

Julia

Lukas Koschmieder, Kai Neumann

July 22, 2014

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Language classification

- ▶ high-level dynamic language
- ▶ designed for technical computing
- ▶ emphasis on performance

Motivation

“The most challenging areas of technical programming benefit the least from technical languages.”

- ▶ dynamic languages
 - ▶ convenient and highly productive
 - ▶ lack in speed
- ▶ compromise (two-tier approach)
 - ▶ high-level logic in dynamic language
 - ▶ computationally-intensive routines in C or Fortran
- ▶ Julia
 - ▶ as fast as a statically compiled language
 - ▶ as interactive/productive as a dynamic language

Overview

- 1 Introduction
- 2 Layers**
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Code representations

1. AST after parsing
2. AST after lowering
3. AST after type inference
4. LLVM IR
5. Assembly code

Layers 1: AST

- ▶ abstract syntax tree
- ▶ internal representation of Julia code
- ▶ specific to Julia
- ▶ produced by parsing Julia code
- ▶ operators become function calls
- ▶ representation is not stored

Layers 1: AST - Example

Infix operator removal

```
julia> 1+2
3
julia> :(1+2)
:(+(1,2))
julia> +(1,2)
3
julia> :(+((1,2)))
:(+(1,2))
julia> :(1+2+3+4+5)
:(+(1,2,3,4,5))
julia> :(1+2-3-4+5)
:(+(-(-+(1,2),3),4),5))
```

Layers 2: Lowered AST

- ▶ lowering
 - ▶ in general: process of moving from surface syntax (highest) to machine code (lowest)
 - ▶ here: transformation of the AST into a simpler version
 - ▶ unnesting expressions
 - ▶ desugaring syntax
 - ▶ simplifying control flow
 - ▶ reduction of the instruction set
 - ▶ etc.

Layers 2: Lowered AST - Example

for loop

```
julia> function loop(x::Int)
    y = 0
    for i = 1:x
        y += x
    end
    y
end
```

Layers 2: Lowered AST - Example

for loop lowered to unless/goto

```

1-element Array{Any,1}:
:($ (Expr(:lambda, {:x}, [{:y, :#s6, :#s5, :i}], [{:x, :Any, 0}, {:y, :Any, 2}, {:#s6, :Any, 2}, {:#s5, :Any, 18}, {:i, :Any, 18}], {}), quote #
none, line 2:
  y = 0 # line 3:
  #s6 = 1
  #s5 = x
  1:
  unless top(<=) (#s6, #s5) goto 2
  i = #s6 # line 4:
  y = +(y, x)
  3:
  #s6 = top(convert)(top(typeof) (#s6), top(+)(1, #s6))
  goto 1
  2:
  0: # line 6:
  return y

```

Layers 3: Type-inferred AST

- ▶ Type inference
 - ▶ automatic deduction of the type of an expression at compile time
 - ▶ also: optimization according to types

Layers 3: Type-inferred AST - Ex. 1

integer definition

```
julia> function foo()  
    x = 1  
end
```

string definition

```
julia> function bar()  
    x = "bar"  
end
```

Layers 3: Type-inferred AST - Ex. 1

lowered integer definition

```
1-element Array{Any,1}:
 :($ (Expr(:lambda, {}, {{:x}}, {{:x, Any, 18}}), {}),
 quote # none, line 2:
     x = 1
     return 1
 End)))
```

lowered string definition

```
1-element Array{Any,1}:
 :($ (Expr(:lambda, {}, {{:x}}, {{:x, Any, 18}}), {}),
 quote # none, line 2:
     x = "bar"
     return "bar"
 end)))
```

Layers 3: Type-inferred AST - Ex. 1

type-inferred integer definition

```
1-element Array{Any,1}:
 :($ (Expr(:lambda, {}, {:{:x}}, {:{:x, Int32, 18}}), {}),
 quote # none, line 2:
     x = 1
     return 1
 End)))
```

type-inferred string definition

```
1-element Array{Any,1}:
 :($ (Expr(:lambda, {}, {:{:x}}, {:{:x, ASCIIString, 18}}), {}),
 quote # none, line 2:
     x = "bar"
     return "bar"
 end)))
```


