

High-performance numerical simulation on distributed heterogeneous hardware platforms

Matthias Möller

Delft University of Technology, Department of Applied Mathematics

8th July 2019, SURFsara, Amsterdam


Overview

- Design-Through-Analysis
- Overview of Spline technologies
- DTA for twin-screw compressor
- Implementation aspects
- Recent updates and collaboration plans

Motivation

DESIGN-THROUGH-ANALYSIS

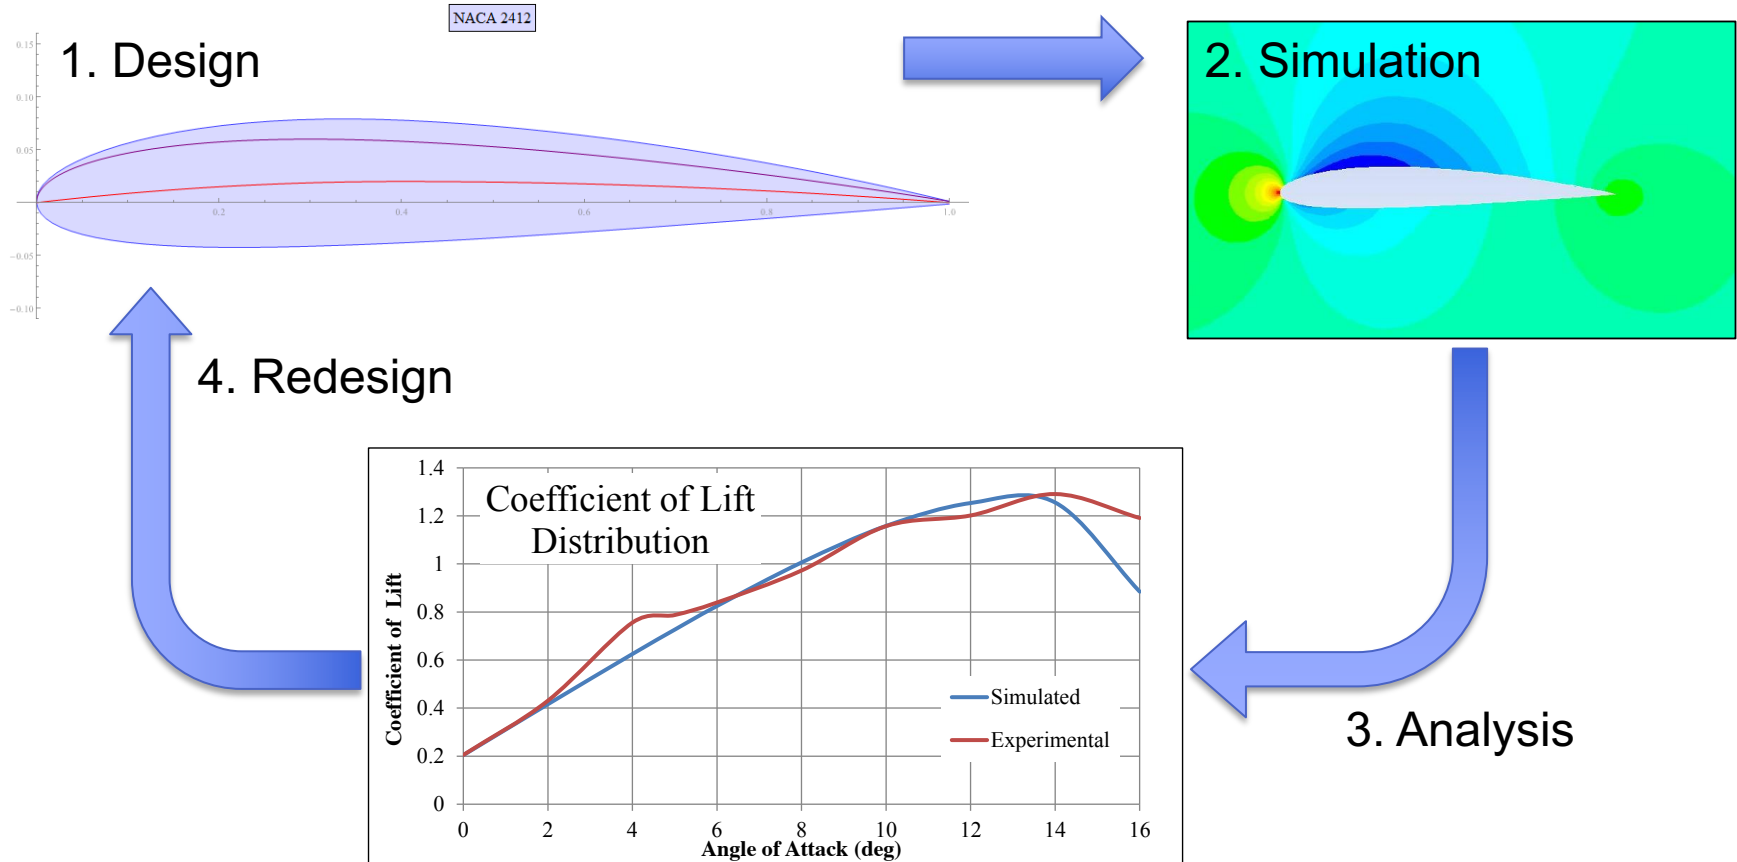
Example: Airfoil design




CFInotebook.net



Design-through-analysis cycle



DTA: 1. Design $D(\mathbf{p})$



- Design parameters

$$\mathbf{p} = (p_1, \dots, p_{12})$$

- Admissible design space

$$\mathcal{S} = [p_1^{min}, p_1^{max}] \times \dots \times [p_{12}^{min}, p_{12}^{max}]$$

DTA: 2. Simulation

- Mathematical model



Navier-Stokes Equations

3 - dimensional - unsteady

Glenn
Research
Center

Coordinates: (x,y,z)	Time: t	Pressure: p	Heat Flux: q
Velocity Components: (u,v,w)	Density: ρ	Stress: τ	Reynolds Number: Re
	Total Energy: Et		Prandtl Number: Pr

Continuity:
$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u)}{\partial x} + \frac{\partial(\rho v)}{\partial y} + \frac{\partial(\rho w)}{\partial z} = 0$$

X - Momentum:
$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial(\rho u^2)}{\partial x} + \frac{\partial(\rho uv)}{\partial y} + \frac{\partial(\rho uw)}{\partial z} = -\frac{\partial p}{\partial x} + \frac{1}{Re_r} \left[\frac{\partial \tau_{xx}}{\partial x} + \frac{\partial \tau_{xy}}{\partial y} + \frac{\partial \tau_{xz}}{\partial z} \right]$$


