithmetic*
 - *Conversions*
 - *Comparisons*
 - *Conditional move*
 - *Phi*
 - *Aggregate operations*
 - *Memory*
 - *Function call*
 - *Branches*
 - *Exception handling*
 - *Inline assembler*
 - *Miscellaneous*

IRBuilder is the workhorse of LLVM *IR* generation. It allows you to fill the *basic blocks* of your functions with LLVM instructions.

A *IRBuilder* internally maintains a current basic block, and a pointer inside the block's list of instructions. When adding a new instruction, it is inserted at that point and the pointer is then advanced after the new instruction.

A *IRBuilder* also maintains a reference to metadata describing the current source location, which will be attached to all inserted instructions.

Instantiation

`class llvmlite.ir.IRBuilder (block=None)`

Create a new IR builder. If *block* (a *Block*) is given, the builder starts right at the end of this basic block.

Properties

IRBuilder has the following attributes:

`IRBuilder.block`

The basic block the builder is operating on.

`IRBuilder.function`

The function the builder is operating on.

`IRBuilder.module`

The module the builder's function is defined in.

`IRBuilder.debug_metadata`

If not *None*, the metadata that will be attached to any inserted instructions as *!dbg*, unless the instruction already has *!dbg* set.

Utilities

`IRBuilder.append_basic_block(name='')`

Append a basic block, with the given optional *name*, to the current function. The current block is not changed. A *Block* is returned.

Positioning

The following *IRBuilder* methods help you move the current instruction pointer around:

`IRBuilder.position_before(instruction)`

Position immediatly before the given *instruction*. The current block is also changed to the instruction's basic block.

`IRBuilder.position_after(instruction)`

Position immediatly after the given *instruction*. The current block is also changed to the instruction's basic block.

`IRBuilder.position_at_start(block)`

Position at the start of the basic *block*.

`IRBuilder.position_at_end(block)`

Position at the end of the basic *block*.

The following context managers allow you to temporarily switch to another basic block, then go back where you were:

`IRBuilder.goto_block(block)`

A context manager which positions the builder either at the end of the basic *block*, if it is not terminated, or just before the *block*'s terminator:

```
new_block = builder.append_basic_block('foo')
with builder.goto_block(new_block):
    # Now the builder is at the end of *new_block*
    # ... add instructions

# Now the builder has returned to its previous position
```

`IRBuilder.goto_entry_block()`

Just like *goto_block()*, but with the current function's entry block.

Flow control helpers

The following context managers make it easier to create conditional code.

`IRBuilder.if_then(pred, likely=None)`

A context manager which creates a basic block whose execution is conditioned on predicate *pred* (a value of type `IntType(1)`). Another basic block is created for instructions after the conditional block. The current basic block is terminated with a conditional branch based on *pred*.

When the context manager is entered, the builder positions at the end of the conditional block. When the context manager is exited, the builder positions at the start of the continuation block.

If *likely* is not `None`, it indicates whether *pred* is likely to be true, and metadata is emitted to specify branch weights in accordance.

`IRBuilder.if_else(pred, likely=None)`

A context manager which sets up two basic blocks whose execution is condition on predicate *pred* (a value of type `IntType(1)`). *likely* has the same meaning as in `if_then()`.

A pair of context managers is yield'ed. Each of them acts as a `if_then()` context manager: the first one for the block to be executed if *pred* is true, the second one for the block to be executed if *pred* is false.

When the context manager is exited, the builder is positioned on a new continuation block which both conditional blocks jump into.

Typical use:

```
with builder.if_else(pred) as (then, otherwise):
    with then:
        # emit instructions for when the predicate is true
    with otherwise:
        # emit instructions for when the predicate is false
# emit instructions following the if-else block
```

Instruction building

The following methods all insert a new instruction (a `Instruction` instance) at the current index in the current block. The new instruction is returned.

An instruction's operands are most always *values*.

Many of these methods also take an optional *name* argument, specifying the local *name* of the result value. If not given, a unique name is automatically generated.

Arithmetic

The *flags* argument in the methods below is an optional sequence of strings serving as modifiers of the instruction's semantics. Examples include the fast-math flags for floating-point operations, or whether wraparound on overflow can be ignored on integer operations.

Integer

`IRBuilder.shl(lhs, rhs, name='', flags=())`

Left-shift *lhs* by *rhs* bits.

`IRBuilder.lshr(lhs, rhs, name='', flags=())`

Logical right-shift *lhs* by *rhs* bits.

`IRBuilder.ashr(lhs, rhs, name='', flags=())`

Arithmetic (signed) right-shift *lhs* by *rhs* bits.

`IRBuilder.add(lhs, rhs, name='', flags=())`

Integer add *lhs* and *rhs*.

`IRBuilder.sadd_with_overflow(lhs, rhs, name='', flags=())`

Integer add *lhs* and *rhs*. A { result, overflow bit } structure is returned.

`IRBuilder.sub(lhs, rhs, name='', flags=())`

Integer subtract *rhs* from *lhs*.

`IRBuilder.sadd_with_overflow(lhs, rhs, name='', flags=())`

Integer subtract *rhs* from *lhs*. A { result, overflow bit } structure is returned.

`IRBuilder.mul (lhs, rhs, name='', flags=())`
Integer multiply *lhs* with *rhs*.

`IRBuilder.smul_with_overflow (lhs, rhs, name='', flags=())`
Integer multiply *lhs* with *rhs*. A { `result`, `overflow bit` } structure is returned.

`IRBuilder.sdiv (lhs, rhs, name='', flags=())`
Signed integer divide *lhs* by *rhs*.

`IRBuilder.udiv (lhs, rhs, name='', flags=())`
Unsigned integer divide *lhs* by *rhs*.

`IRBuilder.srem (lhs, rhs, name='', flags=())`
Signed integer remainder of *lhs* divided by *rhs*.

`IRBuilder.urem (lhs, rhs, name='', flags=())`
Unsigned integer remainder of *lhs* divided by *rhs*.

`IRBuilder.and_ (lhs, rhs, name='', flags=())`
Bitwise AND *lhs* with *rhs*.

`IRBuilder.or_ (lhs, rhs, name='', flags=())`
Bitwise OR *lhs* with *rhs*.

`IRBuilder.xor (lhs, rhs, name='', flags=())`
Bitwise XOR *lhs* with *rhs*.

`IRBuilder.not_ (value, name='')`
Bitwise complement *value*.

`IRBuilder.neg (value, name='')`
Negate *value*.

Floating-point

`IRBuilder.fadd (lhs, rhs, name='', flags=())`
Floating-point add *lhs* and *rhs*.

`IRBuilder.fsub (lhs, rhs, name='', flags=())`
Floating-point subtract *rhs* from *lhs*.

`IRBuilder.fmul (lhs, rhs, name='', flags=())`
Floating-point multiply *lhs* by *rhs*.

`IRBuilder.fdiv (lhs, rhs, name='', flags=())`
Floating-point divide *lhs* by *rhs*.

`IRBuilder.frem (lhs, rhs, name='', flags=())`
Floating-point remainder of *lhs* divided by *rhs*.

Conversions

`IRBuilder.trunc (value, typ, name='')`
Truncate integer *value* to integer type *typ*.

`IRBuilder.zext (value, typ, name='')`
Zero-extend integer *value* to integer type *typ*.

`IRBuilder.sext (value, typ, name='')`
Sign-extend integer *value* to integer type *typ*.

`IRBuilder.fptrunc` (*value*, *typ*, *name*='')
 Truncate (approximate) floating-point *value* to floating-point type *typ*.

`IRBuilder.fpext` (*value*, *typ*, *name*='')
 Extend floating-point *value* to floating-point type *typ*.

`IRBuilder.fptosi` (*value*, *typ*, *name*='')
 Convert floating-point *value* to signed integer type *typ*.

`IRBuilder.fptoui` (*value*, *typ*, *name*='')
 Convert floating-point *value* to unsigned integer type *typ*.

`IRBuilder.sitofp` (*value*, *typ*, *name*='')
 Convert signed integer *value* to floating-point type *typ*.

`IRBuilder.uitofp` (*value*, *typ*, *name*='')
 Convert unsigned integer *value* to floating-point type *typ*.

`IRBuilder.ptrtoint` (*value*, *typ*, *name*='')
 Convert pointer *value* to integer type *typ*.

`IRBuilder.inttoptr` (*value*, *typ*, *name*='')
 Convert integer *value* to pointer type *typ*.

