

University of Stuttgart
Institute of Aerodynamics and Gas Dynamics



**LUMI Hackathon:
Porting GALÆXI To AMD Accelerators**




- Member of the code framework **FLEXI** with growing feature set
- Developed by the Numerics Research Group IAG, University of Stuttgart, Germany
- OpenSource HPC solver for unsteady compressible Navier–Stokes eq.
- High Order Discontinuous Galerkin Spectral Element Method




<https://numericsresearchgroup.org>



<https://github.com/flexi-framework/galaexi>



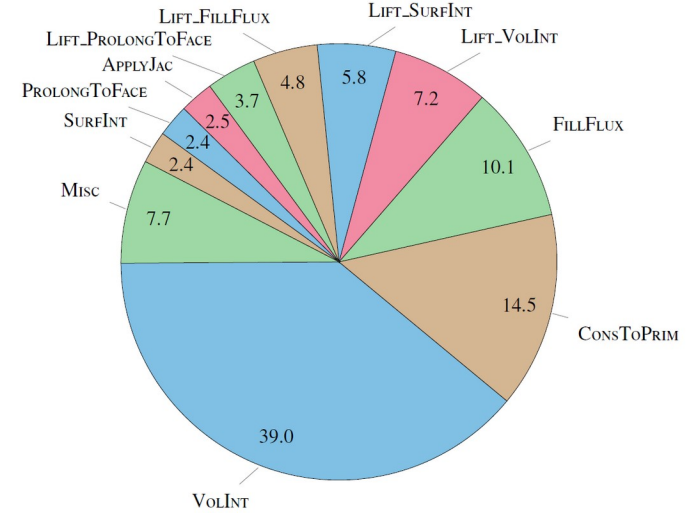
- Member of the code framework **FLEXI** with growing feature set
- Developed by the **Numerics Research Group** IAG, University of Stuttgart, Germany
- **OpenSource** HPC solver for **unsteady compressible** Navier–Stokes eq.
- High Order **Discontinuous Galerkin** Spectral Element Method
- Current feature set
 - Supported: NVIDIA (CUDA Fortran)
 - Running CI/CD pipeline
 - Gauss / Gauss-Lobatto points
 - Variable polynomial degree
 - Several Riemann solv. (comp. flag)
 - Split DG (several types, comp. flag)
 - Boundary conditions
 - MPI parallelized
 - Shock capturing
 - Sponge zones
 - Testcases (TGV, channel, ...)
 - ...



GALÆXI: DGSEM

- Semi-discrete formulation of the DGSEM (weak form):

