

Julia - a programming language

Uwe Hernandez Acosta (CASUS/HZDR) - January 9th, 2025



Introduction

Uwe Hernandez Acosta

- Particle physicist by training
- PhD in Physics 2021 at TU Dresden/HZDR
 - Topic: Strong field QED
- Affiliation: Center of Advanced Systems Understanding
- Research interests:
 - Theoretical particle physics/Quantum Field Theory
 - Strong field physics
 - X-ray diagnostics in matter under extreme conditions
 - Monte-Carlo Event Generation
 - Julia programming language

Survey says!

Survey says!

- I am able to write something useful in a scripting/interpreted programming language. (something like Python, R, Octave, Matlab, Shell, Javascript, Ruby, etc.)

Survey says!

- I am able to write something useful in a scripting/interpreted programming language. (something like Python, R, Octave, Matlab, Shell, Javascript, Ruby, etc.)
- I can write competitive programs in one these scripting languages. (e.g. in Python, competitive means, you know what a meta-class is)

Survey says!

- I am able to write something useful in a scripting/interpreted programming language. (something like Python, R, Octave, Matlab, Shell, Javascript, Ruby, etc.)
- I can write competitive programs in one these scripting languages. (e.g. in Python, competitive means, you know what a meta-class is)
- I can write something useful in a systems-level programming language. (e.g. C, C++, Fortran, Rust, Go, Kotlin, Java, etc.)

Survey says!


- I am able to write something useful in a scripting/interpreted programming language. (something like Python, R, Octave, Matlab, Shell, Javascript, Ruby, etc.)
- I can write competitive programs in one these scripting languages. (e.g. in Python, competitive means, you know what a meta-class is)
- I can write something useful in a systems-level programming language. (e.g. C, C++, Fortran, Rust, Go, Kotlin, Java, etc.)
- I can write competitive code in one of these systems-level languages. (e.g. in C++, competitive means, you can write a template class without getting mad)

Survey says!

- I am able to write something useful in a scripting/interpreted programming language. (something like Python, R, Octave, Matlab, Shell, Javascript, Ruby, etc.)
- I can write competitive programs in one these scripting languages. (e.g. in Python, competitive means, you know what a meta-class is)
- I can write something useful in a systems-level programming language. (e.g. C, C++, Fortran, Rust, Go, Kotlin, Java, etc.)
- I can write competitive code in one of these systems-level languages. (e.g. in C++, competitive means, you can write a template class without getting mad)
- I can write library code in one of these systems-level languages. (e.g. you are one of the authors of Boost)

Software requirements in HEP

- Efficiency
 - Fast execution
 - High data throughput
 - Scalability
- Developer-friendly
 - Quick bug fixes
 - Newest algorithms implemented
 - Good tooling
- User-friendly
 - Rapid development cycles
 - Low entry points
 - Interactivity




" [I propose that] you should use *two* languages for large software system: one, such as C or C++, for manipulating the complex internal data structures where performance is key and another, such as Tcl, for writing small-ish scripts that tie together the C pieces and are used for extensions."

[Ousterhout. "Re: Why you should not use Tcl" 1994] [Ousterhout. IEEE Computer magazine 31.3 (1998)]

Why is this problematic?

The two languages ~~problem~~ annoyance

- Rewriting parts == refactoring
- Different languages == different logics
- Need for glue code
- Extending is a mess
- Debugging is a mess
- Scientists need to be polyglot
- Multithreading? Anyone?




Established Solutions*?

*personal opinion

Why not use ...

... only systems-level languages?

- Take years to learn...
- ...decades to master
- Boilerplate code
- Hardware specific
- Mostly non-interactive
- Missing tools/libraries



Why not use third-party libraries?

- “Use C/C++ under the hood”
- Valid in their scope
- Hard to do something outside the box
- Interoperability? Anyone?
- The vendor decides what is performance-critical



theano

PyTorch

The TensorFlow logo features an orange stylized "TF" icon above the text "TensorFlow" in a bold, dark blue sans-serif font.

TensorFlow

The pandas logo features a dark blue icon of a bar chart with a yellow and pink square above the text "pandas" in a dark blue sans-serif font.

pandas


The dask logo features a colorful icon of three overlapping rectangles (yellow, pink, and red) to the left of the text "dask" in a bold, dark blue sans-serif font.

dask

Why not use ...

... Numba, PyPy, Pythran, etc?

- Sufficient for small code pieces
- These *are* second languages
 - Support only a subset of the host language(s) ...
 - ... and/or add new commands/ logic/concepts
- Usually not a systems-level language
 - e.g Numba is neither Python nor C

The Numba logo features a blue lightning bolt icon to the left of the word "Numba" in a bold, blue, sans-serif font.The PyPy logo consists of a stylized penguin icon on the left and the word "pypy" in a dark blue, lowercase, sans-serif font on the right.The CuPy logo features a 3D cube made of green and grey squares on the left and the word "CuPy" in a grey, sans-serif font on the right.The Cython logo shows a large grey letter 'C' on the left, with a blue and yellow Python logo inside it, followed by the word "ython" in a grey, sans-serif font.The Nuitka logo features a square icon with a blue and yellow geometric pattern on the left and the word "Nuitka" in a blue, sans-serif font on the right.