Y - Momentum:
$$\frac{\partial(\rho v)}{\partial t} + \frac{\partial(\rho uv)}{\partial x} + \frac{\partial(\rho v^2)}{\partial y} + \frac{\partial(\rho vw)}{\partial z} = -\frac{\partial p}{\partial y} + \frac{1}{Re_r} \left[\frac{\partial \tau_{xy}}{\partial x} + \frac{\partial \tau_{yy}}{\partial y} + \frac{\partial \tau_{yz}}{\partial z} \right]$$

Z - Momentum:
$$\frac{\partial(\rho w)}{\partial t} + \frac{\partial(\rho uw)}{\partial x} + \frac{\partial(\rho vw)}{\partial y} + \frac{\partial(\rho w^2)}{\partial z} = -\frac{\partial p}{\partial z} + \frac{1}{Re_r} \left[\frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial \tau_{zz}}{\partial z} \right]$$

Energy:
$$\frac{\partial(E_T)}{\partial t} + \frac{\partial(uE_T)}{\partial x} + \frac{\partial(vE_T)}{\partial y} + \frac{\partial(wE_T)}{\partial z} = -\frac{\partial(u p)}{\partial x} - \frac{\partial(v p)}{\partial y} - \frac{\partial(w p)}{\partial z} - \frac{1}{Re_r Pr_r} \left[\frac{\partial q_x}{\partial x} + \frac{\partial q_y}{\partial y} + \frac{\partial q_z}{\partial z} \right] + \frac{1}{Re_r} \left[\frac{\partial}{\partial x} (u \tau_{xx} + v \tau_{xy} + w \tau_{xz}) + \frac{\partial}{\partial y} (u \tau_{xy} + v \tau_{yy} + w \tau_{yz}) + \frac{\partial}{\partial z} (u \tau_{xz} + v \tau_{yz} + w \tau_{zz}) \right]$$

