



# Optimizing SOD2D for LUMI

Mohammad Umair<sup>†</sup> and Panagiotis-Eleftherios Eleftherakis<sup>§</sup>

<sup>†</sup>FLOW, Engineering Mechanics, KTH Royal Institute of Technology, Sweden

<sup>§</sup> Microprocessors and Digital Systems Laboratory, National Technical University of Athens, Greece

**Hackathon: Optimizing for AMD GPUs 2024**  
14-18 October 2024, Brussels (Belgium)



# SOD2D: Spectral high-Order code 2 solve partial Differential equations

A new Continuous Galerkin High-Order Spectral Element Method (CG-SEM) code designed to perform numerical simulations of both turbulent **compressible** and **incompressible** flows.

Developed with the aim to target large scale Computational Fluid Dynamics simulations for engineering applications by exploiting the capabilities of the leading-edge extreme-scale HPC architectures.

Developed at 



Oriol LEHMKUHL  
Leading Researcher



Jordi MUELA CASTRO  
Recognized Researcher



Lucas GASPARINO  
First Stage Researcher



Mohammad UMAIR  
Postdoctoral Researcher



Ricardo VINUESA  
Principal Investigator



Fran. ALCÁNTARA-ÁVILA  
Postdoctoral Researcher

## Users & Developers at KTH

## Optimization Team at NTUA



Sotirios Xydis  
Assistant Professor



G. Anagnostopoulos  
PhD Student



K. Iliakis  
Postdoc Researcher




P.-E. Eleftherakis  
PhD Student



Marcial SANCHIS  
PhD Student





Cristiano PIMENTA  
PhD Student (Volvo & KTH)



PoI SUÁREZ  
PhD Student

# SOD2D Scaling

NASA High-Lift Common Research Model (*CRM-HL*)



Carried out in **MareNostrum 5 ACC** (Accelerated Partition)

NVIDIA **Hopper H100** 64GB HBM2

Source: Development team at



# SOD2D LUMI Porting and Optimization Project



KTH



Mohammad UMAIR  
Postdoctoral Researcher



Ricardo VINUESA  
Principal Investigator




Fran. ALCÁNTARA-ÁVILA  
Postdoctoral Researcher



Sotirios Xydis  
Assistant Professor



G. Anagnostopoulos  
PhD Student



K. Iliakis  
Postdoc Researcher




P.-E. Eleftherakis  
PhD Student

NTUA

With help from Jing Gong and Jonathan Vincent from PDC Center for High Performance Computing, and Jean-Yves Vet from Hewlett Packard Enterprise (HPE).