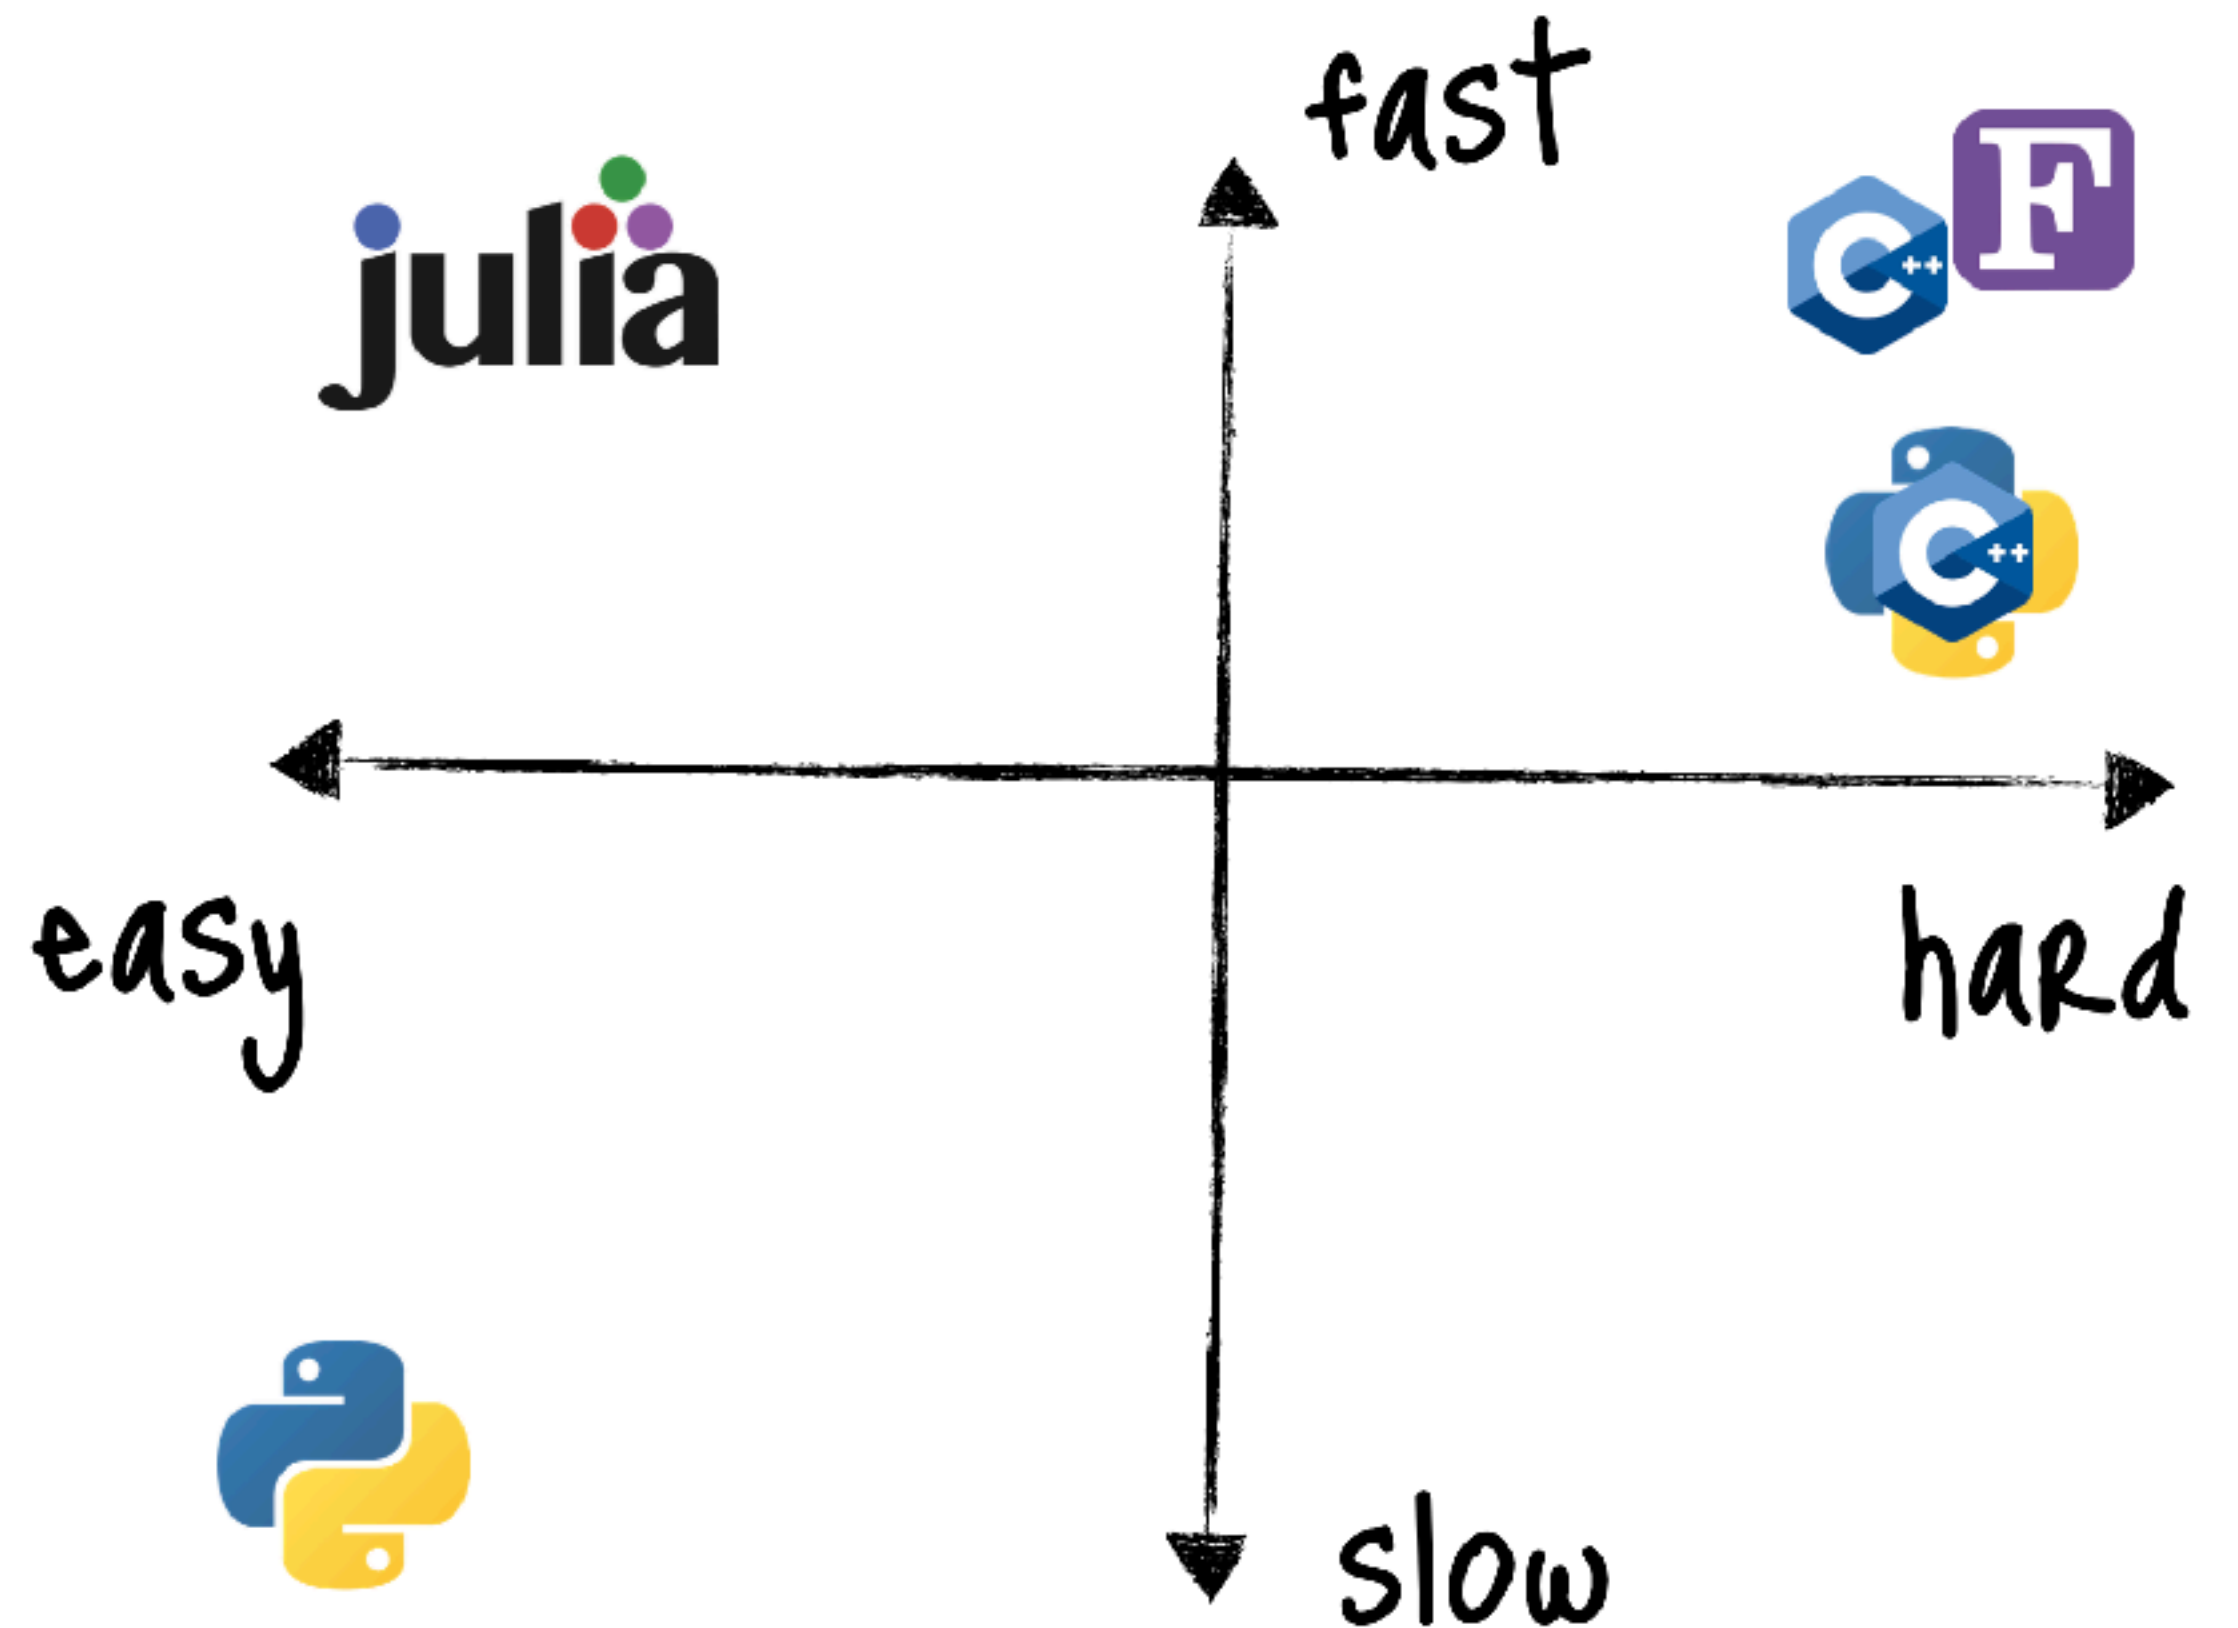
Proposal of a solution

Introduction



The Julia programming language

- Invented 2012 at MIT (mostly)
- Jeff Bezanson, Stefan Karpinski, Viral B. Shah, Alan Edelman
- Design goals
 - Open source
 - Speed like C, dynamic like Ruby
 - Obvious mathematical notation
 - General purpose like Python
 - As easy for statistics as R
 - Powerful linear algebra like in Matlab
 - Good for gluing programs together like the shell



"Something that is dirt simple to learn, yet keeps the most serious hackers happy."

Julia is easy

Ease of use

- Dynamically typed
- Powerful type system
- Garbage collection
- Extensive standard library
 - Mostly written in Julia
 - Math included
 - Performant
- Multiple dispatch for the win!

You can write Julia code as far away from the metal as you want!

```
using DifferentialEquations, Measurements, Plots

g = 9.79 ± 0.02; # Gravitational constants
L = 1.00 ± 0.01; # Length of the pendulum


#Initial Conditions
u₀ = [0 ± 0, π / 60 ± 0.01] # Initial speed and initial angle
tspan = (0.0, 6.3)

#Define the problem
function pendulum(du,u,p,t)
    θ = u[1]
    dθ = u[2]
    du[1] = dθ
    du[2] = -(g/L)*θ
end

#Pass to solvers
prob = ODEProblem(pendulum, u₀, tspan)
sol = solve(prob, Tsit5(), reltol = 1e-6)

# Analytic solution
u = u₀[2] .* cos.(sqrt(g / L) .* sol.t)

plot(sol.t, getindex.(sol.u, 2), label = "Numerical")
plot!(sol.t, u, label = "Analytic")
```




Julia is fast

Not an interpreter

- Just-ahead-of-time compiler
- LLVM empowered
- Statically sizes arrays
- Built-in vector/matrix types
- Arbitrary optimization
- Compiler reflections available
- Native thread support


You can write Julia code as close to the metal as you want!

CPU performance



Data taken from [<https://julialang.org/benchmarks/>]

GPU performance



Taken from [Besard et al. IEEE Trans. Parallel Distrib. Syst. 30.4 (2018)]

Julia is a modern language

Development tooling

```
QEDcore.jl
├── CHANGELOG.md
├── LICENSE
├── Manifest.toml
├── Project.toml
├── README.md
├── docs
├── src
└── test
```

Packaging system

```
name = "QEDcore"
uuid = "35dc0263-cb5f-4c33-a114-1d7f54ab753e"
authors = [
  "Uwe Hernandez Acosta <u.hernandez@hzdr.de>",
  "Anton Reinhard <a.reinhard@hzdr.de>",
]
version = "0.1.1"

[deps]
DocStringExtensions = "ffbed154-4ef7-542d-bbb7-c09d3a79fcae"
QEDbase = "10e22c08-3ccb-4172-bfcf-7d7aa3d04d93"
Reexport = "189a3867-3050-52da-a836-e530ba90ab69"
SimpleTraits = "699a6c99-e7fa-54fc-8d76-47d257e15c1d"
StaticArrays = "90137ffa-7385-5640-81b9-e52037218182"

[compat]
DocStringExtensions = "^0.9"
QEDbase = "0.2.2"
Reexport = "^1.2"
SimpleTraits = "^0.9"
StaticArrays = "^1.9"
julia = "1.6"
```

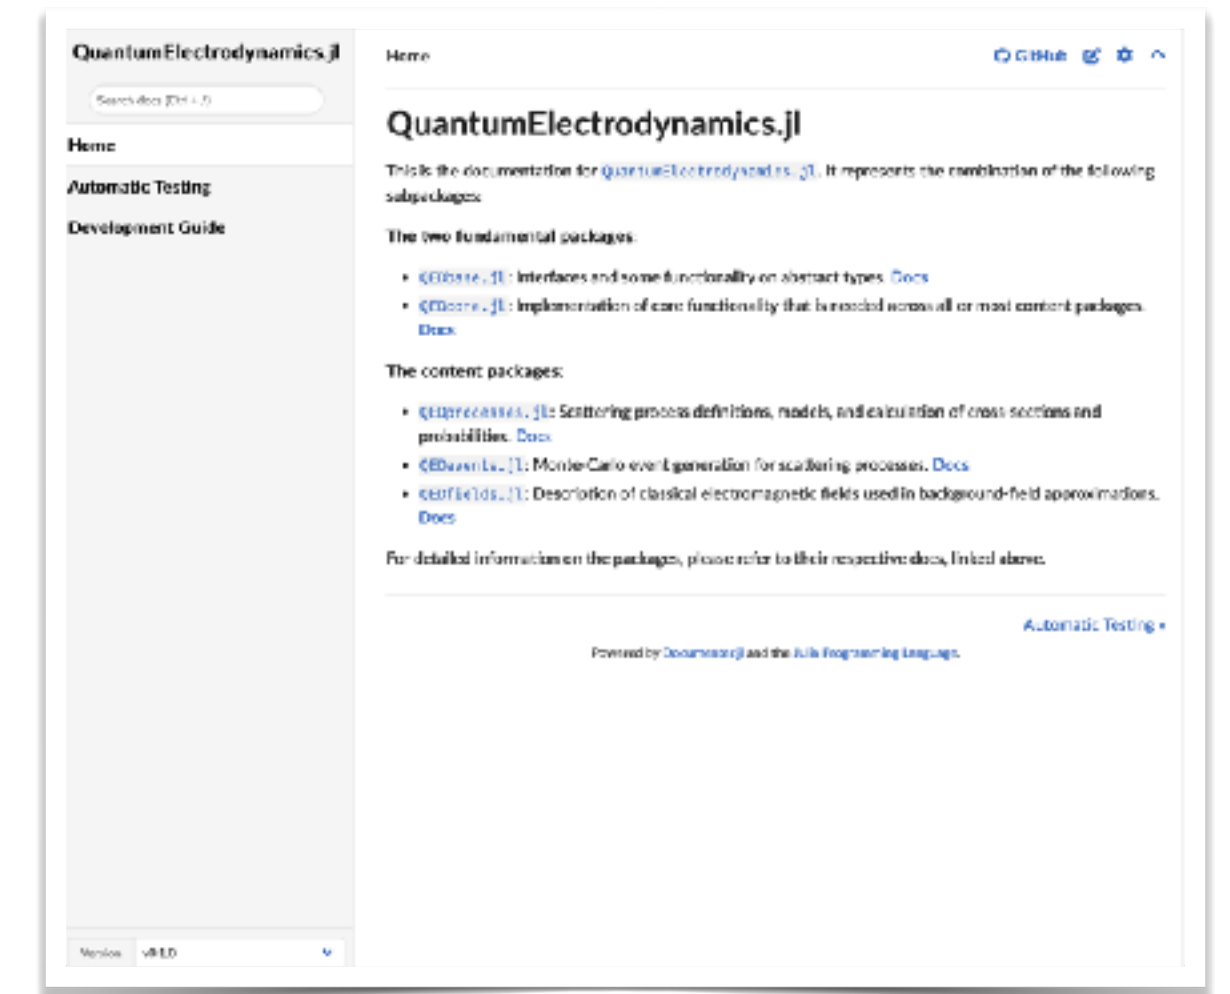
Project.toml

```
(@v1.11) pkg> add QEDcore
Resolving package versions...
Installed QEDcore - v0.1.1
Updating `~/julia/environments/v1.11/Project.toml`
 [35dc0263] + QEDcore v0.1.1
Updating `~/julia/environments/v1.11/Manifest.toml`
 [7d9f7c33] + Accessors v0.1.38
 [dce04be8] + ArgCheck v2.3.0
 [49dc2e85] + Calculus v0.5.1
 [38540f10] + CommonSolve v0.2.4
 [a33af91c] + CompositionsBase v0.1.2
 [187b0558] + ConstructionBase v1.5.8
 [3587e190] + InverseFunctions v0.1.17
 [eff96d63] + Measurements v2.11.0
 [5ad8b20f] + PhysicalConstants v0.2.3
 [10e22c08] + QEDbase v0.2.2
 [35dc0263] + QEDcore v0.1.1
 [f2b01f46] + Roots v2.2.1
 [699a6c99] + SimpleTraits v0.9.4
 [90137ffa] + StaticArrays v1.9.7
 [1e83bf80] + StaticArraysCore v1.4.3
Precompiling project...
4 dependencies successfully precompiled in 6 seconds.
```

Package manager (Pkg.jl)

```
Testing Running tests...
Test Summary: | Pass Total Time
phase spaces | 152 152 3.1s
Test Summary: | Pass Total Time
four momentum | 400 400 2.5s
Test Summary: | Pass Total Time
gamma matrices | 92 92 1.4s
Test Summary: | Pass Total Time
Lorentz vector | 69 69 1.4s
Test Summary: | Pass Total Time
Dirac tensors | 51 51 1.5s
Test Summary: | Pass Total Time
particle types | 35 35 0.1s
Test Summary: | Pass Total Time
particle states | 4367 4367 0.9s
Test Summary: | Pass Total Time
particle spinors | 84 84 0.7s
Test Summary: | Pass Total Time
particle base states | 4367 4367 0.4s
Test Summary: | Pass Total Time
particle propagators | 3 3 0.2s
Test Summary: | Pass Total Time
process interface | 148 148 1.3s
Testing QEDcore tests passed
```

Testing (integrates with Pkg.jl)

