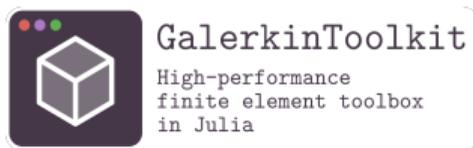


# GalerkinToolkit.jl:

Towards a multi-platform finite element API with deep integration with the Julia package ecosystem

Francesc Verdugo (and collaborators)



NUM-SCARS workshop · 2026-01-23

# The team



## Main developer

- Francesc Verdugo (VU)

## GalerkinToolkit form compiler

- Xiaowei Ouyang (VU)
- Tiziano de Matteis (VU)

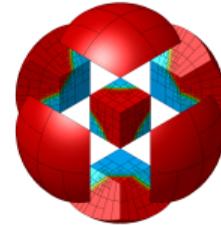
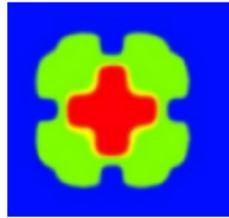


## Software sustainability & GPU

- Robin Richardson (NLeSc)
- Stijn Heldens (NLeSc)
- Alessio Sclocco (NLeSc)

# Why another FEM library?

- Deal.ii
- GalerkinToolkit
- Gridap
- FEniCS
- Firedrake
- MFEM
- ...



# Today's agenda



- Which are the benefits of our API?
- What is new in our code generation strategy?
- Towards GPU support

# Simple FEM in math notation



Let  $f, g : \mathbb{R}^d \rightarrow \mathbb{R}$  be two functions.

Let  $\Omega \subset \mathbb{R}^d$  be a computational domain.

Let  $\mathcal{T}$  be a *triangulation* of  $\Omega$

Let  $V := \{v \in H^1(\Omega) : v(x) \text{ is a polynomial for all } x \in T \in \mathcal{T}\}$

Let  $V_0 := \{v \in V : v(x) = 0 \text{ for } x \in \partial\Omega\}$

Let  $V_g := \{v \in V : v(x) = g(x) \text{ for } x \in \partial\Omega\}$

Find  $u \in V_g$  such that:

$$\int_{\Omega} \nabla u(x) \cdot \nabla v(x) \, dx = \int_{\Omega} f(x)v(x) \, dx \text{ for all } v \in V_0$$

# Implementation with Deal.II



```
1 #include <deal.II/grid/tria.h>
2 #include <deal.II/dofs/dof_handler.h>
3 #include <deal.II/grid/grid_generator.h>
4 #include <deal.II/grid/tria_accessor.h>
5 #include <deal.II/grid/tria_iterator.h>
6 #include <deal.II/dofs/dof_accessor.h>
7 #include <deal.II/fe/fe_q.h>
8 #include <deal.II/dofs/dof_tools.h>
9 #include <deal.II/fe/fe_values.h>
10 #include <deal.II/base/quadature_lib.h>
11 #include <deal.II/base/function.h>
12 #include <deal.II/numerics/vector_tools.h>
13 #include <deal.II/numerics/matrix_tools.h>
14 #include <deal.II/lac/vector.h>
15 #include <deal.II/lac/full_matrix.h>
16 #include <deal.II/lac/sparse_matrix.h>
17 #include <deal.II/lac/dynamic_sparsity_pattern.h>
18 #include <deal.II/lac/solver_cg.h>
19 #include <deal.II/lac/precondition.h>
20 #include <deal.II/numerics/data_out.h>
21 #include <fstream>
22 #include <iostream>
23 using namespace dealii;
24 class Step3
25 {
26 public:
27   Step3();
28   void run();
29 private:
30   void make_grid();
31   void setup_system();
32   void assemble_system();
33   void solve();
34   void output_results() const;
35   Triangulation<2> triangulation;
36   FE_Q<2> fe;
37   DoFHandler<2> dof_handler;
38   SparsityPattern sparsity_pattern;
39   SparseMatrix<double> system_matrix;
40   Vector<double> solution;
41   Vector<double> system_rhs;
42 };
43 Step3::Step3()
44 : fe(1)
45 , dof_handler(triangulation)
46 {}
47 void Step3::make_grid()
48 {
49   GridGenerator::hyper_cube(triangulation, -1, 1);
50   triangulation.refine_global(5);
51   std::cout << "Number of active cells: " << triangulation.n_active_cells()
52             << std::endl;
53 }
54 void Step3::setup_system()
55 {
56   dof_handler.distribute_dofs(fe);
57   std::cout << "Number of degrees of freedom: " << dof_handler.n_dofs()
58             << std::endl;
59   DynamicSparsityPattern dsp(dof_handler.n_dofs());
60   DoFTools::make_sparsity_pattern(dof_handler, dsp);
61   sparsity_pattern.copy_from(dsp);
62   system_matrix.reinit(sparsity_pattern);
63   solution.reinit(dof_handler.n_dofs());
64   system_rhs.reinit(dof_handler.n_dofs());
65 }
66 void Step3::assemble_system()
67 {
68   OGAuss<2> quadrature_formula(fe.degree + 1);
69   FEValues<2> fe_values(fe,
70                       quadrature_formula,
71                       update_values | update_gradients | update_JxW_values);
72   const unsigned int dofs_per_cell = fe.dofs_per_cell;
73   FullMatrix<double> cell_matrix(dofs_per_cell, dofs_per_cell);
74   Vector<double> cell_rhs(dofs_per_cell);
75   std::vector<types::global_dof_index> local_dof_indices(dofs_per_cell);
76   for (const auto &cell : dof_handler.active_cell_iterators())
77   {
78     fe_values.reinit(cell);
79     cell_matrix = 0;
80     cell_rhs = 0;
81     for (const unsigned int q_index : fe_values.quadrature_point_indices())
82     {
83       for (const unsigned int i : fe_values.dof_indices())
84         for (const unsigned int j : fe_values.dof_indices())
85           cell_matrix(i, j) +=
86             (fe_values.shape_grad(i, q_index) * // grad phi_i(x,q)
87              fe_values.shape_grad(j, q_index) * // grad phi_j(x,q)
88              fe_values.JxW(q_index)); // dx
89       for (const unsigned int i : fe_values.dof_indices())
90         cell_rhs(i) += (fe_values.shape_value(1, q_index) * // phi_1(x,q)
91                       1 * // f(x,q)
92                       fe_values.JxW(q_index)); // dx
93     }
94     cell->set_dof_indices(local_dof_indices);
95
96     for (const unsigned int i : fe_values.dof_indices())
97       for (const unsigned int j : fe_values.dof_indices())
98         system_matrix.add(local_dof_indices[i],
99                          local_dof_indices[j],
100                          cell_matrix(i, j));
101     for (const unsigned int i : fe_values.dof_indices())
102       system_rhs(local_dof_indices[i]) += cell_rhs(i);
103   }
104   std::map<types::global_dof_index, double> boundary_values;
105   VectorTools::interpolate_boundary_values(dof_handler,
106                                           0,
107                                           Functions::ZeroFunction<2>(),
108                                           boundary_values);
109   MatrixTools::apply_boundary_values(boundary_values,
110                                       system_matrix,
111                                       solution,
112                                       system_rhs);
113 }
114 void Step3::solve()
115 {
116   SolverControl solver_control(1000, 1e-12);
117   SolverCG<Vector<double>> solver(solver_control);
118   solver.solve(system_matrix, solution, system_rhs, PreconditionIdentity());
119 }
120 void Step3::output_results() const
121 {
122   DataOut<2> data_out;
123   data_out.attach_dof_handler(dof_handler);
124   data_out.add_data_vector(solution, "solution");
125   data_out.build_patches();
126   std::ofstream output("solution.vtk");
127   data_out.write_vtk(output);
128 }
129 void Step3::run()
130 {
131   make_grid();
132   setup_system();
133   assemble_system();
134   solve();
135   output_results();
136 }
137 int main()
138 {
139   deallog.depth_console(2);
140   Step3 laplace_problem;
141   laplace_problem.run();
142   return 0;
143 }
```