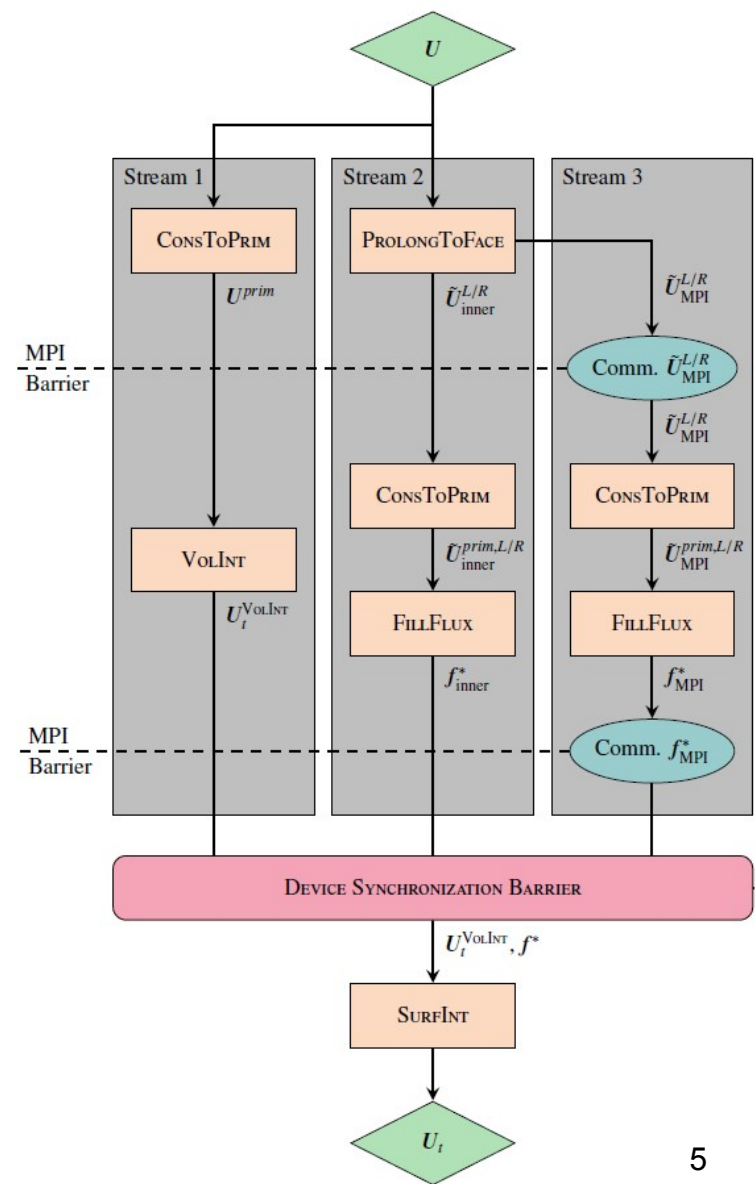
$$\frac{\partial \hat{U}_{ijk}}{\partial t} = \underbrace{\frac{1}{\mathcal{J}_{ijk}}}_{\text{APPLYJAC}} \left[\underbrace{\sum_{\alpha=0}^N \mathcal{F}_{\alpha jk}^1 \hat{D}_{i\alpha}}_{\text{VOLINT}} + \underbrace{\sum_{\beta=0}^N \mathcal{F}_{i\beta k}^2 \hat{D}_{j\beta}}_{\text{VOLINT}} + \underbrace{\sum_{\gamma=0}^N \mathcal{F}_{ij\gamma}^3 \hat{D}_{k\gamma}}_{\text{VOLINT}} \right] + \underbrace{\left(\underbrace{([f^* \hat{s}]_{jk}^{\xi^+} \hat{\ell}_i^+)}_{\text{FILLFLUX}} - \underbrace{([f^* \hat{s}]_{jk}^{\xi^-} \hat{\ell}_i^-)}_{\text{FILLFLUX}} \right)}_{\text{SURFINT}} + \underbrace{\left(\underbrace{([f^* \hat{s}]_{ik}^{\eta^+} \hat{\ell}_j^+)}_{\text{FILLFLUX}} - \underbrace{([f^* \hat{s}]_{ik}^{\eta^-} \hat{\ell}_j^-)}_{\text{FILLFLUX}} \right)}_{\text{SURFINT}} + \underbrace{\left(\underbrace{([f^* \hat{s}]_{ij}^{\zeta^+} \hat{\ell}_k^+)}_{\text{FILLFLUX}} - \underbrace{([f^* \hat{s}]_{ij}^{\zeta^-} \hat{\ell}_k^-)}_{\text{FILLFLUX}} \right)}_{\text{SURFINT}}$$



Routine	Vol/Surf	DOF-local	LIFT_*	Operations	Explanation
CONSTO PRIM	Surf, Vol	YES	NO	$O(N^{2,3})$	Computes primitive variables U^{prim} from state U .
VOLINT	Vol	NO	YES	$O(N^4)$	Evaluates volume fluxes \mathcal{F} and multiplies with \hat{D} .
PROLONGTOFACE	Vol \rightarrow Surf	NO	YES	$O(N^2)$	Evaluates solution at element faces $U^{L/R}$ to compute f^* .
FILLFLUX	Surf	YES	YES	$O(N^2)$	Computes common flux f^* on faces with Riemann solver.
SURFINT	Vol \leftarrow Surf	NO	YES	$O(N^2)$	Computes surface integral with f^* and $\hat{\ell}^\pm$.
APPLYJAC	Vol	YES	YES	$O(N^3)$	Applies Jacobian \mathcal{J} to \hat{U}_t .


