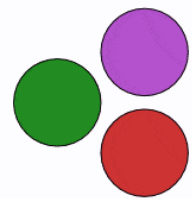


WaterLily.jl

An **incompressible flow solver** for dynamic bodies
running seamlessly on **CPU/GPU**

Bernat Font, Gabe Weymouth



b.font@tudelft.nl

g.d.weymouth@tudelft.nl

Julia4PDEs 2026

Behind WaterLily

Paella Appreciator



Seafood Appreciator



Paella Aficionado



The requirements: leveraging the Julia ecosystem

1. Fast:

- Compiled: Julia uses **LLVM**
- Multi-threading on CPU and GPU of any vendor: **KernelAbstractions.jl**
- Memory-distributed computing: **MPI.jl**

2. Flexible

1. Support of STL-like geometries via **MeshIO.jl** and **GeometryBasics.jl**
2. **ForwardDiff.jl** to compute body velocity from position and time-mapping function

3. Natively linked to ML

- Can be coupled with **Lux.jl**

The requirements: leveraging the Julia ecosystem

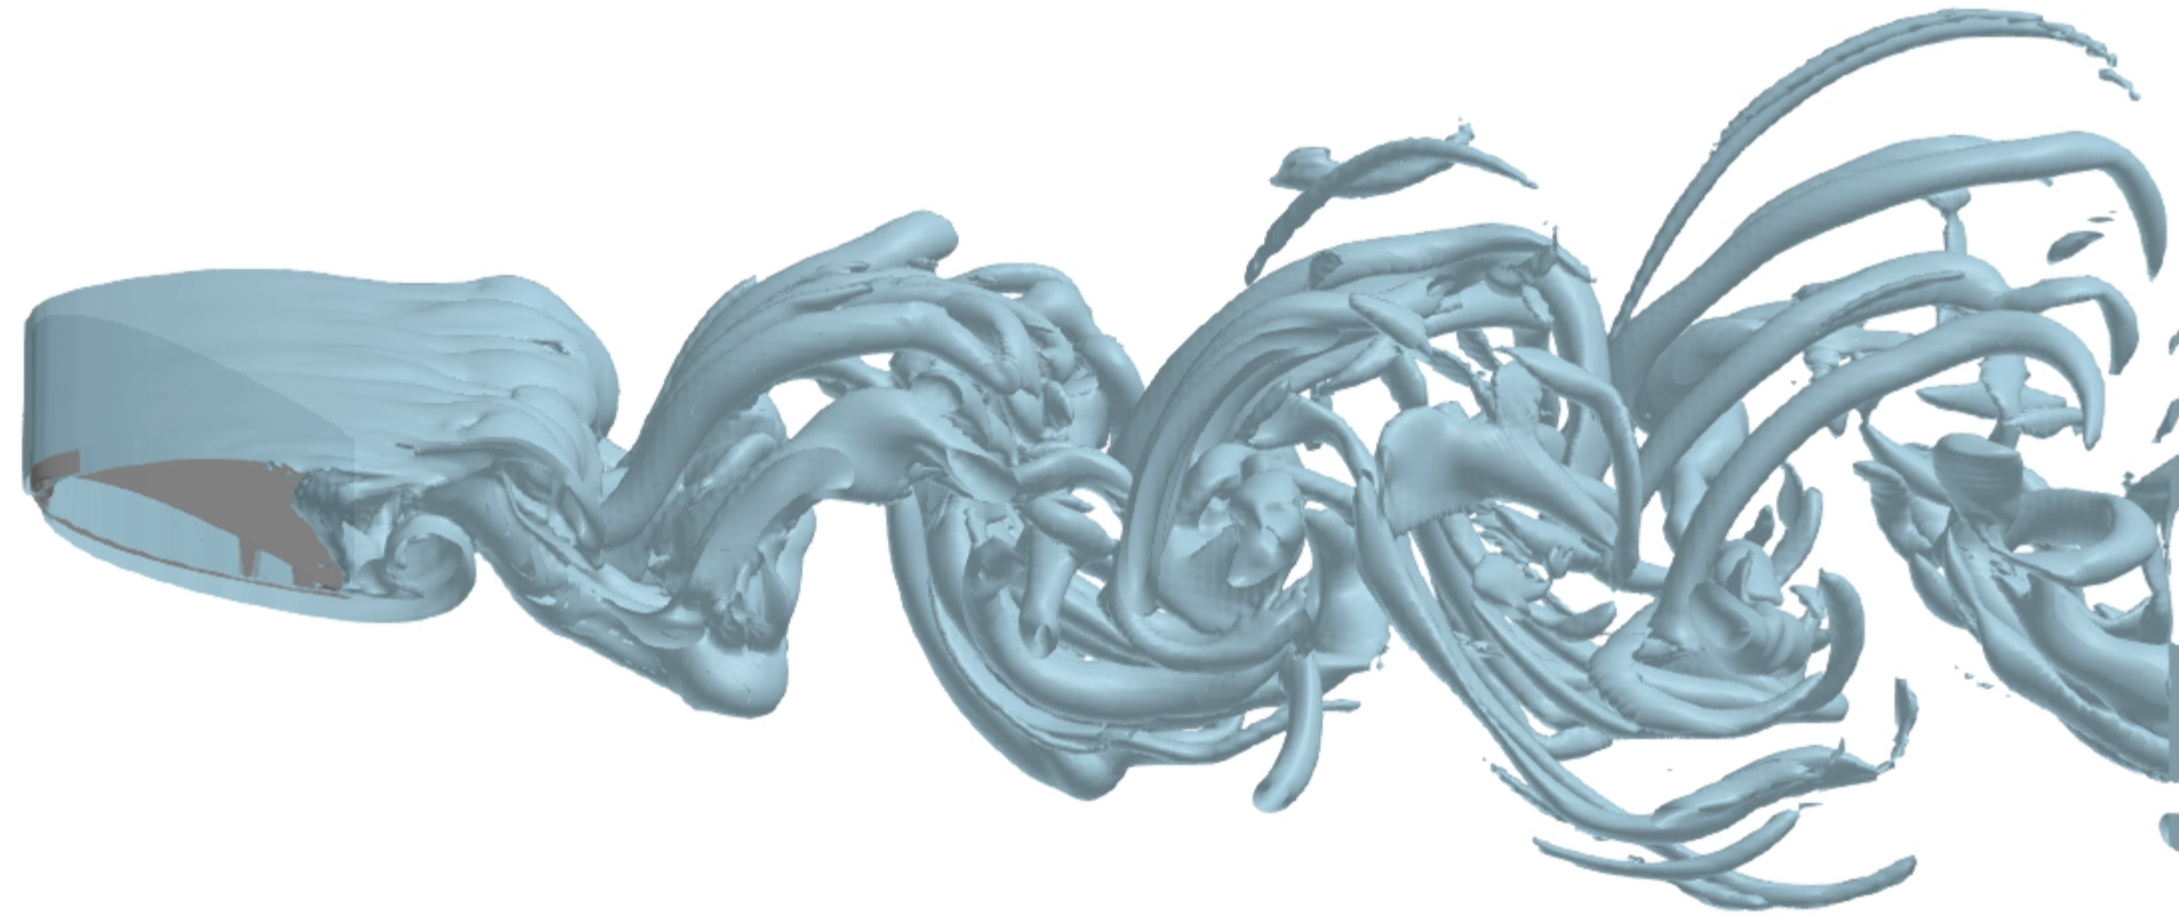
4. Differentiable

- Full solver differentiability via **ForwardDiff.jl**

5. Visualization

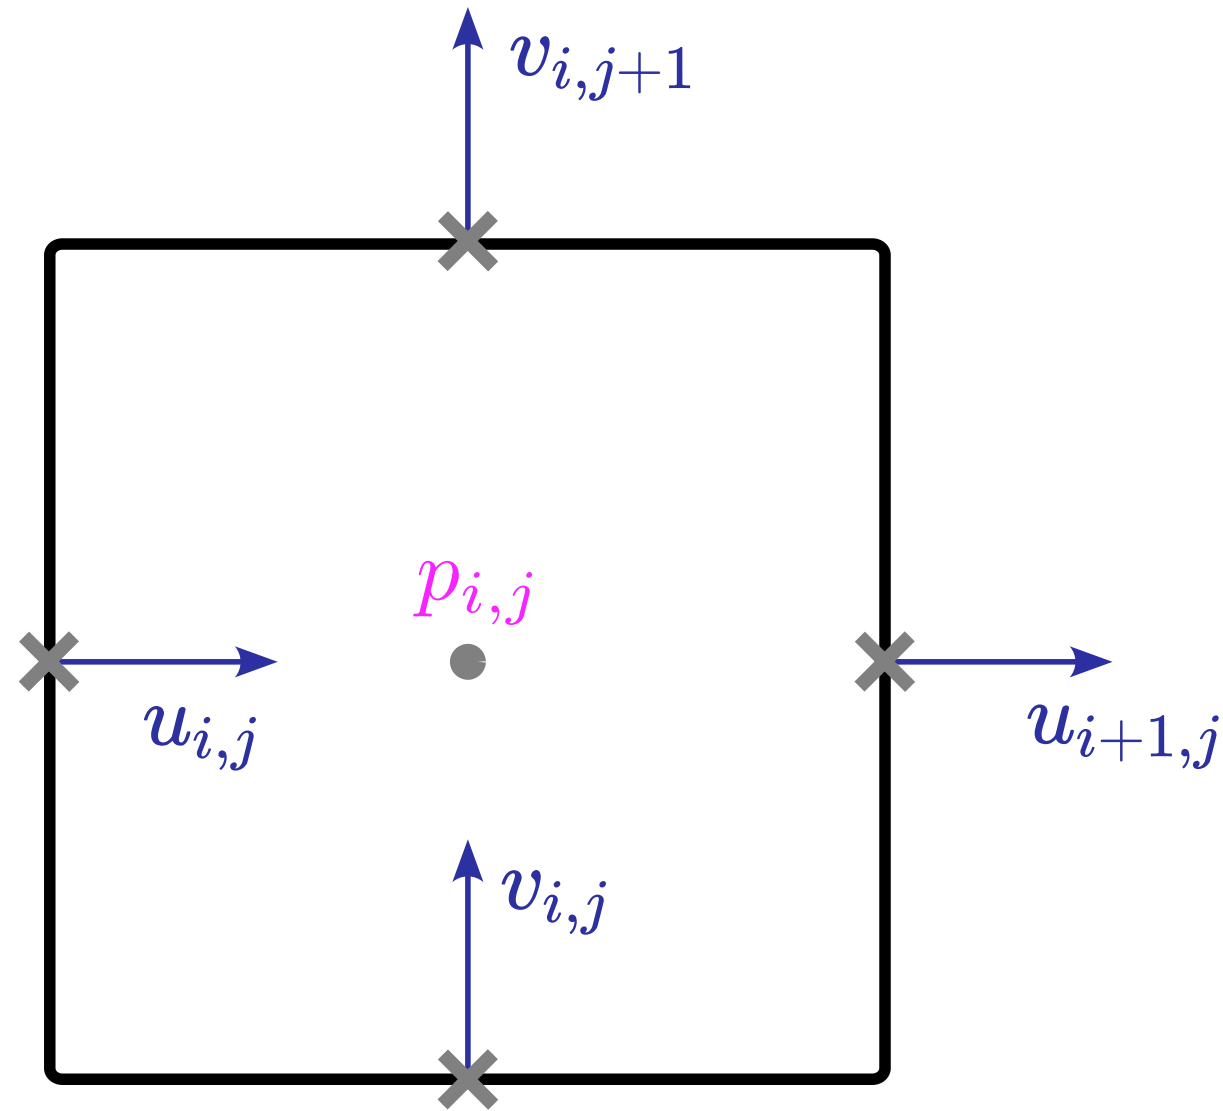
- Fast rendering of 2D or 3D flow fields using **Makie.jl**

WaterLily | Numerical methods



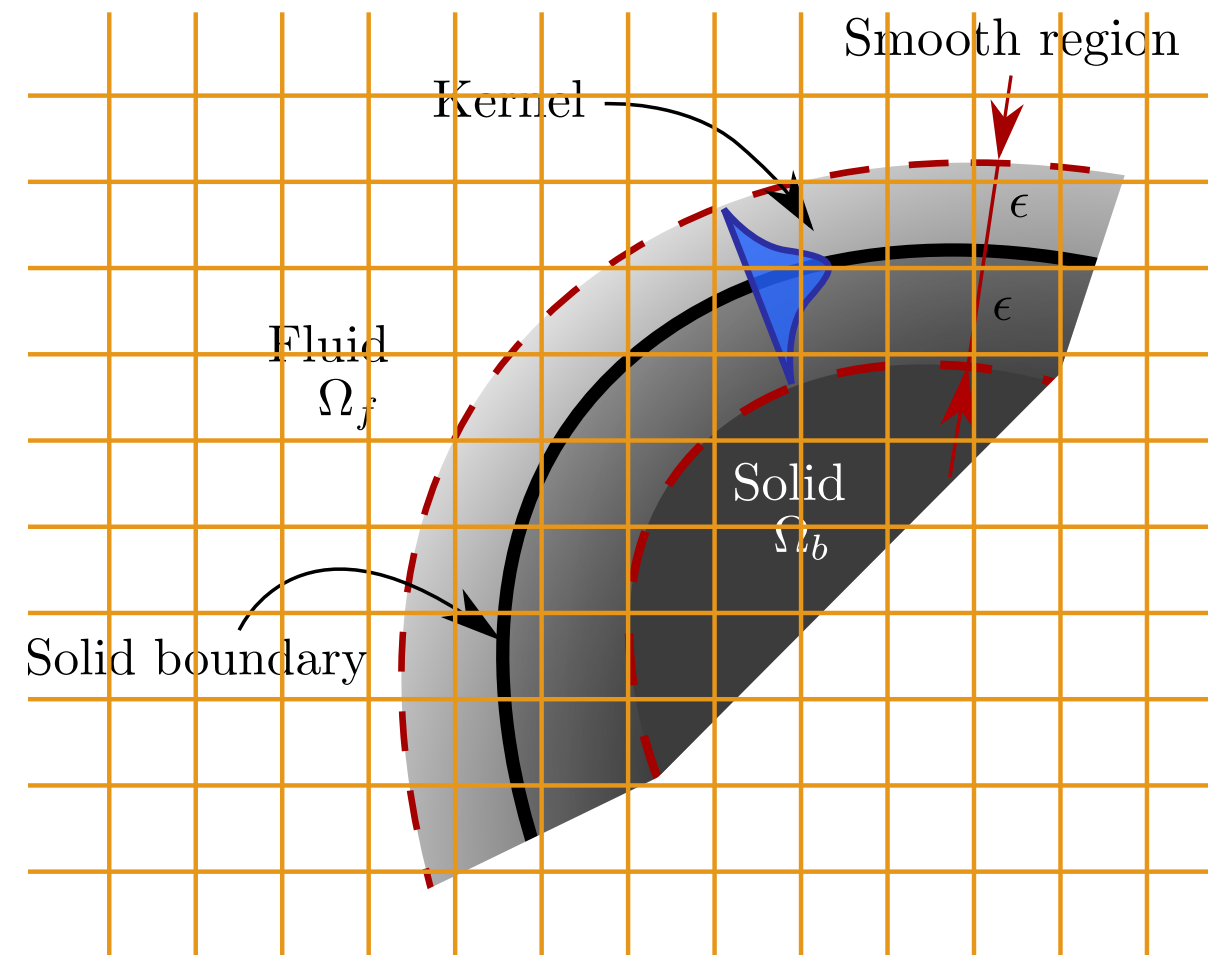
WaterLily | Numerical methods

- Incompressible flow
- Finite volume in a staggered grid
 - QUICK reconstruction scheme in space, predictor-corrector in time (2nd order)