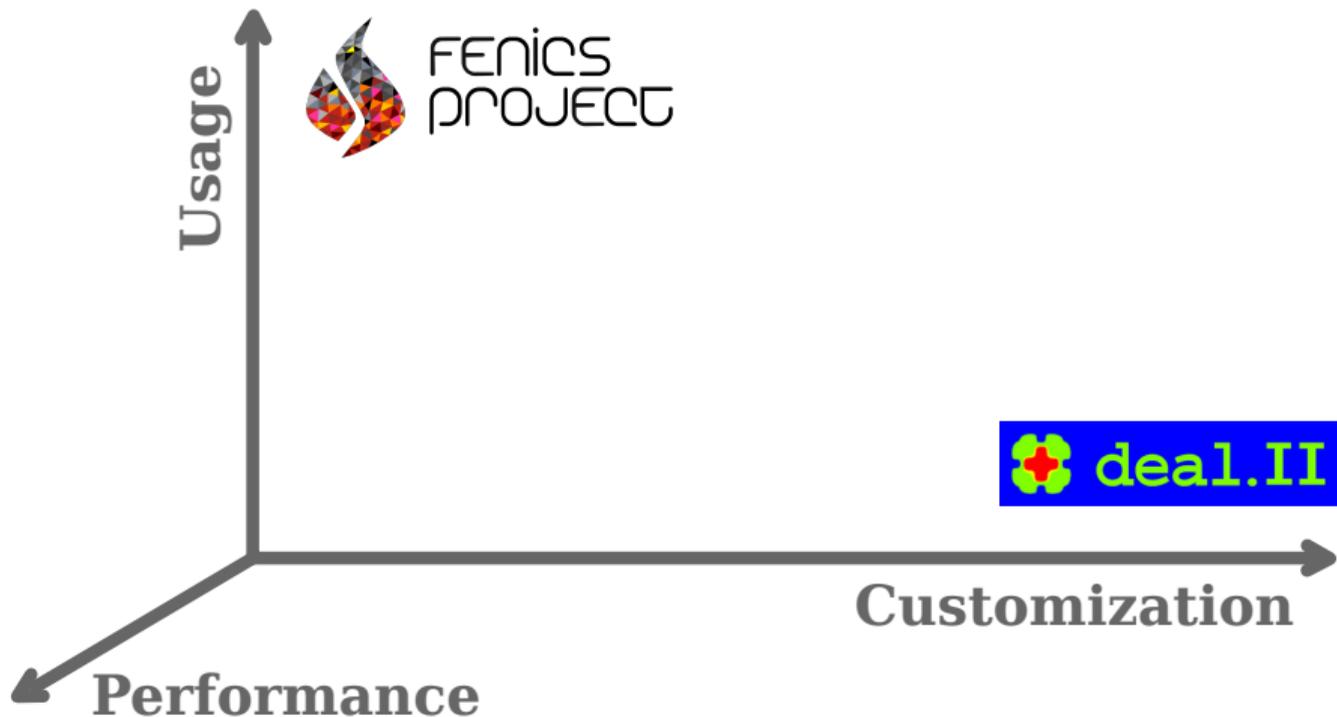
# Implementation with FEniCS

```
1 from fenics import *
2 mesh = UnitSquareMesh(8, 8)
3 V = FunctionSpace(mesh, 'P', 1)
4 u_D = Expression('1 + x[0]*x[0] + 2*x[1]*x[1]', degree=2)
5 def boundary(x, on_boundary):
6     return on_boundary
7 bc = DirichletBC(V, u_D, boundary)
8 u = TrialFunction(V)
9 v = TestFunction(V)
10 f = Constant(-6.0)
11 a = dot(grad(u), grad(v))*dx
12 L = f*v*dx
13 u = Function(V)
14 solve(a == L, u, bc)
15 vtkfile = File('poisson/solution.pvd')
16 vtkfile << u
```