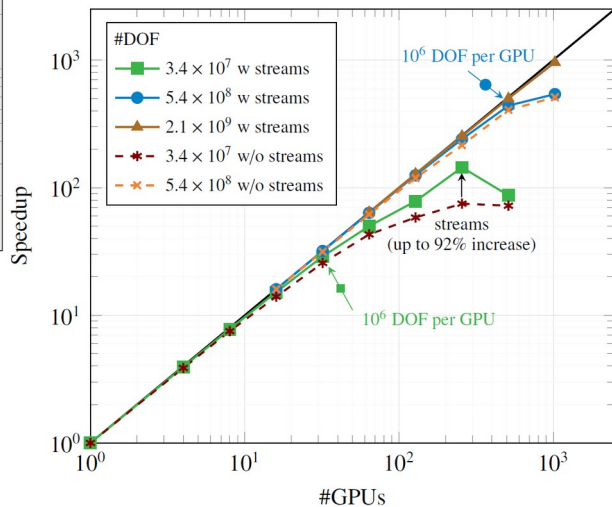

GALÆXI v1.0: Implementation

- CUDA Fortran
- Design principles
 1. Retain the general data structure and parallelization strategy of FLEXI for unstructured meshes
 2. Retain the majority of the codebase and the associated features of the original Implementation
 3. All routines called during the time-stepping are executed on the accelerator without the need of data transfers
- **Inter-GPU parallelization:** Distributed memory approach using *GPU-aware MPI* and systematic *hiding of communication latency*
- **Intra-GPU parallelization:** Launch multiple kernels concurrently and asynchronously to *maximize the occupancy* using e.g. *streams*
- **Kernel implementation:** Maximize the utilization of the available parallel resources, individual parallelization strategies for *pointwise* and *volume surface* operations



GALÆXI v1.0: Performance Evaluation

- Performance index:
$$PID = \frac{\text{Walltime} \times \#\text{Ranks}}{\#\text{RK-stages} \times \#\text{DOF}}$$
- PID describes the walltime required by a single rank to advance a single DOF for one stage of the explicit Runge–Kutta time-stepping





GALÆXI

GALÆXI v2.0

GALÆXI v2.0

- Retain design principles from GALÆXI v1.0
- Allow computation on ALL architectures (including CPU)
 - Rewrite all CUDA Fortran kernels in C.
- Governing philosophy of multi-backend approach
 - Fortran = host code
 - C = device code
- Host code (Fortran) calls device code (C) through interfaces
 - Interfaces wrap CUDA/HIP API calls and Kernels
 - Hides GPU code for those who don't want/need to interact with it
 - Easily extendable to multiple backends/vendors

```
ALLOCATE(U( PP_nVar,0:PP_N,0:PP_N,0:PP_NZ,nElems))
CALL AllocateDeviceMemory("d_U", SIZE_C_DOUBLE, SIZE(U))

CALL CopyToDevice("d_U", U, SIZE(U))

CALL CopyFromDevice(U, "d_U", SIZE(U))
```

```
void AllocateDeviceMemory(char* dVarKey, size_t typeSize_bytes, int arraySize)
{
    void* d_arr;
    #if (USE_ACCEL == ACCEL_CUDA)
        DEVICE_ERR_CHECK( cudaMalloc(&d_arr, typeSize_bytes*arraySize) );
    #elif (USE_ACCEL == ACCEL_HIP)
        DEVICE_ERR_CHECK( hipMalloc(&d_arr, typeSize_bytes*arraySize) );
    #elif (USE_ACCEL == ACCEL_HYBRID)
        d_arr = NULL;
    #endif

    DeviceVars[dVarKey] = d_arr;
}
```


GALÆXI v2.0: Development Plan

- Phase 0 (Finished)
 - Updates to supporting function surrounding code (CI/CD, unit testing, build, etc.)
 - Incorporated device memory management methods
- Phase 1 (In Progress - Due 13.12.2024)
 - Port core kernels from CUDA Fortran to CUDA/HIP C++.
 - Code able to run inviscid, incompressible test problem on a single GPU (all architectures)
- **Current Status**
 - Core kernels rewritten in C
 - Computation of inviscid test problem is possible on a single NVIDIA GPU with new kernels
 - No support for AMD
- Phase 2 and Phase 3 (Planned)
 - Reintroduce multi-GPU support
 - Support for more physics and features (viscous flow, shock capturing, WMLES, etc.)




GALÆXI

Goals For Hackathon

LUMI Hackathon Goals

- Easy
 - Extend device support code (memory management, kernel launches, etc) to support HIP
 - Build with HIP compiler
 - Familiarize with AMD profiling tools
- Better
 - Complete adjustments to kernels (if needed)
 - Code can run test problem on LUMI-G nodes
- Best
 - Tuning and optimization
 - Start looking at hybrid architectures



University of Stuttgart
Institute of Aerodynamics and Gas Dynamics

THANK YOU!