- Solution for one particular design and one particular angle of attach

$$U = U(D(\mathbf{p}); AoA)$$


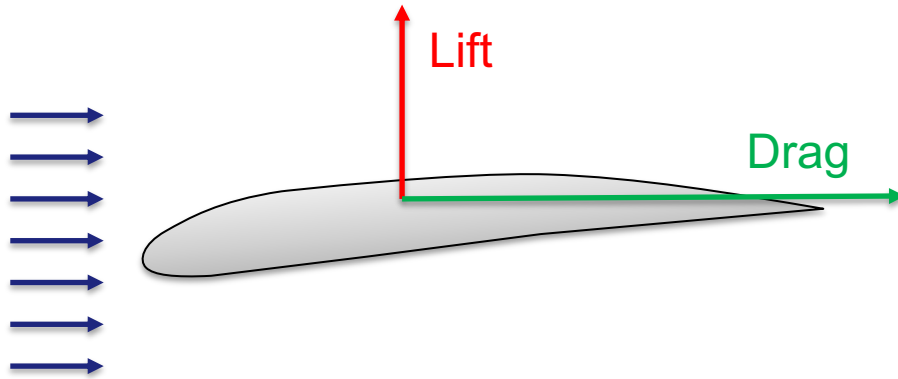


Linné FLOW Centre and SeRC, KTH, Sweden

DTA: 3. Analysis

- Cost functional

$$C(U; D)$$




Example: Operation conditions



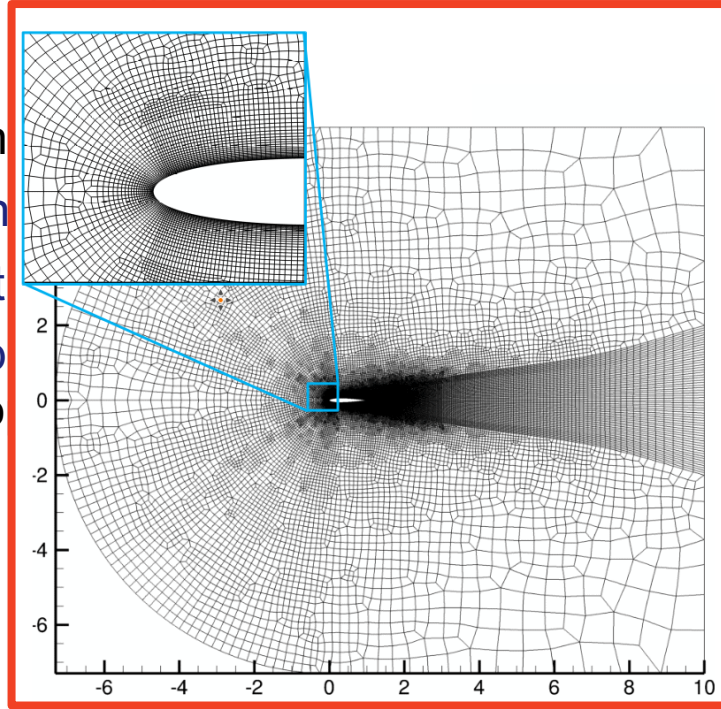
Design-through-analysis

1. Find a set of admissible **design parameters** \mathbf{p} and generate the **design** $D(\mathbf{p})$
2. Compute **solutions** $U(D(\mathbf{p}); A \circ A)$ to the **mathematical model** $\mathcal{M}(U, D(\mathbf{p}))$
3. Evaluate the **cost functional** $\mathcal{C}(U, D(\mathbf{p}))$ for all solutions/operating conditions
4. Vary the **design parameters** \mathbf{p} to optimize the **cost functional** $\mathcal{C}(U, D(\mathbf{p}))$ for a wide range of operating conditions and repeat the DTA cycle at step 1