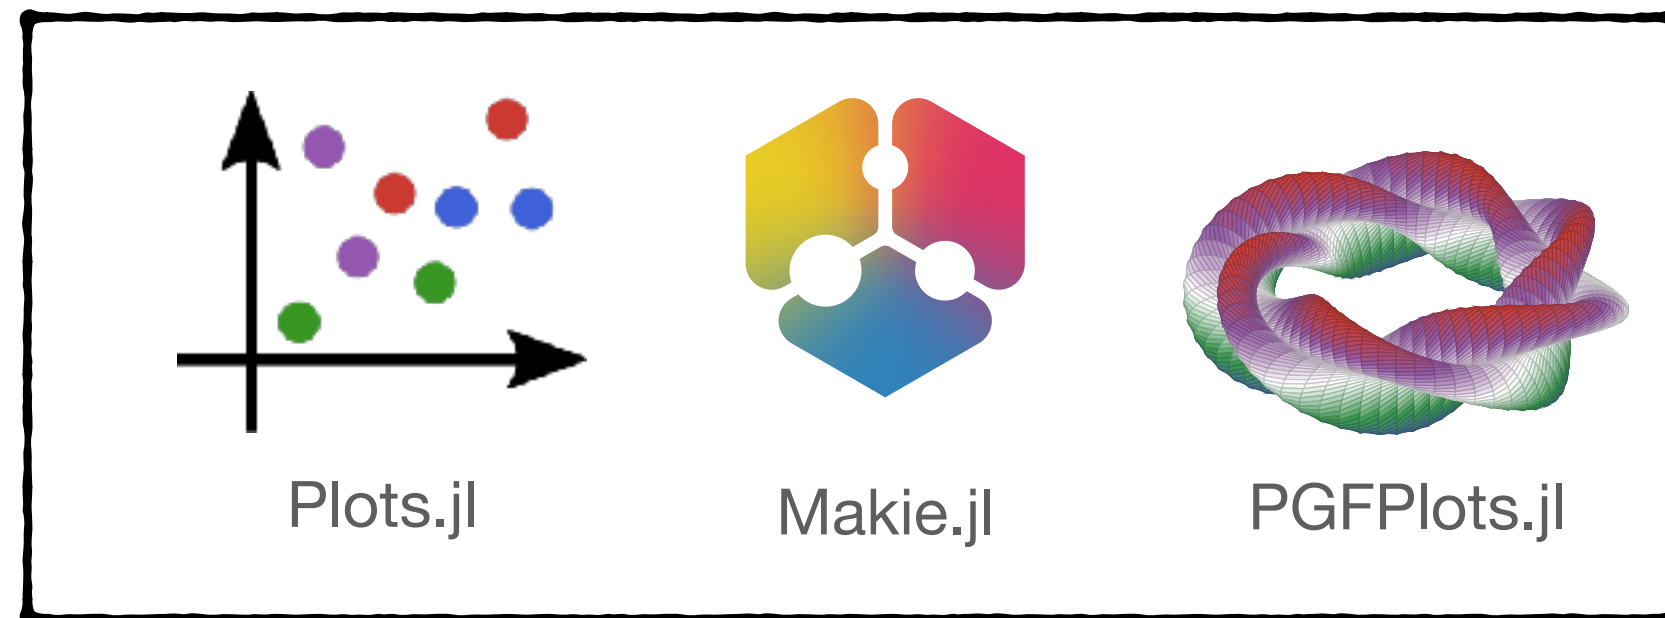


Docmenter.jl

Rich eco-system

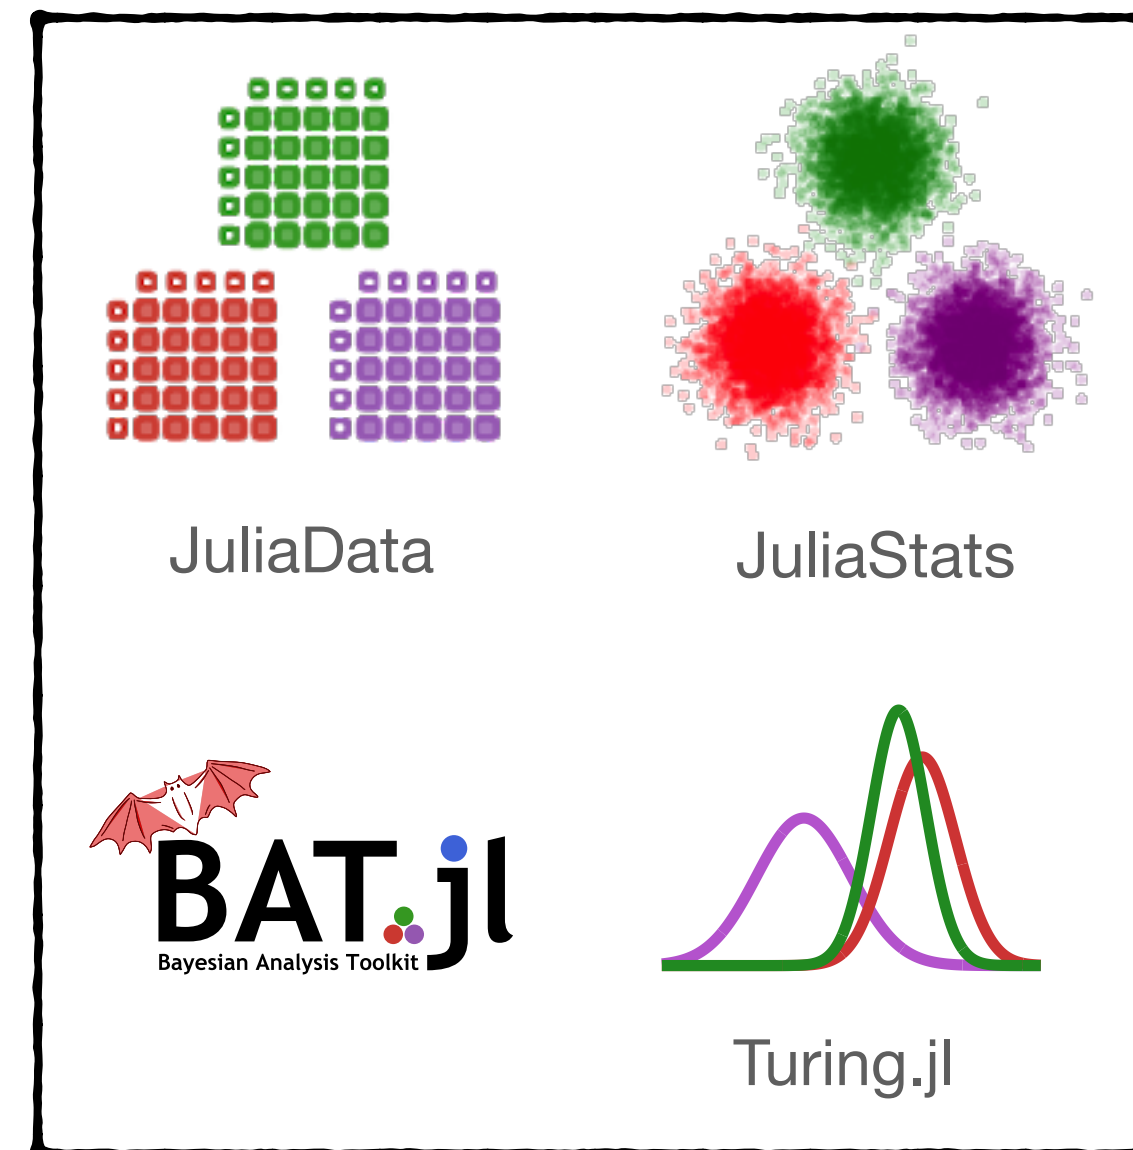
>10k packages

Visualization



Plots.jl Makie.jl PGFPlots.jl

Data and Statistics

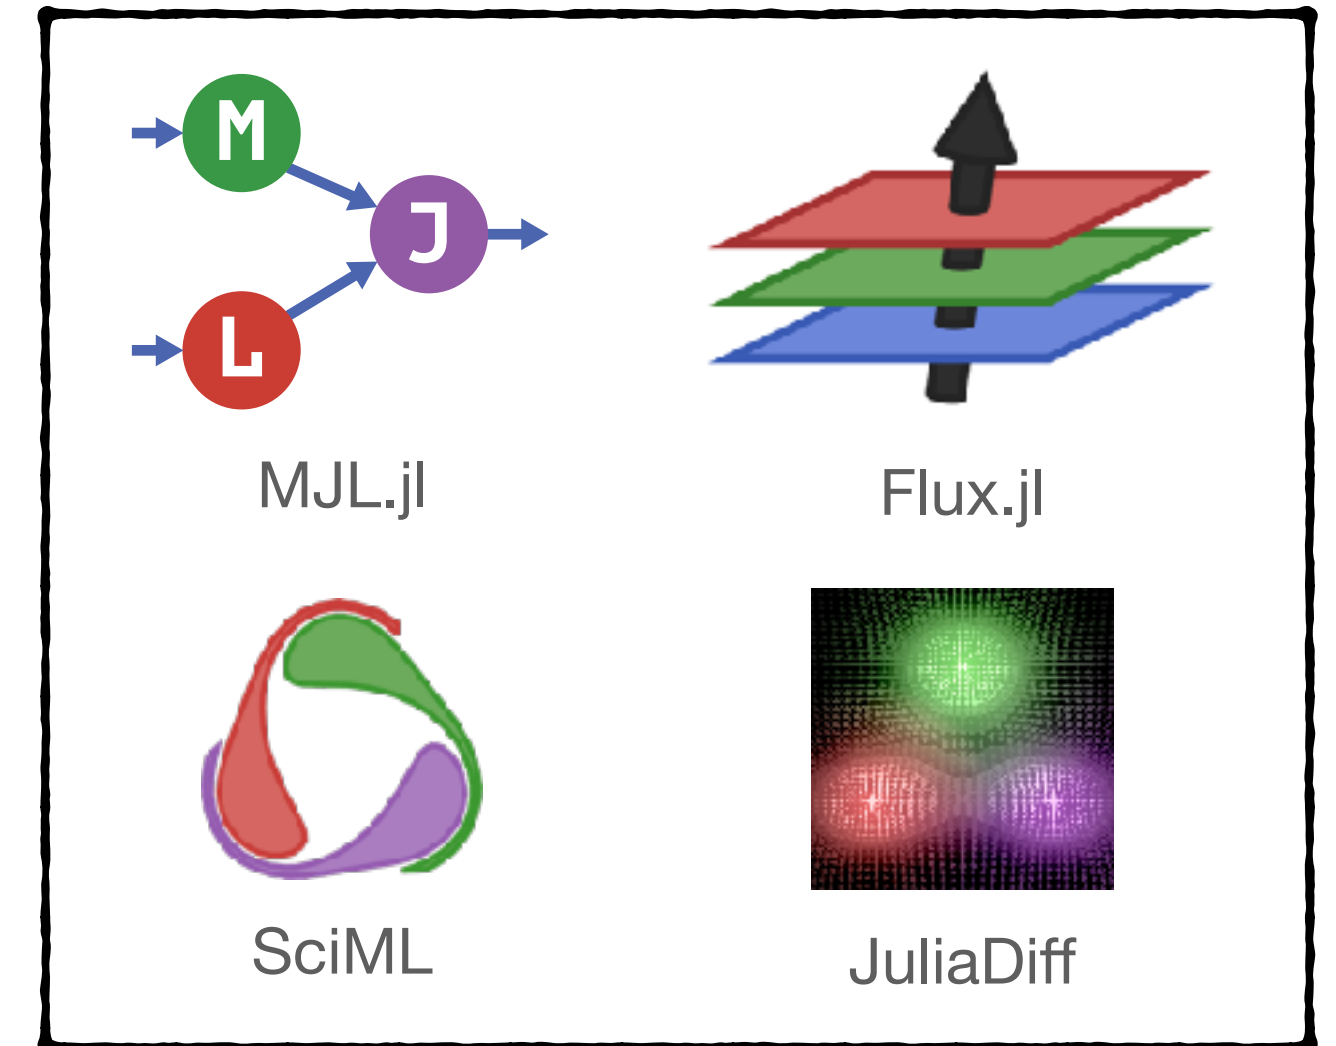


JuliaData JuliaStats

BAT.jl
Bayesian Analysis Toolkit

Turing.jl

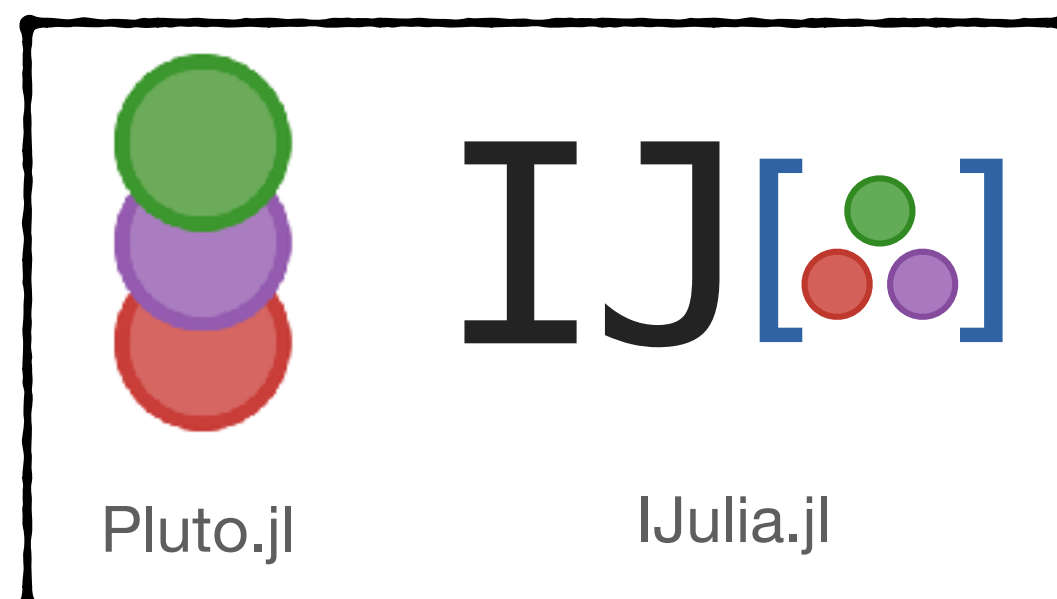
Machine learning



MJL.jl Flux.jl

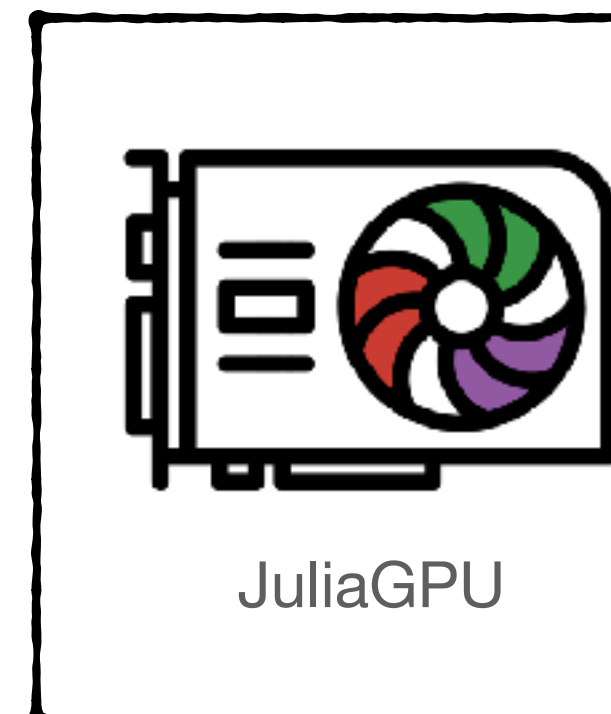
SciML JuliaDiff

Notebooks



Pluto.jl IJulia.jl

GPU support

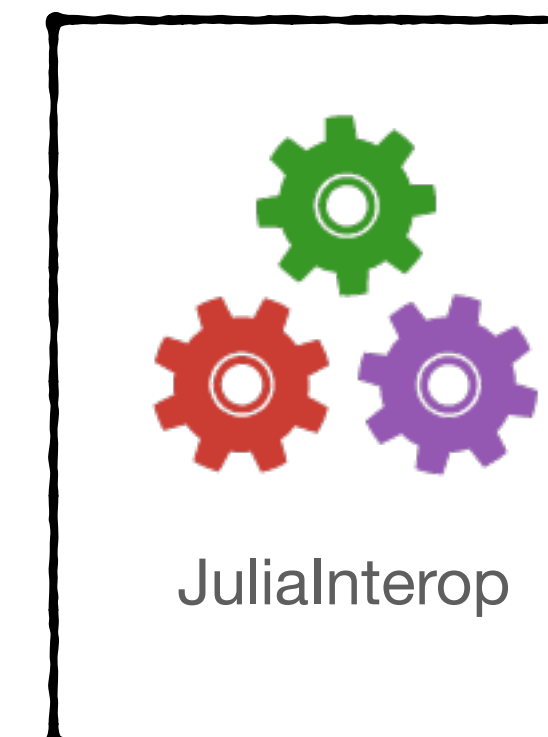


JuliaGPU

CUDA.jl
AMDGPU.jl
oneAPI.jl
Metal.jl

KernelAbstractions.jl

Interoperability



JuliaInterop

CxxWrap.jl
PyCall.jl
RCall.jl
MathLink.jl

Loading data

HEP data formats

UnROOT: an I/O library for the CERN ROOT file format written in Julia

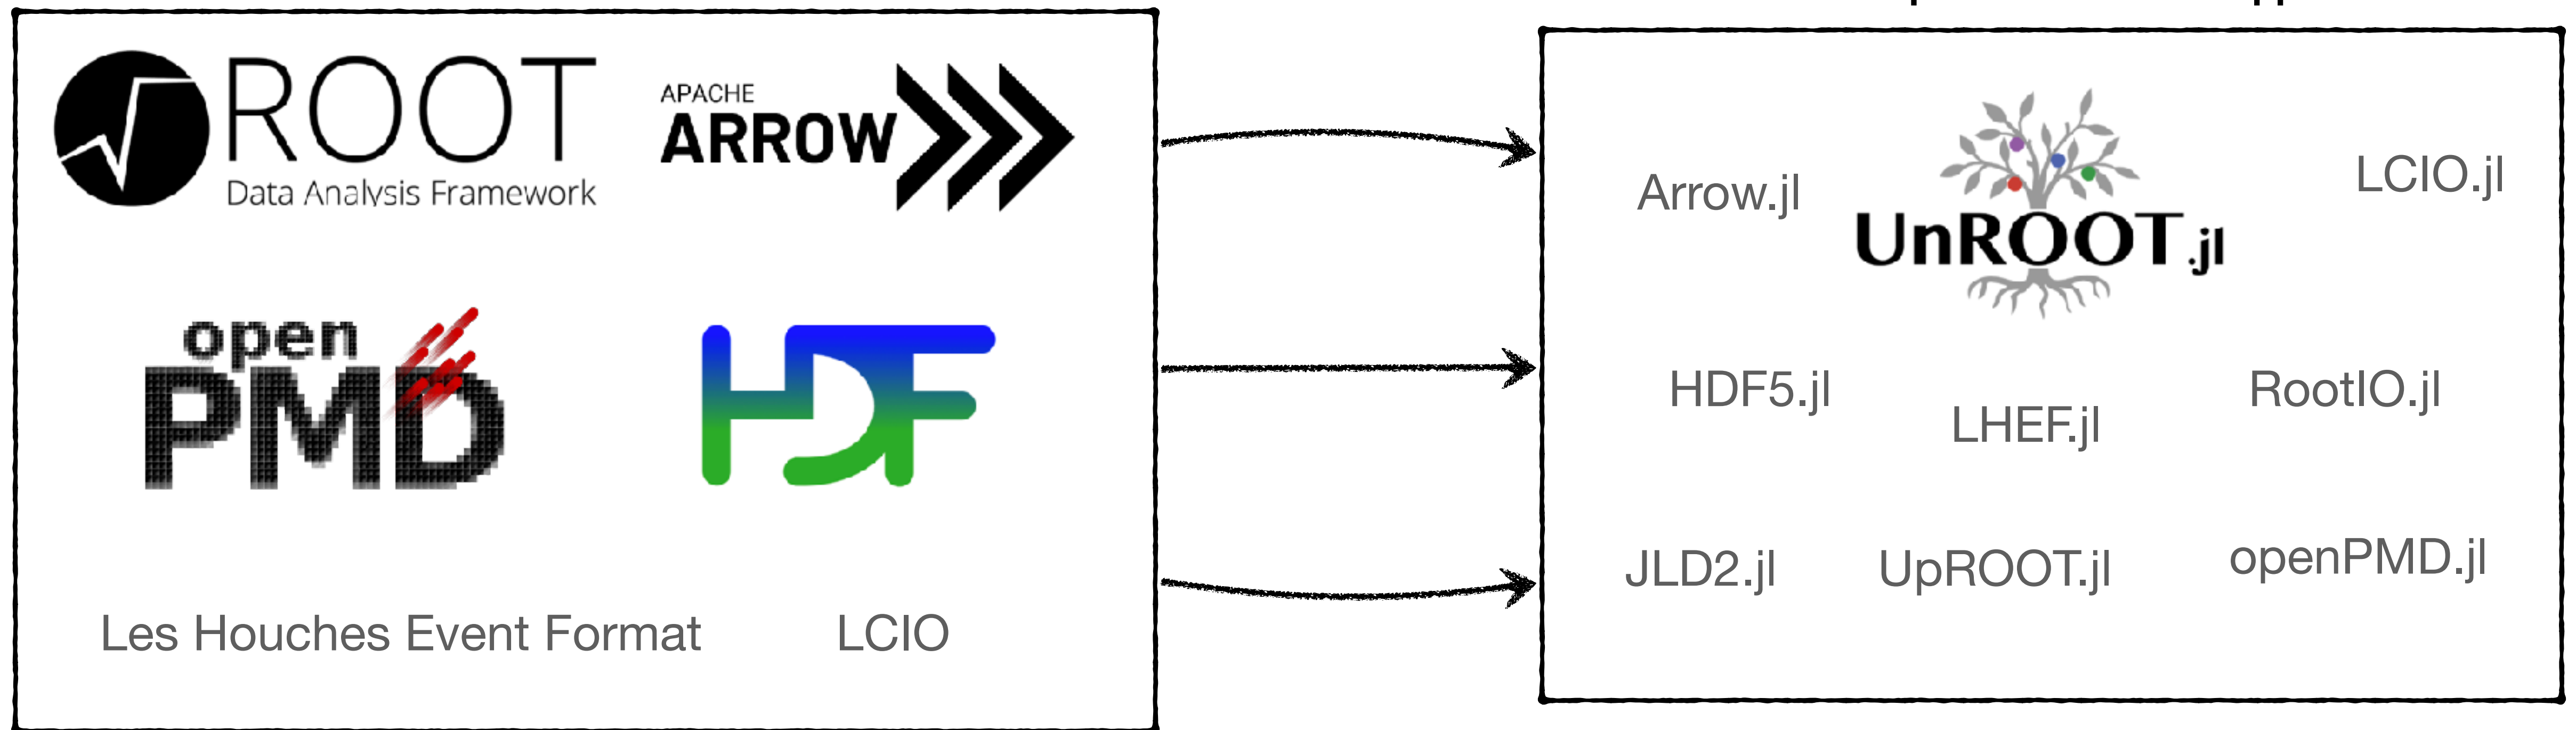