FENICS  
PROJECT

# API trade-offs



# Implementation with GalerkinToolkit



```
1 using LinearAlgebra
2 import Gmsh
3 import GalerkinToolkit as GT
4 import LinearSolve
5 import ForwardDiff
6 import GLMakie
7 mesh = GT.mesh_from_msh("model.msh")
8 Ω = GT.interior(mesh)
9 Γd = GT.boundary(mesh)
10 f = GT.analytical_field(x->1.0,Ω)
11 g = GT.analytical_field(x->2.0,Ω)
12 k = 1
13 V = GT.lagrange_space(Ω,k;dirichlet_boundary=Γd)

14 T = Float64
15 uhd = GT.zero_dirichlet_field(T,V)
16 GT.interpolate_dirichlet!(g,uhd)
17 degree = 2*k
18 dΩ = GT.measure(Ω,degree)
19 ∇ = ForwardDiff.gradient
20 a = (u,v) -> GT.∫( x->∇(u,x)·∇(v,x), dΩ)
21 l = v -> GT.∫( x->v(x)*f(x), dΩ)
22 p = GT.SciMLBase_LinearProblem(uhd,a,l)
23 sol = LinearSolve.solve(p)
24 uh = GT.solution_field(uhd,sol)
25 GT.makie_surfaces(Ω;color=uh)
```

A small version of the GalerkinToolkit logo, which is a dark square with a white cube icon, is positioned in the center of the slide between the two columns of code.

```
20 a = (u,v) -> GT.∫( x->∇(u,x)·∇(v,x), dΩ)
21 l = v -> GT.∫( x->v(x)*f(x), dΩ)
22 p = GT.SciMLBase_LinearProblem(uhd,a,l)
```

# Implementation with GalerkinToolkit (v2)



```
1 using LinearAlgebra
2 import Gmsh
3 import GalerkinToolkit as GT
4 import LinearSolve
5 import ForwardDiff
6 import GLMakie as Makie
7 mesh = GT.mesh_from_msh("model.msh")
8 Ω = GT.interior(mesh)
9 Γd = GT.boundary(mesh)
10 f = GT.analytical_field(x->1.0,Ω)
11 g = GT.analytical_field(x->2.0,Ω)
12 k = 1
13 V = GT.lagrange_space(Ω,k;dirichlet_boundary=Γd)
14 T = Float64
15 uhd = GT.zero_dirichlet_field(T,V)
16 GT.interpolate_dirichlet!(g,uhd)
17 degree = 2*k
18 dΩ = GT.measure(Ω,degree)
19 A_alloc = GT.allocate_matrix(T,V,V,Ω)
20 free_or_dirichlet=(GT.FREE,GT.DIRICHLET)
21 Ad_alloc = GT.allocate_matrix(T,V,V,Ω;free_or_dirichlet)
22 b_alloc = GT.allocate_vector(T,V,Ω,Γn)
23 function assemble_in_Ω!(A_alloc,Ad_alloc,b_alloc,V,f,dΩ)
24     tabulate = (GT.value,∇)
25     compute = (GT.coordinate,)
26     V_faces = GT.each_face(V,dΩ;tabulate,compute)
27     n = GT.max_num_reference_dofs(V)
28     T = Float64
29     Auu = zeros(T,n,n)
30     bu = zeros(T,n)
31     for V_face in V_faces
32         dofs = GT.dofs(V_face)
33         fill!(Auu,zero(T))
34         fill!(bu,zero(T))
35         for V_point in GT.each_point(V_face)
36             x = GT.coordinate(V_point)
37             dV = GT.weight(V_point)
38             dof_s = GT.shape_functions(GT.value,V_point)
39             dof_Vs = GT.shape_functions(∇,V_point)
40             for (i,dofi) in enumerate(dofs)
41                 v = dof_s[i]
42                 ∇v = dof_Vs[i]
43                 bu[i] += f.definition(x)*v*dV
44                 for (j,dofj) in enumerate(dofs)
45                     ∇u = dof_Vs[j]
46                     Auu[i,j] += ∇v·∇u*dV
47                 end
48             end
49         end
50         GT.contribute!(A_alloc,Auu,dofs,dofs)
51         GT.contribute!(Ad_alloc,Auu,dofs,dofs)
52         GT.contribute!(b_alloc,bu,dofs)
53     end
54 end
55 assemble_in_Ω!(A_alloc,Ad_alloc,b_alloc,V,f,dΩ)
56 A = GT.compress(A_alloc)
57 Ad = GT.compress(Ad_alloc)
58 b = GT.compress(b_alloc)
59 xd = GT.dirichlet_values(uhd)
60 b .= b .- Ad*xd
61 p = LinearSolve.LinearProblem(A,b)
62 sol = LinearSolve.solve(p)
63 uh = GT.solution_field(uhd,sol)
64 GT.makie_surfaces(Ω;color=uh)
```