`IRBuilder.bitcast` (*value*, *typ*, *name*='')
 Convert pointer *value* to pointer type *typ*.

`IRBuilder.addrspacecast` (*value*, *typ*, *name*='')
 Convert pointer *value* to pointer type *typ* of different address space.

Comparisons

`IRBuilder.icmp_signed` (*cmpop*, *lhs*, *rhs*, *name*='')
 Signed integer compare *lhs* with *rhs*. *cmpop*, a string, can be one of <, <=, ==, !=, >=, >.

`IRBuilder.icmp_unsigned` (*cmpop*, *lhs*, *rhs*, *name*='')
 Unsigned integer compare *lhs* with *rhs*. *cmpop*, a string, can be one of <, <=, ==, !=, >=, >.

`IRBuilder.fcmp_ordered` (*cmpop*, *lhs*, *rhs*, *name*='', *flags*=[])
 Floating-point ordered compare *lhs* with *rhs*. *cmpop*, a string, can be one of <, <=, ==, !=, >=, >, ord, uno. *flags*, a list, can include any of nnan, ninf, nsz, arcp, and fast (which implies all previous flags).

`IRBuilder.fcmp_unordered` (*cmpop*, *lhs*, *rhs*, *name*='', *flags*=[])
 Floating-point unordered compare *lhs* with *rhs*. *cmpop*, a string, can be one of <, <=, ==, !=, >=, >, ord, uno. *flags*, a list, can include any of nnan, ninf, nsz, arcp, and fast (which implies all previous flags).

Conditional move

`IRBuilder.select` (*cond*, *lhs*, *rhs*, *name*='')
 Two-way select: *lhs* if *cond* else *rhs*.

Phi

`IRBuilder.phi` (*typ*, *name*='')
 Create a phi node. Add incoming edges and their values using the `add_incoming()` method on the return value.

Aggregate operations

`IRBuilder.extract_value (agg, index, name='')`

Extract the element at *index* of the *aggregate value* *agg*. *index* may be an integer or a sequence of integers. If *agg* is of an array type, indices can be arbitrary values; if *agg* is of a struct type, indices have to be constant.

`IRBuilder.insert_value (agg, value, index, name='')`

Build a copy of *aggregate value* *agg* by setting the new *value* at *index*. *index* can be of the same types as in `extract_value()`.

Memory

`IRBuilder.alloca (typ, size=None, name='')`

Statically allocate a stack slot for *size* values of type *typ*. If *size* is not given, a stack slot for one value is allocated.

`IRBuilder.load (ptr, name='', align=None)`

Load value from pointer *ptr*. If *align* is passed, it should be a Python integer specifying the guaranteed pointer alignment.

`IRBuilder.store (value, ptr, align=None)`

Store *value* to pointer *ptr*. If *align* is passed, it should be a Python integer specifying the guaranteed pointer alignment.

`IRBuilder.gep (ptr, indices, inbounds=False, name='')`

The *getelementptr* instruction. Given a pointer *ptr* to an aggregate value, compute the address of the inner element given by the sequence of *indices*.

`llvmlite.ir.cmpxchg (ptr, cmp, val, ordering, failordering=None, name='')`

Atomic compare-and-swap at address *ptr*. *cmp* is the value to compare the contents with, *val* the new value to be swapped into. Optional *ordering* and *failordering* specify the memory model for this instruction.

`llvmlite.ir.atomic_rmw (op, ptr, val, ordering, name='')`

Atomic in-memory operation *op* at address *ptr*, with operand *val*. *op* is a string specifying the operation (e.g. add or sub). The optional *ordering* specifies the memory model for this instruction.

Function call

`IRBuilder.call (fn, args, name='', cconv=None, tail=False, fastmath=())`

Call function *fn* with arguments *args* (a sequence of values). *cconv* is the optional calling convention. *tail*, if true, is a hint for the optimizer to perform tail-call optimization. *fastmath* is a string or a sequence of strings of names for *fast-math* flags.

Branches

These instructions are all *terminators*.

`IRBuilder.branch (target)`

Unconditional jump to the *target* (a *Block*).

`IRBuilder.cbranch (cond, truebr, falsebr)`

Conditional jump to either *truebr* or *falsebr* (both *Block* instances), depending on *cond* (a value of type `IntType(1)`). This instruction is a *PredictableInstr*.

`IRBuilder.ret(value)`

Return the *value* from the current function.

`IRBuilder.ret_void()`

Return from the current function without a value.

`IRBuilder.switch(value, default)`

Switch to different blocks based on the *value*. *default* is the block to switch to if no other block is matched.

Add non-default targets using the `add_case()` method on the return value.

`IRBuilder.indirectbr(address)`

Jump to basic block with address *address* (a value of type `IntType(8).as_pointer()`). A block address can be obtained using the `BlockAddress` constant.

Add all possible jump destinations using the `add_destination()` method on the return value.

Exception handling

`IRBuilder.invoke(self, fn, args, normal_to, unwind_to, name='', cconv=None, tail=False)`

Call function *fn* with arguments *args* (a sequence of values). *cconv* is the optional calling convention. *tail*, if true, is a hint for the optimizer to perform tail-call optimization.

If the function *fn* returns normally, control is transferred to *normal_to*. Otherwise, it is transferred to *unwind_to*, the first non-phi instruction of which must be `LandingPad`.

`IRBuilder.landingpad(typ, personality, name='', cleanup=False)`

Describe which exceptions this basic block can handle.

typ specifies the return type of the landing pad. It is a structure with two pointer-sized fields. *personality* specifies an exception personality function. *cleanup* specifies whether control should be always transferred to this landing pad, even when no matching exception is caught.

Add landing pad clauses using the `add_clause()` method on the return value.

There are two kinds of landing pad clauses:

- A `CatchClause`, which specifies a typeinfo for a single exception to be caught. The typeinfo is a value of type `IntType(8).as_pointer().as_pointer()`;
- A `FilterClause`, which specifies an array of typeinfos.

Every landing pad must either contain at least one clause, or be marked for cleanup.

The semantics of a landing pad are entirely determined by the personality function. See [Exception handling in LLVM](#) for details on the way LLVM handles landing pads in the optimizer, and [Itanium exception handling ABI](#) for details on the implementation of personality functions.

`IRBuilder.resume(landingpad)`

Resume an exception caught by landing pad *landingpad*. Used to indicate that the landing pad did not catch the exception after all (perhaps because it only performed cleanup).

Inline assembler

`IRBuilder.asm(ftype, asm, constraint, args, side_effect, name='')`

Add an inline assembler call instruction. This is used for instance in `load_reg()` and `store_reg()`.

Arguments:

- *fty* is a function type specifying the inputs and output of the inline assembler call;
- *asm* is the inline assembler snippet, e.g. "mov \$2, \$0\nadd \$1, \$0" (x86 inline ASM uses the AT&T syntax);
- *constraint* defines the input/output constraints, e.g. =r, r, r;
- *args* is the list of inputs, as IR values;
- *side_effect* (boolean), whether this instruction has side effects not visible in the constraint list;
- *name* optional name of the returned LLVM value.

For more information about these parameters, see the official LLVM documentation here: <http://llvm.org/docs/LangRef.html#inline-asm-constraint-string>.

Example that adds two 64-bit values on x86:

```
fty = FunctionType(IntType(64), [IntType(64), IntType(64)])
add = builder.asm(fty, "mov $2, $0\nadd $1, $0", "=r,r,r",
                  (arg_0, arg_1), name="asm_add")
```

`IRBuilder.load_reg(reg_type, reg_name, name='')`
Load a register value into an LLVM value.

Example to get the value of the rax register:

```
builder.load_reg(IntType(64), "rax")
```

`IRBuilder.store_reg(value, reg_type, reg_name, name='')`
Store an LLVM value inside a register

Example to store 0xAAAAAAAAAAAAAAAA into the rax register:

```
builder.store_reg(Constant(IntType(64), 0xAAAAAAAAAAAAAAAA), IntType(64), "rax")
```

Miscellaneous

`IRBuilder.assume(cond)`
Let the LLVM optimizer assume that *cond* (a value of type `IntType(1)`) is true.

`IRBuilder.unreachable()`
Mark an unreachable point in the code.