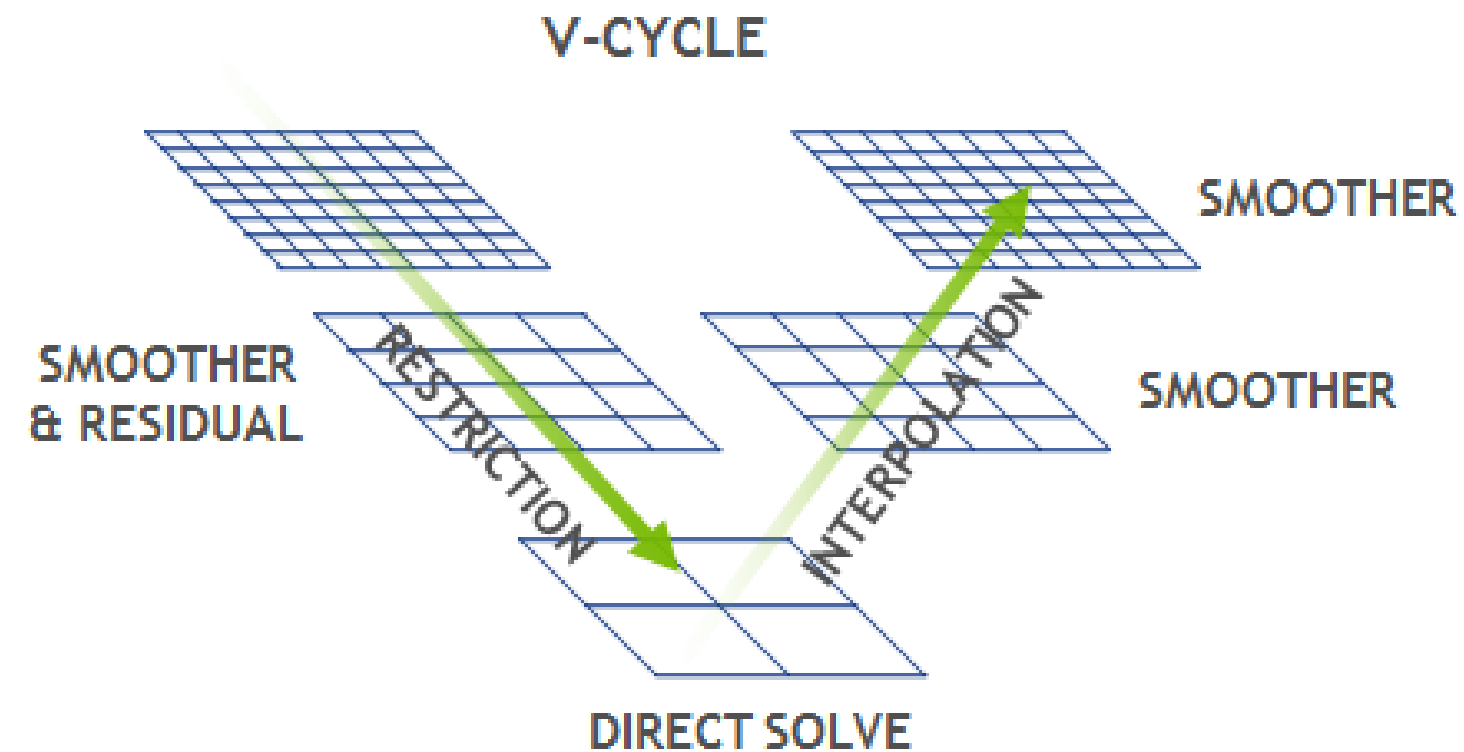
WaterLily | Numerical methods

- Incompressible flow
- Finite volume in a staggered grid
 - QUICK reconstruction scheme in space, predictor-corrector in time (2nd order)
- Boundary data immersion method for solid geometries
 - Signed distance function, **ForwardDiff.jl** (normal, curvature, velocity)



WaterLily | Numerical methods

- Incompressible flow
- Finite volume in a staggered grid
 - QUICK reconstruction scheme in space, predictor-corrector in time (2nd order)
- Boundary data immersion method for solid geometries
 - Signed distance function, **ForwardDiff.jl** (normal, curvature, velocity)
- Geometric multigrid (GMG) for the pressure solver + projection method



Porting to GPU using `KernelAbstractions.jl`

- Consider the divergence of a 2D vector field \vec{u} in a staggered and Cartesian grid where $\Delta x = 1$

$$\sigma = \iiint (\nabla \cdot \vec{u}) dV = \iint \vec{u} \cdot \hat{n} dS \rightarrow \sigma_{i,j} = (u_{i+1,j} - u_{i,j}) + (v_{i,j+1} - v_{i,j})$$

- This can be written as

```
1 N = (10, 10) # domain size
2 σ = zeros(N) # scalar field
3 u = rand(N..., length(N)) # 2D vector field with ghost cells
4
5 for d ∈ 1:ndims(σ), I ∈ inside(σ)
6     σ[I] += ∂(d, I, u)
7 end
```

Porting to GPU using `KernelAbstractions.jl`

- `KernelAbstractions.jl` allows to write a simple kernel function that is then **specialized** for different GPU backend:
 - `CUDA.jl`, `AMDGPU.jl`, `oneAPI.jl`, and `Metal.jl`

```
1 using KernelAbstractions: get_backend, @index, @kernel
2
3 @kernel function _divergence!( $\sigma$ , u)
4     I = @index(Global, Cartesian) # define a parallel Cartesian index in Global memory
5      $\sigma$ _sum = zero(eltype( $\sigma$ )) # initialize accumulation
6     for d  $\in$  1:ndims( $\sigma$ )
7          $\sigma$ _sum +=  $\partial$ (d, I, u)
8     end
9      $\sigma$ [I] =  $\sigma$ _sum # assign result
10 end
11 function divergence!( $\sigma$ , u)
12     _divergence!(get_backend( $\sigma$ ), 64)( $\sigma$ , u, ndrange=size(inside( $\sigma$ ))) # generate & call the kernel
13 end
```

Porting to GPU using `KernelAbstractions.jl`

```
1 for d ∈ 1:ndims(σ)
2     @loop σ[I] += ∂(d, I, u) over I ∈ inside(σ)
3 end
```

- Generating a kernel for every **loop** in our solver

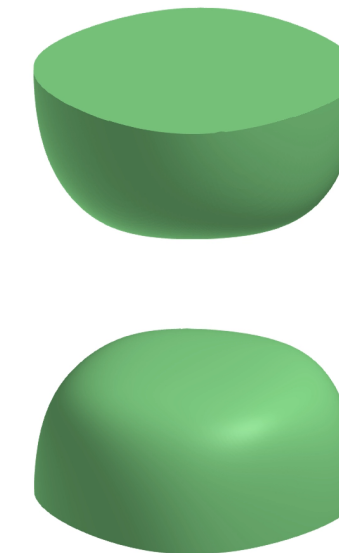
```
1 macro loop(args...)
2     ex, I, R, sym = args, [] # get expression, iterator, range
3     grab!(sym, ex)          # get kernel arguments (symbols)
4     @gensym kern           # generate unique kernel function name
5     return quote
6         @kernel function $kern($(rep.(sym)...)) # define the kernel
7             $I = @index(Global, Cartesian)
8             $ex
9         end
10        $kern(get_backend($(sym[1])), 64)($(sym...), ndrange=size($R)) # generate & call the kernel
11    end |> esc
12 end
```