# Explicit assembly loop in GalerkinToolkit



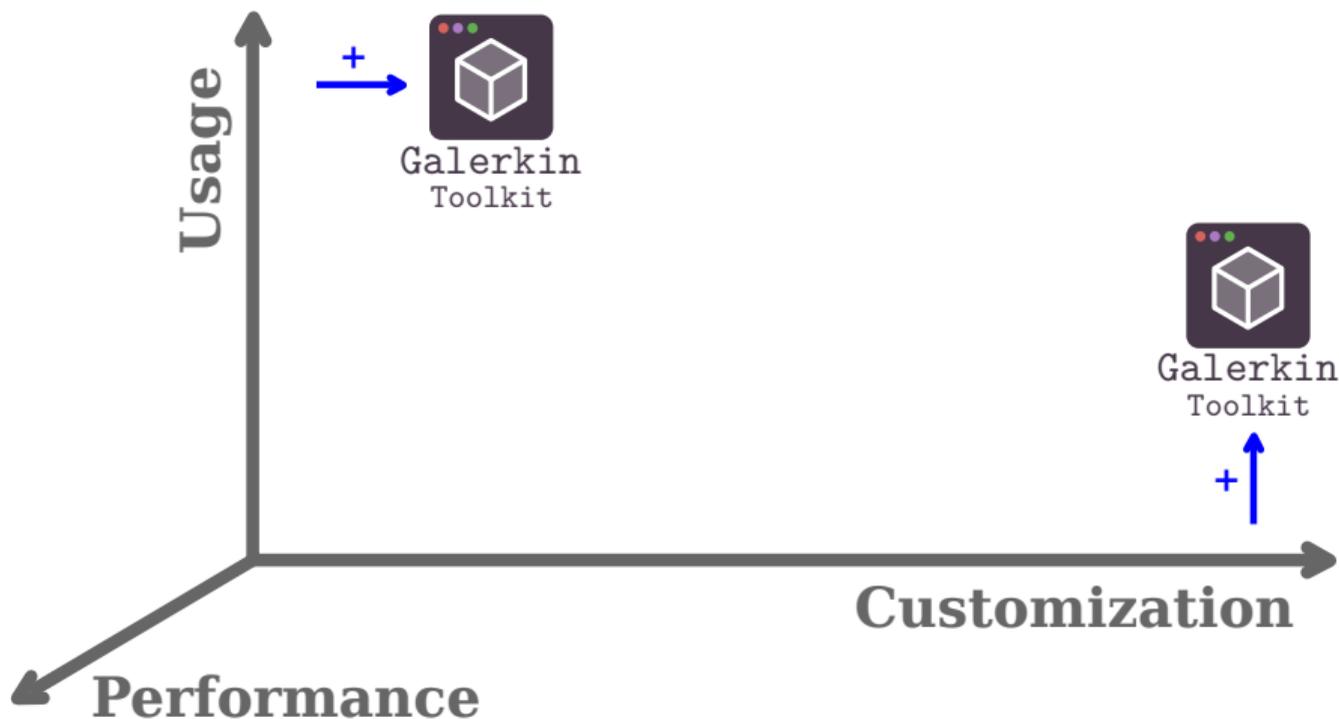
```
23 function assemble_in_Ω!(A_alloc,Ad_alloc,b_alloc,V,f,dΩ) 40
24     tabulate = (GT.value,∇) 41
25     compute = (GT.coordinate,) 42
26     V_faces = GT.each_face(V,dΩ;tabulate,compute) 43
27     n = GT.max_num_reference_dofs(V) 44
28     T = Float64 45
29     Auu = zeros(T,n,n) 46
30     bu = zeros(T,n) 47
31     for V_face in V_faces 48
32         dofs = GT.dofs(V_face) 49
33         fill!(Auu,zero(T)) 50
34         fill!(bu,zero(T)) 51
35         for V_point in GT.each_point(V_face) 52
36             x = GT.coordinate(V_point) 53
37             dV = GT.weight(V_point) 54
38             dof_s = GT.shape_functions(GT.value,V_point)
39             dof_∇s = GT.shape_functions(∇,V_point)
```



```
        for (i,dofi) in enumerate(dofs)
            v = dof_s[i]
            ∇v = dof_∇s[i]
            bu[i] += f.definition(x)*v*dV
            for (j,dofj) in enumerate(dofs)
                ∇u = dof_∇s[j]
                Auu[i,j] += ∇v·∇u*dV
            end
        end
    end
end
GT.contribute!(A_alloc,Auu,dofs,dofs)
GT.contribute!(Ad_alloc,Auu,dofs,dofs)
GT.contribute!(b_alloc,bu,dofs)
end
end
```

```
26     V_faces = GT.each_face(V,dΩ;tabulate,compute)
35         for V_point in GT.each_point(V_face)
```

# Enhanced API trade-offs



# GalerkinToolkit form compiler (GTFC)



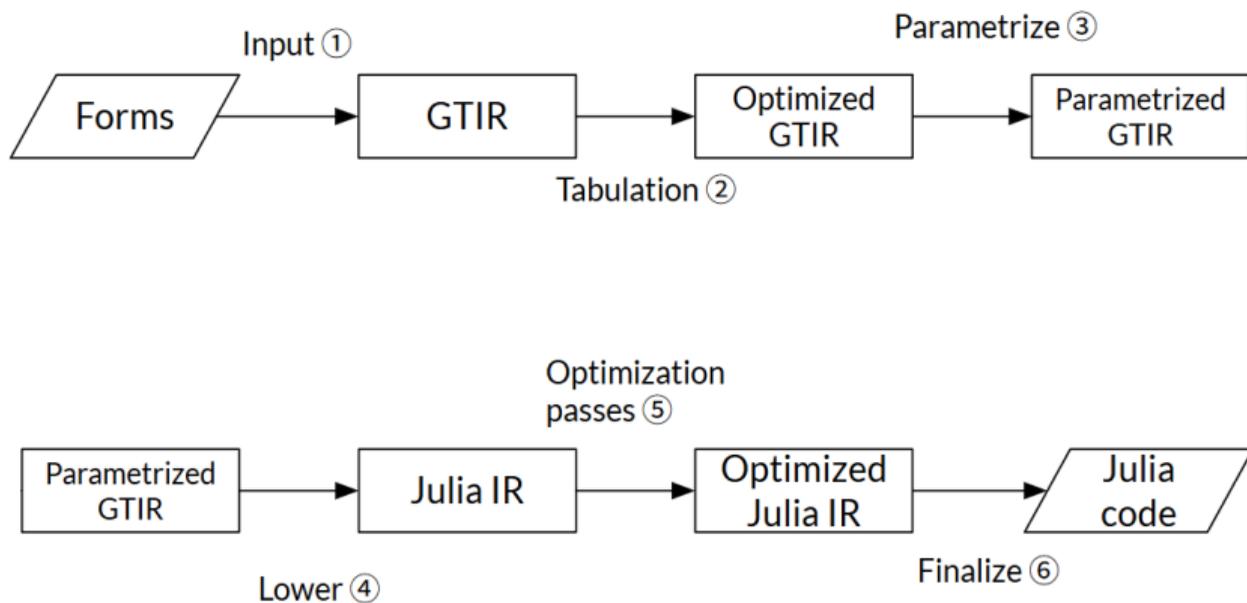
$GT.\int( x \rightarrow \nabla(u,x) \cdot \nabla(v,x), d\Omega)$



```
function assemble_in_Ω!(A_alloc,V,dΩ)
  n = GT.max_num_reference_dofs(V)
  Auu = zeros(Float64,n,n)
  for V_face in GT.each_face(V,dΩ;tabulate=(∇,))
    dofs = GT.dofs(V_face)
    fill!(Auu,zero(Float64))
    for V_point in GT.each_point(V_face)
      dV = GT.weight(V_point)
      dof_∇s = GT.shape_functions(∇,V_point)
      for (i,dofi) in enumerate(dofs)
        ∇v = dof_∇s[i]
        for (j,dofj) in enumerate(dofs)
          ∇u = dof_∇s[j]
          Auu[i,j] += ∇v·∇u*dV
        end
      end
    end
  end
  GT.contribute!(A_alloc,Auu,dofs,dofs)
end
```