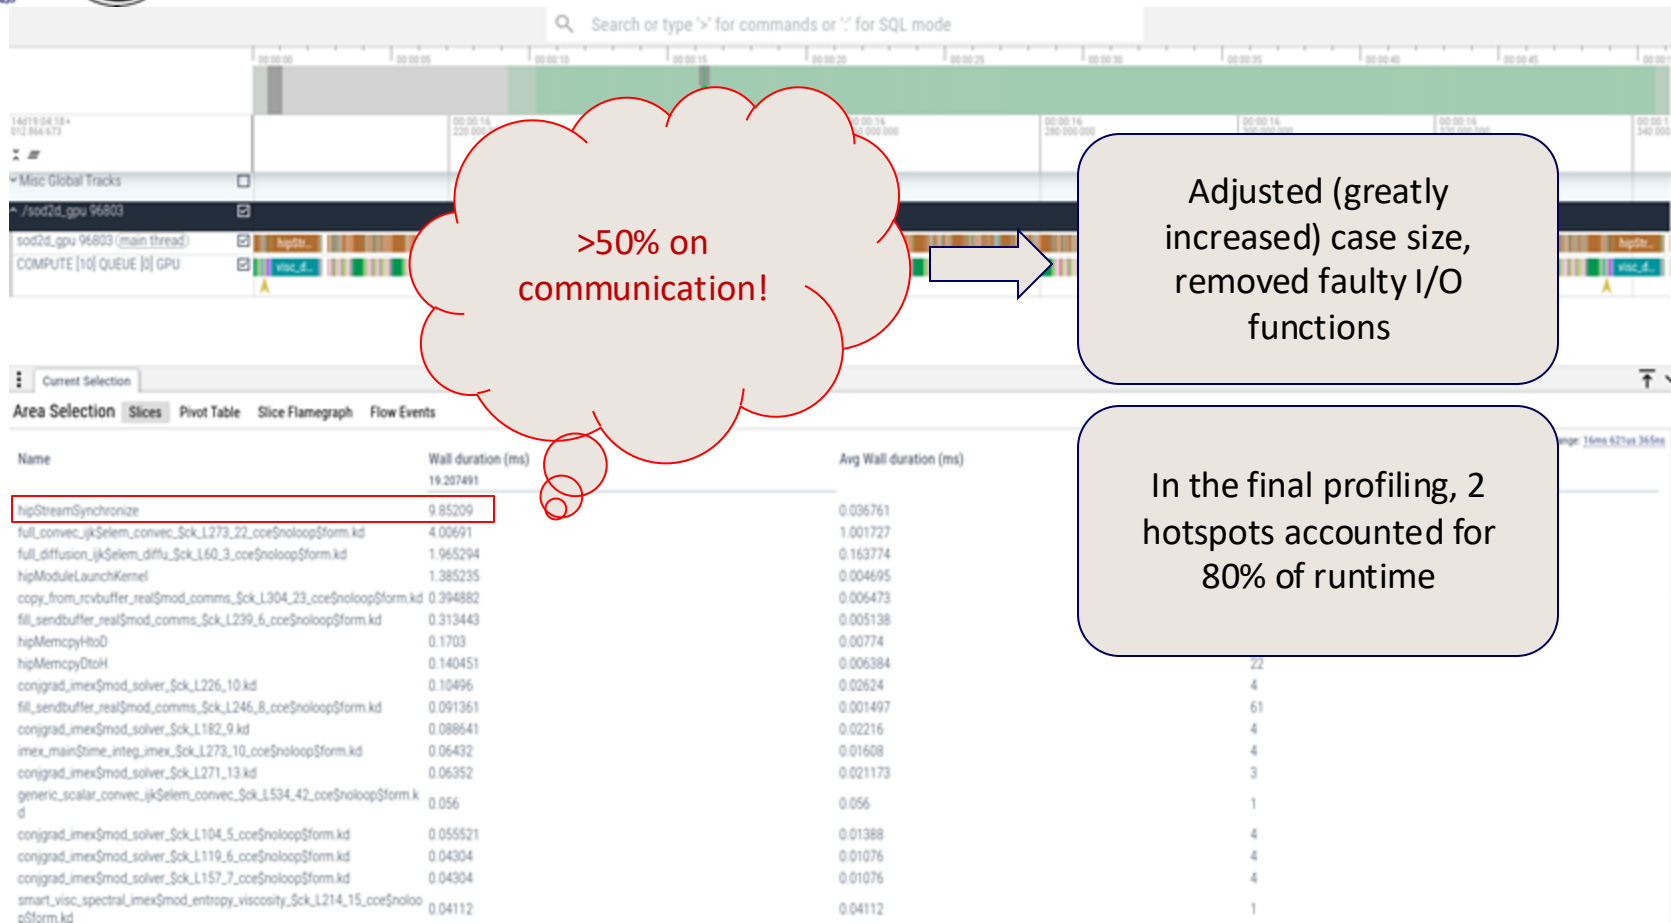


# Accomplished goals



Indicative Specification	AMD MI250X (LUMI)	NVIDIA H100	NVIDIA DGX A100 (per GPU)
Peak FP64 FLOPS	47.9 TFLOPS	30 TFLOPS	19.5 TFLOPS
<b>Peak FP32 FLOPS</b>	<b>47.9 TFLOPS</b>	<b>60 TFLOPS</b>	<b>39 TFLOPS</b>
Memory Bandwidth	3.28 TB/s	2.0 TB/s	1.6 TB/s
Peak Power	500W (LUMI)	700W	500W
SOD2D norm SLOWDOWN	~6x	1x	2x
SOD2D PPW deterioration	~4.3x	1x	~1.4x