Design-through-analysis

1. Find a set of adm
2. Compute solution
3. Evaluate the cost
4. Vary the design p
a wide range of o




Generate the design $D(\mathbf{p})$
al model $\mathcal{M}(U, D(\mathbf{p}))$
ons/operating conditions
unctional $\mathcal{C}(U, D(\mathbf{p}))$ for
DTA cycle at step 1




Introduction to

SPLINE TECHNOLOGIES

Basis functions: $\hat{B}_i(\xi) \quad i = 1, \dots, N$

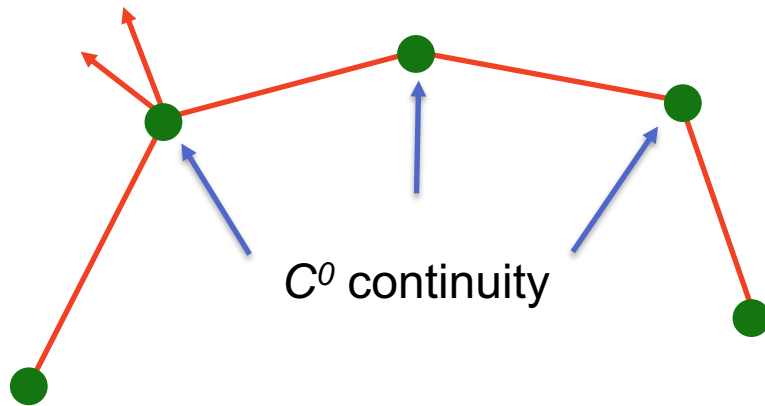


Linear 'tent' basis functions




Quadratic B-spline basis functions

Curves: $\mathbf{C}(\xi) = \sum_{i=1}^N \mathbf{c}_i \hat{B}_i(\xi) : [0,1] \rightarrow \mathbb{R}^d$




Linear 'tent' basis functions




Quadratic B-spline basis functions

Surfaces: $\mathbf{S}(\xi, \eta) = \sum_{i,j=1}^{N,M} \mathbf{c}_{i,j} \hat{B}_i(\xi) \hat{B}_j(\eta) : [0,1] \times [0,1] \rightarrow \mathbb{R}^d$