- Automatically **specialized for any computing backend**

WaterLily | Easy simulation setup

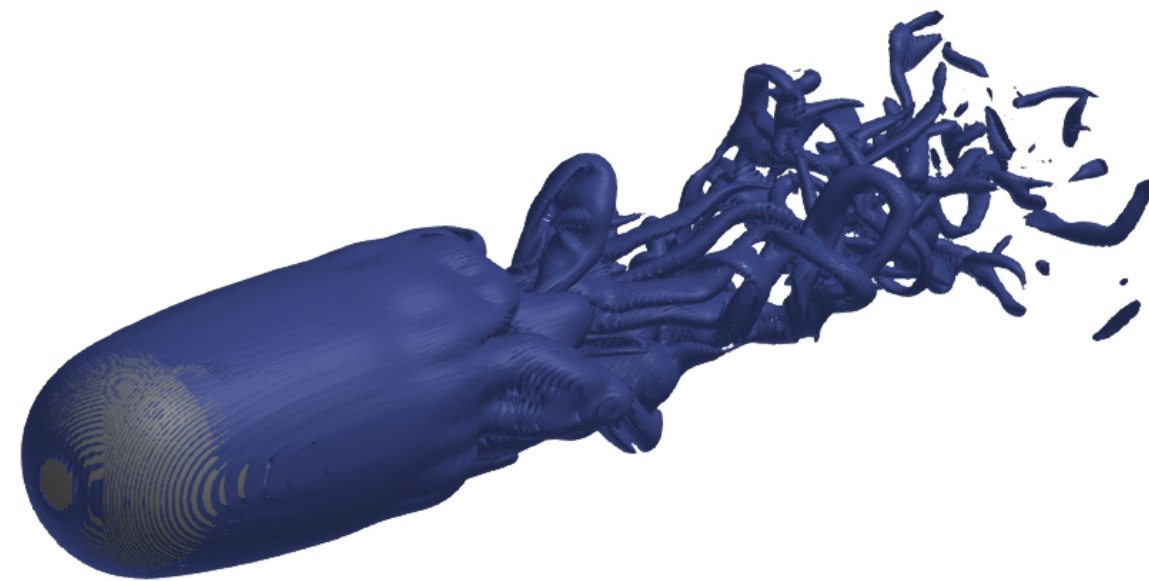
- The **Simulation** type holds information of **Fluid** and **Body** types
- User defines the IC, BC, simulation size, and the body SDF (if any)
- **Simulation** can be run on CPU/GPU changing the array type: **Array**, **CuArray**, **ROCArray**

```
1 using WaterLily, CUDA
2
3 function TGV(; p=6, Re=1600, T=Float32, mem=Array)
4     # Define vortex size, velocity, viscosity
5     L = 2^p; U = T(1); κ=T(π/L); ν = T(1/(κ*Re))
6     function uλ(i,xyz) # initial velocity field
7         x,y,z = @. xyz*π/L
8         i==1 && return -U*sin(x)*cos(y)*cos(z) # u_x
9         i==2 && return U*cos(x)*sin(y)*cos(z) # u_y
10        return 0*U # u_z
11    end
12    return Simulation((L, L, L), (0, 0, 0), L; U, uλ, ν, T, mem)
13 end
14 sim = TGV(mem=CuArray)
15 sim_step!(sim, 20)
```



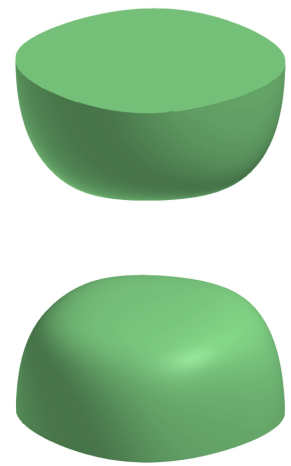
WaterLily | Easy simulation setup

```
1 using WaterLily, CUDA, StaticArrays
2
3 function sphere(D; Re=3700, U=1, L=(7,3,3), center=(1.5,1.5,1.5), mem=CuArray, T=Float32)
4     center = SA{T}[center...]
5     body = AutoBody((x,t)-> √sum(abs2, x .- (center .* D)) - D÷2)
6     Simulation(L.*D, (U, 0, 0), D; U, v=U*D/Re, body, mem, T, exitBC=true)
7 end
8
9 sim = sphere(128)
10 sim_step!(sim, 10)
```

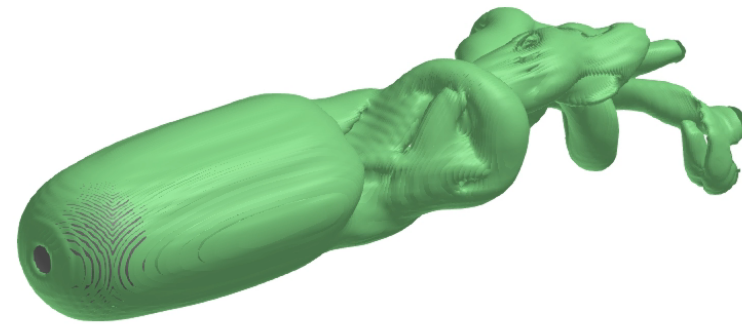


Benchmarking | Not only cool, but fast too!