Layers 3: Type-inferred AST - Ex. 2

star operator

```
julia> function foobar(x)
    x = x * x
end
```

star operator on integers

```
julia> foobar(256)
65536
```

star operator on strings

```
julia> foobar("256")
"256256"
```

Layers 3: Type-inferred AST - Ex. 2

lowered star operator

```
1-element Array{Any,1}:
 :($ (Expr(:lambda, {x}, {[:#s5]}, {[:x, :Any, 18]}, {[:#s5, :Any, 18]}}, {}), quote # none, line 2:
     #s5 = *(x,x)
     x = #s5
     return #s5
 end)))
```

Layers 3: Type-inferred AST - Ex. 2

type-inferred star operator on integers

```
(...)  
  #s5 = top(box)(Int,top(mul_int)(x::Int32,x::Int32))::Int32  
  x = #s5::Int32  
  return #s5::Int32  
(...)
```

type-inferred star operator on strings

```
(...)  
  #s5 = string(x::String,x::String)::Union(UTF8String,  
      ASCIIString)  
  x = #s5::Union(UTF8String,ASCIIString)  
  return #s5::Union(UTF8String,ASCIIString)  
(...)
```

Layers 4: LLVM IR

- ▶ LLVM IR
 - ▶ LLVM intermediate representation
 - ▶ not specific to Julia
 - ▶ passed to LLVM compiler
 - ▶ requires types

Layers 4: LLVM IR - Example

star operator on integers

```
define i32 @julia_foobar926(i32) {
top:
%1 = mul i32 %0, %0, !dbg !4746
ret i32 %1, !dbg !4746
}
```

star operator on strings

```
define %jl_value_t* @julia_foobar927(%jl_value_t*, %jl_value_t**,
    i32) {
top:
%3 = alloca [6 x %jl_value_t*], align 4
%.sub = getelementptr inbounds [6 x %jl_value_t*]* %3, i32 0, i32
(51 more lines)
}
```

Layers 5: Assembler code

- ▶ Assembler code
 - ▶ optimized machine-specific code
 - ▶ provided by LLVM compiler

star operator on integers

```
.text
Filename: none
Source line: 2
push EBP
mov EBP, ESP
mov EAX, 1
pop EBP
ret
```

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)**
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Multiple dispatch

“Multiple dispatch is the ability to choose which version of a function to call based on the run-time type of the arguments passed to the function call.”

- ▶ object-oriented programming
 - ▶ single dispatch
 - ▶ dispatch based only on 1st argument (object type)
- ▶ Java/C++ (single dispatch)
 - ▶ `object.function(arg1, arg2, ...)`
- ▶ Julia (multiple dispatch)
 - ▶ `function(object, arg1, arg2, ...)`

Multiple dispatch

Asteroids collision system in Julia

```
function collide_with(x::Asteroid, y::Asteroid) { ... }  
function collide_with(x::Asteroid, y::Spaceship) { ... }  
function collide_with(x::Spaceship, y::Asteroid) { ... }  
function collide_with(x::Spaceship, y::Spaceship) { ... }
```

- ▶ useful for mathematic code from the programmer's point of view

a+b in Julia

```
function +(a::Int, b::Int) = {  
    ...  
}  
+(a,b)
```

The key ingredients of performance

1. Rich type information provided naturally by multiple dispatch

The key ingredients of performance

1. Rich type information provided naturally by multiple dispatch
2. Aggressive code specialization against run-time types

The key ingredients of performance

1. Rich type information provided naturally by multiple dispatch
2. Aggressive code specialization against run-time types
3. JIT compilation using the LLVM compiler framework

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation**
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Julia Multiple Dispatch

- ▶ Julia chooses which function is called depending on its parameters' types
- ▶ the `@which` macro can be used to see which function julia will call on which parameters

Integer Multiplication

```
julia> @which 1*2  
*(x::Int64,y::Int64) at int.jl:55
```

Vector Multiplication

```
julia> A=ones(1,10); @which A*2  
*(A::AbstractArray{T,N},B::Number) at abstractarray.jl:340
```

Julia Multiple Dispatch: IR

Scalar Multiplication: $Int64 \times Int64$

```
define i64 @"julia_*1396"(i64, i64) {
top:
  %2 = mul i64 %1, %0, !dbg !6861
  ret i64 %2, !dbg !6861
}
```

Vector Multiplication: $Vector \times Int64$

```
define %jl_value_t* @"julia_*1378"(%jl_value_t*, i64) {
top:
  %2 = call %jl_value_t* @"julia_.*1380"(%jl_value_t* %0, i64 %1)
    , !dbg !6861
  ret %jl_value_t* %2, !dbg !6861
}
```