# Initial Profiling



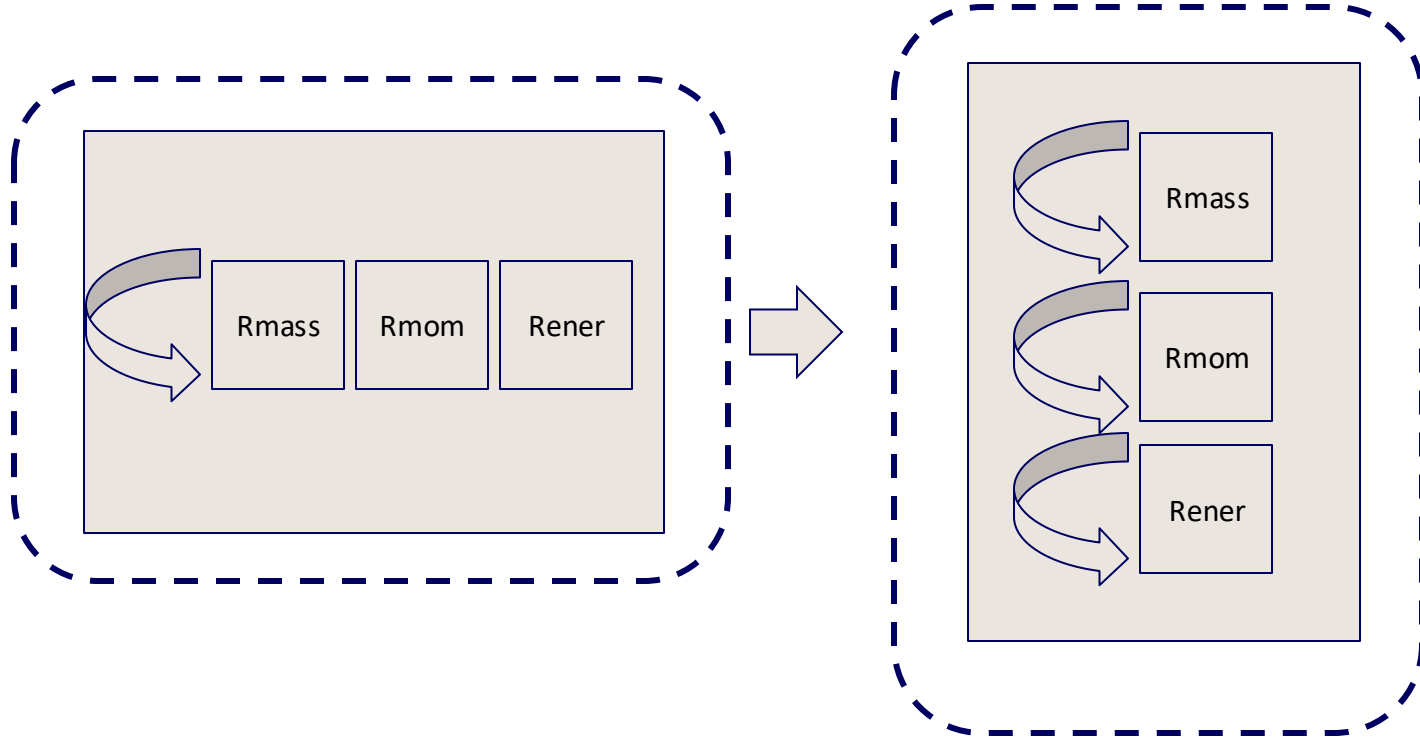


# Register spilling of the hotspots

## 7.1 Wavefront Launch Stats

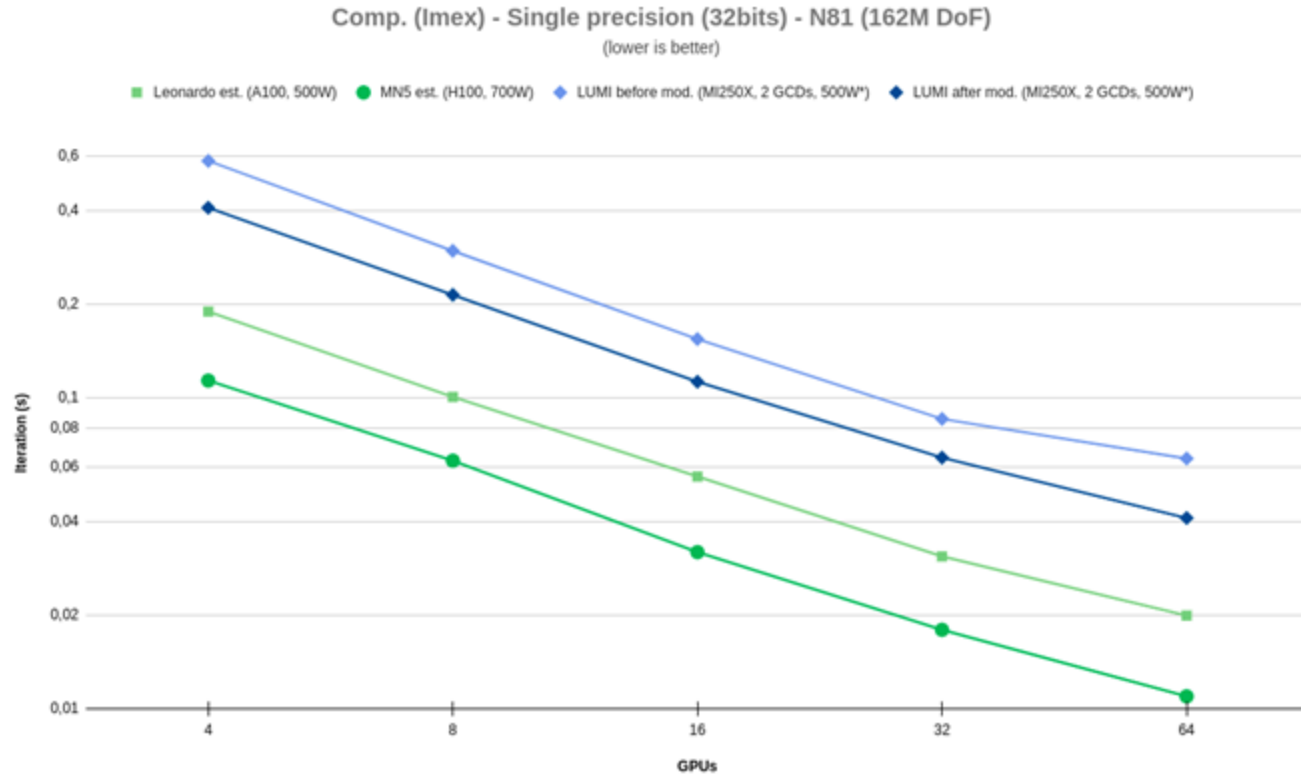
Metric_ID	Metric	Avg	Min	Max	Unit
7.1.0	Grid Size	16384000.00	16384000.00	16384000.00	Work items
7.1.1	Workgroup Size	256.00	256.00	256.00	Work items
7.1.2	Total Wavefronts	0.00	0.00	0.00	Wavefronts
7.1.3	Saved Wavefronts	0.00	0.00	0.00	Wavefronts
7.1.4	Restored Wavefronts	0.00	0.00	0.00	Wavefronts
7.1.5	VGPRs	128.00	128.00	128.00	Registers
7.1.6	AGPRs	0.00	0.00	0.00	Registers
7.1.7	SGPRs	96.00	96.00	96.00	Registers
7.1.8	LDS Allocation	4608.00	4608.00	4608.00	Bytes
7.1.9	Scratch Allocation	624.00	624.00	624.00	Bytes/workitem