Bilinear basis functions





Biquadratic B-spline basis functions

Matrix structure for single-patch domain




Biquadratic B-spline basis functions

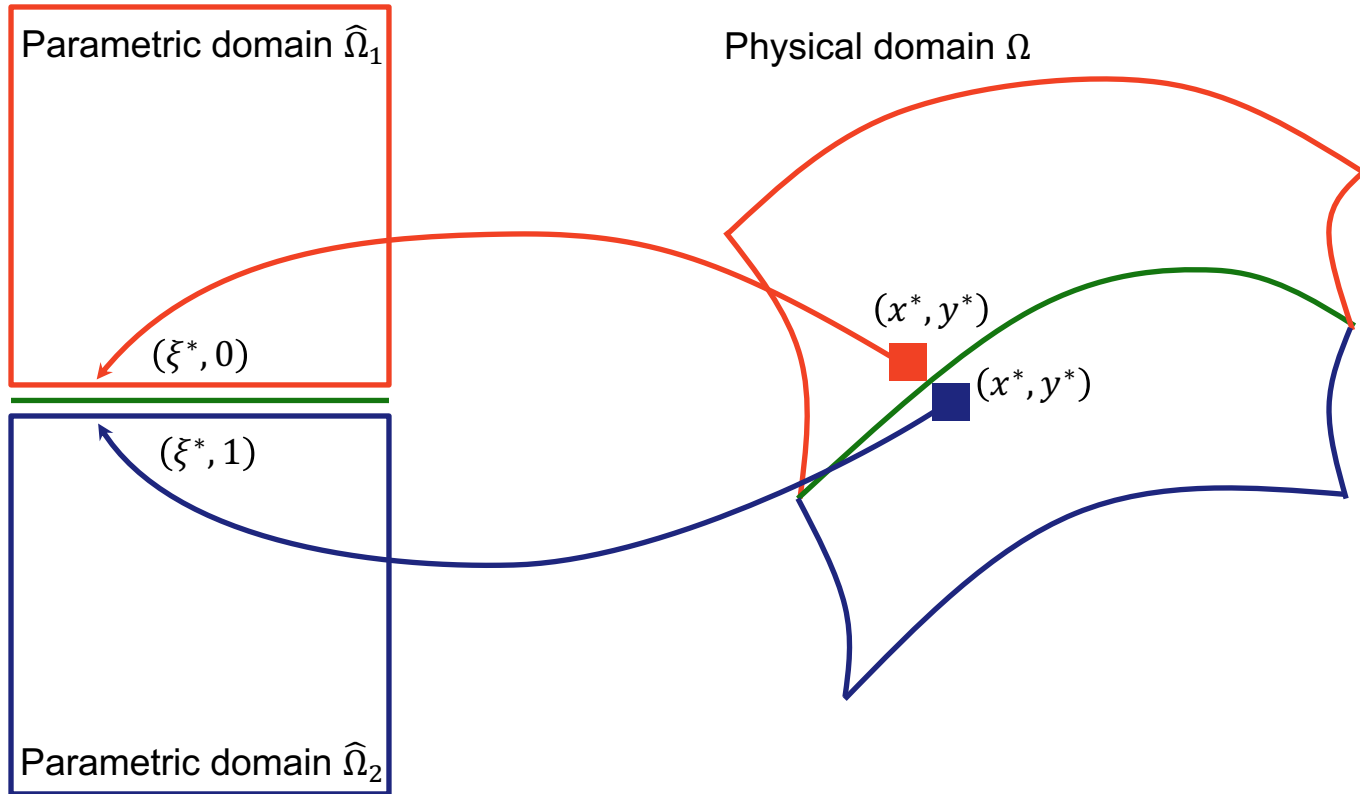
Matrix structure for multi-patch domain





Isogeometric Analysis in a nutshell




Isogeometric Analysis in a nutshell





Advanced spline technologies: THB splines





Application of THB splines



B-spline Basis



THB-spline Basis





Design-Through-Analysis for

TWIN-SCREW COMPRESSORS

Long-term vision

- Develop a computer code for the efficient **geometry** modelling, **simulation** and **optimization** of rotary twin-screw compressors and expanders