Tamás Gál^{1,2}, Jerry (Jiahong) Ling³, and Nick Amin⁴

High-performance end-user analysis in pure Julia programming language

Jerry Ling^{1,*} and Tamás Gál^{2,**}

File formats/standards

Julia implementations/wrapper



Drawbacks of using Julia?

Julia should be better or shouldn't it?

- Formatter/Linter/LSP could be better
- Little scripts*
- Startup time*
- Vendor lock
 - Only LLVM and Clang
 - Only one reference implementation
- Building binaries*
- Calling Juila from other Languages*
- Context-based programming*
- Cumbersome static performance prediction
- Cumbersome static analysis/checking*

*solved (kinda)

Julia in the wild

Data throughput

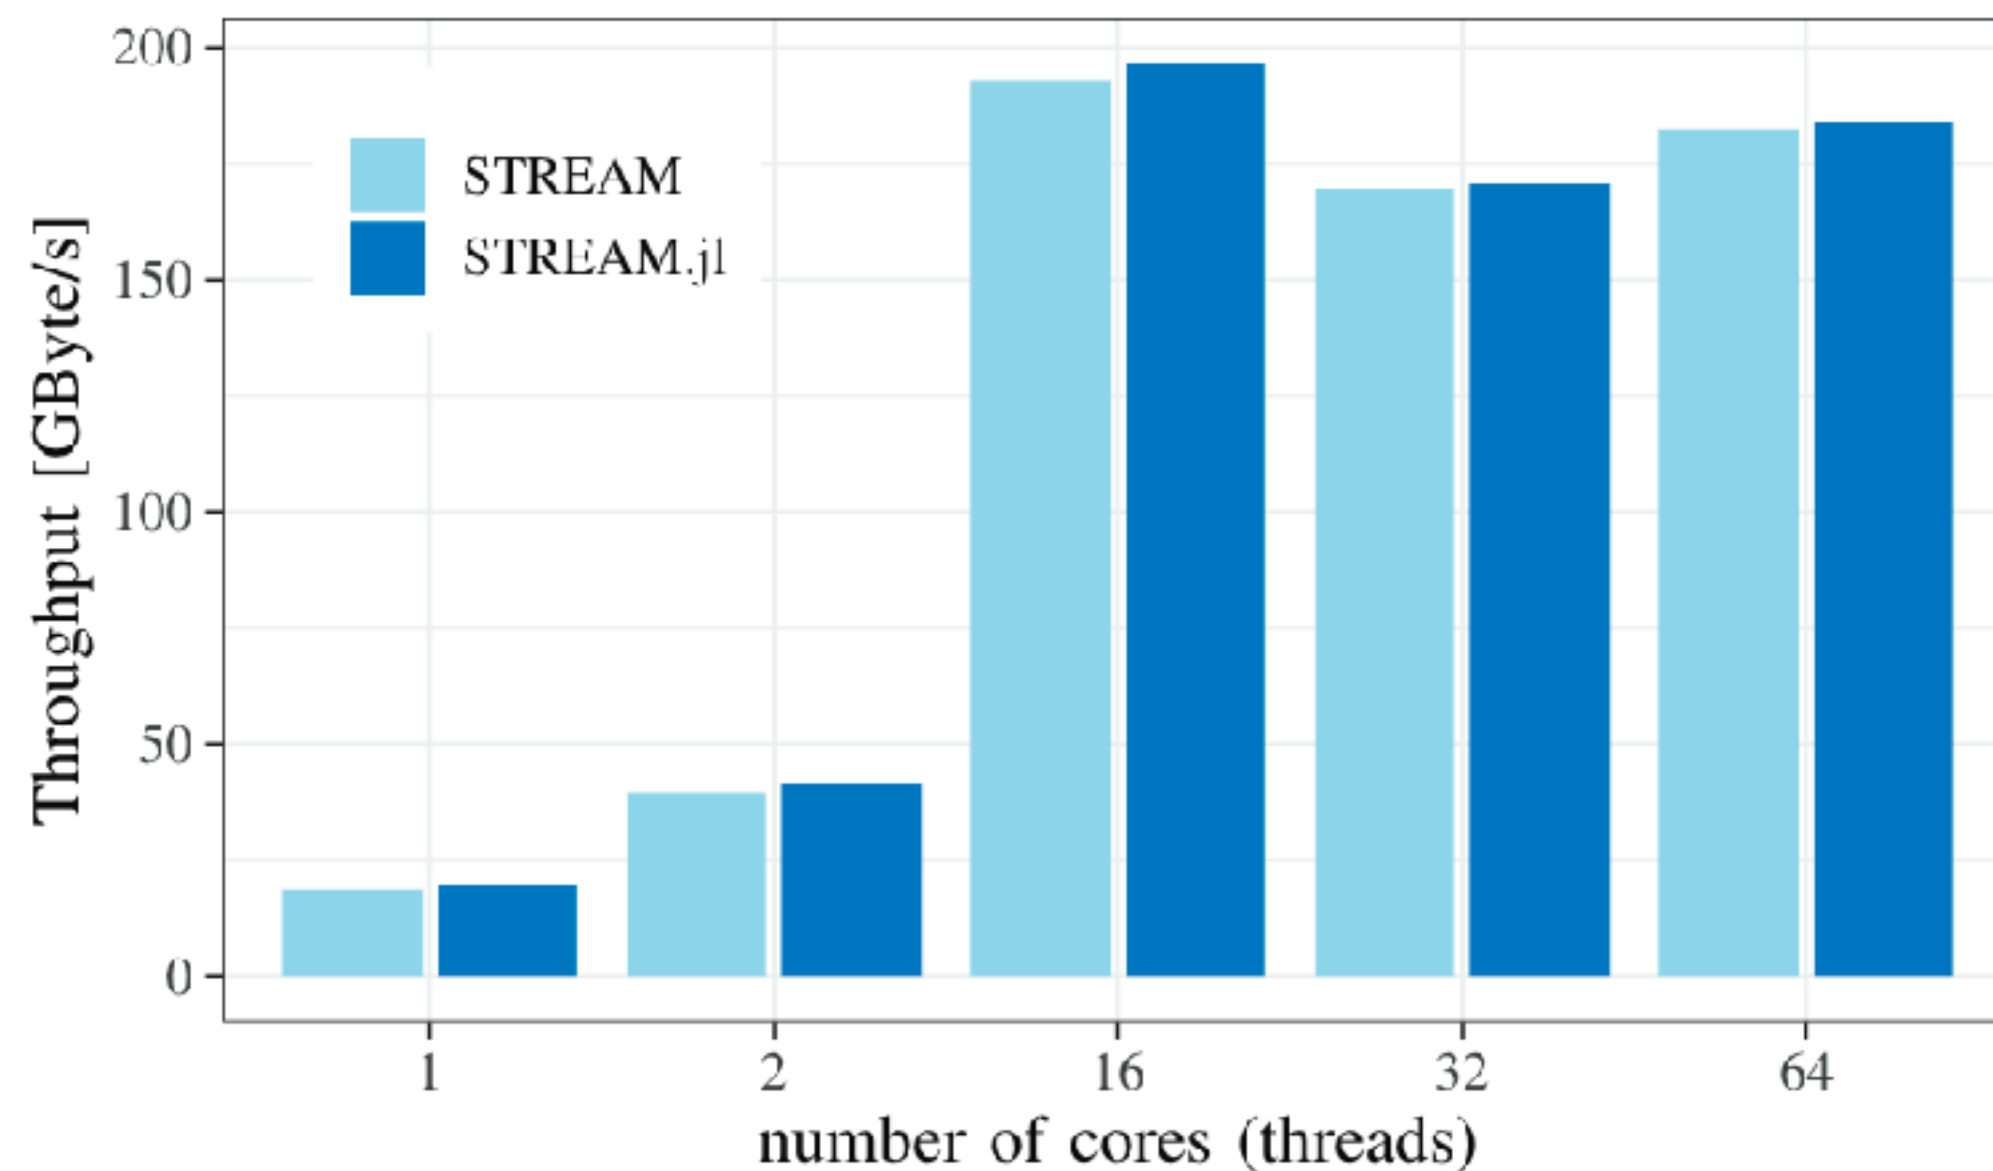
Memory bandwidth benchmarks

Benchmarking Julia's Communication Performance:
Is Julia HPC ready or Full HPC?