\* The actual generated code is slightly different

# GalerkinToolkit form compiler (GTFC)



# User-defined functions and types

Let  $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^d$  be a user-defined flux

The weak form transforms into

$$\int_{\Omega} \sigma(\nabla u(x)) \cdot \nabla v(x) \, dx = \int_{\Omega} v(x) \, dx$$

```
function  $\sigma(\nabla u)$ 
    # You can implement here
    # anything that
    # returns a vector of
    # the same length as  $\nabla u$ 
end
GT. $\int(d\Omega)$  do x
     $\sigma\_u = \text{GT.external}(\sigma, \nabla(u, x))$ 
     $\sigma\_u \cdot \nabla(v, x)$ 
end
```

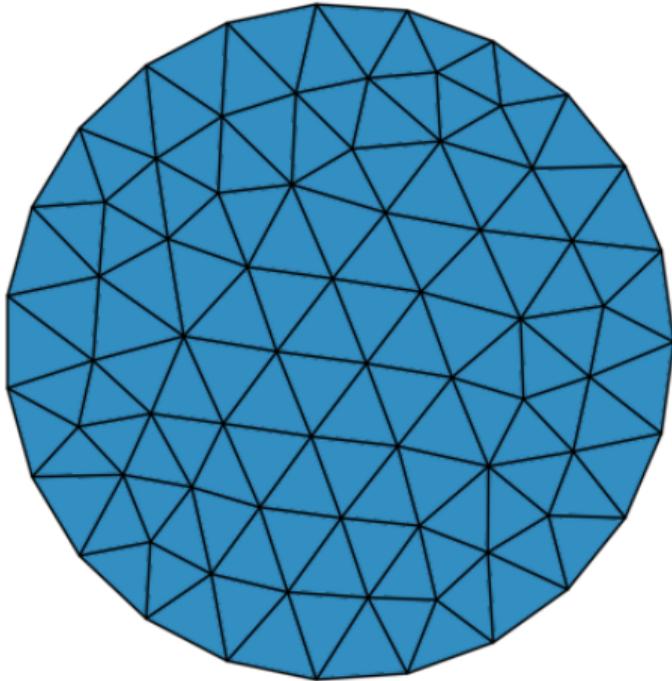
# User-defined functions and types

```
function assemble_in_Ω!(A_alloc,V,dΩ)
  n = GT.max_num_reference_dofs(V)
  Auu = zeros(Float64,n,n)
  σ_u = # allocate σ_u
  for V_face in GT.each_face(V,dΩ;tabulate=(∇,))
    dofs = GT.dofs(V_face)
    fill!(Auu,zero(Float64))
    for V_point in GT.each_point(V_face)
      dV = GT.weight(V_point)
      dof_∇s = GT.shape_functions(∇,V_point)
      for (j,dofj) in enumerate(dofs)
        ∇u = dof_∇s[j]
        σ_u[j] = σ(∇u)
      end
      for (i,dofi) in enumerate(dofs)
        ∇v = dof_∇s[i]
        for (j,dofj) in enumerate(dofs)
          Auu[i,j] += σ_u[j]·∇u*dV
        end
      end
    end
  end
  GT.contribute!(A_alloc,Auu,dofs,dofs)
end
```

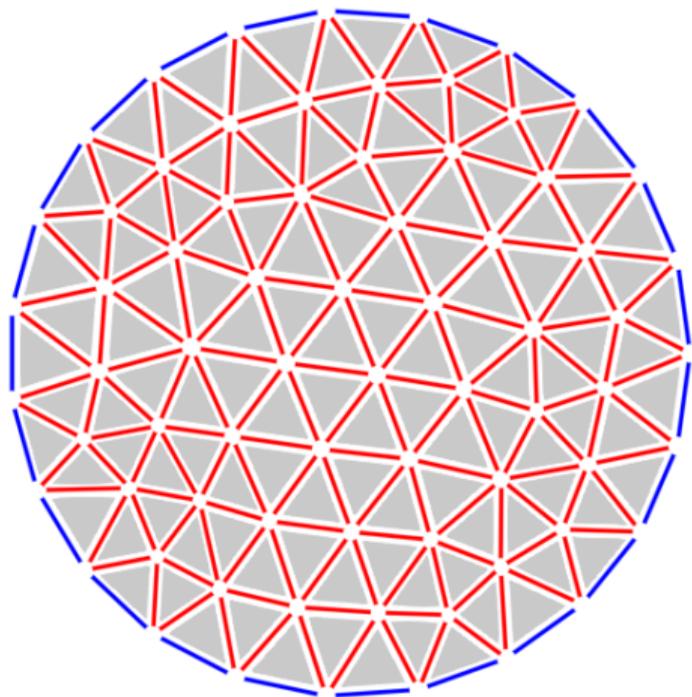
```
dof_∇s = GT.shape_functions(∇,V_point)
for (j,dofj) in enumerate(dofs)
  ∇u = dof_∇s[j]
  σ_u[j] = σ(∇u) ←
end
for (i,dofi) in enumerate(dofs)
  ∇v = dof_∇s[i]
  for (j,dofj) in enumerate(dofs)
    Auu[i,j] += σ_u[j]·∇u*dV
  end
end
```

