

# A TASK-BASED PARALLELIZATION OF A FINITE VOLUME CODE FOR THE EULER FLOWS

**Dr. Sangeeth Simon**

- Team CAGIRE, INRIA Bordeaux Sud-Ouest &  
Laboratoire de Mathématiques et de leurs applications, Université de Pau et des Pays  
de l'Adour

December 8, 2020

# What is this research about?

- ✦ **The overall aim:** To have a highly scalable, hardware-agnostic, Higher Order, Discontinuous Galerkin - based, Navier-Stokes solver.
- ✦ **The problem:** How to *easily* develop scalable and portable applications for heterogeneous HPC clusters?
- ✦ **Our approach:** Use a task-based parallelization paradigm to benefit from coexisting multi-core & GPU architectures.
- ✦ **Our findings so far:** Decent scalability obtained using StarPU on multi-core architectures.

## What is this research about?

- ✦ **The overall aim:** To have a highly scalable, hardware-agnostic, Higher Order, Discontinuous Galerkin - based, Navier-Stokes solver.
- ✦ **The problem:** How to *easily* develop scalable and portable applications for heterogeneous HPC clusters?
- ✦ **Our approach:** Use a task-based parallelization paradigm to benefit from coexisting multi-core & GPU architectures.
- ✦ **Our findings so far:** Decent scalability obtained using StarPU on multi-core architectures.

## What is this research about?

- ✦ **The overall aim:** To have a highly scalable, hardware-agnostic, Higher Order, Discontinuous Galerkin - based, Navier-Stokes solver.
- ✦ **The problem:** How to *easily* develop scalable and portable applications for heterogeneous HPC clusters?
- ✦ **Our approach:** Use a task-based parallelization paradigm to benefit from coexisting multi-core & GPU architectures.
- ✦ **Our findings so far:** Decent scalability obtained using StarPU on multi-core architectures.

## What is this research about?

- ✦ **The overall aim:** To have a highly scalable, hardware-agnostic, Higher Order, Discontinuous Galerkin - based, Navier-Stokes solver.
- ✦ **The problem:** How to *easily* develop scalable and portable applications for heterogeneous HPC clusters?
- ✦ **Our approach:** Use a task-based parallelization paradigm to benefit from coexisting multi-core & GPU architectures.
- ✦ **Our findings so far:** Decent scalability obtained using StarPU on multi-core architectures.

## What is this research about?

- ✦ **The overall aim:** To have a highly scalable, hardware-agnostic, Higher Order, Discontinuous Galerkin - based, Navier-Stokes solver.
- ✦ **The problem:** How to *easily* develop scalable and portable applications for heterogeneous HPC clusters?
- ✦ **Our approach:** Use a task-based parallelization paradigm to benefit from coexisting multi-core & GPU architectures.
- ✦ **Our findings so far:** Decent scalability obtained using StarPU on multi-core architectures.

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work



# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Outline of the presentation

- ⊙ Higher order methods for Nonlinear balance laws & Discontinuous Galerkin approach
- ⊙ Task Based Parallelization & StarPU
- ⊙ Important findings from Project HODINS so far
- ⊙ Taking HODINS ahead→ Approach 1: Extensions to Three Dimensions
- ⊙ Taking HODINS ahead→ Approach 2: Increasing order of accuracy
- ⊙ Conclusions
- ⊙ Future work

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach *India*

## Linear Balance laws

- ⊙ Linear balance laws have **solutions that are superpositions** and hence comparatively easier to obtain.
- ⊙ Computationally, they reduce to **global matrix operations** that can be optimized for an HPC application.


# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach

## Nonlinear Balance laws

- ⊙ Nonlinear balance laws tend to **develop discontinuities** as part of their solution.
- ⊙ Obtaining interface fluxes at the element boundaries involves **nonlinear function evaluations**.
- ⊙ Relatively **higher computational complexity** per time step.


# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach *Anria*

## ⊙ Why need higher order at all? <sup>1</sup>



<sup>1</sup> Advection of a vortex on 64x64 grid



Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin  
approach *Inria*⊙ Why need higher order at all? <sup>1</sup>

First order is unacceptably more 'diffusive' than second order.

<sup>1</sup>Results for advection of a vortex on  $64 \times 64$  grid for 5000 iterations. Second order accuracy using SSPRK-MUSCL.

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach *India*

## ⊙ Why need higher order at all?

Objective of any higher order method

Obtaining better resolution at cheaper cost compared to a lower order method!

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach



- ① **How does Discontinuous Galerkin approach help?**

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach

## How does Discontinuous Galerkin approach help?

- ⊙ **Best of both frameworks** → Finite Element (polynomial representation of solution on elements & matrix operations) & Finite Volume (conservation, interface fluxes, limiters).