Sascha Hunold
TU Wien, Faculty of Informatics
Vienna, Austria

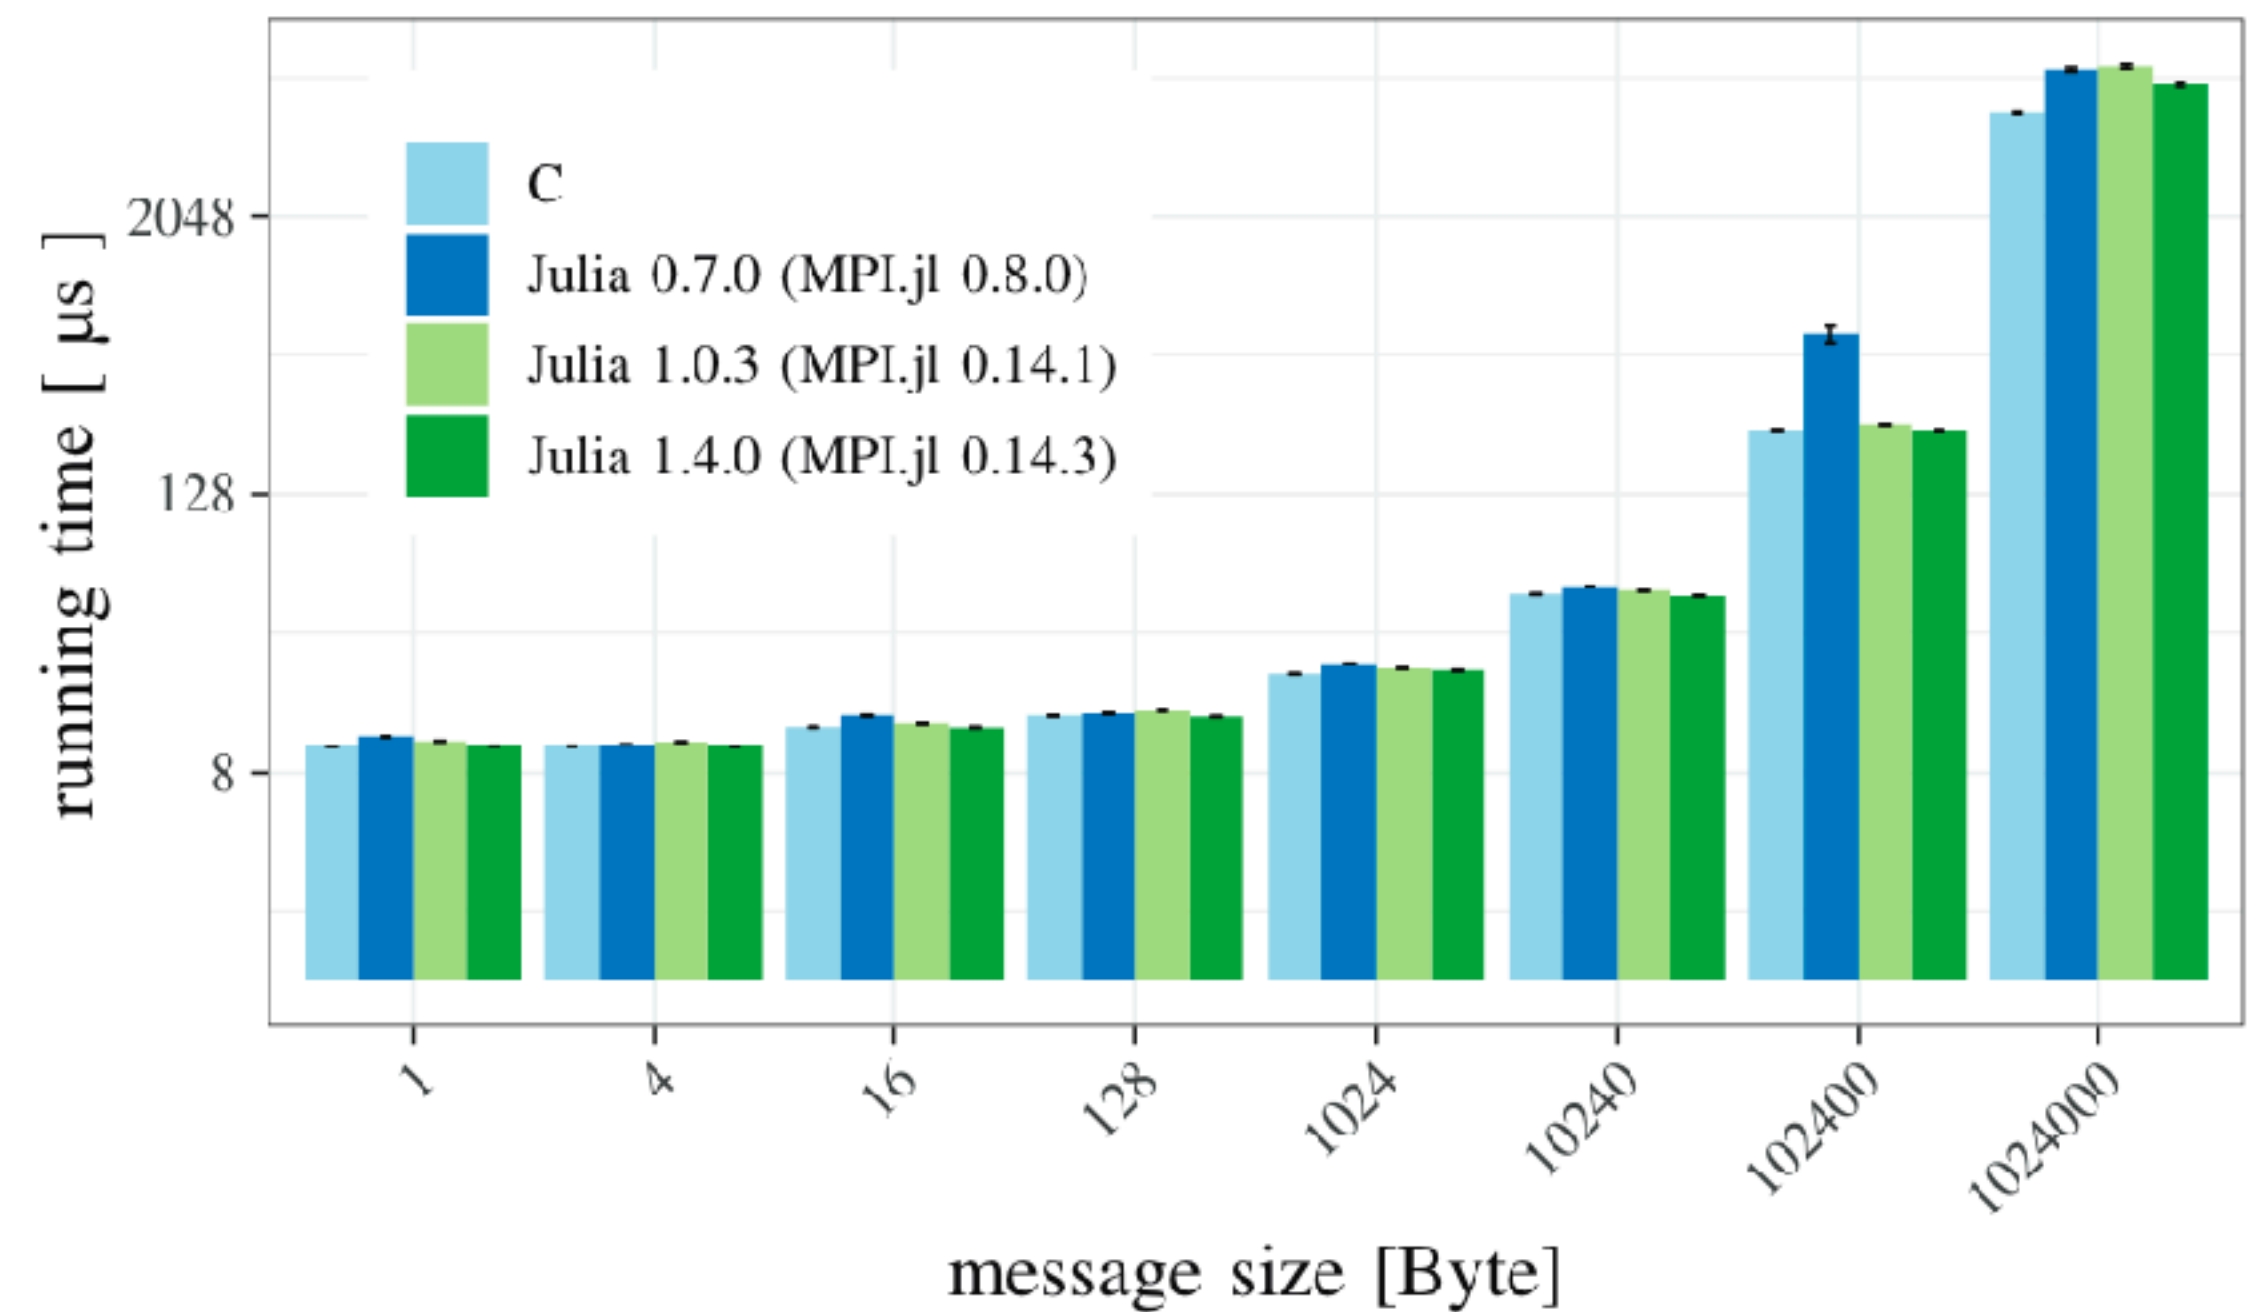
Sebastian Steiner
TU Wien, Faculty of Informatics
Vienna, Austria

Intra-node performance



STREAM benchmark up to 64 AMD CPU cores
LoC: 378 (C) vs 156 (Julia)

Inter-node performance



MPI broadcasting benchmark: 36 × 32 processes

Single-node performance

Single-thread axpy benchmarks on Fugaku (A64FX)