\* The actual generated code is slightly different

# Polyhedral complexes



# Polyhedral complexes



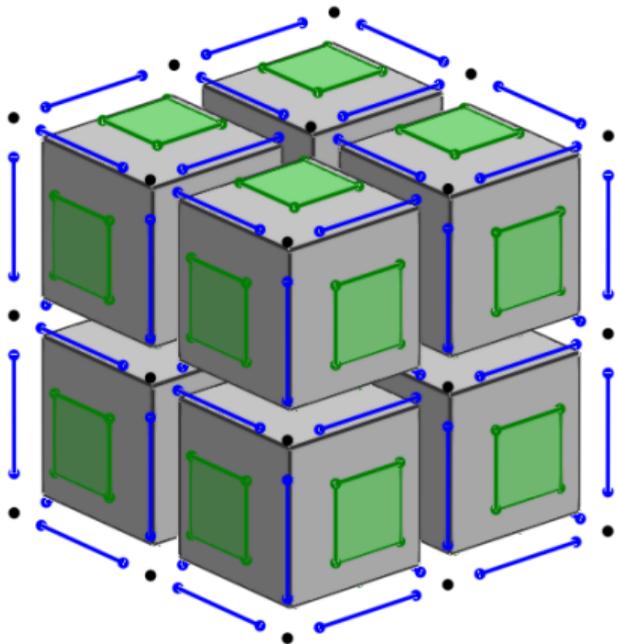
$$GT . \int (x \rightarrow v(x), d\Omega)$$

$$GT . \int (x \rightarrow v(x), d\Gamma)$$

$$GT . \int (x \rightarrow v[1](x) - v[2](x), d\Lambda)$$

$$GT . \int (x \rightarrow v(x) + q(x), d\Gamma)$$

# Polyhedral complexes



```
# 3D
shrink = 0.8
linewidth = 3
GT.makie_surfaces!(ax,mesh;shrink,dim=3,color="gray90")
GT.makie_edges!(ax,mesh;shrink,dim=3,color="gray20",linewidth)

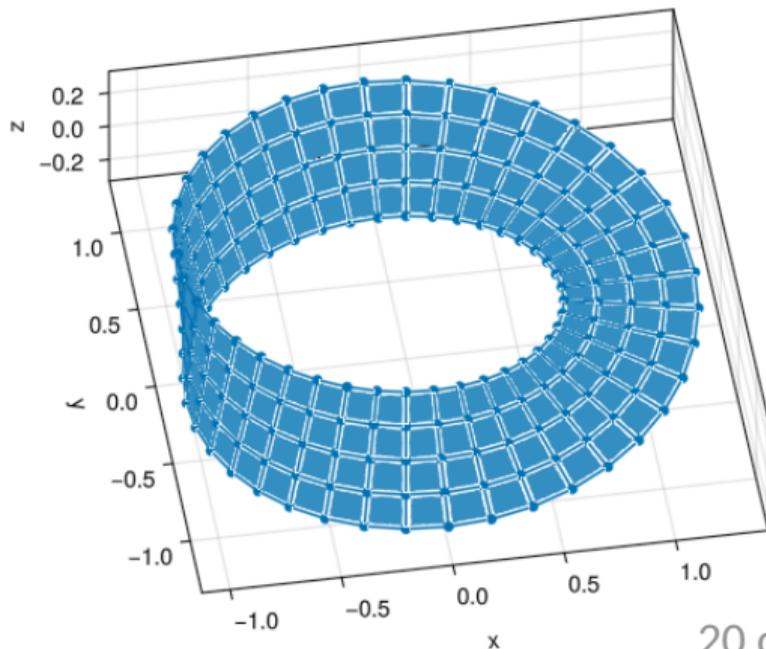
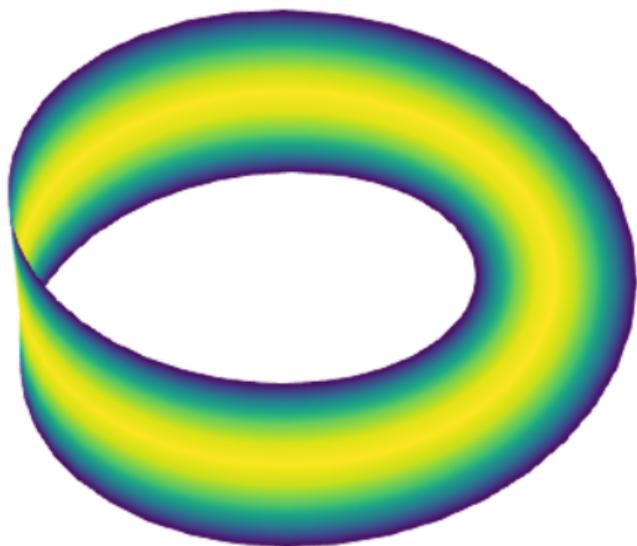
# 2D
shrink = 0.4
markersize = 12
GT.makie_surfaces!(ax,mesh;shrink,dim=2,color="palegreen")
GT.makie_edges!(ax,mesh;shrink,dim=2,color="green",linewidth)
GT.makie_vertices!(ax,mesh;shrink,dim=2,color="green",markersize)

# 1D
shrink = 0.6
GT.makie_edges!(ax,mesh;shrink,dim=1,color="blue",linewidth)
GT.makie_vertices!(ax,mesh;shrink,dim=1,color="blue",markersize)

# 0D
GT.makie_vertices!(ax,mesh;dim=0,color="black",markersize)
```

# Manifolds

Example: Laplace-Beltrami on a Möbius strip



# Parametrizable code generation



Example: Residual of a non-linear problem ( $p$ -Laplacian).

$$r(v) = \int_{\Omega} \nabla v \cdot (|\nabla u_h|^{p-2} \nabla u_h) \, d\Omega$$

```
uh = # Define your FE field
flux(∇u) = norm(∇u)^(p-2) * ∇u
r = v -> GT.∫( x-> ∇(v,x) · GT.external(flux,∇(uh,x)), dΩ)
# Assembly (code gets generated here)
b,cache = GT.assemble_vector(r,Float64,V;parameters=(uh,))
```

```
uh2 = # Another FE field
# Reuse the code generated above
GT.update_vector!(b,cache;parameters=(uh2,))
```

# Building SciML problems



## # Linear problem

```
a = (u,v) -> GT.∫( x->∇(u,x)·∇(v,x), dΩ)
l = v -> GT.∫( x->v(x)*f(x), dΩ)
p = GT.SciMLBase_LinearProblem(uhd,a,l)
sol = LinearSolve.solve(p)
uh = GT.solution_field(uhd,sol)
```

## # Nonlinear problem

```
r = u -> v -> GT.∫( x-> ∇(v,x)·GT.external(flux,∇(u,x)),dΩ)
j = u -> (du,v) -> GT.∫( x-> ∇(v,x)·GT.external(dflux,∇(du,x),∇(u,x)),dΩ)
p = GT.SciMLBase_NonlinearProblem(uh,r,j)
sol = NonlinearSolve.solve(p)
uh = GT.solution_field(uh,sol)
```