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach


## How does Discontinuous Galerkin approach help?

- ⊙ **Good data locality & compact stencil** → Each element is equipped with enough data (except interface fluxes) to compute its own residuals. Each element interacts only with its immediate neighbours.

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach *India*


## Elemental polynomial order $k = 4$

DG



pointers of size  $12 \times 15 = 180$

CG



pointer of size  $12 \times 1 = 12$



pointer of size  $22 \times 3 = 66$



pointer of size  $12 \times 3 = 36$

# Higher Order Methods for Nonlinear Balance laws & Discontinuous Galerkin approach

## How does Discontinuous Galerkin approach help?

- ⊙ **Local matrices instead of global matrices** → In DG we completely avoid global matrix operations.
- ⊙ **Diffusion operator adds more computations** → Evaluation of second order gradients using same DOF.
- ⊙ **Enhanced arithmetic intensity** → Each element ideally does more computation with local data compared to communication.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- Developing new applications/ porting legacy codes etc. requires substantial effort.
- Poor portability across different clusters.
- Performance gains depend on how memory management is achieved.
- Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- A framework that enables swift HPC code development.
- Minimal memory management efforts.
- Offers good scalability with little effort.
- Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- Offers a unified execution model (unifies all computing units and their embedded memory).
- Automatically takes care of low-level memory management between host and accelerators.
- StarPU efficiently maps tasks to available computing units.
- Offers scope for development/tuning of scheduling policies.



# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
  - 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
  - 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
  - 4 Poor portability across different clusters.
  - 5 Performance gains depend on how memory management is achieved.
- Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- A framework that enables swift HPC code development.
- Minimal memory management efforts.
- Offers good scalability with little effort.
- Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- Offers a unified execution model (unifies all computing units and their embedded memory).
- Automatically takes care of low-level memory management between host and accelerators.
- StarPU efficiently maps tasks to available computing units.
- Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.



# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

# Task-Based Parallelization & StarPU

## ⊙ Heterogeneous clusters entail:

- 1 Several multiprocessor architectures coexisting on a single node (CPU cores, specialized coprocessors, GPGPUs).
- 2 Need for using different hardware specific API's to interact with them (CUDA, OpenMP, OpenACC, openCL).
- 3 Developing new applications/ porting legacy codes etc. requires substantial effort.
- 4 Poor portability across different clusters.
- 5 Performance gains depend on how memory management is achieved.
- 6 Domain experts often need to double-up as HPC experts.

## ⊙ The ideal dream of today's CFD developer:

- 1 A framework that enables swift HPC code development.
- 2 Minimal memory management efforts.
- 3 Offers good scalability with little effort.
- 4 Same code runs on CPU/GPU/COProcessors depending on what is available.

## ⊙ StarPU: A dynamic runtime scheduler supporting task based parallelism.

- 1 Offers a unified execution model (unifies all computing units and their embedded memory).
- 2 Automatically takes care of low-level memory management between host and accelerators.
- 3 StarPU efficiently maps tasks to available computing units.
- 4 Offers scope for development/tuning of scheduling policies.

- HODINS → **H**igh **O**rders **D**iscontinuous methods with **r**untime **S**cheduler.




- HODINS-2D → a first order accurate Finite Volume code for the Euler equations with source terms using StarPU. <sup>1</sup>

---

<sup>1</sup>Essadki, M. Jung, J. Larat, A. Pelletier, M. Perrier, V. A Task-Driven Implementation of a Simple Numerical Solver for Hyperbolic Conservation Laws. *ESAIM: ProcS* (2018) **63**:228–247.


- Experiments performed on a 2 Dodeca-core Haswell Intel Xeon E5-2680, a Xeon Phi KNL and Nvidia K40-M GPU.

## ⊙ Importance of tasksizeOverhead <sup>1</sup>



(A) 2 dodeca-core Haswell Intel Xeon E5-2680.

(a) Haswell




(B) Xeon Phi KNL.

(b) KNL

Tasks must have a minimum execution time (~1ms) for scalability from StarPU!


<sup>1</sup> Taken from Essadki et al, 2018. Legends are in ( $\mu$  s)

## Scalability results<sup>1</sup>



(A) 2 dodeca-core Haswell Intel Xeon E5-2680.

(a) Haswell



(B) Xeon Phi KNL.

(b) KNL

For scalability from StarPU: Numerous tasks ( $O(\text{No of Cores})$ ) that are individually large enough ( $>\text{tasksizeoverhead}$ )

<sup>1</sup> Taken from Essadki et al, 2018. Results for  $1024 \times 1024$  grid.


- ⊙ How do we improve arithmetic intensity of kernels?

## HODINS-3D

- Extend HODINS-2D to HODINS-3D.