Example

A trivial function

Define a function adding two double-precision floating-point numbers.

```
"""
This file demonstrates a trivial function "fpadd" returning the sum of
two floating-point numbers.
"""

from llvmlite import ir

# Create some useful types
```

```
double = ir.DoubleType()
fnty = ir.FunctionType(double, (double, double))

# Create an empty module...
module = ir.Module(name=__file__)
# and declare a function named "fpadd" inside it
func = ir.Function(module, fnty, name="fpadd")

# Now implement the function
block = func.append_basic_block(name="entry")
builder = ir.IRBuilder(block)
a, b = func.args
result = builder.fadd(a, b, name="res")
builder.ret(result)

# Print the module IR
print(module)
```

The generated LLVM *IR* is printed out at the end, and should look like this:

```
; ModuleID = "examples/ir_fpadd.py"
target triple = "unknown-unknown-unknown"
target datalayout = ""

define double @"fpadd"(double %.1", double %.2")
{
entry:
    %.res" = fadd double %.1", %.2"
    ret double %.res"
}
```

To learn how to compile and execute this function, refer to the [binding layer](#) documentation.

llvmlite.binding – The LLVM binding layer

The `llvmlite.binding` module provides classes to interact with functionalities of the LLVM library. They generally mirror concepts of the C++ API closely. A small subset of the LLVM API is mirrored, though: only those parts that have proven useful to implement Numba's JIT compiler.

Initialization and finalization

These functions need only be called once per process invocation.

```
llvmlite.binding.initialize()
    Initialize the LLVM core.
```

```
llvmlite.binding.initialize_all_targets()
    Initialize all targets. Necessary before targets can be looked up via the Target class.
```

```
llvmlite.binding.initialize_all_asmprinters()
    Initialize all code generators. Necessary before generating any assembly or machine code via the
    TargetMachine.emit_object() and TargetMachine.emit_assembly() methods.
```

```
llvmlite.binding.initialize_native_target()
    Initialize the native (host) target. Calling this function once is necessary before doing any code generation.
```

```
llvmlite.binding.initialize_native_asmprinter()
```

Initialize the native assembly printer.

```
llvmlite.binding.initialize_native_asmparser()
```

Initialize the native assembly parser. This is necessary for inline assembly to work.

```
llvmlite.binding.shutdown()
```

Shutdown the LLVM core.

```
llvmlite.binding.llvm_version_info
```

A three-integer tuple representing the LLVM version number, for example (3, 7, 1). Since LLVM is statically linked into the llvmlite DLL, this is guaranteed to represent the true LLVM version in use.

Dynamic libraries and symbols

These functions tell LLVM how to resolve external symbols referred from compiled LLVM code.

```
llvmlite.binding.add_symbol(name, address)
```

Register the *address* of global symbol *name*, for use from LLVM-compiled functions.

```
llvmlite.binding.address_of_symbol(name)
```

Get the in-process address of symbol named *name*. An integer is returned, or None if the symbol isn't found.

```
llvmlite.binding.load_library_permanently(filename)
```

Load an external shared library. *filename* should be the path to the shared library file.

Target information

Target information allows you to inspect and modify aspects of the code generation, such as which CPU is targetted or what optimization level is desired.

Minimal use of this module would be to create a *TargetMachine* for later use in code generation:

```
from llvmlite import binding
target = binding.Target.from_default_triple()
target_machine = target.create_target_machine()
```

Functions

```
llvmlite.binding.get_default_triple()
```

Return the default target triple LLVM is configured to produce code for, as a string. This represents the host's architecture and platform.

```
llvmlite.binding.get_process_triple()
```

Return a target triple suitable for generating code for the current process. An example when the default triple from `get_default_triple()` is not be suitable is when LLVM is compiled for 32-bit but the process is executing in 64-bit mode.

```
llvmlite.binding.get_object_format(triple=None)
```

Get the object format for the given *triple* string (or the default triple if None). A string is returned such as "ELF", "COFF" or "MachO".

```
llvmlite.binding.get_host_cpu_name()
```

Get the name of the host's CPU as a string. You can use the return value with *Target.create_target_machine()*.

```
llvmlite.binding.get_host_cpu_features()
```

Returns a dictionary-like object indicating the CPU features for current architecture and whether they are enabled for this CPU. The key-value pairs are the feature name as string and a boolean indicating whether the feature is available. The returned value is an instance of `FeatureMap` class, which adds a new method `.flatten()` for returning a string suitable for use as the “features” argument to `Target.create_target_machine()`.

If LLVM has not implemented this feature or it fails to get the information, this function will raise a `RuntimeError` exception.

```
llvmlite.binding.create_target_data(data_layout)
```

Create a `TargetData` representing the given `data_layout` (a string).

Classes

```
class llvmlite.binding.TargetData
```

A class providing functionality around a given data layout. It specifies how the different types are to be represented in memory. Use `create_target_data()` to instantiate.

```
get_abi_size(type)
```

Get the ABI-mandated size of the LLVM `type` (as returned by `ValueRef.type`). An integer is returned.

```
get_pointee_abi_size(type)
```

Similar to `get_abi_size()`, but assumes `type` is a LLVM pointer type, and returns the ABI-mandated size of the type pointed to by. This is useful for global variables (whose type is really a pointer to the declared type).

```
get_pointee_abi_alignment(type)
```

Similar to `get_pointee_abi_size()`, but return the ABI-mandated alignment rather than the ABI size.

```
class llvmlite.binding.Target
```

A class representing a compilation target. The following factories are provided:

```
classmethod from_triple(triple)
```

Create a new `Target` instance for the given `triple` string denoting the target platform.

```
classmethod from_default_triple()
```

Create a new `Target` instance for the default platform LLVM is configured to produce code for. This is equivalent to calling `Target.from_triple(get_default_triple())`

The following methods and attributes are available:

```
description
```

A description of the target.

```
name
```

The name of the target.

```
triple
```

The triple (a string) uniquely identifying the target, for example “x86_64-pc-linux-gnu”.

```
create_target_machine(cpu='', features='', opt=2, reloc='default', codemodel='jitdefault')
```

Create a new `TargetMachine` instance for this target and with the given options. `cpu` is an optional CPU name to specialize for. `features` is a comma-separated list of target-specific features to enable or disable. `opt` is the optimization level, from 0 to 3. `reloc` is the relocation model. `codemodel` is the code model. The defaults for `reloc` and `codemodel` are appropriate for JIT compilation.

Tip: To list the available CPUs and features for a target, execute `llc -mcpu=help` on the command line.

`class llvmlite.binding.TargetMachine`

A class holding all the settings necessary for proper code generation (including target information and compiler options). Instantiate using `Target.create_target_machine()`.

`add_analysis_passes(pm)`

Register analysis passes for this target machine with the `PassManager` instance `pm`.

`emit_object(module)`

Represent the compiled `module` (a `ModuleRef` instance) as a code object, suitable for use with the platform's linker. A bytestring is returned.

`set_asm_verbosity(is_verbose)`

Set whether this target machine will emit assembly with human-readable comments describing control flow, debug information, and so on.

`emit_assembly(module)`

Return the compiled `module`'s native assembler, as a string.

`initialize_native_asmprinter()` must have been called first.

`target_data`

The `TargetData` associated with this target machine.

`class llvmlite.binding.FeatureMap`

For storing processor feature information in a dictionary-like object. This class extends `dict` and only adds the `.flatten()` method.

`flatten(sort=True)`

Returns a string representation of the stored information that is suitable for use in the “features” argument of `Target.create_target_machine()`. If `sort` keyword argument is `True` (the default), the features are sorted by name to give a stable ordering between python session.

Modules

While they conceptually represent the same thing, modules in the *IR layer* and modules in the *binding layer* don't have the same roles and don't expose the same API.

While modules in the IR layer allow to build and group functions together, modules in the binding layer give access to compilation, linking and execution of code. To distinguish between the two, the module class in the binding layer is called `ModuleRef` as opposed to `llvmlite.ir.Module`.

To go from one realm to the other, you must use the `parse_assembly()` function.

Factory functions

A module can be created from the following factory functions:

`llvmlite.binding.parse_assembly(llvmir)`

Parse the given `llvmir`, a string containing some LLVM IR code. If parsing is successful, a new `ModuleRef` instance is returned.

`llvmir` can be obtained, for example, by calling `str()` on a `llvmlite.ir.Module` object.

```
llvmlite.binding.parse_bitcode(bitcode)
```

Parse the given *bitcode*, a bytestring containing the LLVM bitcode of a module. If parsing is successful, a new *ModuleRef* instance is returned.

bitcode can be obtained, for example, by calling *ModuleRef.as_bitcode()*.