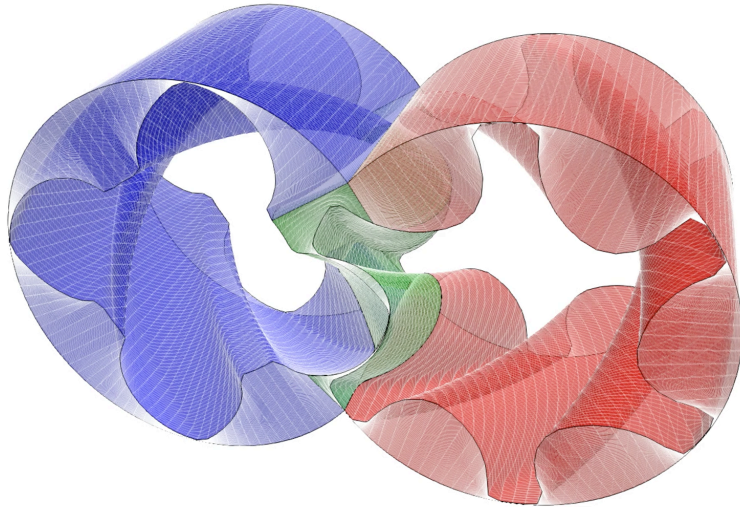


Source: Chair of Fluidics, TU Dortmund University, DE

A challenging industrial application

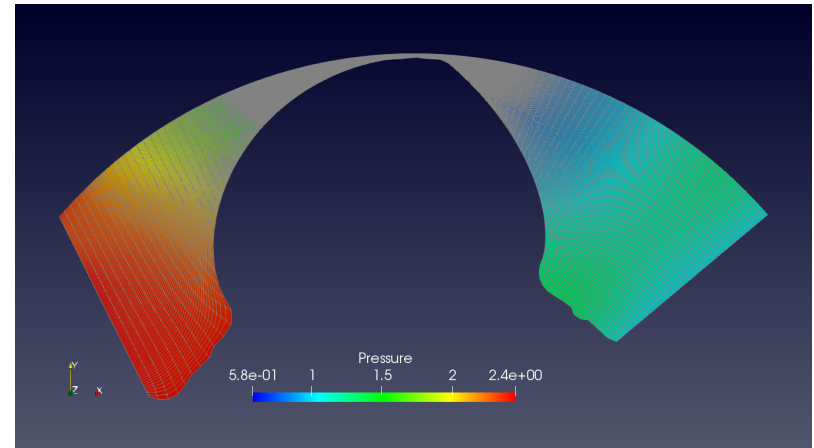
Geometry modelling

- Counter-rotating helical rotors
- Narrow clearances (<0.4mm)
- Complex deforming fluid domain




Multi-physics simulation


- Compressible high-speed flow
- Thermal expansion of solids
- Extension: injection of oil, ...



Co-design of geometry model and simulation pipeline




Multi-patch topology



Multi-block matrix



Multi-device computer




Co-design principles

- No topology changes (casing-to-rotor)
- Exploit block structure of matrices
- Keep design simple and extendible
- Support heterogeneous hardware


Design-Through-Analysis for twin-screw compressors

GEOMETRY MODELLING



Test cases




Test case #1: rotor-casing passage




Test case #2: separator with CUSP points



Test case #3: rotating twin rotors



Test case #4: rotating twin screws



Design-Through-Analysis for twin-screw compressors

ISOGEOMETRIC FLOW SOLVER

Compressible Euler equations

$$\frac{\partial}{\partial t} \underbrace{\begin{bmatrix} \rho \\ \rho \mathbf{u} \\ \rho E \end{bmatrix}}_{U = \begin{bmatrix} u_1 \\ \vdots \\ u_{d+2} \end{bmatrix}} + \nabla \cdot \underbrace{\begin{bmatrix} \rho \mathbf{u} \\ \rho \mathbf{u} \otimes \mathbf{u} + p \mathcal{J} \\ (\rho E + p) \mathbf{u} \end{bmatrix}}_{\mathbf{F}(U) = \begin{bmatrix} f_1^1 & \cdots & f_1^d \\ \vdots & \ddots & \vdots \\ f_{d+2}^1 & \cdots & f_{d+2}^d \end{bmatrix}} = 0$$

- Equation of state for an ideal gas (isentropic index $\gamma = 1.4$ for dry air)

$$p(\rho, e) = (\gamma - 1)\rho e, \quad \rho e = (\rho E - 0.5\rho\|\mathbf{u}\|^2)$$

- Flux-Jacobian matrix $A(U) = \frac{\partial \mathbf{F}(U)}{\partial U}$, homogeneity property $\mathbf{F}(\lambda U) = \lambda \mathbf{F}(U)$

Variational formulation

- Find solution $U(\cdot, t)$ at fixed time t such that for all test functions W

$$\int_{\Omega} W \partial_t U - \nabla W \cdot \mathbf{F}(U) d\mathbf{x} + \int_{\Gamma} W G(U, U^*) ds = 0$$

- Boundary fluxes

$$G(U, U^*) = \begin{cases} [0, p\mathbf{n}, 0]^T & \text{at solid walls} \\ 0.5(\mathbf{F}(U) + \mathbf{F}(U^*)) \cdot \mathbf{n} & \\ -0.5|\mathbf{A}(\text{Roe}(U, U^*))| \cdot \mathbf{n} & \text{otherwise} \end{cases}$$

Discretization

- Find solution $U_h(\cdot, t)$ at fixed time t such that for all test functions W_h

$$\int_{\Omega} W_h \partial_t U_h - \nabla W_h \cdot \mathbf{F}_h(U) d\mathbf{x} + \int_{\Gamma} W_h G_h(U, U^*) ds = 0$$

- Fletcher's group approximation (CMAME '83)

$$U_h(\mathbf{x}, t) = \sum_{j=1}^N B_j(\mathbf{x}) U_j(t), \quad B_j = \hat{B}_j \circ \mathbf{x}^{-1}$$
$$\mathbf{F}_h(U(\mathbf{x}, t)) = \sum_{j=1}^N B_j(\mathbf{x}) \mathbf{F}_j(t), \quad \mathbf{F}_j = \mathbf{F}(U_j)$$