LLVM IR: Julia vs clang++

Increment: Julia

```
define i64 @julia_increment(i64) {  
top:  
  %1 = add i64 %0, 1, !dbg !6898  
  ret i64 %1, !dbg !6898  
}
```

Increment: Clang++

```
define i32 @_Z9incrementi(i32 %a) #0 {  
entry:  
  %a.addr = alloca i32, align 4  
  store i32 %a, i32* %a.addr, align 4  
  %0 = load i32* %a.addr, align 4  
  %add = add nsw i32 %0, 1  
  ret i32 %add  
}
```


Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance**
- 6 Parallelism
- 7 Language Features
- 8 Conclusion

Performance: Microbenchmarks

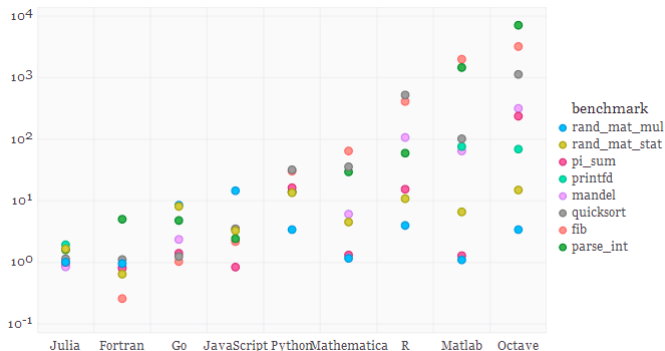


Figure : Benchmark times relative to C, <http://julialang.org>,
<https://github.com/JuliaLang/julia/tree/master/test/perf/micro>

Performance (2)

- ▶ Julia performance within factor of 2 compared to native C
- ▶ Julia Startup-Time not included
- ▶ Julia caches native code, first call to function is slower, but only min is measured
- ▶ Also: Other languages often faster in special cases

Performance (3)

CumProd in MATLAB 2014a

```
A=ones(100,1000000);b=0.9;tic;cumprod(A*b);toc  
Elapsed time is 0.679127 seconds.
```

CumProd in Julia

```
A=ones(100,1000000);b=0.9;@elapsed cumprod(A*b)  
2.988972363
```

⇒ Speedup = 4,4

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism**
- 7 Language Features
- 8 Conclusion

Parallelism

- ▶ Message-Passing Model
- ▶ Unlike MPI, Communication is 'one-sided'
- ▶ ⇒ View only from Master thread, call external functions and fetch results
- ▶ Macros for distributed loops

Parallel for loop

```
nheads = @parallel (+) for i=1:100000000
    int(randbool())
end
```

Parallelism(2)

Spawning functions on other processes

```
julia> r = @spawn rand(2,2)
RemoteRef(1,1,0)

julia> s = @spawn 1 .+ fetch(r)
RemoteRef(1,1,1)

julia> fetch(s)
1.10824216411304866 1.13798233877923116
1.12376292706355074 1.18750497916607167
```

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features**
- 8 Conclusion

Language Features

- ▶ Package Manager (Packages for Plotting, Statistics, Websockets, etc)
- ▶ Integration (C, Fortran Calls)
- ▶ Free & Open Source
- ▶ Active Community
- ▶ High-Level syntax

Overview

- 1 Introduction
- 2 Layers
- 3 Performance (theory)
- 4 Demo: LLVM IR Code Generation
- 5 Performance
- 6 Parallelism
- 7 Language Features
- 8 Conclusion**

Conclusion

- ▶ Julia is a dynamic high-level language using LLVM as JIT compiler
- ▶ 5 stages/layers to transform julia code into assembly code
- ▶ Julia aims on the gap between dynamic and statically compiled languages
- ▶ Even with nice performance, still a tradeoff between productivity and performance
- ▶ Does not outperform every other language in every category

More Information

- ▶ Interactive web-based Julia (including Tutorial):
<http://forio.com/julia/repl/>
- ▶ YouTube: JuliaLanguage
- ▶ <http://julialang.org/>

References I



Jeff Bezanson and Stefan Karpinski and Viral B. Shah and Alan Edelman

Julia: A Fast Dynamic Language for Technical Computing

CoRR, abs/1209.5145, 2012.



<http://julialang.org/>



<https://github.com/JuliaLang/julia>



<http://media.readthedocs.org/pdf/julia/latest/julia.pdf>



<http://blog.leahhanson.us/julia-introspects.html>



<http://scientopia.org/blogs/goodmath/2014/02/04/everyone-stop-implementing-programming-languages-right-now-its-been-solved/>