The ModuleRef class

```
class llvmlite.binding.ModuleRef
```

A wrapper around a LLVM module object. The following methods and properties are available:

```
as_bitcode()
```

Return the bitcode of this module as a bytes object.

```
get_function(name)
```

Get the function with the given *name* in this module. If found, a *ValueRef* is returned. Otherwise, *NameError* is raised.

```
get_global_variable(name)
```

Get the global variable with the given *name* in this module. If found, a *ValueRef* is returned. Otherwise, *NameError* is raised.

```
link_in(other, preserve=False)
```

Link the *other* module into this module, resolving references wherever possible. If *preserve* is true, the other module is first copied in order to preserve its contents; if *preserve* is false, the other module is not usable after this call anymore.

```
verify()
```

Verify the module's correctness. *RuntimeError* is raised on error.

```
data_layout
```

The data layout string for this module. This attribute is settable.

```
functions
```

An iterator over the functions defined in this module. Each function is a *ValueRef* instance.

```
global_variables
```

An iterator over the global variables defined in this module. Each global variable is a *ValueRef* instance.

```
name
```

The module's identifier, as a string. This attribute is settable.

```
triple
```

The platform "triple" string for this module. This attribute is settable.

Value references

A value reference is a wrapper around a LLVM value for you to inspect. You can't create one yourself; instead, you'll get them from methods of the *ModuleRef* class.

Enumerations

```
class llvmlite.binding.Linkage
```

The different linkage types allowed for global values. The following values are provided:

```
external
```

```
available_externally
```

`linkonce_any`
`linkonce_odr`
`linkonce_odr_autohide`
`weak_any`
`weak_odr`
`Appending`
`internal`
`private`
`dllimport`
`dllexport`
`external_weak`
`ghost`
`common`
`linker_private`
`linker_private_weak`

class `llvmlite.binding.Visibility`

The different visibility styles allowed for global values. The following values are provided:

`default`
`hidden`
`protected`

class `llvmlite.binding.StorageClass`

The different storage classes allowed for global values. The following values are provided:

`default`
`dllimport`
`dllexport`

The ValueRef class

class `llvmlite.binding.ValueRef`

A wrapper around a LLVM value. The following properties are available:

is_declaration
True if the global value is a mere declaration, False if it is defined in the given module.

linkage
The linkage type (a *Linkage* instance) for this value. This attribute is settable.

module
The module (a *ModuleRef* instance) this value is defined in.

name
This value's name, as a string. This attribute is settable.

type

This value's LLVM type. An opaque object is returned. It can be used with e.g. `TargetData.get_abi_size()`.

storage_class

The storage class (a `StorageClass` instance) for this value. This attribute is settable.

visibility

The visibility style (a `Visibility` instance) for this value. This attribute is settable.

Execution engine

The execution engine is where actual code generation and execution happens. The currently supported LLVM version (LLVM 3.8) exposes one execution engine, named MCJIT.

Functions

`llvmlite.binding.create_mcjit_compiler(module, target_machine)`

Create a MCJIT-powered engine from the given *module* and *target_machine*. A `ExecutionEngine` instance is returned. The *module* need not contain any code.

`llvmlite.binding.check_jit_execution()`

Ensure the system allows creation of executable memory ranges for JIT-compiled code. If some security mechanism (such as SELinux) prevents it, an exception is raised. Otherwise the function returns silently.

Calling this function early can help diagnose system configuration issues, instead of letting JIT-compiled functions crash mysteriously.

The ExecutionEngine class

class `llvmlite.binding.ExecutionEngine`

A wrapper around a LLVM execution engine. The following methods and properties are available:

add_module (*module*)

Add the *module* (a `ModuleRef` instance) for code generation. When this method is called, ownership of the module is transferred to the execution engine.

finalize_object ()

Make sure all modules owned by the execution engine are fully processed and “usable” for execution.

get_pointer_to_function (*gv*)

Warning: This function is deprecated. User should use `ExecutionEngine.get_function_address()` and `ExecutionEngine.get_global_value_address()` instead

Return the address of the function value *gv*, as an integer. The value should have been looked up on one of the modules owned by the execution engine.

Note: This method may implicitly generate code for the object being looked up.

Note: This function is formerly an alias to `get_pointer_to_global()`, which is now removed because it returns an invalid address in MCJIT when given a non-function global value.

get_function_address (*name*)

Return the address of the function named *name* as an integer.

get_global_value_address (*name*)

Return the address of the global value named *name* as an integer.

remove_module (*module*)

Remove the *module* (a *ModuleRef* instance) from the modules owned by the execution engine. This allows releasing the resources owned by the module without destroying the execution engine.

set_object_cache (*notify_func=None*, *getbuffer_func=None*)

Set the object cache callbacks for this engine.

notify_func, if given, is called whenever the engine has finished compiling a module. It is passed two arguments (*module*, *buffer*). The first argument *module* is a *ModuleRef* instance. The second argument *buffer* is a bytes object of the code generated for the module. The return value is ignored.

getbuffer_func, if given, is called before the engine starts compiling a module. It is passed one argument, *module*, a *ModuleRef* instance of the module being compiled. The function can return `None`, in which case the module will be compiled normally. Or it can return a bytes object of native code for the module, which will bypass compilation entirely.

target_data

The *TargetData* used by the execution engine.

Optimization passes

LLVM gives you the possibility to fine-tune optimization passes. llvmlite exposes several of these parameters. Optimization passes are managed by a pass manager; there are two kinds thereof: *FunctionPassManager*, for optimizations which work on single functions, and *ModulePassManager*, for optimizations which work on whole modules.

To instantiate any of those pass managers, you first have to create and configure a *PassManagerBuilder*.

class llvmlite.binding.PassManagerBuilder

Create a new pass manager builder. This object centralizes optimization settings. The following method is available:

populate (*pm*)

Populate the pass manager *pm* with the optimization passes configured in this pass manager builder.

The following writable properties are also available:

disable_unroll_loops

If true, disable loop unrolling.

inlining_threshold

The integer threshold for inlining a function into another. The higher, the more likely inlining a function is. This attribute is write-only.

loop_vectorize

If true, allow vectorizing loops.

opt_level

The general optimization level as an integer between 0 and 3.

size_level

Whether and how much to optimize for size. An integer between 0 and 2.

slp_vectorize

If true, enable the “SLP vectorizer”, which uses a different algorithm from the loop vectorizer. Both may be enabled at the same time.

class llvmlite.binding.PassManager

The base class for pass managers. Use individual `add_*` methods or `PassManagerBuilder.populate()` to add optimization passes.

add_constant_merge_pass()

See [constmerge pass documentation](#).

add_dead_arg_elimination_pass()

See [deadargelim pass documentation](#).

add_function_attrs_pass()

See [functionattrs pass documentation](#).

add_function_inlining_pass(self)

See [inline pass documentation](#).

add_global_dce_pass()

See [globaldce pass documentation](#).

add_global_optimizer_pass()

See [globalopt pass documentation](#).

add_ipsccp_pass()

See [ipsccp pass documentation](#).

add_dead_code_elimination_pass()

See [dce pass documentation](#).

add_cfg_simplification_pass()

See [simplifycfg pass documentation](#).

add_gvn_pass()

See [gvn pass documentation](#).

add_instruction_combining_pass()

See [instcombine pass documentation](#).

add_licm_pass()

See [licm pass documentation](#).

add_sccp_pass()

See [sccp pass documentation](#).

add_sroa_pass()

See [scalarrepl pass documentation](#).

Note that while the link above describes the transformation performed by the pass added by this function, it corresponds to the `opt -sroa` command-line option and not `opt -scalarrepl`.

add_type_based_alias_analysis_pass()

add_basic_alias_analysis_pass()

See [basicaa pass documentation](#).

class llvmlite.binding.ModulePassManager

Create a new pass manager to run optimization passes on a module.

The following method is available:

run (*module*)

Run optimization passes on the *module* (a *ModuleRef* instance). True is returned if the optimizations made any modification to the module, False instead.

class llvmlite.binding.**FunctionPassManager** (*module*)

Create a new pass manager to run optimization passes on a function of the given *module* (an *ModuleRef* instance).

The following methods are available:

finalize ()

Run all the finalizers of the optimization passes.

initialize ()

Run all the initializers of the optimization passes.

run (*function*)

Run optimization passes on the *function* (a *ValueRef* instance). True is returned if the optimizations made any modification to the module, False instead.