Partition=6	Partition=7	Partition=8
Partition=3	Partition=4	Partition=5
Partition=0	Partition=1	Partition=2


(a) Haswell




(b) KNL

Scalability results<sup>1</sup>


HODINS-3D-scalability study on 2 dodeca-core Haswell Intel Xeon E5-2667



HODINS-3D-scalability study on 2 dodeca-core Haswell Intel Xeon E5-2667



HODINS-3D-scalability study on 2 dodeca-core Haswell Intel Xeon E5-2667



- Same pattern as 2D → Too many/Too few tasks not good for scalability.
- NParts = 64, 512 and 4096 offers some scaling.
- HODINS-3D scalability → Not as great as reported for HODINS-2D.
- Need to investigate the reason for poor scaling.

<sup>1</sup> Results for 512<sup>3</sup> grid.

## HODINS-Higher order

- ⊙ Higher order increases arithmetic intensity per kernel due to additional operations like gradient computation and reconstruction.




## HODINS-Higher order

- ⦿ Higher order achieved using SSPRK-MUSCL combination.

# HODINS-Higher order


## Old overlaps model



- 5 pointers
- 5 `starpu_data_handle`
- 1st step: launch concurrently
  - Internal residual computation[R]
  - copy inside overlaps[W]
- 2nd step: compute the border residual, involving both internal[R] and borders[R] data handles
  - But matching indices between internal and overlaps is not straightforward to implement

# HODINS-Higher order

## New overlaps model



- 1 pointers
- 6 `starp_data_handle`
  - One view with Internal and overlaps
  - One view with full pointer
- 1st step: launch concurrently
  - Internal residual computation
  - copy inside overlaps
- 2nd step: compute the border residual, involving the full view of the pointer
  - Use the same algorithm as for internal residual, but with different indices.

## What lies ahead...

- ⊙ Complete H.O implementation and do performance analyses (Gantt charts, GFLOPS, Memory bandwidth, time taken by codelets etc.).
- ⊙ Assimilate learning to develop HODINS-MUSCL-3D. Performance analyses.
- ⊙ Apply learning to DG through AEROSOL framework.

## What lies ahead...

- ⊙ Complete H.O implementation and do performance analyses (Gantt charts, GFLOPS, Memory bandwidth, time taken by codelets etc.).
- ⊙ Assimilate learning to develop HODINS-MUSCL-3D. Performance analyses.
- ⊙ Apply learning to DG through AEROSOL framework.

## What lies ahead...

- ⊙ Complete H.O implementation and do performance analyses (Gantt charts, GFLOPS, Memory bandwidth, time taken by codelets etc.).
- ⊙ Assimilate learning to develop HODINS-MUSCL-3D. Performance analyses.
- ⊙ Apply learning to DG through AEROSOL framework.

# HODINS-Higher order

- ① We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ② We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ③ We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ④ We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ⑤ To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ⑥ Also **built HODINS-MUSCL-2D with new data management features**.

# HODINS-Higher order

- ① We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ① We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ① We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ① We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ① To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ① Also **built HODINS-MUSCL-2D with new data management features**.



# HODINS-Higher order

- ⊙ We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ⊙ We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ⊙ We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ⊙ We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ⊙ To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ⊙ Also **built HODINS-MUSCL-2D with new data management features**.

# HODINS-Higher order

- ⊙ We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ⊙ We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ⊙ We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ⊙ We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ⊙ To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ⊙ Also **built HODINS-MUSCL-2D with new data management features**.

# HODINS-Higher order

- ⊙ We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ⊙ We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ⊙ We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ⊙ We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ⊙ To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ⊙ Also **built HODINS-MUSCL-2D with new data management features**.

# HODINS-Higher order

- ⊙ We motivated the **need for Higher order methods for Non-Linear equations** (like Navier-Stokes eqns).
- ⊙ We motivated why **Discontinuous Galerkin might be an ideal candidate** for achieving this.
- ⊙ We motivated the **need for using task-based parallelism** in a world of heterogeneous architectures.
- ⊙ We reviewed the work done under HODINS-2D on a first order Euler code with source terms: We realized that **saturation task heap with compute intense tasks is necessary for achieving scalability**.
- ⊙ To improve compute intensity, we built HODINS-3D: **Decent scalability observed** for  $N_{\text{parts}}=64$  for a grid of  $512^3$  cartesian cells on different architectures. Scope for improvement exist.
- ⊙ Also **built HODINS-MUSCL-2D with new data management features**.

My sincere thanks to:

(1cm) **Inria & "Conseil Régional Nouvelle Aquitaine"**.

Team CAGIRE (Vincent Perrier, Matthieu Hafele, Jonathan Jung)

ANR Solharis / Région Nouvelle Aquitaine hpc-scalable-ecosystem Organizers