# Building SciML problems



```
# ODE Problem
```

```
m = (u,v) -> GT.∫(x->C*v(x)*u(x),dΩ)
```

```
r = (uh,t) -> v -> -1*GT.∫(x->∇(uh,x)·∇(v,x),dΩ)
```

```
j = (uh,t) -> (u,v) -> -1*GT.∫(x->∇(u,x)·∇(v,x),dΩ)
```

```
tspan = (0.0,T)
```

```
problem = GT.SciMLBase_ODEProblem(tspan,uh,m,r,j;dirichlet_dynamics!)
```

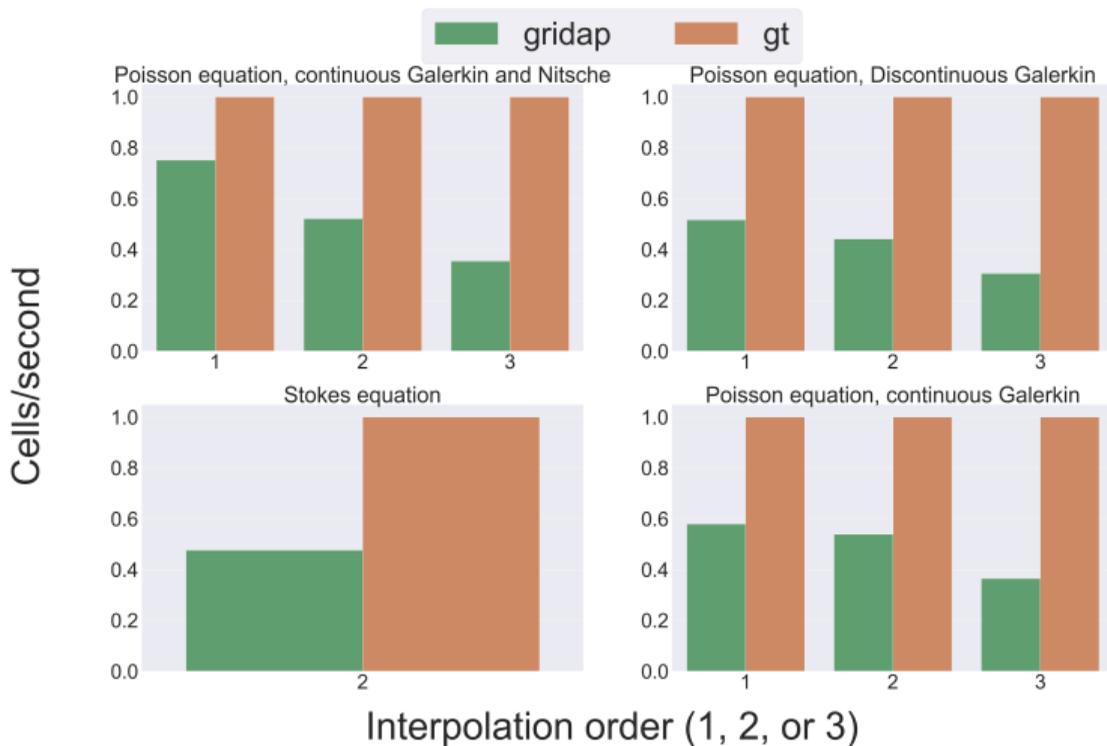
```
integrator = DifferentialEquations.init(problem)
```

```
for integrator in integrator
```

```
    uh = GT.solution_field(integrator)
```

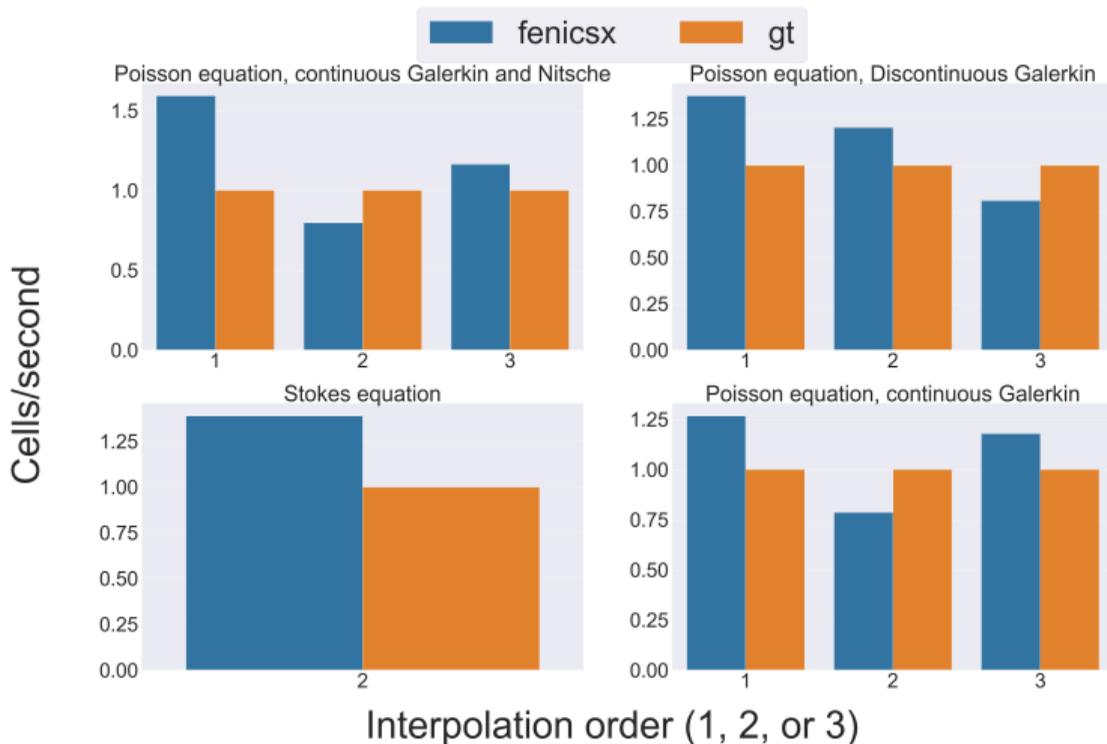
```
end
```

# Performance: Gridap vs GalerkinToolkit



Higher is better. Results for unstructured hexahedral meshes. Internal git branch.