Analysis Utilities

llvmlite.binding.**get_function_cfg** (*func*, *show_inst=True*)

Return a string of the control-flow graph of the function in DOT format. If the input *func* is not a materialized function, the module containing the function is parsed to create an actual LLVM module. The *show_inst* flag controls whether the instructions of each block are printed.

llvmlite.binding.**view_dot_graph** (*graph*, *filename=None*, *view=False*)

View the given DOT source. If *view* is True, the image is rendered and viewed by the default application in the system. The file path of the output is returned. If *view* is False, a *graphviz.Source* object is returned. If *view* is False and the environment is in a IPython session, an IPython image object is returned and can be displayed inline in the notebook.

Note: This function requires the *graphviz* package.

Examples

Compiling a trivial function

Compile and execute the function defined in *A trivial function*. The function is compiled with no specific optimizations.

```
from __future__ import print_function

from ctypes import CFUNCTYPE, c_double

import llvmlite.binding as llvm

# All these initializations are required for code generation!
llvm.initialize()
llvm.initialize_native_target()
llvm.initialize_native_asmprinter() # yes, even this one
```

```

llvm_ir = """
; ModuleID = "examples/ir_fpadd.py"
target triple = "unknown-unknown-unknown"
target datalayout = ""

define double @"fpadd"(double %.1", double %.2")
{
entry:
    %"res" = fadd double %.1", %.2"
    ret double %"res"
}
"""

def create_execution_engine():
    """
    Create an ExecutionEngine suitable for JIT code generation on
    the host CPU. The engine is reusable for an arbitrary number of
    modules.
    """
    # Create a target machine representing the host
    target = llvm.Target.from_default_triple()
    target_machine = target.create_target_machine()
    # And an execution engine with an empty backing module
    backing_mod = llvm.parse_assembly("")
    engine = llvm.create_mcjit_compiler(backing_mod, target_machine)
    return engine

def compile_ir(engine, llvm_ir):
    """
    Compile the LLVM IR string with the given engine.
    The compiled module object is returned.
    """
    # Create a LLVM module object from the IR
    mod = llvm.parse_assembly(llvm_ir)
    mod.verify()
    # Now add the module and make sure it is ready for execution
    engine.add_module(mod)
    engine.finalize_object()
    return mod

engine = create_execution_engine()
mod = compile_ir(engine, llvm_ir)

# Look up the function pointer (a Python int)
func_ptr = engine.get_function_address("fpadd")

# Run the function via ctypes
cfunc = CFUNCTYPE(c_double, c_double, c_double)(func_ptr)
res = cfunc(1.0, 3.5)
print("fpadd(...) =", res)

```

Contributing

llvmlite originated to fulfill the needs of the [Numba](#) project. It is still mostly maintained by the [Numba](#) team. As such, we tend to prioritize the needs and constraints of [Numba](#) over other conflicting desires. However, we welcome any contributions, under the form of *bug reports* or *pull requests*.

Communication

Mailing-list

For now, we use the [Numba](#) public mailing-list, which you can e-mail at numba-users@continuum.io. If you have any questions about contributing to llvmlite, it is ok to ask them on this mailing-list. You can subscribe and read the archives on [Google Groups](#), and there is also a [Gmane mirror](#) allowing NNTP access.

Bug tracker

We use the [Github issue tracker](#) to track both bug reports and feature requests. If you report an issue, please include specifics:

- what you are trying to do;
- which operating system you have and which version of llvmlite you are running;
- how llvmlite is misbehaving, e.g. the full error traceback, or the unexpected results you are getting;
- as far as possible, a code snippet that allows full reproduction of your problem.

Pull requests

If you want to contribute code, we recommend you fork our [Github repository](#), then create a branch representing your work. When your work is ready, you should submit it as a pull request from the Github interface.

Development rules

Coding conventions

All Python code should follow [PEP 8](#). Our C++ code doesn't have a well-defined coding style (would it be nice to follow [PEP 7?](#)). Code and documentation should generally fit within 80 columns, for maximum readability with all existing tools (such as code review UIs).

Platform support

llvmlite is to be kept compatible with Python 2.7, 3.4 and later under at least Linux, OS X and Windows. It only needs to be compatible with the currently supported LLVM version (currently, the 3.8 series).

We don't expect contributors to test their code on all platforms. Pull requests are automatically built and tested using [Travis-CI](#). This takes care of Linux compatibility. Other operating systems are tested on an internal continuous integration platform at Continuum Analytics.

Documentation

This documentation is maintained in the `docs` directory inside the [llvmlite repository](#). It is built using Sphinx.

You can edit the source files under `docs/source/`, after which you can build and check the documentation:

```
$ make html
$ open _build/html/index.html
```

Glossary

basic block A sequence of instructions inside a function. A basic block always starts with a *label* and ends with a *terminator*. No other instruction inside the basic block can transfer control out of the block.

function declaration The specification of a function's prototype (including the argument and return types, and other information such as the calling convention), without an associated implementation. This is like a `extern` function declaration in C.

function definition A function's prototype (like in a *function declaration*) plus a body implementing the function.

getelementptr A LLVM *instruction* allowing to get the address of a subelement of an aggregate data structure.

See also:

Official documentation: [‘getelementptr’ Instruction](#)

global value A named value accessible to all members of a module.

global variable A variable whose value is accessible to all members of a module. Under the hood, it is a constant pointer to a module-allocated slot of the given type.

All global variables are global values. However, the converse is not true: a function declaration or definition is not a global variable; it is only a *global value*.

instruction The fundamental element(s) used in implementing a LLVM function. LLVM instructions define a procedural assembly-like language.

IR

Intermediate Representation A high-level assembly language describing to LLVM the code to be compiled to native code.

label A branch target inside a function. A label always denotes the start of a *basic block*.

metadata Optional ancillary information which can be associated with LLVM instructions, functions, etc. Metadata is used to convey certain information which is not critical to the compiling of LLVM *IR* (such as the likelihood of a condition branch or the source code location corresponding to a given instruction).

module A compilation unit for LLVM *IR*. A module can contain any number of function declarations and definitions, global variables, and metadata.

terminator

terminator instruction A kind of *instruction* which explicitly transfers control to another part of the program (instead of simply going to the next instruction after it is executed). Examples are branches and function returns.

Release Notes

v0.17.0

In this release, we are upgrading to LLVM 4.0. We are also starting to provide wheel packages for 64-bit Linux platforms (manylinux).

Fixes:

- Issue #249, PR #250: Disable static linking of libstdc++ by default.

Enhancements:

- PR #246: Add requirements.txt for pip dependency resolving
- PR #238: LLVM 4.0
- PR #222: Enable wheel builds

v0.16.0

API changes:

- Switched from LLVM 3.8 to 3.9
- `TargetData.add_pass` is removed in LLVM 3.9.

Enhancements:

- PR #239: Enable fastmath flags
- PR #233: Updates for llvm3.9.1
- PR #199: Update for changes in LLVM 3.9

Fixes:

- PR #236: Fix metadata with long value
- PR #231: Fix setup.py for Python2.7 so that pip auto installs dependencies
- PR #226: Fix `get_host_cpu_features()` so that it fails properly

v0.15.0

Enhancements:

- PR #213: Add partial LLVM bindings for ObjectFile.
- PR #215: Add inline assembly helpers in the builder.
- PR #216: Allow specifying alignment in alloca instructions.
- PR #219: Remove unnecessary verify in module linkage.

Fixes:

- PR #209, Issue #208: Fix overly restrictive test for library filenames.

v0.14.0

Enhancements:

- PR #104: Add binding to get and view function control-flow graph.
- PR #210: Improve llvmdrv recipe.
- PR #212: Add initializer for the native assembly parser.

v0.13.0

Enhancements:

- PR #176: Switch from LLVM 3.7 to LLVM 3.8.
- PR #191: Allow setting the alignment of a global variable.
- PR #198: Add missing function attributes.
- PR #160: Escape the contents of metadata strings, to allow embedding any characters.
- PR #162: Add support for creating debug information nodes.
- PR #200: Improve the usability of metadata emission APIs.
- PR #200: Allow calling functions with metadata arguments (such as `@llvm.dbg.declare`).

Fixes:

- PR #190: Suppress optimization remarks printed out in some cases by LLVM.
- PR #200: Allow attaching metadata to a `ret` instruction.

v0.12.1

New release to fix packages on PyPI. Same as v0.12.0.

v0.12.0

Enhancements:

- PR #179: Let llvmlite build on armv7l Linux.
- PR #161: Allow adding metadata to functions.
- PR #163: Allow emitting fast-math `fcmp` instructions.
- PR #159: Allow emitting verbose assembly in TargetMachine.