# Kernel Code refactoring to reduce spilling

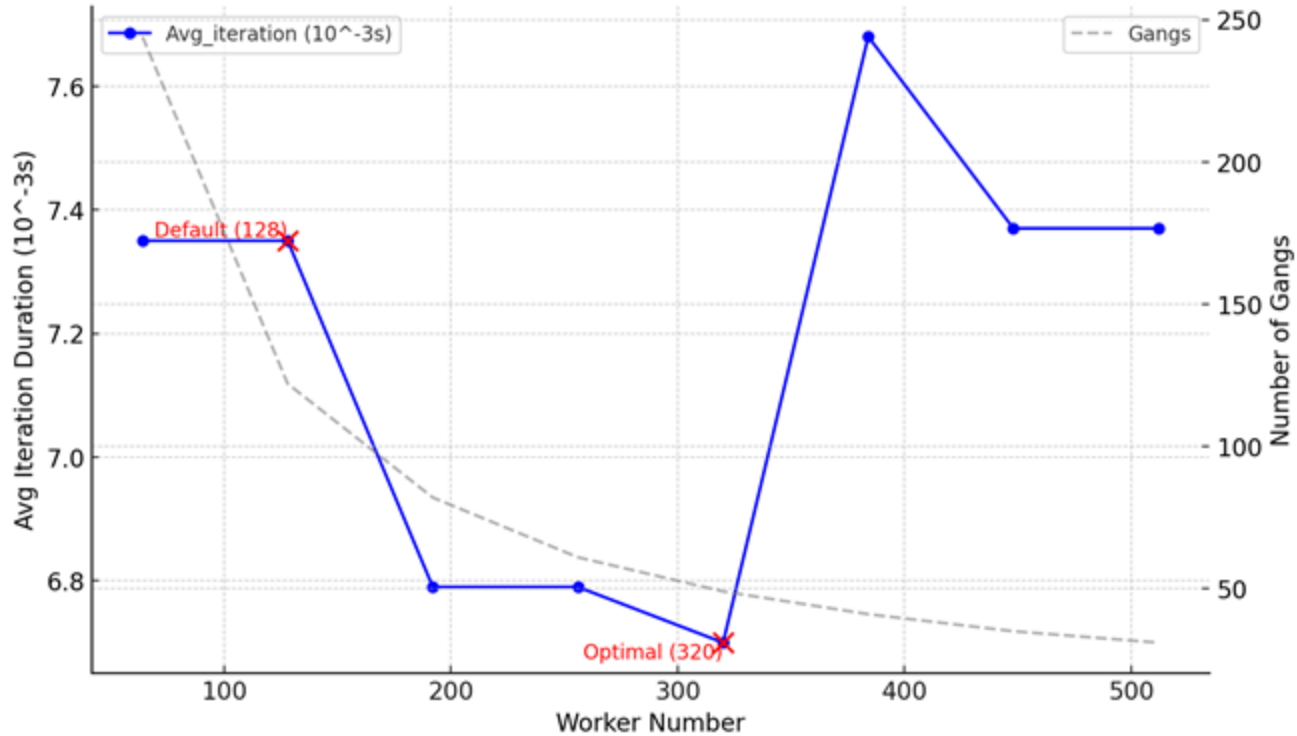




# Scaling runs and refactoring speed-up



Avg Iteration Duration and Gangs vs Worker Number





# Recap

- With  $N=86$ , `full_convvec_ijk` took 53.60% of the iteration time
- Splitting the kernel in 3 sped up that kernel by deducing register spilling.
- Changing the launch bounds led to an other performance improvement on that kernel
- Now the sum of the 3 split kernels takes about  $\sim 33\%$ . which is about the same duration as the next dominant kernel `full_diffusion_ijk`. Overall iteration improvement is about 30%.
- Final slowdown is  $\sim 2x$  compared to NVIDIA A100



# Future work



## We will focus more on this kernel

- The idea is to refactor this kernel and then move on to the next big kernel
- There has to be a compromise between performance and readability
- SOD2D is a continuously evolving code
- Many new features are being implemented almost every month

## Autotuning

- Full integration of the autotuning framework is required:
  - Changing nvtx profiling calls
  - Scripting changes to adjust to slurm and interface with AMD profilers
- More autotuning runs should be included:
  - Different solvers
  - At scale
  - Different examples

# The effort here will reflect in better future

© Benet Eiximeno Franch (UPC) and Cristiano Pimenta Silva (Volvo & KTH)

