

Tasks all the way down

Parallelism in **julia**

Rapid adoption for a young language

20M downloads; 4,000 packages; 10,000 companies; 1,500 universities

IEEE Spectrum Language Rankings

Choose a Ranking

IEEE Spectrum

Trending

Jobs

Open

Custom

Language Types

Web



Enterprise



Mobile



Embedded




(Click to hide)

17	HTML▼		61.4
18	Kotlin▼		57.8
19	Julia▼		56.0
20	Rust▼		55.6
21	Shell▼		52.0
22	Processing▼		49.2

TOOLBOX

JULIA: COME FOR THE SYNTAX, STAY FOR THE SPEED

Researchers often find themselves coding algorithms in one programming language, only to have to rewrite them in a faster one. An up-and-coming language could be the answer.



BY JEFFREY M. PERKEL

When it comes to climate modeling, every computational second counts. Designed to account for air, land, sun and sea, and the complicated physics that links them, these models can run to millions of lines of code, which are executed on the world's most powerful computers. So when the coder-climatologists of the Climate Modeling Alliance (CIMA) — a coalition of US-based scientists, engineers and mathematicians — set out to build a model from the ground up, they opted for a language that could handle their needs. They opted for Julia.

Launched in 2012, Julia is an open-source language that combines the interactivity and syntax of 'scripting' languages, such as Python, Matlab and R, with the speed of 'compiled' languages such as Fortran and C.

Among climate scientists, the lingua franca is Fortran: speedy, but — with roots dating to the 1950s — not terribly exciting. "A lot of people, when they hear 'Fortran', are like, 'Oh, my God, I don't want to program in that,'" says Frank Giraldo, a mathematician at the Naval Postgraduate School in Monterey, California, and a co-principal investigator on the CIMA project. Younger programmers prefer languages that can accommodate the latest trends in software and hardware design,

Giraldo says, and since adopting Julia he has seen an uptick in interest. "Some of them are really interested in climate modelling, but others are intrigued by the idea of using Julia for some large-scale application," he says.

Jane Herriman, who is studying materials science at the California Institute of Technology in Pasadena, says that she has seen ten-

nature
International journal of science

Customers, Partners & Companies Using Julia



10,000+ companies using Julia








Many case studies here:

<https://juliacomputing.com/case-studies/>

Universities Using and Teaching Julia



Books on Julia





Celeste.jl: Julia at Peta-scale

Cori: 650,000 cores, 1.3M threads, 60 TB of data



Most light sources are near the detection limit.

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^{||}

^{*}Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

[†]Parallel Computing Lab, Intel Corporation

[‡]Julia Computing

[§]Computer Science and AI Laboratories, Massachusetts Institute of Technology

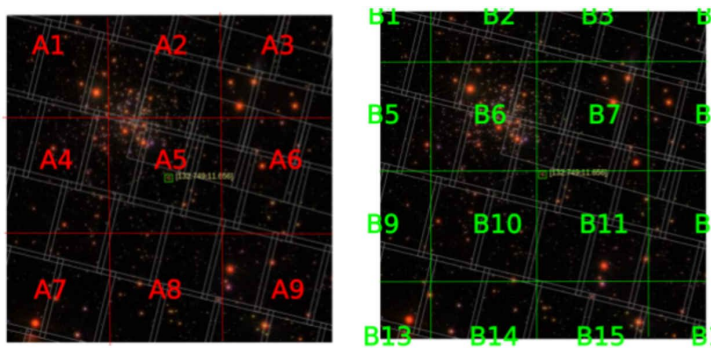
[¶]Department of Statistics, University of California, Berkeley

^{||}Lawrence Berkeley National Laboratory




Irregular, Multi-Scale Parallelism


Nodes



Threads



SIMD



Light sources that do not overlap may be updated concurrently.

Fundamental Schedule Unit: Task

```
function pfib(n::Int)
    if n <= 1
        return n
    end
    t = Threads.@spawn pfib(n-2)
    return pfib(n-1) + fetch(t)::Int
end
```

- Concurrency
 - w/ High Performance I/O Scheduler
- Parallelism
- Low memory footprint (Millions of tasks per node)
- Dynamically serializable

Single Node Schedule: Parallel Depth First


Work Stealing:



Parallel Depth First:



- Highly Cache efficient for regular problems
- Composability/Nested parallelism without sacrificing performance
 - Fearless parallelism for library authors




Jameson
Nash

Jeff
Bezanson

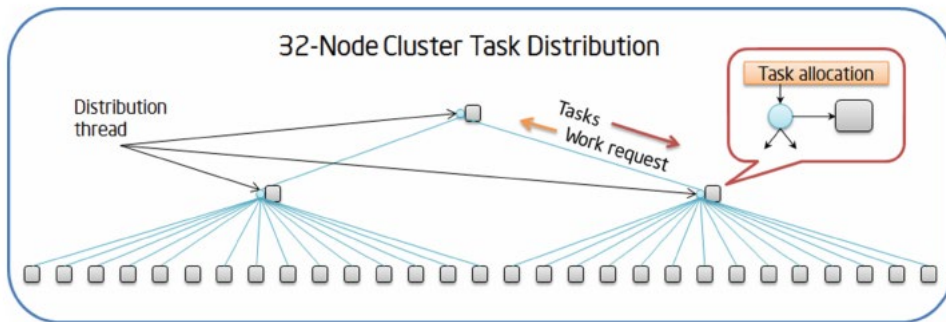
Kiran
Pamnany

Distributed Schedule: User Policy



Dagger.jl: Dask-like global DAG scheduler

- Different Applications need different scheduling approaches
- Common needs available from package repository



Gasp.jl: Dtree load-balancing irregular work scheduler (scaling to millions of concurrent threads)

Active work

- Compiler integration
 - LLVM-level optimization of task states
 - Julia-level semantics for compiler optimized tasks
- SIMT unification (in two directions)
 - Each GPU “thread” should be a task
 - SIMT programming model for regular workloads across the abstraction levels
- Code Loading/Distributed JIT/Code Caching Opportunities