Fixes:

- Issue #177: Make `setup.py` compatible with `pip install`.

v0.11.0

Enhancements:

- PR #175: Check LLVM version at build time
- PR #169: Default initializer for non-external global variable

- PR #168: add `ir.Constant.literal_array()`

v0.10.0

Enhancements:

- PR #146: Improve `setup.py clean` to wipe more leftovers.
- PR #135: Remove some llvmpy compatibility APIs.
- PR #151: Always copy `TargetData` when adding to a pass manager.
- PR #148: Make errors more explicit on loading the binding DLL.
- PR #144: Allow overriding `-flto` in Linux builds.
- PR #136: Remove Python 2.6 and 3.3 compatibility.
- Issue #131: Allow easier creation of constants by making type instances callable.
- Issue #130: The test suite now ensures the runtime DLL dependencies are within a certain expected set.
- Issue #121: Simplify build process on Unix and remove hardcoded linking with `LLVMOPProfileJIT`.
- Issue #125: Speed up formatting of raw array constants.

Fixes:

- PR #155: Properly emit IR for metadata null.
- PR #153: Remove deprecated uses of `TargetMachine::getDataLayout()`.
- PR #156: Move personality from `LandingPadInstr` to `FunctionAttributes`. It was moved in LLVM 3.7.
- PR #149: Implement LLVM scoping correctly.
- PR #141: Ensure no `CMakeCache.txt` file is included in sdist.
- PR #132: Correct constant in `llvmir.py` example.

v0.9.0

Enhancements:

- PR #73: Add `get_process_triple()` and `get_host_cpu_features()`
- Switch from LLVM 3.6 to LLVM 3.7. The generated IR for some memory operations has changed.
- Improved performance of IR serialization.
- Issue #116: improve error message when the operands of a binop have differing types.
- PR #113: Let `Type.get_abi_{size,alignment}` not choke on identified types.
- PR #112: Support non-alphanumeric characters in type names.

Fixes:

- Remove the `libcurses` dependency on Linux.

v0.8.0

- Update LLVM to 3.6.2
- Add an *align* parameter to `IRBuilder.load()` and `IRBuilder.store()`.
- Allow setting visibility, DLL storageclass of `ValueRef`
- Support profiling with `OProfile`

v0.7.0

- PR #88: Provide hooks into the MCJIT object cache
- PR #87: Add indirect branches and exception handling APIs to `ir.Builder`.
- PR #86: Add `ir.Builder` APIs for integer arithmetic with overflow
- Issue #76: Fix non-Windows builds when LLVM was built using CMake
- Deprecate `.get_pointer_to_global()` and add `.get_function_address()` and `.get_global_value_address()` in `ExecutionEngine`.

v0.6.0

Enhancements:

- Switch from LLVM 3.5 to LLVM 3.6. The generated IR for metadata nodes has slightly changed, and the “old JIT” engine has been remove (only MCJIT is now available).
- Add an optional `flags` argument to arithmetic instructions on `IRBuilder`.
- Support attributes on the return type of a function.

v0.5.1

Fixes:

- Fix implicit termination of basic block in nested `if_then()`

v0.5.0

New documentation hosted at <http://llvmlite.pydata.org>

Enhancements:

- Add code-generation helpers from `numba.cgutils`
- Support for `memset`, `memcpy`, `memmove` intrinsics

Fixes:

- Fix string encoding problem when round-triping `parse_assembly()`

v0.4.0

Enhancements:

- Add `Module.get_global()`
- Renamed `Module.global_variables` to `Module.global_values`
- Support loading library permanently
- Add `Type.get_abi_alignment()`

Fixes:

- Expose LLVM version as a tuple

Patched LLVM 3.5.1: Updated to 3.5.1 with the same ELF relocation patched for v0.2.2.

v0.2.2

Enhancements:

- Support for `addrspacecast`
- Support for tail call, calling convention attribute
- Support for `IdentifiedStructType`

Fixes:

- GEP `addrspace` propagation
- Various installation process fixes

Patched LLVM 3.5: The binaries from the numba binstar channel use a patched LLVM3.5 for fixing a LLVM ELF relocation bug that is caused by the use of 32-bit relative offset in 64-bit binaries. The problem appears to occur more often on hardened kernels, like in CentOS. The patched source code is available at: <https://github.com/numba/llvm-mirror/releases/tag/3.5p1>

v0.2.0

This is the first official release. It contains a few feature additions and bug fixes. It meets all requirements to replace `llvmpy` in `numba` and `numbapro`.

v0.1.0

This is the first release. This is released for beta testing `llvmlite` and `numba` before the official release.

CHAPTER 2

Indices

- `genindex`
- `modindex`

I

`llvmlite.binding`, [19](#)
`llvmlite.ir`, [3](#)

Symbols

`__call__()` (llvmlite.ir.Type method), 3

A

`add()` (llvmlite.ir.IRBuilder method), 13
`add()` (llvmlite.ir.NamedMetaData method), 7
`add_analysis_passes()` (llvmlite.binding.TargetMachine method), 22
`add_attribute()` (llvmlite.ir.Argument method), 6
`add_case()` (llvmlite.ir.SwitchInstr method), 8
`add_clause()` (llvmlite.ir.LandingPad method), 9
`add_debug_info()` (llvmlite.ir.Module method), 9
`add_destination()` (llvmlite.ir.IndirectBranch method), 9
`add_global()` (llvmlite.ir.Module method), 10
`add_incoming()` (llvmlite.ir.PhiInstr method), 9
`add_metadata()` (llvmlite.ir.Module method), 10
`add_module()` (llvmlite.binding.ExecutionEngine method), 25
`add_named_metadata()` (llvmlite.ir.Module method), 10
`add_symbol()` (in module llvmlite.binding), 20
`address_of_symbol()` (in module llvmlite.binding), 20
`addrspace` (llvmlite.ir.PointerType attribute), 4
`addrspacecast()` (llvmlite.ir.IRBuilder method), 15
Aggregate (class in llvmlite.ir), 4
`align` (llvmlite.ir.GlobalVariable attribute), 7
`alloca()` (llvmlite.ir.IRBuilder method), 16
`and_()` (llvmlite.ir.IRBuilder method), 14
`append_basic_block()` (llvmlite.ir.Function method), 7
`append_basic_block()` (llvmlite.ir.IRBuilder method), 12
`args` (llvmlite.ir.Function attribute), 8
Argument (class in llvmlite.ir), 6
ArrayType (class in llvmlite.ir), 4
`as_bitcode()` (llvmlite.binding.ModuleRef method), 23
`as_pointer()` (llvmlite.ir.Type method), 3
`ashr()` (llvmlite.ir.IRBuilder method), 13
`asm()` (llvmlite.ir.IRBuilder method), 17
`assume()` (llvmlite.ir.IRBuilder method), 18
`atomic_rmw()` (in module llvmlite.ir), 16
attributes (llvmlite.ir.Function attribute), 8

B

basic block, 31
`basic_block` (llvmlite.ir.BlockAddress attribute), 6
`bitcast()` (llvmlite.ir.Constant method), 5
`bitcast()` (llvmlite.ir.IRBuilder method), 15
Block (class in llvmlite.ir), 6
`block` (llvmlite.ir.IRBuilder attribute), 11
BlockAddress (class in llvmlite.ir), 6
`branch()` (llvmlite.ir.IRBuilder method), 16

C

`call()` (llvmlite.ir.IRBuilder method), 16
`calling_convention` (llvmlite.ir.Function attribute), 8
CatchClause (class in llvmlite.ir), 9
`cbranch()` (llvmlite.ir.IRBuilder method), 16
`check_jit_execution()` (in module llvmlite.binding), 25
`cmpxchg()` (in module llvmlite.ir), 16
Constant (class in llvmlite.ir), 5
`create_mcjit_compiler()` (in module llvmlite.binding), 25
`create_target_data()` (in module llvmlite.binding), 21
`create_target_machine()` (llvmlite.binding.Target method), 21

D

`data_layout` (llvmlite.binding.ModuleRef attribute), 23
`data_layout` (llvmlite.ir.Module attribute), 10
`debug_metadata` (llvmlite.ir.IRBuilder attribute), 11
`description` (llvmlite.binding.Target attribute), 21
`disable_unroll_loops` (llvmlite.binding.PassManagerBuilder attribute), 26
DIToken (class in llvmlite.ir), 6
DIValue (class in llvmlite.ir), 6
DoubleType (class in llvmlite.ir), 4