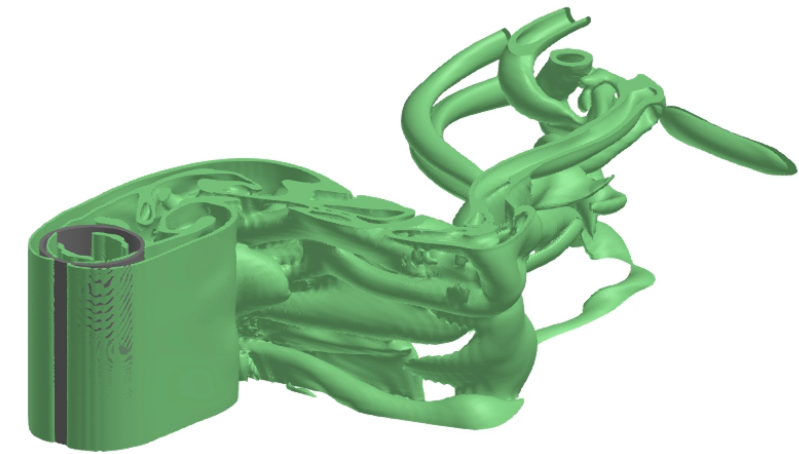
TGV



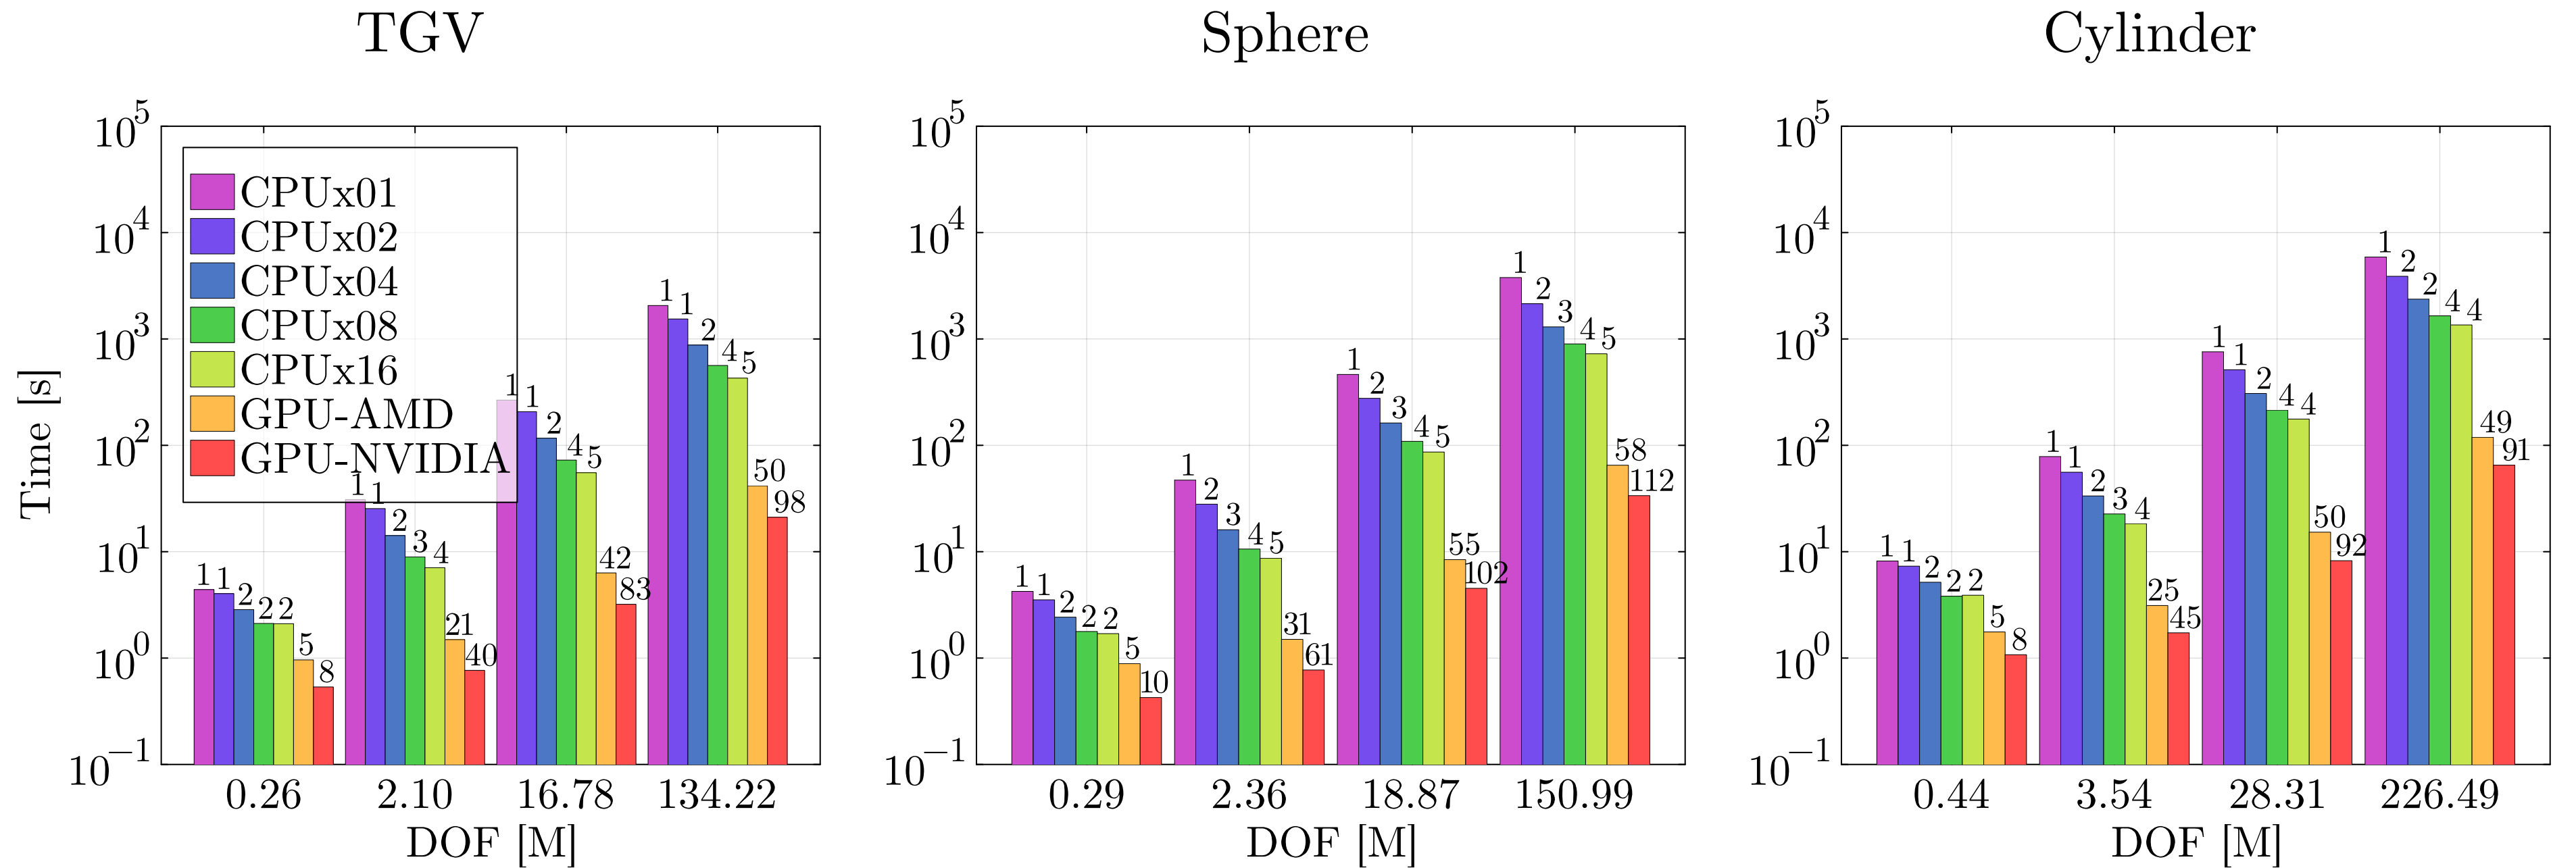
Sphere



Cylinder



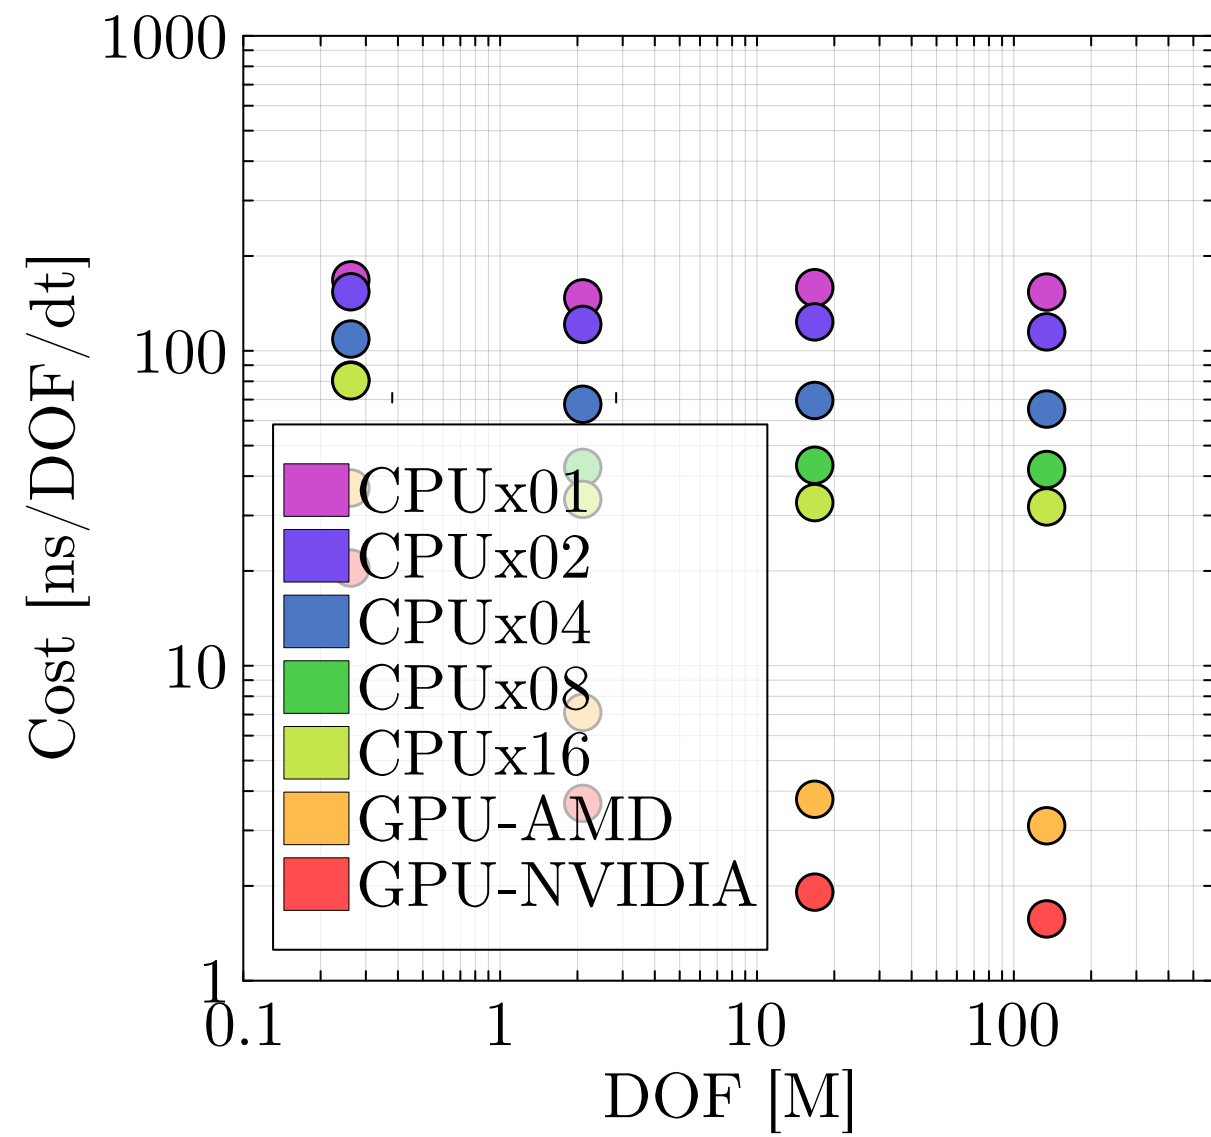
Benchmarking | x100 speedup on modern GPUs vs serial CPU