# Performance: FEniCSx vs GalerkinToolkit



Higher is better. Results for unstructured hexahedral meshes. Internal git branch. 25 of 30

# Porting low-level API to GPUs

Example: Integrate a user-defined function on the CPU

$$\int_{\Omega} f(x) dx$$

```
Ω = # define domain here
f(x) = 2*sin(sum(x))

degree = 5
dΩ = GT.quadrature(Ω,degree)
s = 0.0
for dΩ_face in GT.each_face(dΩ)
    for dΩ_point in GT.each_point(dΩ_face)
        x = GT.coordinate(dΩ_point)
        dx = GT.weight(dΩ_point)
        s += f(x)*dx
    end
end
```

# Porting low-level API to GPUs



Example: Integrate a user-defined function on the GPU

```
function kernel!(contributions,dΩ_faces_gpu)
    face_id = (blockIdx().x-1)*blockDim().x +
              threadIdx().x
    if face_id > length(dΩ_faces_gpu)
        return nothing
    end
    dΩ_face = dΩ_faces_gpu[face_id]
    s = 0.0
    for dΩ_point in GT.each_point(dΩ_face)
        x = GT.coordinate(dΩ_point)
        dx = GT.weight(dΩ_point)
        s += f(x)*dx
    end
    contributions[face_id] = s
    return nothing
end
```

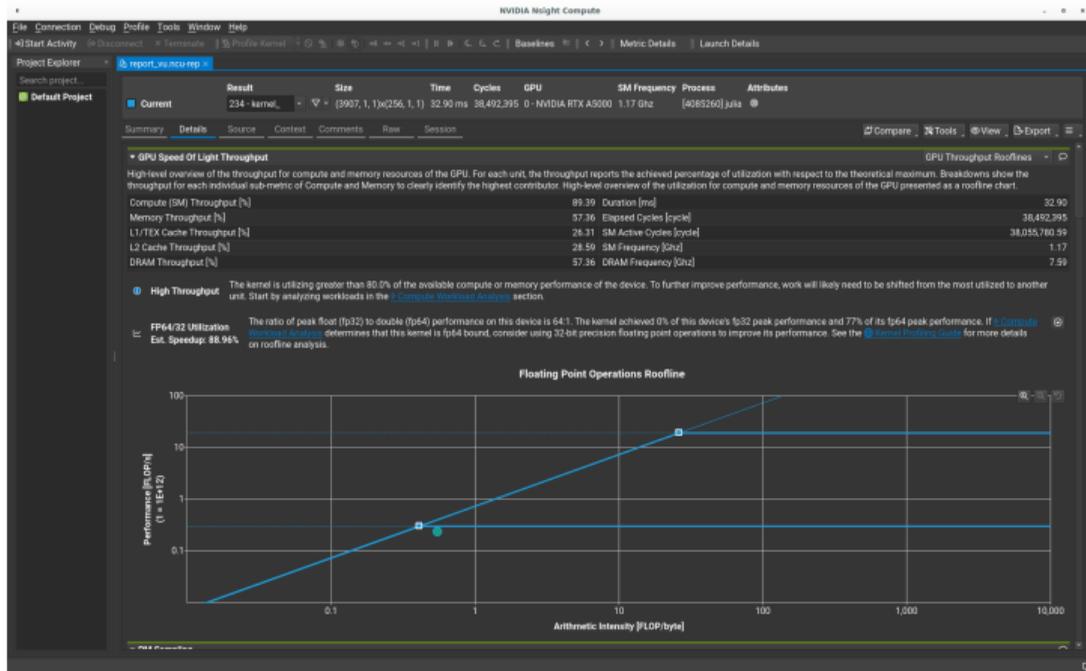
```
dΩ_faces_cpu = GT.each_face(dΩ)
dΩ_faces_gpu = CUDA.cu(dΩ_faces_cpu)
nfaces = length(dΩ_faces_gpu)
contributions = CUDA.zeros(Float64,nfaces)
threads_in_block = 256
blocks_in_grid = ceil{Int, nfaces/256}
@cuda threads=threads_in_block
        blocks=blocks_in_grid
        kernel!(contributions,dΩ_faces_gpu)
s = sum(contributions)
```

Internal git branch.

# Porting low-level API to GPUs



## Example: Integrate a user-defined function on the GPU



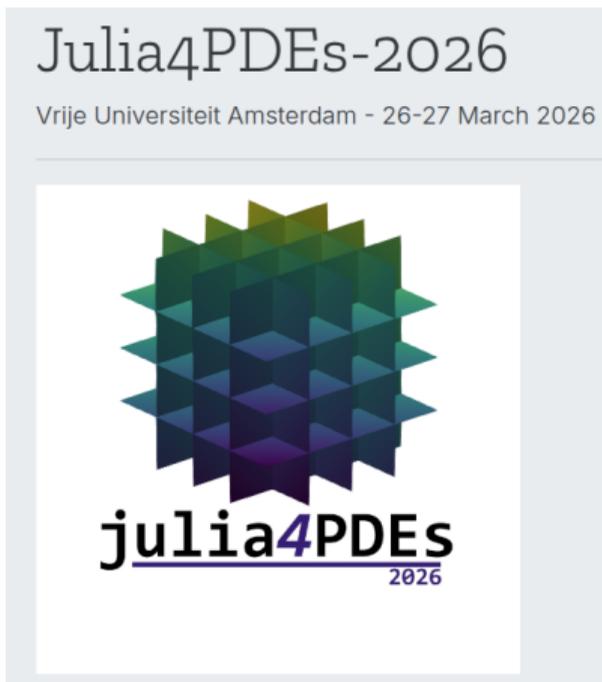
# Summary

## GalerkinToolkit

- Multi-purpose
- Modern
- Easy to use / customizable / performant
- Inter-operable
- Towards multi-platform
- It is FOSS!



Registration and call for abstract are open!



Grant ID: NLESC.SS.2023.008

Co-organized with O. Colomes, A. Palha, R. Richardson, and D. Toshniwal