E

elements (llvmlite.ir.Aggregate attribute), 4
`emit_assembly()` (llvmlite.binding.TargetMachine method), 22

emit_object() (llvmlite.binding.TargetMachine method), 22
 ExecutionEngine (class in llvmlite.binding), 25
 extract_value() (llvmlite.ir.IRBuilder method), 16

F

fadd() (llvmlite.ir.IRBuilder method), 14
 fcmp_ordered() (llvmlite.ir.IRBuilder method), 15
 fcmp_unordered() (llvmlite.ir.IRBuilder method), 15
 fdiv() (llvmlite.ir.IRBuilder method), 14
 FeatureMap (class in llvmlite.binding), 22
 FilterClause (class in llvmlite.ir), 9
 finalize() (llvmlite.binding.FunctionPassManager method), 28
 finalize_object() (llvmlite.binding.ExecutionEngine method), 25
 flatten() (llvmlite.binding.FeatureMap method), 22
 FloatType (class in llvmlite.ir), 4
 fmul() (llvmlite.ir.IRBuilder method), 14
 fpext() (llvmlite.ir.IRBuilder method), 15
 fptosi() (llvmlite.ir.IRBuilder method), 15
 fptoui() (llvmlite.ir.IRBuilder method), 15
 fptrunc() (llvmlite.ir.IRBuilder method), 14
 frem() (llvmlite.ir.IRBuilder method), 14
 from_default_triple() (llvmlite.binding.Target class method), 21
 from_triple() (llvmlite.binding.Target class method), 21
 fsub() (llvmlite.ir.IRBuilder method), 14
 Function (class in llvmlite.ir), 7
 function (llvmlite.ir.Block attribute), 6
 function (llvmlite.ir.BlockAddress attribute), 6
 function (llvmlite.ir.Instruction attribute), 8
 function (llvmlite.ir.IRBuilder attribute), 11
 function declaration, 31
 function definition, 31
 FunctionPassManager (class in llvmlite.binding), 28
 functions (llvmlite.binding.ModuleRef attribute), 23
 functions (llvmlite.ir.Module attribute), 10
 FunctionType (class in llvmlite.ir), 4

G

gep() (llvmlite.ir.Constant method), 5
 gep() (llvmlite.ir.IRBuilder method), 16
 get_abi_alignment() (llvmlite.ir.Type method), 3
 get_abi_size() (llvmlite.binding.TargetData method), 21
 get_abi_size() (llvmlite.ir.Type method), 3
 get_default_triple() (in module llvmlite.binding), 20
 get_function() (llvmlite.binding.ModuleRef method), 23
 get_function_address() (llvmlite.binding.ExecutionEngine method), 26
 get_function_cfg() (in module llvmlite.binding), 28
 get_global() (llvmlite.ir.Module method), 10
 get_global_value_address() (llvmlite.binding.ExecutionEngine method), 26

get_global_variable() (llvmlite.binding.ModuleRef method), 23
 get_host_cpu_features() (in module llvmlite.binding), 20
 get_host_cpu_name() (in module llvmlite.binding), 20
 get_named_metadata() (llvmlite.ir.Module method), 10
 get_object_format() (in module llvmlite.binding), 20
 get_pointee_abi_alignment() (llvmlite.binding.TargetData method), 21
 get_pointee_abi_size() (llvmlite.binding.TargetData method), 21
 get_pointer_to_function() (llvmlite.binding.ExecutionEngine method), 25
 get_process_triple() (in module llvmlite.binding), 20
 get_unique_name() (llvmlite.ir.Module method), 10
 getelementptr, 31
 global value, 31
 global variable, 31
 global_constant (llvmlite.ir.GlobalVariable attribute), 7
 global_values (llvmlite.ir.Module attribute), 10
 global_variables (llvmlite.binding.ModuleRef attribute), 23
 GlobalValue (class in llvmlite.ir), 7
 GlobalVariable (class in llvmlite.ir), 7
 goto_block() (llvmlite.ir.IRBuilder method), 12
 goto_entry_block() (llvmlite.ir.IRBuilder method), 12

I

icmp_signed() (llvmlite.ir.IRBuilder method), 15
 icmp_unsigned() (llvmlite.ir.IRBuilder method), 15
 if_else() (llvmlite.ir.IRBuilder method), 12
 if_then() (llvmlite.ir.IRBuilder method), 12
 indirectbr() (llvmlite.ir.IRBuilder method), 17
 IndirectBranch (class in llvmlite.ir), 8
 initialize() (in module llvmlite.binding), 19
 initialize() (llvmlite.binding.FunctionPassManager method), 28
 initialize_all_asmprinters() (in module llvmlite.binding), 19
 initialize_all_targets() (in module llvmlite.binding), 19
 initialize_native_asmparser() (in module llvmlite.binding), 20
 initialize_native_asmprinter() (in module llvmlite.binding), 19
 initialize_native_target() (in module llvmlite.binding), 19
 initializer (llvmlite.ir.GlobalVariable attribute), 7
 inlining_threshold (llvmlite.binding.PassManagerBuilder attribute), 26
 insert_basic_block() (llvmlite.ir.Function method), 7
 insert_value() (llvmlite.ir.IRBuilder method), 16
 instruction, 31
 Instruction (class in llvmlite.ir), 8
 Intermediate Representation, 31
 inttoptr() (llvmlite.ir.Constant method), 5
 inttoptr() (llvmlite.ir.IRBuilder method), 15

IntType (class in llvmlite.ir), 4
 IR, 31
 IRBuilder (class in llvmlite.ir), 11
 is_declaration (llvmlite.binding.ValueRef attribute), 24
 is_declaration (llvmlite.ir.Function attribute), 8
 is_terminated (llvmlite.ir.Block attribute), 6

L

label, 31
 LabelType (class in llvmlite.ir), 5
 LandingPad (class in llvmlite.ir), 9
 landingpad() (llvmlite.ir.IRBuilder method), 17
 link_in() (llvmlite.binding.ModuleRef method), 23
 Linkage (class in llvmlite.binding), 23
 linkage (llvmlite.binding.ValueRef attribute), 24
 linkage (llvmlite.ir.GlobalValue attribute), 7
 Linkage.appending (in module llvmlite.binding), 24
 Linkage.available_externally (in module llvmlite.binding), 23
 Linkage.common (in module llvmlite.binding), 24
 Linkage.dllexport (in module llvmlite.binding), 24
 Linkage.dllimport (in module llvmlite.binding), 24
 Linkage.external (in module llvmlite.binding), 23
 Linkage.external_weak (in module llvmlite.binding), 24
 Linkage.ghost (in module llvmlite.binding), 24
 Linkage.internal (in module llvmlite.binding), 24
 Linkage.linker_private (in module llvmlite.binding), 24
 Linkage.linker_private_weak (in module llvmlite.binding), 24
 Linkage.linkonce_any (in module llvmlite.binding), 23
 Linkage.linkonce_odr (in module llvmlite.binding), 24
 Linkage.linkonce_odr_autohide (in module llvmlite.binding), 24
 Linkage.private (in module llvmlite.binding), 24
 Linkage.weak_any (in module llvmlite.binding), 24
 Linkage.weak_odr (in module llvmlite.binding), 24
 literal_array() (llvmlite.ir.Constant class method), 5
 literal_struct() (llvmlite.ir.Constant class method), 5
 LiteralStructType (class in llvmlite.ir), 4
 llvm_version_info (in module llvmlite.binding), 20
 llvmlite.binding (module), 19
 llvmlite.ir (module), 3
 load() (llvmlite.ir.IRBuilder method), 16
 load_library_permanently() (in module llvmlite.binding), 20
 load_reg() (llvmlite.ir.IRBuilder method), 18
 loop_vectorize (llvmlite.binding.PassManagerBuilder attribute), 26
 lshr() (llvmlite.ir.IRBuilder method), 13

M

MDValue (class in llvmlite.ir), 6
 metadata, 31
 MetadataString (class in llvmlite.ir), 6

MetadataType (class in llvmlite.ir), 5
 module, 31
 Module (class in llvmlite.ir), 9
 module (llvmlite.binding.ValueRef attribute), 24
 module (llvmlite.ir.Instruction attribute), 8
 module (llvmlite.ir.IRBuilder attribute), 11
 ModulePassManager (class in llvmlite.binding), 27
 ModuleRef (class in llvmlite.binding), 23
 mul() (llvmlite.ir.IRBuilder method), 13