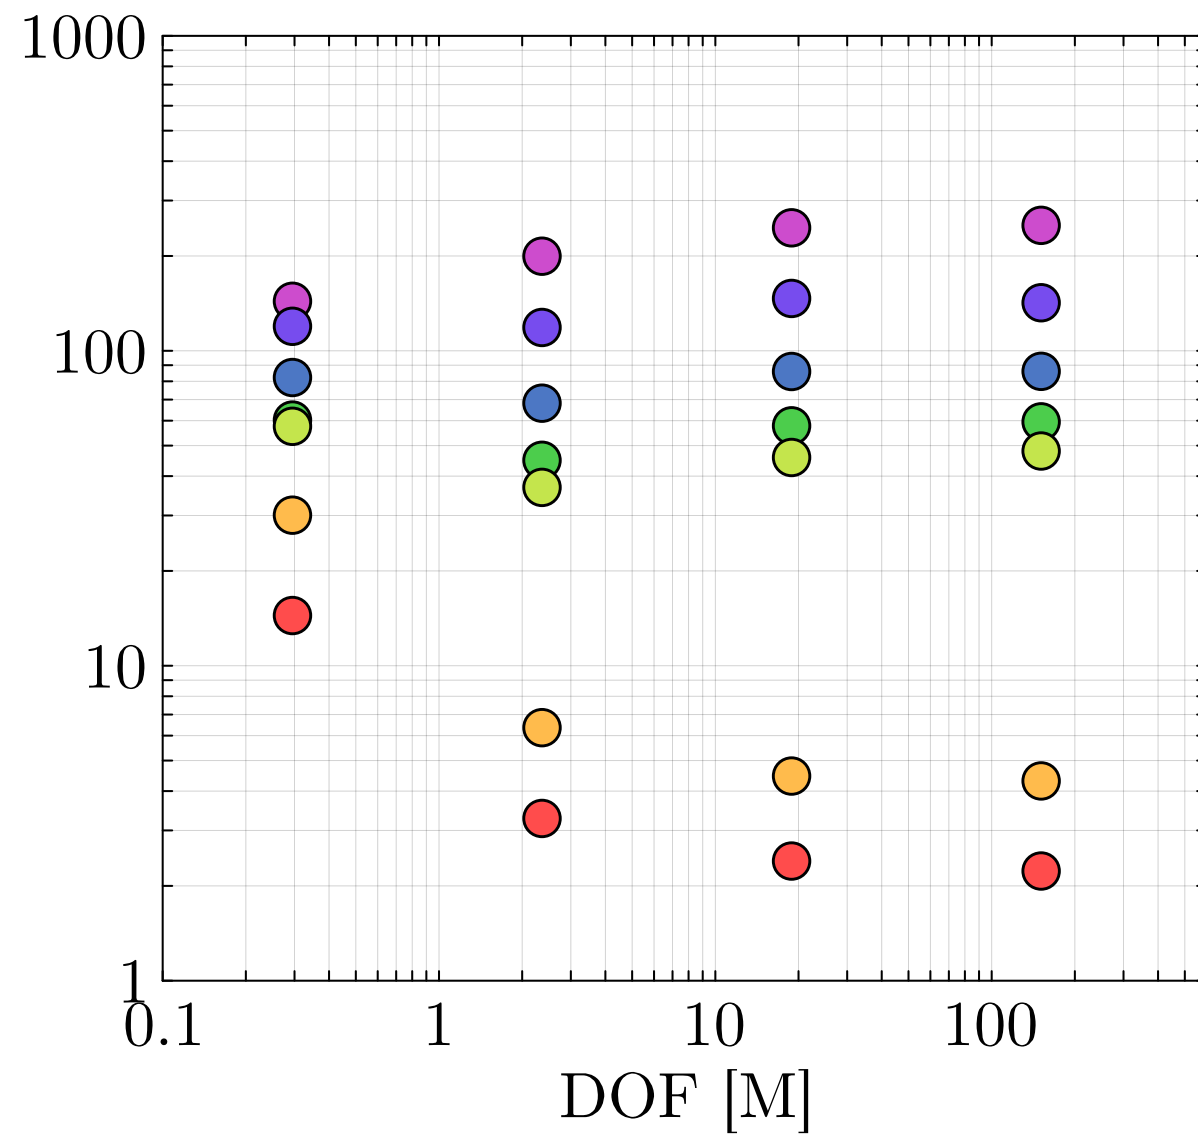
- FP32, 100 dts, NVIDIA H100 GPU, AMD MI250X GPU
- 302M DOF consume 39GB VRAM