Productivity meets Performance: Julia on A64FX

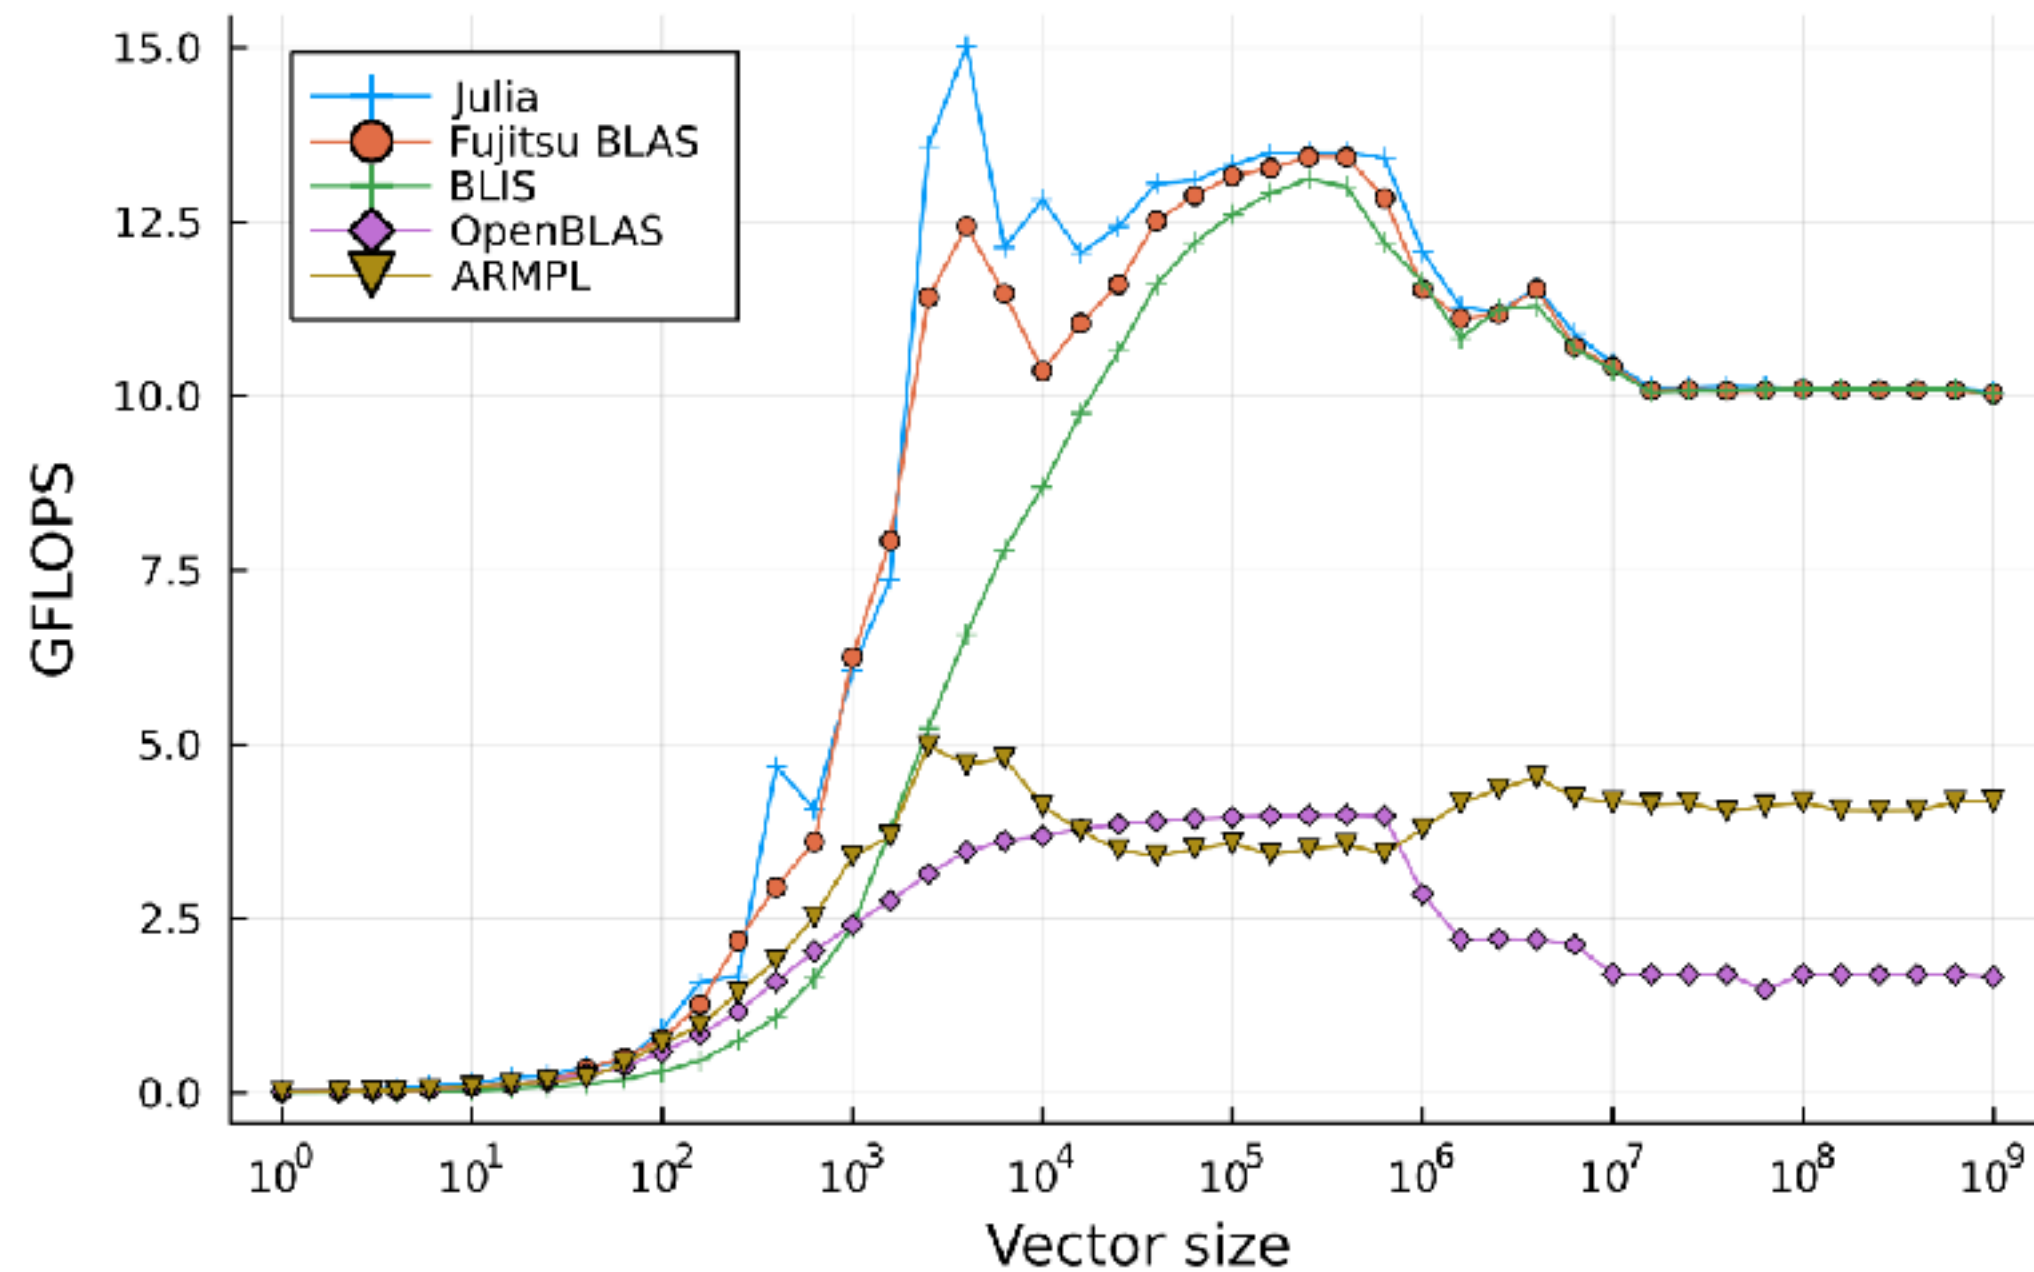
1st Mosè Giordano
Advanced Research Computing
UCL
London, United Kingdom
m.giordano@ucl.ac.uk

2nd Milan Klöwer
Atmospheric, Oceanic and Planetary Physics
University of Oxford
Oxford, United Kingdom
milan.kloewer@physics.ox.ac.uk

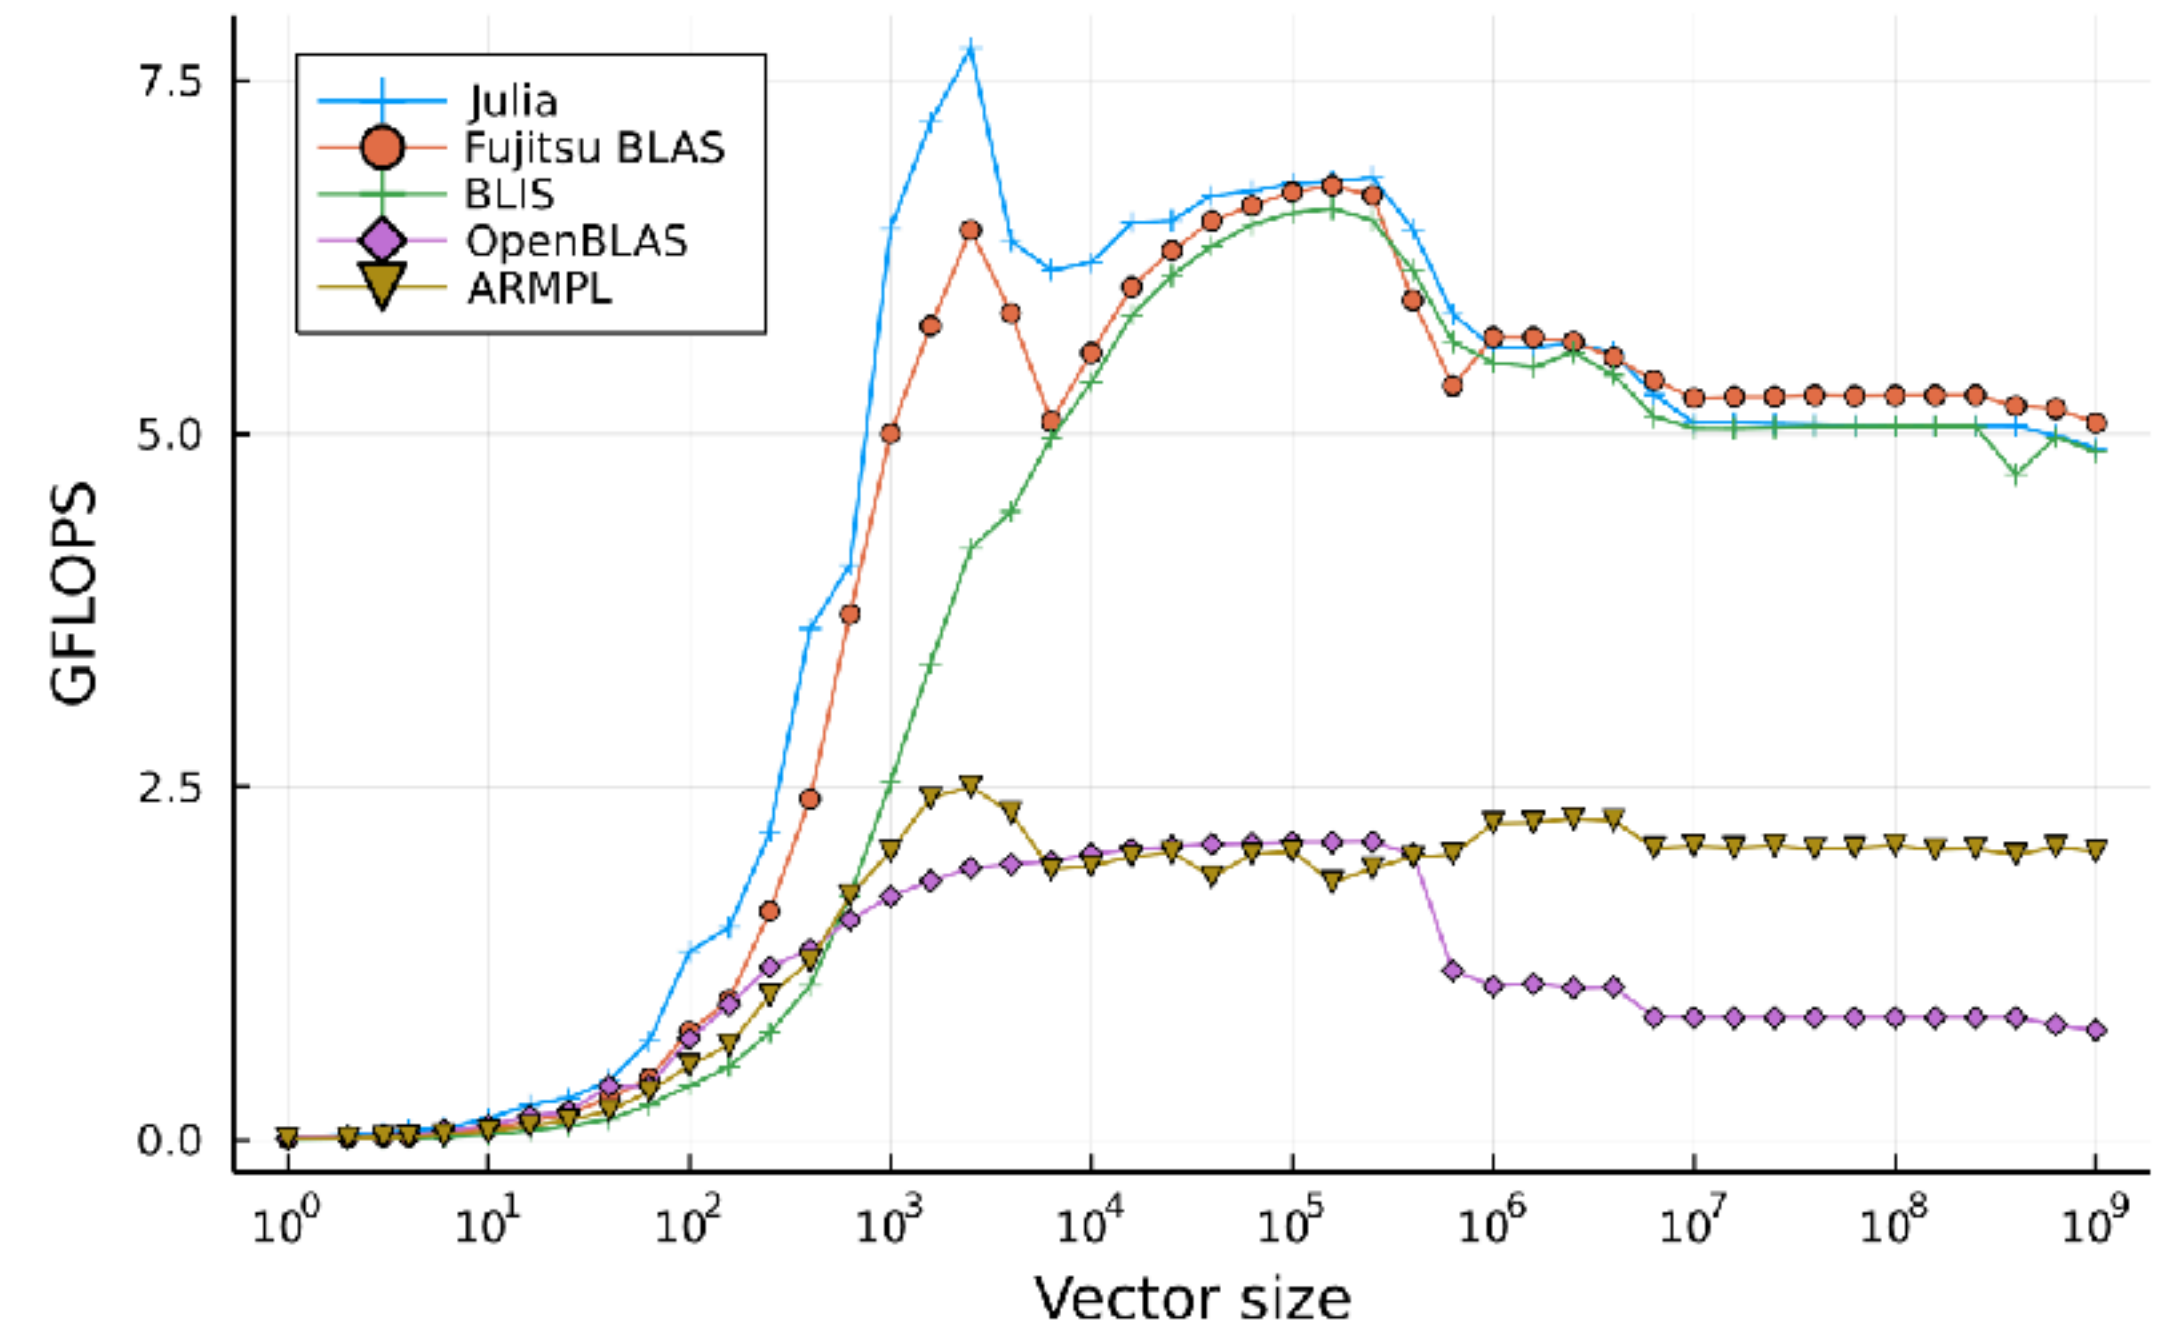
3rd Valentin Churavy
CSAIL, EECS
Massachusetts Institute of Technology
Cambridge, United States of America
vchuravy@mit.edu

```
function axpy!(a::T, x::Vector{T}, y::Vector{T}) where {T<:Number}
    @simd for i in eachindex(x, y)
        @inbounds y[i] = muladd(a, x[i], y[i])
    end
    return y
end
```

Single precision



Double precision




Julia on scale

Celeste.jl project

Cataloging the Visible Universe through Bayesian Inference at Petascale

Jeffrey Regier^{*}, Kiran Pamnany[†], Keno Fischer[‡], Andreas Noack[§], Maximilian Lam^{*}, Jarrett Revels[§],
Steve Howard[¶], Ryan Giordano[¶], David Schlegel^{||}, Jon McAuliffe[¶], Rollin Thomas^{||}, Prabhat^{||}

- 2017 at NERSC (Berkeley)
 - Analysis of 178 TB telescope data
 - Inferred parameters of 1.88×10^8 stars
 - Done in 14.6 min
 - 1.3×10^6 threads on 650,000 Intel Xeon Phi cores
 - 1.54 PFLOPS peak performance



Community efforts

JuliaHEP working group in the HEP software foundation

HEP software foundation

- Founded around 2014
- facilitate coordination and common efforts in HEP software and computing
- Objectives
 - Share expertise
 - Raise awareness
 - Catalyse new common projects
 - Promote collaborations in new developments
 - Provide training
 - ...