N

name (llvmlite.binding.ModuleRef attribute), 23
 name (llvmlite.binding.Target attribute), 21
 name (llvmlite.binding.ValueRef attribute), 24
 NamedMetadata (class in llvmlite.ir), 6
 neg() (llvmlite.ir.IRBuilder method), 14
 not_() (llvmlite.ir.IRBuilder method), 14

O

opt_level (llvmlite.binding.PassManagerBuilder attribute), 26
 or_() (llvmlite.ir.IRBuilder method), 14

P

parse_assembly() (in module llvmlite.binding), 22
 parse_bitcode() (in module llvmlite.binding), 22
 PassManager (class in llvmlite.binding), 27
 PassManager.add_basic_alias_analysis_pass() (in module llvmlite.binding), 27
 PassManager.add_cfg_simplification_pass() (in module llvmlite.binding), 27
 PassManager.add_constant_merge_pass() (in module llvmlite.binding), 27
 PassManager.add_dead_arg_elimination_pass() (in module llvmlite.binding), 27
 PassManager.add_dead_code_elimination_pass() (in module llvmlite.binding), 27
 PassManager.add_function_attrs_pass() (in module llvmlite.binding), 27
 PassManager.add_function_inlining_pass() (in module llvmlite.binding), 27
 PassManager.add_global_dce_pass() (in module llvmlite.binding), 27
 PassManager.add_global_optimizer_pass() (in module llvmlite.binding), 27
 PassManager.add_gvn_pass() (in module llvmlite.binding), 27
 PassManager.add_instruction_combining_pass() (in module llvmlite.binding), 27
 PassManager.add_ipscpp_pass() (in module llvmlite.binding), 27
 PassManager.add LICM_pass() (in module llvmlite.binding), 27

[PassManager.add_sccp_pass\(\)](#) (in module `llvmlite.binding`), [27](#)
[PassManager.add_sroa_pass\(\)](#) (in module `llvmlite.binding`), [27](#)
[PassManager.add_type_based_alias_analysis_pass\(\)](#) (in module `llvmlite.binding`), [27](#)
[PassManagerBuilder](#) (class in `llvmlite.binding`), [26](#)
[phi\(\)](#) (`llvmlite.ir.IRBuilder` method), [15](#)
[PhiInstr](#) (class in `llvmlite.ir`), [9](#)
[pointee](#) (`llvmlite.ir.PointerType` attribute), [4](#)
[PointerType](#) (class in `llvmlite.ir`), [4](#)
[populate\(\)](#) (`llvmlite.binding.PassManagerBuilder` method), [26](#)
[position_after\(\)](#) (`llvmlite.ir.IRBuilder` method), [12](#)
[position_at_end\(\)](#) (`llvmlite.ir.IRBuilder` method), [12](#)
[position_at_start\(\)](#) (`llvmlite.ir.IRBuilder` method), [12](#)
[position_before\(\)](#) (`llvmlite.ir.IRBuilder` method), [12](#)
[PredictableInstr](#) (class in `llvmlite.ir`), [8](#)
[ptrtoint\(\)](#) (`llvmlite.ir.IRBuilder` method), [15](#)
[Python Enhancement Proposals](#)
 [PEP 7](#), [30](#)
 [PEP 8](#), [30](#)

R

[remove_module\(\)](#) (`llvmlite.binding.ExecutionEngine` method), [26](#)
[replace\(\)](#) (`llvmlite.ir.Block` method), [6](#)
[resume\(\)](#) (`llvmlite.ir.IRBuilder` method), [17](#)
[ret\(\)](#) (`llvmlite.ir.IRBuilder` method), [16](#)
[ret_void\(\)](#) (`llvmlite.ir.IRBuilder` method), [17](#)
[run\(\)](#) (`llvmlite.binding.FunctionPassManager` method), [28](#)
[run\(\)](#) (`llvmlite.binding.ModulePassManager` method), [28](#)

S

[sadd_with_overflow\(\)](#) (`llvmlite.ir.IRBuilder` method), [13](#)
[sdiv\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)
[select\(\)](#) (`llvmlite.ir.IRBuilder` method), [15](#)
[set_asm_verbosity\(\)](#) (`llvmlite.binding.TargetMachine` method), [22](#)
[set_metadata\(\)](#) (`llvmlite.ir.Function` method), [8](#)
[set_metadata\(\)](#) (`llvmlite.ir.Instruction` method), [8](#)
[set_object_cache\(\)](#) (`llvmlite.binding.ExecutionEngine` method), [26](#)
[set_weights\(\)](#) (`llvmlite.ir.PredictableInstr` method), [8](#)
[sext\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)
[shl\(\)](#) (`llvmlite.ir.IRBuilder` method), [13](#)
[shutdown\(\)](#) (in module `llvmlite.binding`), [20](#)
[sitofp\(\)](#) (`llvmlite.ir.IRBuilder` method), [15](#)
[size_level](#) (`llvmlite.binding.PassManagerBuilder` attribute), [26](#)
[slp_vectorize](#) (`llvmlite.binding.PassManagerBuilder` attribute), [27](#)
[smul_with_overflow\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)

[srem\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)
[storage_class](#) (`llvmlite.binding.ValueRef` attribute), [25](#)
[storage_class](#) (`llvmlite.ir.GlobalValue` attribute), [7](#)
[StorageClass](#) (class in `llvmlite.binding`), [24](#)
[StorageClass.default](#) (in module `llvmlite.binding`), [24](#)
[StorageClass.dllexport](#) (in module `llvmlite.binding`), [24](#)
[StorageClass.dllimport](#) (in module `llvmlite.binding`), [24](#)
[store\(\)](#) (`llvmlite.ir.IRBuilder` method), [16](#)
[store_reg\(\)](#) (`llvmlite.ir.IRBuilder` method), [18](#)
[sub\(\)](#) (`llvmlite.ir.IRBuilder` method), [13](#)
[switch\(\)](#) (`llvmlite.ir.IRBuilder` method), [17](#)
[SwitchInstr](#) (class in `llvmlite.ir`), [8](#)

T

[Target](#) (class in `llvmlite.binding`), [21](#)
[target_data](#) (`llvmlite.binding.ExecutionEngine` attribute), [26](#)
[target_data](#) (`llvmlite.binding.TargetMachine` attribute), [22](#)
[TargetData](#) (class in `llvmlite.binding`), [21](#)
[TargetMachine](#) (class in `llvmlite.binding`), [22](#)
[terminator](#), [31](#)
[terminator](#) (`llvmlite.ir.Block` attribute), [6](#)
[terminator instruction](#), [31](#)
[triple](#) (`llvmlite.binding.ModuleRef` attribute), [23](#)
[triple](#) (`llvmlite.binding.Target` attribute), [21](#)
[triple](#) (`llvmlite.ir.Module` attribute), [10](#)
[trunc\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)
[Type](#) (class in `llvmlite.ir`), [3](#)
[type](#) (`llvmlite.binding.ValueRef` attribute), [24](#)

U

[udiv\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)
[uitofp\(\)](#) (`llvmlite.ir.IRBuilder` method), [15](#)
[Undefined](#) (in module `llvmlite.ir`), [5](#)
[unnamed_addr](#) (`llvmlite.ir.GlobalVariable` attribute), [7](#)
[unreachable\(\)](#) (`llvmlite.ir.IRBuilder` method), [18](#)
[urem\(\)](#) (`llvmlite.ir.IRBuilder` method), [14](#)

V

[Value](#) (class in `llvmlite.ir`), [5](#)
[ValueRef](#) (class in `llvmlite.binding`), [24](#)
[verify\(\)](#) (`llvmlite.binding.ModuleRef` method), [23](#)
[view_dot_graph\(\)](#) (in module `llvmlite.binding`), [28](#)
[Visibility](#) (class in `llvmlite.binding`), [24](#)
[visibility](#) (`llvmlite.binding.ValueRef` attribute), [25](#)
[Visibility.default](#) (in module `llvmlite.binding`), [24](#)
[Visibility.hidden](#) (in module `llvmlite.binding`), [24](#)
[Visibility.protected](#) (in module `llvmlite.binding`), [24](#)
[VoidType](#) (class in `llvmlite.ir`), [4](#)

W

[width](#) (`llvmlite.ir.IntType` attribute), [4](#)

X

`xor()` (Ilvmlite.ir.IRBuilder method), [14](#)

Z

`zext()` (Ilvmlite.ir.IRBuilder method), [14](#)