Semi-discrete formulation

$$\begin{bmatrix} \mathbf{M} & & \\ & \ddots & \\ & & \mathbf{M} \end{bmatrix} \begin{bmatrix} \dot{u}_1 \\ \vdots \\ \dot{u}_{d+2} \end{bmatrix} - \begin{bmatrix} f_1^1 & \cdots & f_1^d \\ \vdots & \ddots & \vdots \\ f_{d+2}^1 & \cdots & f_{d+2}^d \end{bmatrix} \begin{bmatrix} \mathbf{C}^1 \\ \vdots \\ \mathbf{C}^d \end{bmatrix} + \begin{bmatrix} g_1^1 & \cdots & g_1^d \\ \vdots & \ddots & \vdots \\ g_{d+2}^1 & \cdots & g_{d+2}^d \end{bmatrix} \begin{bmatrix} \mathbf{S}^1 \\ \vdots \\ \mathbf{S}^d \end{bmatrix} = 0$$

- Constant coefficient matrices

$$\mathbf{M} = \left\{ \int_{\Omega} B_i B_j d\mathbf{x} \right\}_{i,j=1}^N, \quad \mathbf{C}^k = \left\{ \int_{\Omega} \partial_k (B_i) B_j d\mathbf{x} \right\}_{i,j=1}^N, \quad \mathbf{S}^k = \left\{ \int_{\Omega} B_i B_j \mathbf{n} ds \right\}_{i,j=1}^N$$

are pre-assembled using Gauss quadrature and stored to efficiently form the divergence term as SpMV-operation when it is required

From the programmer's perspective


$$M[\dot{u}_1 \cdots \dot{u}_{d+2}] - \underbrace{[C^1 \cdots C^d] \begin{bmatrix} f_1^1 & \cdots & f_1^{d+2} \\ \vdots & \ddots & \vdots \\ f_d^1 & \cdots & f_d^{d+2} \end{bmatrix}}_{\text{block vector of sparse matrices}} + [S^1 \cdots S^d] \begin{bmatrix} g_1^1 & \cdots & g_1^{d+2} \\ \vdots & \ddots & \vdots \\ g_d^1 & \cdots & g_d^{d+2} \end{bmatrix} = 0$$



=



*



1 x d+2 block vector
of dense vectors

1 x d block vector
of sparse matrices

d x d+2 block matrix
of function expressions

FDBB: Fluid dynamics building blocks

Low-level API

- **Unified function wrappers** to the core functionality of established HPC linear algebra libraries, e.g. ArrayFire, Blaze, Eigen, VexCL
- **Compile-time** block linear algebra backend with full support for
 - **Dense vectors**
 - **Sparse matrices**
 - **Function expressions**

Example


```
vex::vector<double>      x,y;  
vex::sparse::matrix<double> Cx,Cy;  
  
BlockMatrix<...,1,2> Mat(Cx,Cy);  
  
auto Expr = make_BlockExpr<2,1>(cos(x)+y,  
                                sin(x)-y );  
  
BlockRowVector<...,1> Vec = Mat * Expr;
```

Loops are unrolled at compile time
and fused in a single compute kernel

FDBB: Fluid dynamics building blocks

High-level API

- C++ expression templates for
 - Variables & Riemann invariants
 - Fluxes with 'generic' pressure
 - Equations of state