Benchmarking | Top speed: 1.5 ns/DOF/dt

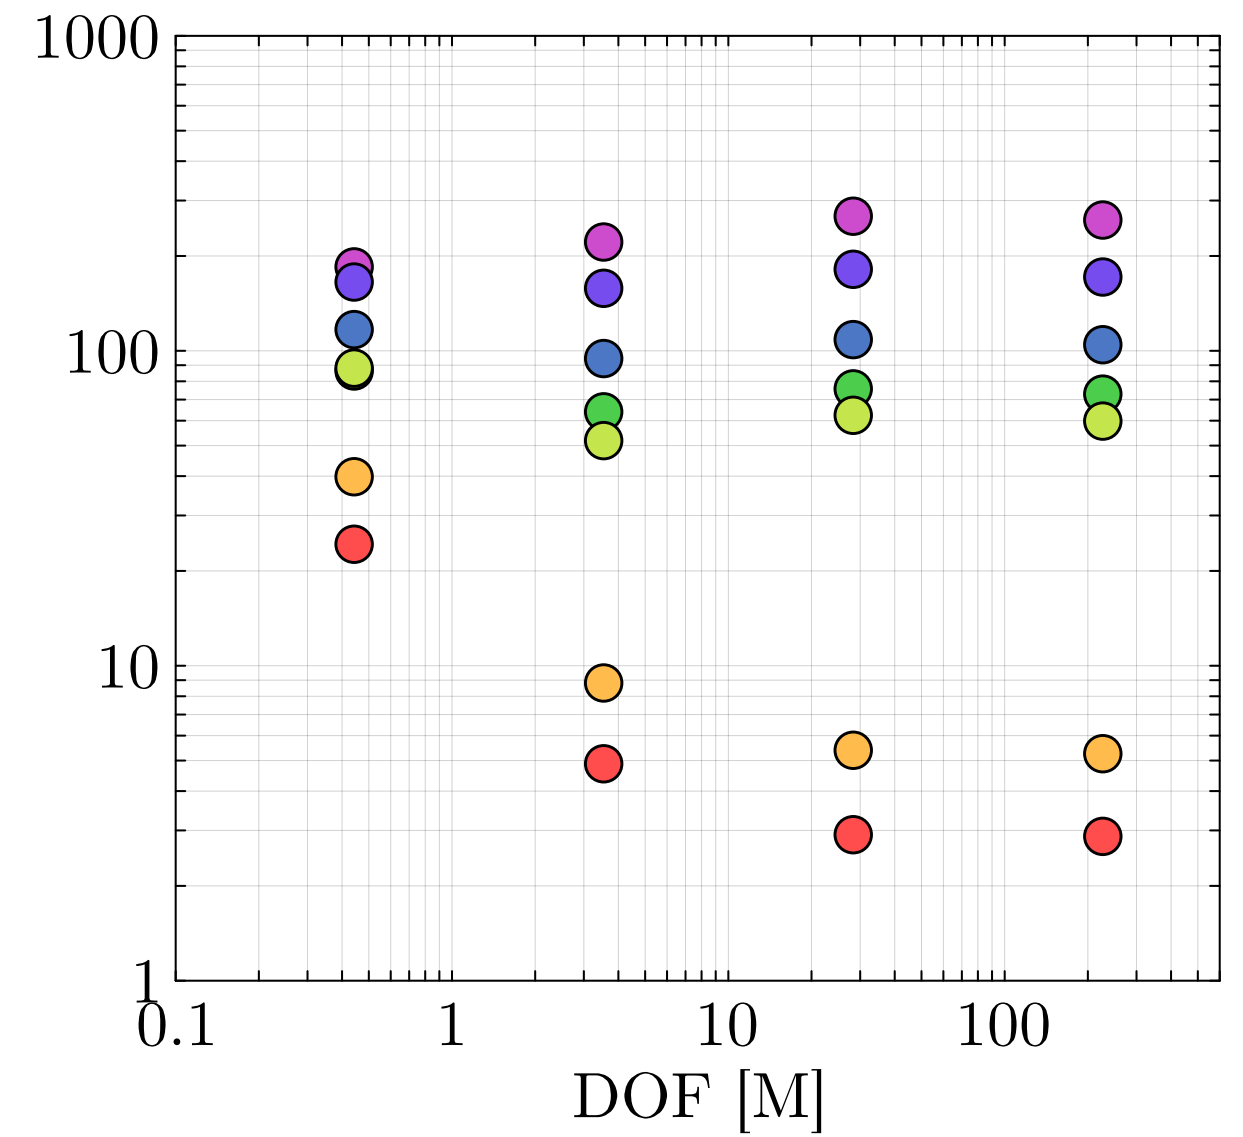
TGV



Sphere

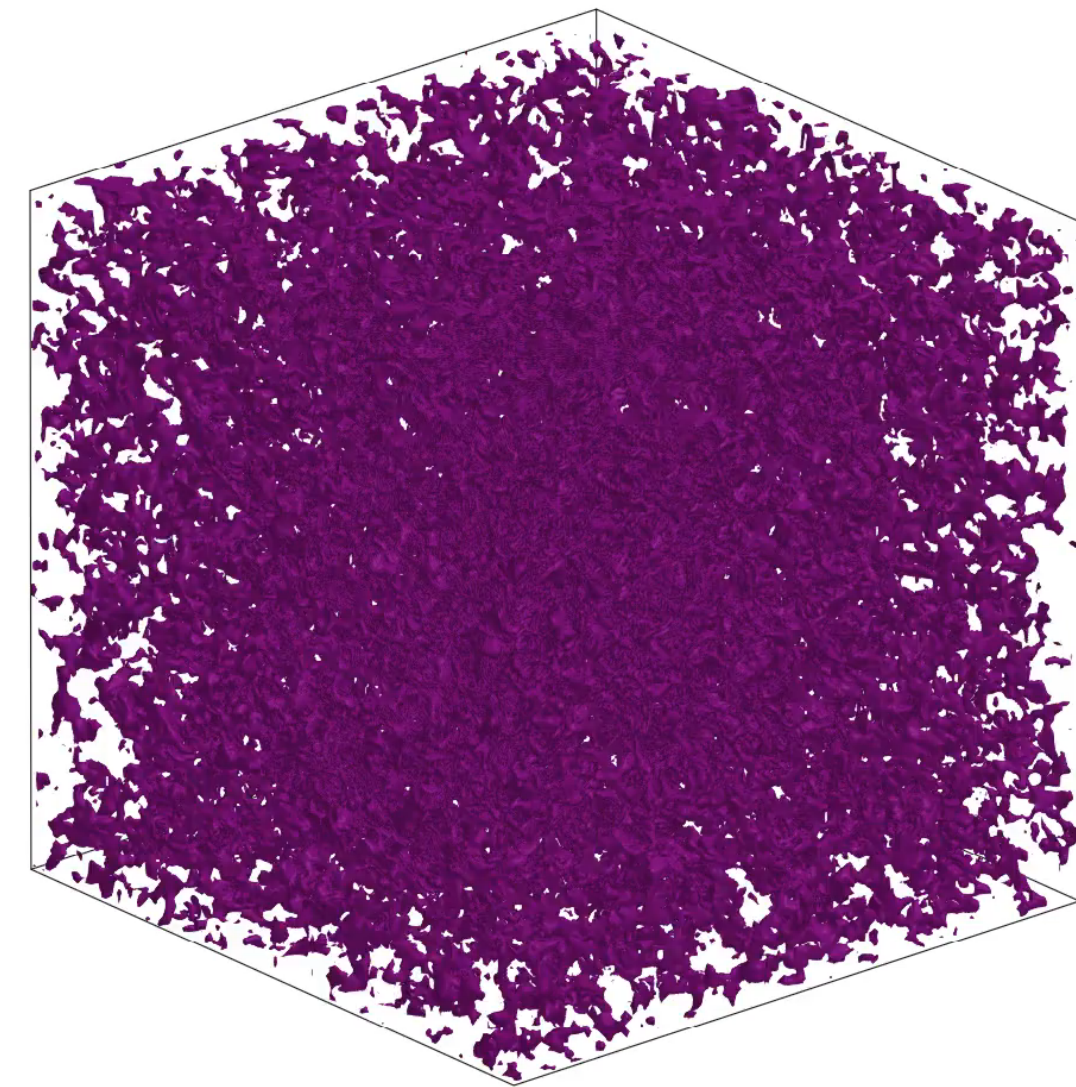
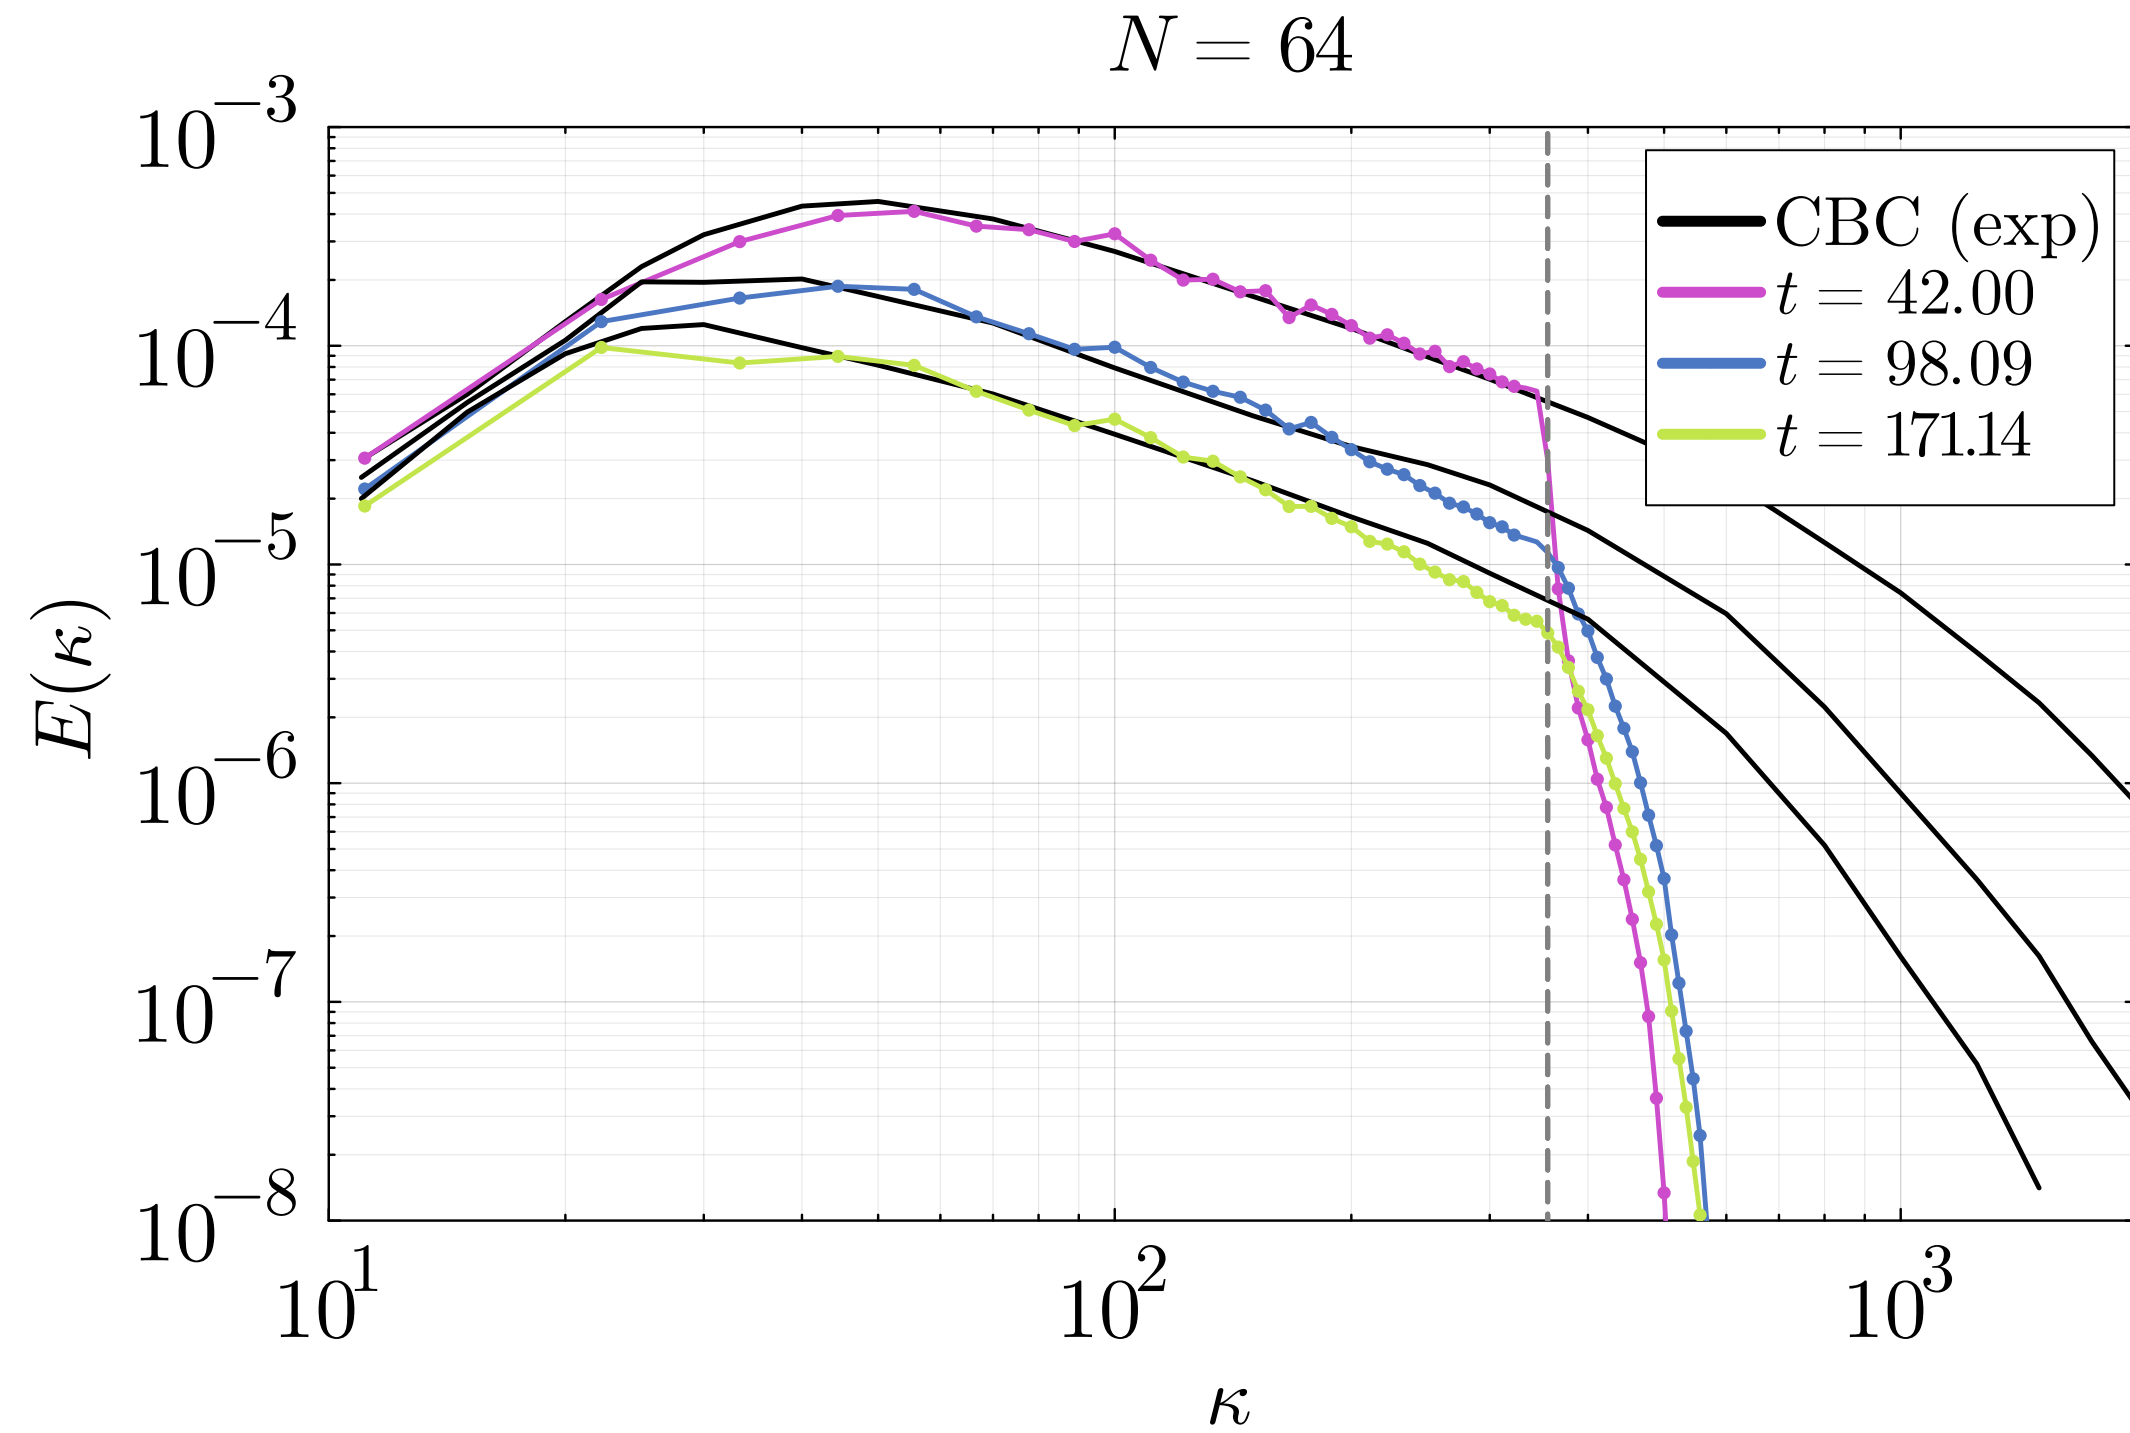


Cylinder



Validation | Homogeneous isotropic turbulence

- $Re = 34k$, $N = 64^3$ (periodic), CDS + Smagorinsky model $C_s = 0.17$



Validation | Sphere

- $Re = 3700$, $D = \{88, 128, 128\}$, $L = (7, \{6, 3\}, \{6, 3\})D$
DOFs = 300M, QUICK (iLES)

