

Interactive Computing with Julia in JupyterLab


CSCS User Lab Day – Meet the Swiss National Supercomputing Centre

Samuel Omlin

September 1st 2020

Prototype

```
P = rand(4,3)
```



2

Production code

```
float* P;  
P = malloc(...);  
rand(P,...);
```

The two language problem

Prototype (MATLAB / Python / ...)

simple & high-level

interactive

low development cost

slow


Production code (C / C++ / Fortran / ...)

complex & low-level


not interactive

high development cost


fast




The two language problem



The two language problem



The two language problem



Solution

A language that can be used for both
Prototype & Production code

Solution

A language that can be used for both
Prototype & Production code



Solution



simple & high-level

interactive

low development cost

fast

Solution



simple & high-level

interactive


low development cost

fast

Fast and interactive???

Julia code is compiled, yet only shortly before you use it **the first time.**

Solution



CUDA.jl

Native Julia Code for GPUs!




simple & high-level

interactive

low development cost

fast




MPI.jl

Julia suitable for GPU supercomputing

Single GPU performance:

93% of the the CUDA C code



Agenda

- Introduction: the two language problem ✓
- Julia on Piz Daint
- Julia in JupyterLab at CSCS
- Julia Notebook examples
- Conclusions & Outlook

Julia on Piz Daint

Julia modules:

```
$> module load daint-gpu # or daint-mc
$> module load Julia
$> module load JuliaExtensions
```

<- includes MPI + CUDA packages
<- Plots, PyCall & HDF5 packages...

Available packages:

```
julia> versioninfo()
```

Note on the Julia package manager manager:

julia> Pkg.status() shows only the packages installed by the user by default, but you can load the above packages normally, e.g.:

```
julia> using MPI
```

Start an interactive Julia session with GPU:

```
$> srun -C gpu --time=04:00:00 --pty bash
$> julia
```

Julia on Piz Daint

Julia modules:

```
$> module load daint-gpu # or daint-mc  
$> module load Julia  
$> module load JuliaExtensions
```

stacked environment:
user installed packages have precedence!

Available packages:

```
julia> versioninfo()
```

Note on the Julia package manager manager:

julia> Pkg.status() shows only the packages installed by the user by default, but you can load the above packages normally, e.g.:

```
julia> using MPI
```

Start an interactive Julia session with GPU:

```
$> srun -C gpu --time=04:00:00 --pty bash  
$> julia
```

Julia on Piz Daint

Julia modules:

```
$> module load daint-gpu # or daint-mc
$> module load Julia
$> module load JuliaExtensions
```

<- includes MPI + CUDA packages
<- Plots, PyCall & HDF5 packages...

Available packages:

```
julia> versioninfo()
```

Note on the Julia package manager manager:

julia> Pkg.status() shows only the packages installed by the user by default, but you can load the above packages normally, e.g.:

```
julia> using MPI
```

More information: <https://user.cscs.ch/tools/interactive/julia/>

Agenda

- Introduction: the two language problem ✓
- Julia on Piz Daint ✓
- Julia in JupyterLab at CSCS
- Julia Notebook examples
- Conclusions & Outlook

Julia in JupyterLab at CSCS

- Accesses the same stacked environment
- The modules **Julia** and **JuliaExtensions** are automatically loaded.
- Currently not set up for usage with MPI (not yet straightforward and well supported): use a single node.

Installing a package from the command line or from JupyterLab gives the exact same result!

Agenda

- Introduction: the two language problem ✓
- Julia on Piz Daint ✓
- Julia in JupyterLab at CSCS ✓
- Julia Notebook examples
- Conclusions & Outlook

Notebook 1: using the stacked environment

<https://user.cscs.ch/tools/interactive/jupyterlab/#ijulia>

Notebook 2: glacier flow using GPU

2-D Shallow ice equations

$$\frac{\partial H}{\partial t} = -\nabla_i(qH_i)$$

$$qH_i = -\frac{H^3 g}{3\mu} \nabla_i(H + B)$$

Notebook 2: glacier flow using GPU

2-D Shallow ice equations

$$\frac{\partial H}{\partial t} = -\nabla_i(qH_i)$$

$$qH_i = -\frac{\mathbf{H}^3 g}{3\mu} \nabla_i(H + B)$$

Nonlinear diffusion!

Notebook 2: glacier flow using GPU

Numerics

- Iterative algorithm with implicit time stepping
- Pseudo-transient method
- Numerical damping for convergence acceleration

Notebook 2: glacier flow using GPU

Demo...

JupyterHub - Chromium

JupyterHub

omlins.jupyter.csccs.ch/hub/spawn/omlins

CSCS ETHzürich

Home Token User: omlins


Node Type: GPU Nodes: 1 Duration (hr): 1

Advanced options


Launch JupyterLab

Help | Privacy | Terms 2019 © CSCS | www.csccs.ch

3-D OpenGL visualization in Julia (different topography)



3-D OpenGL visualization in Julia (different topography)



Uses Makie.jl

Done on Laptop.
We will see if Makie.jl
can be installed after Piz
Daint upgrade.

Agenda

- Introduction: the two language problem ✓
- Julia on Piz Daint ✓
- Julia in JupyterLab at CSCS ✓
- Julia Notebook examples ✓
- Conclusions & Outlook

Conclusions & outlook

- same stacked environment in JupyterLab as when using Julia from command line
- `CUDA.jl` enables writing native Julia code for GPUs
- We will see if `Makie.jl` can be installed after Piz Daint upgrade.

Conclusions & outlook

- same stacked environment in JupyterLab as when using Julia from command line
- `CUDA.jl` enables writing native Julia code for GPUs
- We will see if `Makie.jl` can be installed after Piz Daint upgrade.

Questions / advice / feedback / ...

I am the responsible for Julia computing – get in touch with me!

help@cscs.ch

|

Samuel.Omlin@cscs.ch

Thank you for your kind attention