

Towards New Solutions for Scientific Computing: The Case of Julia

Maurizio Tomasi¹ and Mosé Giordano²

¹*Università degli Studi, Milano, Italy; maurizio.tomasi@unimi.it*

²*Università del Salento, Lecce, Italy; mose.giordano@le.infn.it*

Abstract. This year marks the consolidation of Julia (<https://julialang.org/>), a programming language designed for scientific computing, as the first stable version (1.0) has been released, in August 2018. Among its main features, expressiveness and high execution speeds are the most prominent: the performance of Julia code is similar to statically compiled languages, yet Julia provides a nice interactive shell and fully supports Jupyter. Moreover, it can transparently call external codes written in C, Fortran, and even Python and R without the need of wrappers. The usage of Julia in the astronomical community is growing, and a GitHub organization named JuliaAstro takes care of coordinating the development of packages. In this paper we present the features and shortcomings of this language, and discuss its application in astronomy and astrophysics.

1. Introduction

Julia (Bezanson et al. 2017) is a programming language that has recently reached its first stable milestone: version 1.0 has been released in August 2018, and the language specification has been frozen. Julia provides a number of features that makes it extremely interesting for astrophysics, astronomy, and scientific applications in general:

- Nice and simple syntax, similar to Matlab's.
- The speed of Julia codes often matches the speed of other languages used for High Performance Computing (HPC), namely C, C++, and Fortran, thanks to a number of features: type inference, Just-In-Time compilation, use of LLVM to produce optimized machine code;
- Native support for vectors, matrices, and tensors;
- Support for missing values (using the keyword `missing`), useful when dealing with data acquired using real-world experiments;
- First-class support for many numeric types, apart from integers and floating-point numbers: rationals, complex numbers, arbitrary-precision numbers.
- Symbolic computation (e.g., estimation of analytical derivatives) is easy to implement;
- Easy to call functions defined in dynamic libraries, using the `ccall` function;
- Ability to import packages written in Python or R; several wrappers to well-known Python libraries are available (e.g., `PyPlot.jl` wraps `Matplotlib`).

2. Features of Julia

2.1. Compilation model

Julia compiles functions the first time they are executed. The compilation depends on the type of the function parameters, as shown in this example:

```
f(x) = 2x + 1  # Define a function
f(1)          # Compile f assuming an integer argument
f(1.0)       # Compile again f assuming a float argument
f(3)         # No compilation is necessary, as 3 is an int
```

2.2. Operations on arrays, matrices, and tensors

Julia’s arrays are similar to Fortran’s:

1. Indices start from 1;
2. Arrays are stored in column-major order;
3. The compiler is able to propagate operators and functions to arrays, performing loop fusion.

The latter point is particularly important. If `a`, `b`, `c`, and `result` are arrays of the same size, the statement `result = a + b + c` in Fortran corresponds to one `do` loop. On the other side, the same code in Python applied on NumPy arrays is equivalent to the application of *three* `for`-loop cycles, because NumPy is not able to perform¹ *loop fusion*, i.e., the combination of several `for` loops into one.

Loop fusion is an important feature for HPC languages. Julia provides loop fusion through the so-called *dotted operators*: if `#` is a two-argument operator, `.#` applies the operator to all the elements of the two arrays. Therefore, in Julia the code `result .= a .+ b .+ c` is equivalent to the Fortran code `result = a + b + c`. Julia’s approach is more general, as this applies to custom operators and functions as well:

```
++(a::Real, b::Real) = 2a + b      # Custom operator
3 ++ 4                             # Result: 10
[3, 4] .++ [4, 7]                  # Result: [10, 15]
f(x::Real) = 3x^2                  # Custom function
f.([3, 6, 5])                      # Result: [27, 108, 75]
```

2.3. Homoiconicity

Julia provides the syntax for manipulating its own code with the same syntax used to manipulate variables. This feature, called *homoiconicity* (“same representation”), is inspired by LISP-like languages, and it has several applications in the domain of symbolic analysis (e.g., automatic computation of analytical derivatives). An interesting applications of homoiconicity in Julia is provided by the Zygote package (Innes 2018), which is able to perform automatic symbolic differentiation at compile time:

¹This limitation can be circumvented by other libraries, like WeldNumPy (<https://www.weld.rs/weldnumpy/>), Numba (<https://numba.pydata.org/>), or Cython (<https://cython.org/>).

```
julia> using Zygote
julia> f(x) = 2x + 1
julia> @code_llvm f'(0)
; Function #68
; Location: /somewhere/interface.jl:49
define i64 @"julia_#68_37159"(i64) {
top:
ret i64 2 # Return 2 immediately (the derivative is a constant)
}
```