HEP computing collaborations for the challenges of the next decade

Contacts: Simone Campana (Simone.Campana@cern.ch), Zach Marshall (ZLMarshall@lbl.gov), Alessandro Di Girolamo (Alessandro.Di.Girolamo@cern.ch), Heidi Schellman (Heidi.Schellman@cern.ch), Graeme A. Stewart (Graeme.A.Stewart@cern.ch)

A Roadmap for HEP Software and Computing R&D for the 2020s


The HEP Software Foundation⁵ · Johannes Albrecht⁶⁹ · Antonio Augusto Alves Jr⁸¹ · Guilherme Amadio⁵ · Giuseppe Andronico²⁷ · Nguyen Anh-Ky¹²² · Laurent Aphecetche⁶⁶ · John Apostolakis⁵ · Makoto Asai⁶³ · Luca Atzori⁵ · Marian Babik⁵ · Giuseppe Bagliesi³² · Marilena Bandieramonte⁵ · Sunanda Banerjee¹⁶ · Martin Barisits⁵ · Lothar A. T. Bauerdick¹⁶ · Stefano Belforte³⁵ · Douglas Benjamin⁸² · Catrin Bernius⁶³ · Wahid Bhimji⁴⁶ · Riccardo Maria Bianchi¹⁰⁵ · Ian Bird⁵ · Catherine Biscarat⁵² · Jakob Blomer⁵ · Kenneth Bloom⁹⁷ · Tommaso Boccali³² · Concezio Bozzi²⁸ · Ma

Challenges in Monte Carlo Event Generator Software for High-Luminosity LHC

The HSF Physics Event Generator WG · Andrea Valassi¹ · Efe Yazgan² · Josh McFayden^{1,3,4} · Simone Amoroso⁵ · Joshua Bendavid¹ · Andy Buckley⁶ · Matteo Cacciari^{7,8} · Taylor Childers⁹ · Vitaliano Ciulli¹⁰ · Rikkert Frederix¹¹ · Stefano Frixione¹² · Francesco Giuliani¹³ · Alexander Grohsjean⁵ · Christian Gütschow¹⁴ · Stefan Höche¹⁵ · Walter Hopkins⁹ · Philip Ilten^{16,17} · Dmitri Konstantinov¹⁸ · Frank Krauss¹⁹ · Qiang Li²⁰ · Leif Lönnblad¹¹ · Fabio Maltoni^{21,22} · Michelangelo Mangano¹ · Zach Marshall³ · Olivier Mattelaer²² · Javier Fernandez Menendez²³ · Stephen Mrenna¹⁵ · Servesh Muralidharan^{1,9} · Tobias Neumann^{14,24} · Simon Plätzer²⁵ · Stefan Prestel¹¹ · Stefan Roiser¹ · Marek Schönherr¹⁹ · Holger Schulz¹⁷ · Markus Schulz¹ · Elizabeth Sexton-Kennedy¹⁵ · Frank Siegert²⁶ · Andrzej Siódmok²⁷ · Graeme A. Stewart¹


JuliaHEP @ HSF

- HSF working group founded in 2022
- JuliaHEP annual workshop
 - 2023: ECAP in Erlangen
 - 2024: CERN
- Monthly community calls
- Monitoring/Supporting development
- Tutorial material + example project



Potential of the Julia Programming Language for High Energy Physics Computing

Jonas Eschle¹  · Tamás Gál²  · Mosè Giordano³  · Philippe Gras⁴  · Benedikt Hegner⁵ · Lukas Heinrich⁶  · Uwe Hernandez Acosta^{7,8}  · Stefan Kluth⁶  · Jerry Ling⁹  · Pere Mato⁵  · Mikhail Mikhasenko^{10,11}  · Alexander Moreno Briceño¹²  · Jim Pivarski¹³  · Konstantinos Samaras-Tsakiris⁵  · Oliver Schulz⁶  · Graeme Andrew Stewart⁵  · Jan Strube^{14,15}  · Vassil Vassilev¹³



github.com/JuliaHEP



#hep

Next JuliaHEP workshop

July 28 - 31, 2025

Princeton University



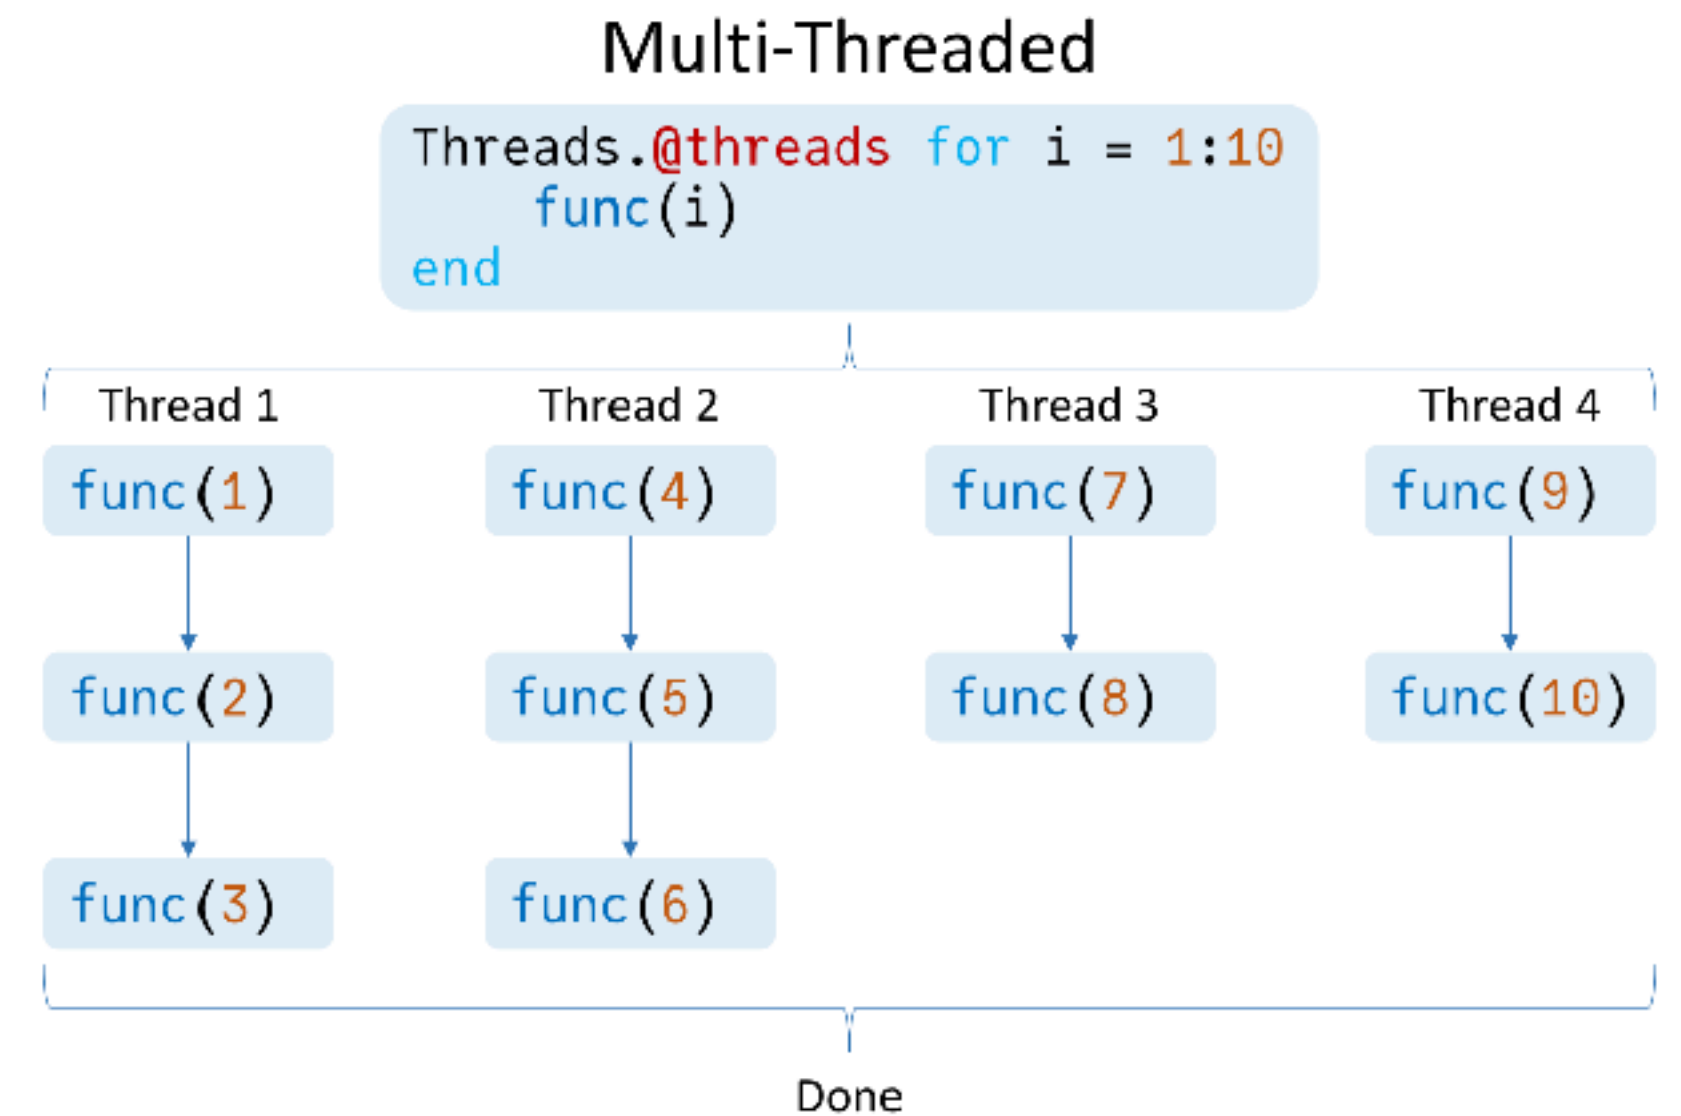

Abstract submissions are now open!

Backup

Parallel computing





Native Threading support

- Support for OpenMP-like models
 - Parallelization of loops
- Support for M:N threading
 - M user threads are mapped onto N kernel threads
- Support for task migration
 - Tasks can be started, suspended, and resumed again

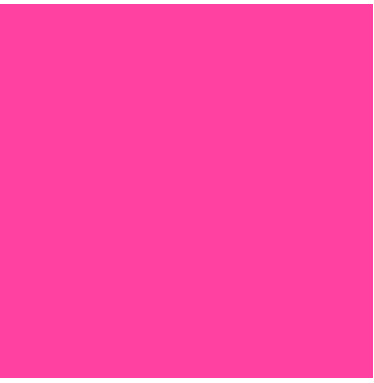

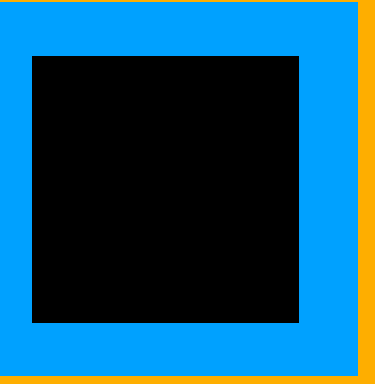






Multiple dispatch

Function and methods

	<code>f(::Any, ::Number)</code>
	<code>f(::T, ::T) where {T<:Number}</code>
	<code>f(::Int64, ::Int64)</code>
	<code>f(::String, ::Any)</code>

Float64<:AbstractFloat<:Real<:Number<:Any

	String	Int64	Float64
String			
Int64			
Float64			

Multiple dispatch II

Expressiveness

Dispatch degree	Syntax	Dispatched on	Selection power
None	$f(x, y, z)$	$\{ \}$	1
Single	$x.f(y, z)$	$\{x\}$	$ X $
Multiple	$f(x::X, y::Y, z::Z)$	$\{x, y, z\}$	$ X \cdot Y \cdot Z $

Multiple dispatch III

Unreasonable effectiveness

- Allows generic code based on abstract types
- Allows arbitrary optimization
- Orthogonal development
- Solves the expression problem

```
using DifferentialEquations, Plots

g = 9.79          # Gravitational constants
L = 1.00         # Length of the pendulum


#Initial Conditions
u₀ = [0, π / 60] # Initial speed and initial angle
tspan = (0.0, 6.3)

#Define the problem
function pendulum(du,u,p,t)
    θ = u[1]
    dθ = u[2]
    du[1] = dθ
    du[2] = -(g/L)*θ
end

#Pass to solvers
prob = ODEProblem(pendulum, u₀, tspan)
sol = solve(prob, Tsit5(), reltol = 1e-6)

# Analytic solution
u = u₀[2] .* cos.(sqrt(g / L) .* sol.t)

plot(sol.t, getindex.(sol.u, 2), label = "Numerical")
plot!(sol.t, u, label = "Analytic")
```