Example

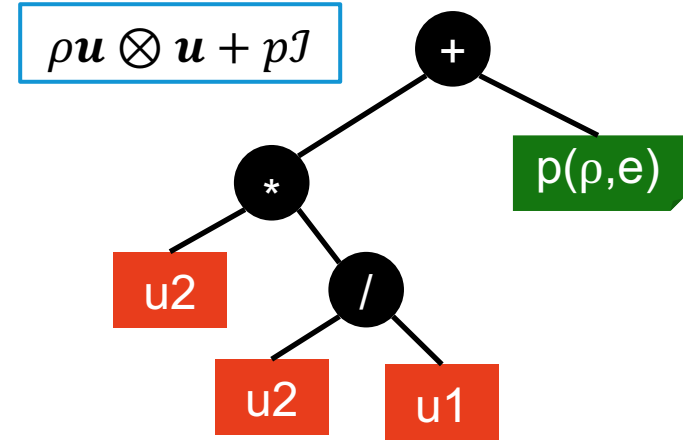
```
using eos = EOS::idealGas<double,  
                        ratio<7,2>,  
                        ratio<5,2>>;  
  
// 1x4 dim conservative state vector  
using var = Variables<eos,2,Conservative>;  
auto U = create<vex::vector<double>,4>(N);  
  
// 2x4 dim inviscid flux tensor  
using flx = Fluxes<var>;  
auto F = flx::inviscid(U);  
  
// Explicit solution update  
U += dt * Mat * F
```

Auto-generation of device-optimized CFD kernels

```

double rhs_6_sum = 0;
{
  for(size_t j = 0; j < rhs_6_ell_width; ++j)
  {
    int nnz_idx = idx + j * rhs_6_ell_pitch;
    int c = rhs_6_ell_col[nnz_idx];
    if (c != (int)(-1))
    {
      int idx = c;
      rhs_6_sum += rhs_6_ell_val[nnz_idx] * ((( prm_tag_2_1[idx] * prm_tag_2_1[idx] ) /
      prm_tag_0_1[idx] ) + ( rhs_6_x_4 * ( prm_tag_0_1[idx] * ( ( prm_tag_3_1[idx] / prm_tag_0_1[idx] ) - (
      ( 5.0000000000000000e-01 ) * ( ( prm_tag_1_1[idx] * prm_tag_1_1[idx] ) + ( prm_tag_2_1[idx] *
      prm_tag_2_1[idx] ) ) / ( prm_tag_0_1[idx] * prm_tag_0_1[idx] ) ) ) ) ) );
    } else break;
  }
  if (rhs_6_csr_ptr)
  ...

```



Heterogeneous HPC support


	CPU	CUDA	OpenCL	Intel MIC	FPGA
ArrayFire	X	X	X		
Armadillo	(X)				
Blaze	X				
Eigen	X				
MTL4	(X)				
uBLAS	(X)				
VexCL ¹	X	X	X	(X)	(X ²)
ViennaCL	X	X	X		

¹ source code is generated and JIT-compiled; ² Maxeler DFEs


Example: Computational performance



- Intel 2xE5-2670 @2.6 GHz
- IBM Power8NVL @4.02 GHz
- IBM + NVIDIA GP100GL



Test case #1: Inviscid compressible flow, $p_{in}:p_{out} = 2:1$



Conclusions

Hardware-oriented Numerics with IGA: **co-design** of geometry and simulation

- IGA package G+Smo: <https://github.com/gismo/gismo>
- Open-source FDBB: <https://gitlab.com/mmoelle1/FDBB>

Ongoing and future work

- Extension to turbulent flows and ALE formulation for rotating geometries
- Automatic compute resource scheduling and dynamic load balancing

References

- MM, A. Jaeschke: High-order IGA for compressible flows I+II, arXiv: [1809.10893](https://arxiv.org/abs/1809.10893) and [1809.10896](https://arxiv.org/abs/1809.10896)
- MM, A. Jaeschke: FDBB: Fluid Dynamics Building Blocks, arXiv: [1809.09851](https://arxiv.org/abs/1809.09851)

Financial support by EC under GA No 672787 is acknowledged

Recent updates

- **MPI-communication and dynamic load-balancing tool framework** : *joint student or master project (UvA, TUD); topic 2*
 - Cluster inventory tool implemented with support for x86, x86_64, ppc64le, ARM, CUDA, OpenCL, SDAccel, ...
- **Sparse band matrix implementation** : *student project TUD*
- **DSL-based algorithmic differentiation** : *PhD project by A. Jaeschke*
 - Gradient-based shape optimization in turbomachinery applications
- **PyG+Smo** : *hobby project of mine*
 - Python wrappers of G+Smo library based on PyBind11