3. Julia in Astronomy

3.1. JuliaAstro

The JuliaAstro GitHub organization (<https://github.com/JuliaAstro>) collects all the packages related to astronomy developed for Julia. At the time of writing (November 2018), the packages are the following:

- `AstroImages.jl`: Visualization of astronomical images;
- `AstroLib.jl`: Bundle of small astronomical and astrophysical routines;
- `AstroTime.jl`: Astronomical time keeping;
- `Cosmology.jl`: Library of cosmological functions;
- `DustExtinction.jl`: Models for the interstellar extinction due to dust;
- `ERFA.jl`: Wrapper to `liberfa`²;
- `EarthOrientation.jl`: Earth orientation parameters from IERS tables;
- `FITSIO.jl`: Flexible Image Transport System (FITS) file support;
- `LombScargle.jl`: Compute Lomb-Scargle periodogram;
- `SPICE.jl`: Julia wrapper for NASA NAIF's SPICE toolkit;
- `SkyCoords.jl`: Support for astronomical coordinate systems;
- `UnitfulAstro.jl`: An extension of `Unitful.jl` (a package to attach measure units to variables) for astronomers;
- `WCS.jl`: Astronomical World Coordinate Systems library.

3.2. Simulating a CMB space mission

One of us (MT) has had the opportunity to use Julia in a few studies involving the design of a CMB space mission (CORE, PICO, and LiteBIRD). These studies involved the simulation of the operations needed to observe the sky, and they required the generation of simulated noisy data timelines acquired by instruments mounted onboard the spacecraft. The quantity of data was of the order of hundreds of GB, and the exploratory nature of the study made existing codes (developed in C++ for the Planck experiment) cumbersome to use, as they were conceived as large monolithic programs meant to be ran end-to-end. A rewrite of some modules in Julia provided similar performance (within 10%) with the existing C++ codes; moreover, the Julia codes were runnable in Jupyter notebooks, thus allowing to interactively explore the parameter space and ease data analysis.

²<https://github.com/liberfa/erfa>. This is a BSD-licensed replica of the SOFA library (<http://www.iausofa.org/>).

4. Conclusions

Julia has several features that make it an interesting solution for astronomical and astrophysical projects. It can achieve performance similar to compiled languages, like C and Fortran, but it is considerably more expressive and easy to use.

Notwithstanding the long list of interesting features, we believe it would not be fair to omit some of Julia's most important shortcomings:

- Compilation times can be significant. Since compilation happens at runtime, a Julia script that calls several short functions can be noticeably slower than a similar script written in other compiled or interpreted languages.
- The language is new, and there are not as many libraries as for other languages. Python, R, C, and Fortran library are easy to import; however, if a code heavily relies only on a few libraries, it is usually easier to just use the language for which these libraries were developed than wrapping everything in Julia.
- It is still not possible to produce stand-alone executables. This makes code deployment more difficult.
- As any new language, it is necessary to grasp a number of concepts before being fully productive with it. For instance, a programmer experienced in NumPy might find surprising that explicit `for` loop can be more performant than expressions involving broadcasting. (The repository <https://github.com/ziotom78/python-julia-c-> provides an example.)

In the opinion of the authors, there are two contexts in astrophysical data analysis where Julia can provide a significant advantage over existing solutions:

- Analysis of large amounts of data, where no existing codes are available and the amount of calculations is significant. In this case, Julia codes can be as performant as other codes written using multiple libraries and languages: the typical case uses Python for most of the code and some optimized library (Numba, Fortran codes wrapped using `f2py`) for the most performance-critical routines. As an application of this use case we mention the Celeste project, which was able to load and process 178 TB of data from the SDSS catalogue in 14.6 minutes across 8192 nodes (Regier et al. 2018).
- Existing codes are monolithic and difficult to use interactively, and the expense of rewriting code in Julia can be rewarded by the possibility to run the code interactively, either in Julia's command line or in Jupyter notebooks.

Acknowledgments. We thank the Julia community at for many useful discussions, <https://discourse.julialang.org/>

References

- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, *SIAM Review*, 59, 65
 Innes, M. 2018, ArXiv e-prints. 1810.07951
 Regier, J., Pamnany, K., Fischer, K., et al. 2018, ArXiv e-prints. 1801.10277