Multiple dispatch III

Unreasonable effectiveness

- Allows generic code based on abstract types
- Allows arbitrary optimization
- Orthogonal development
- Solves the expression problem

```
using DifferentialEquations, Measurements, Plots

g = 9.79 ± 0.02; # Gravitational constants
L = 1.00 ± 0.01; # Length of the pendulum


#Initial Conditions
u₀ = [0 ± 0, π / 60 ± 0.01] # Initial speed and initial angle
tspan = (0.0, 6.3)

#Define the problem
function pendulum(du,u,p,t)
    θ = u[1]
    dθ = u[2]
    du[1] = dθ
    du[2] = -(g/L)*θ
end

#Pass to solvers
prob = ODEProblem(pendulum, u₀, tspan)
sol = solve(prob, Tsit5(), reltol = 1e-6)

# Analytic solution
u = u₀[2] .* cos(sqrt(g / L) .* sol.t)

plot(sol.t, getindex.(sol.u, 2), label = "Numerical")
plot!(sol.t, u, label = "Analytic")
```

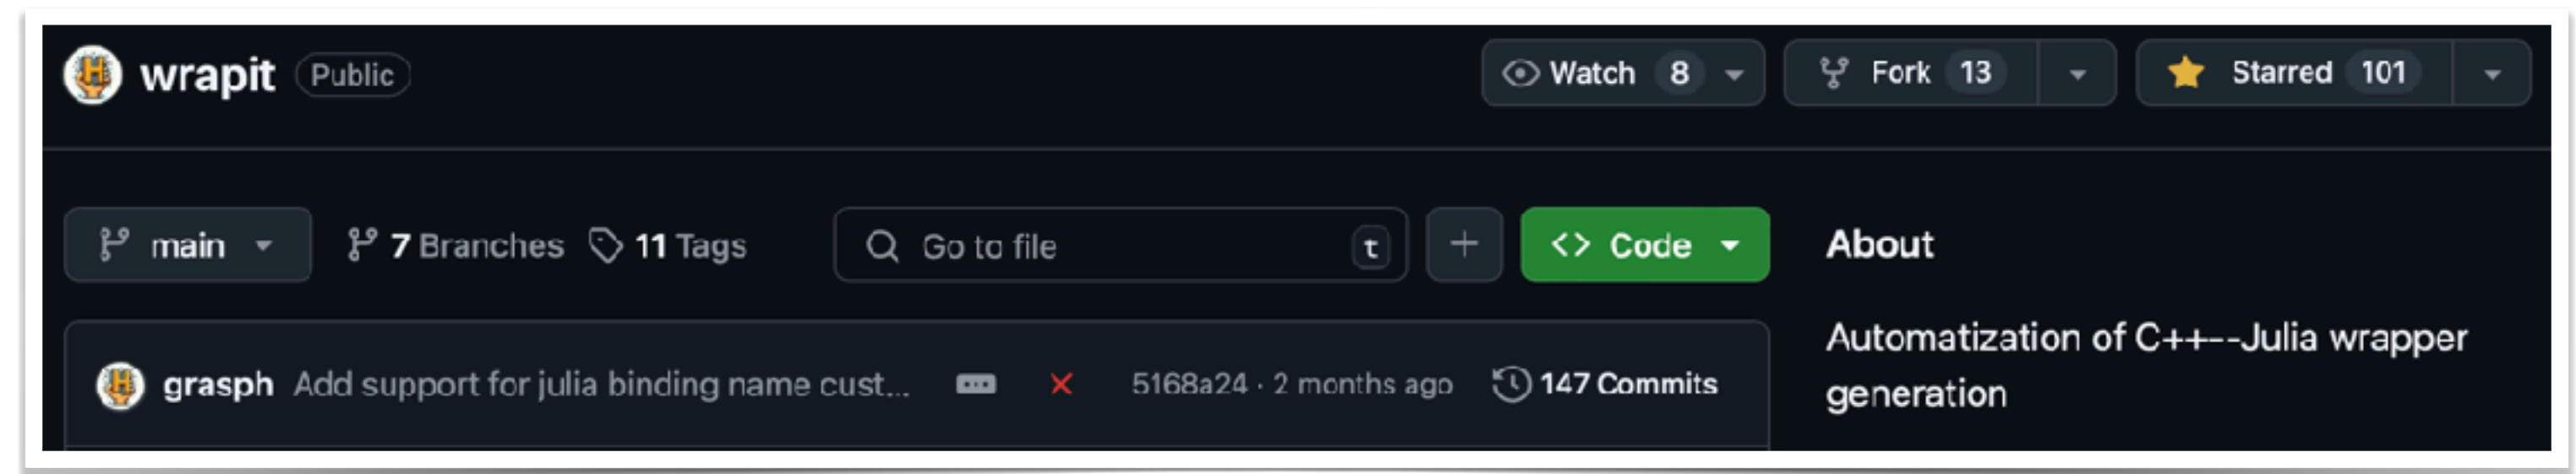



Interoperability and Legacy code

Everything is wrapped

- Use foreign code from Julia
- Wrapit and CxxWrap.jl for (semi-) automatic building of bindings
- non-exhausted list of wrapped libraries
 - Geant4.jl
 - ROOT.jl
 - XRootD.jl
 - Pythia8.jl
 - FastJet.jl
 - UpROOT.jl
 - Etc.

Interoperability



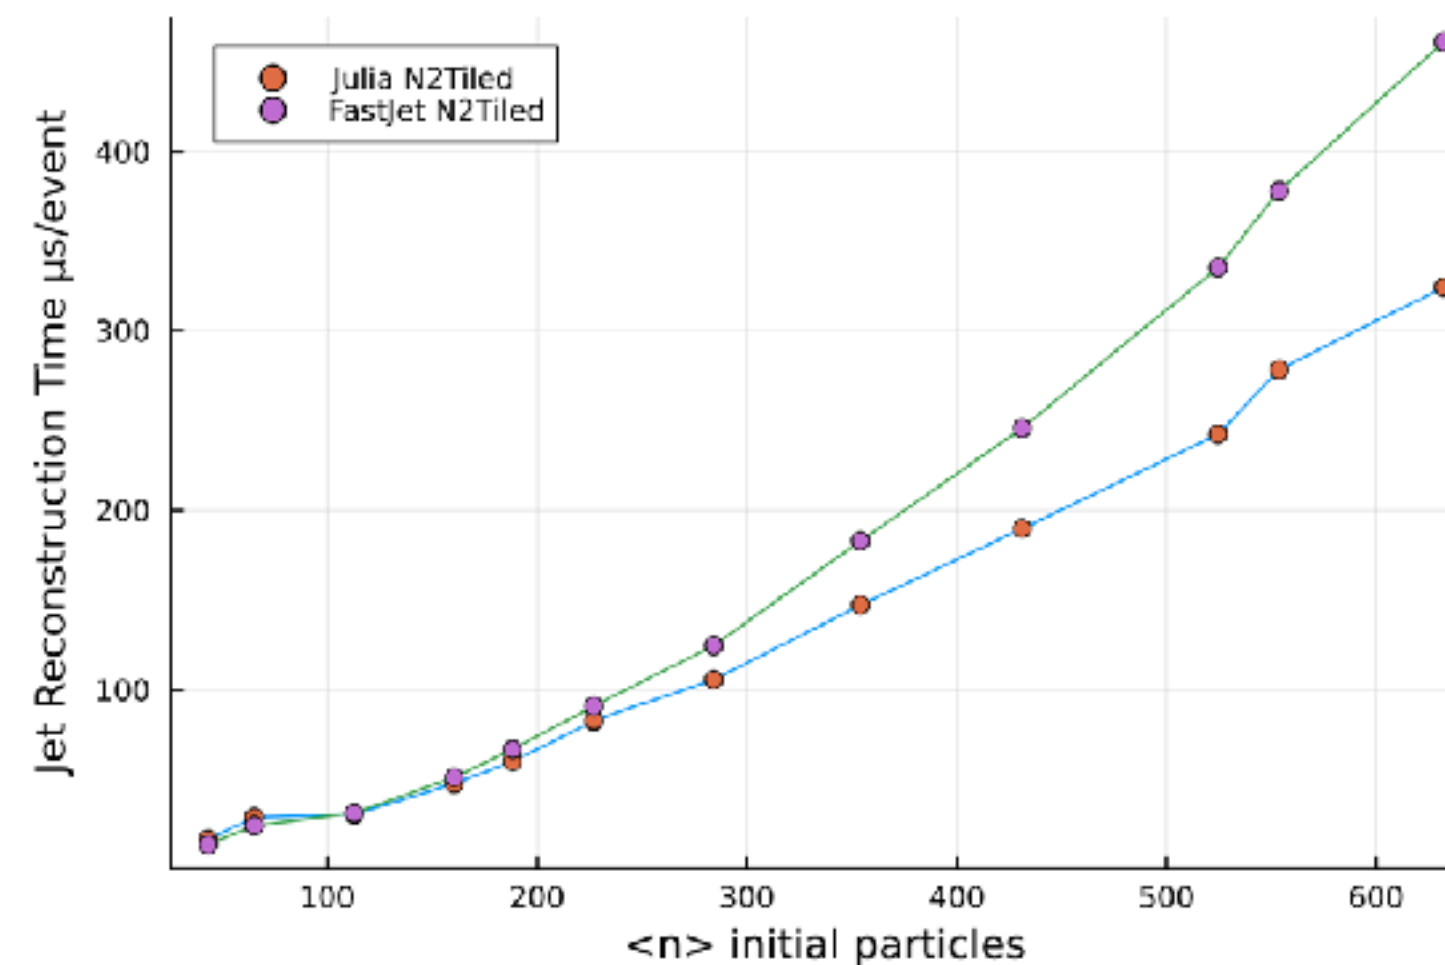
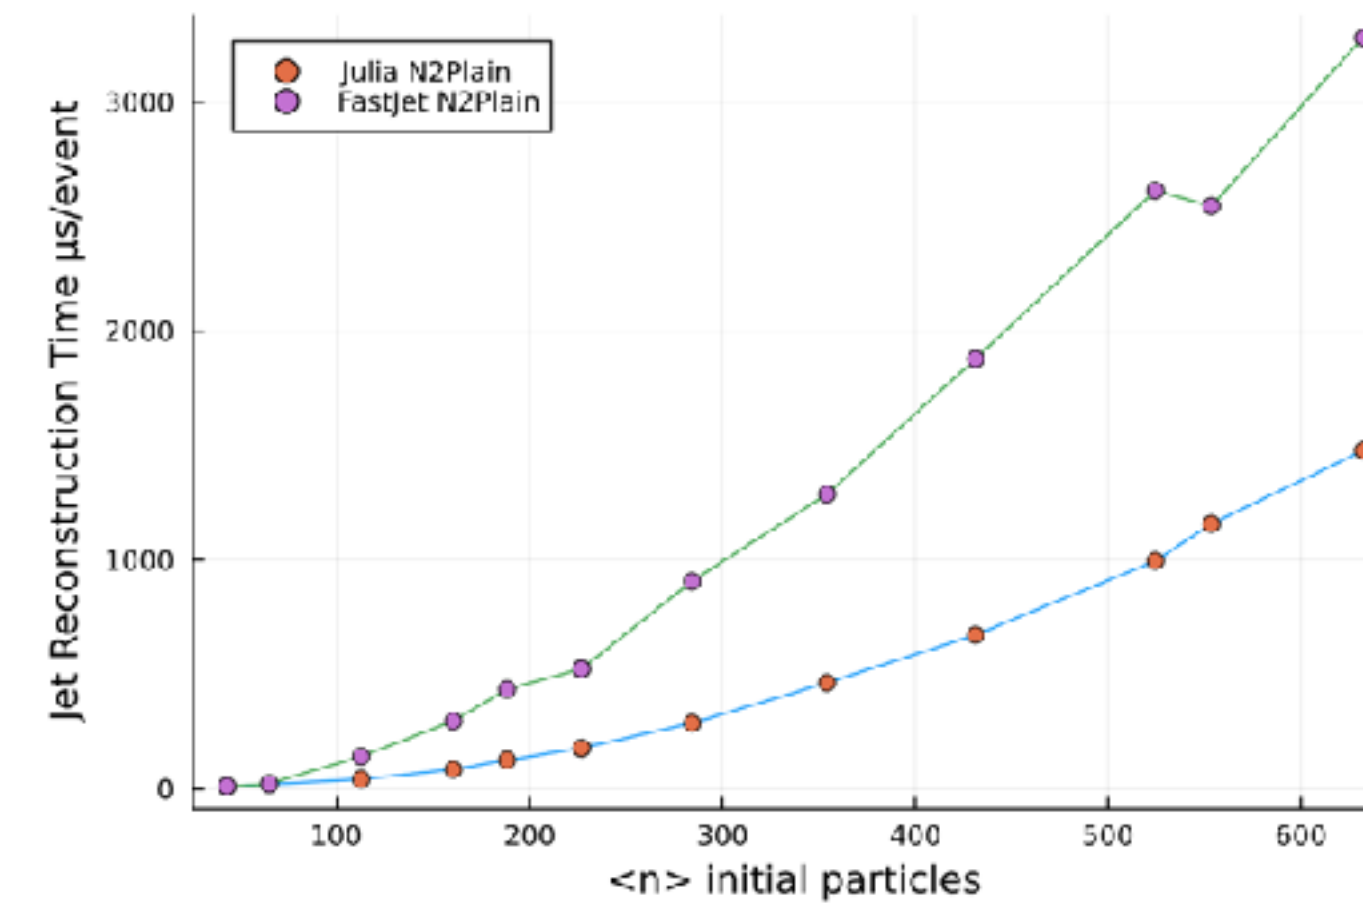
JetReconstruction.jl

Example for rewriting


- Sequential jet clustering
 - Algorithms from FastJet
 - Fully written in Julia
 - Visualization included
- Lesson learned
 - Better ergonomics
 - Better tooling
 - Neat visualization
 - More flexible usage

Polyglot Jet Finding

Graeme Andrew Stewart^{1,*}, Philippe Gras^{2*}, Benedikt Hegner^{1*}, and Atell Krasnopolski³



Anti- k_T Jet Reconstruction, 13TeV pp collision




JetReconstruction.jl

QuantumElectrodynamics.jl


Interfaces and tools available

- Particles
- Lorentz Vectors
- Phase space points
- Computational models
- Scattering processes
- Particle distributions
- Laser fields
- Event generation

$$e^- + \text{laser} \rightarrow e^- + (e^+e^-)$$



$$e^- + \text{laser} \rightarrow e^- + \gamma$$



QuantumElectrodynamics.jl