On the Impact of Number Representation for High-Order LES

F.D. Witherden and A. Jameson

Department of Aeronautics & Astronautics
Stanford University

Motivation

- LES is expensive…
- …really expensive.

Computer Arithmetic

- Binary floating point following **IEEE 754**

- $x = sign \cdot mantissa \cdot 2^{exponent}$

	sign	exponent	mantissa
binary32	1	8	23
binary64	1	11	52

Computer Arithmetic

- Complicated!
- If you think you understand floating point arithmetic—you don't!

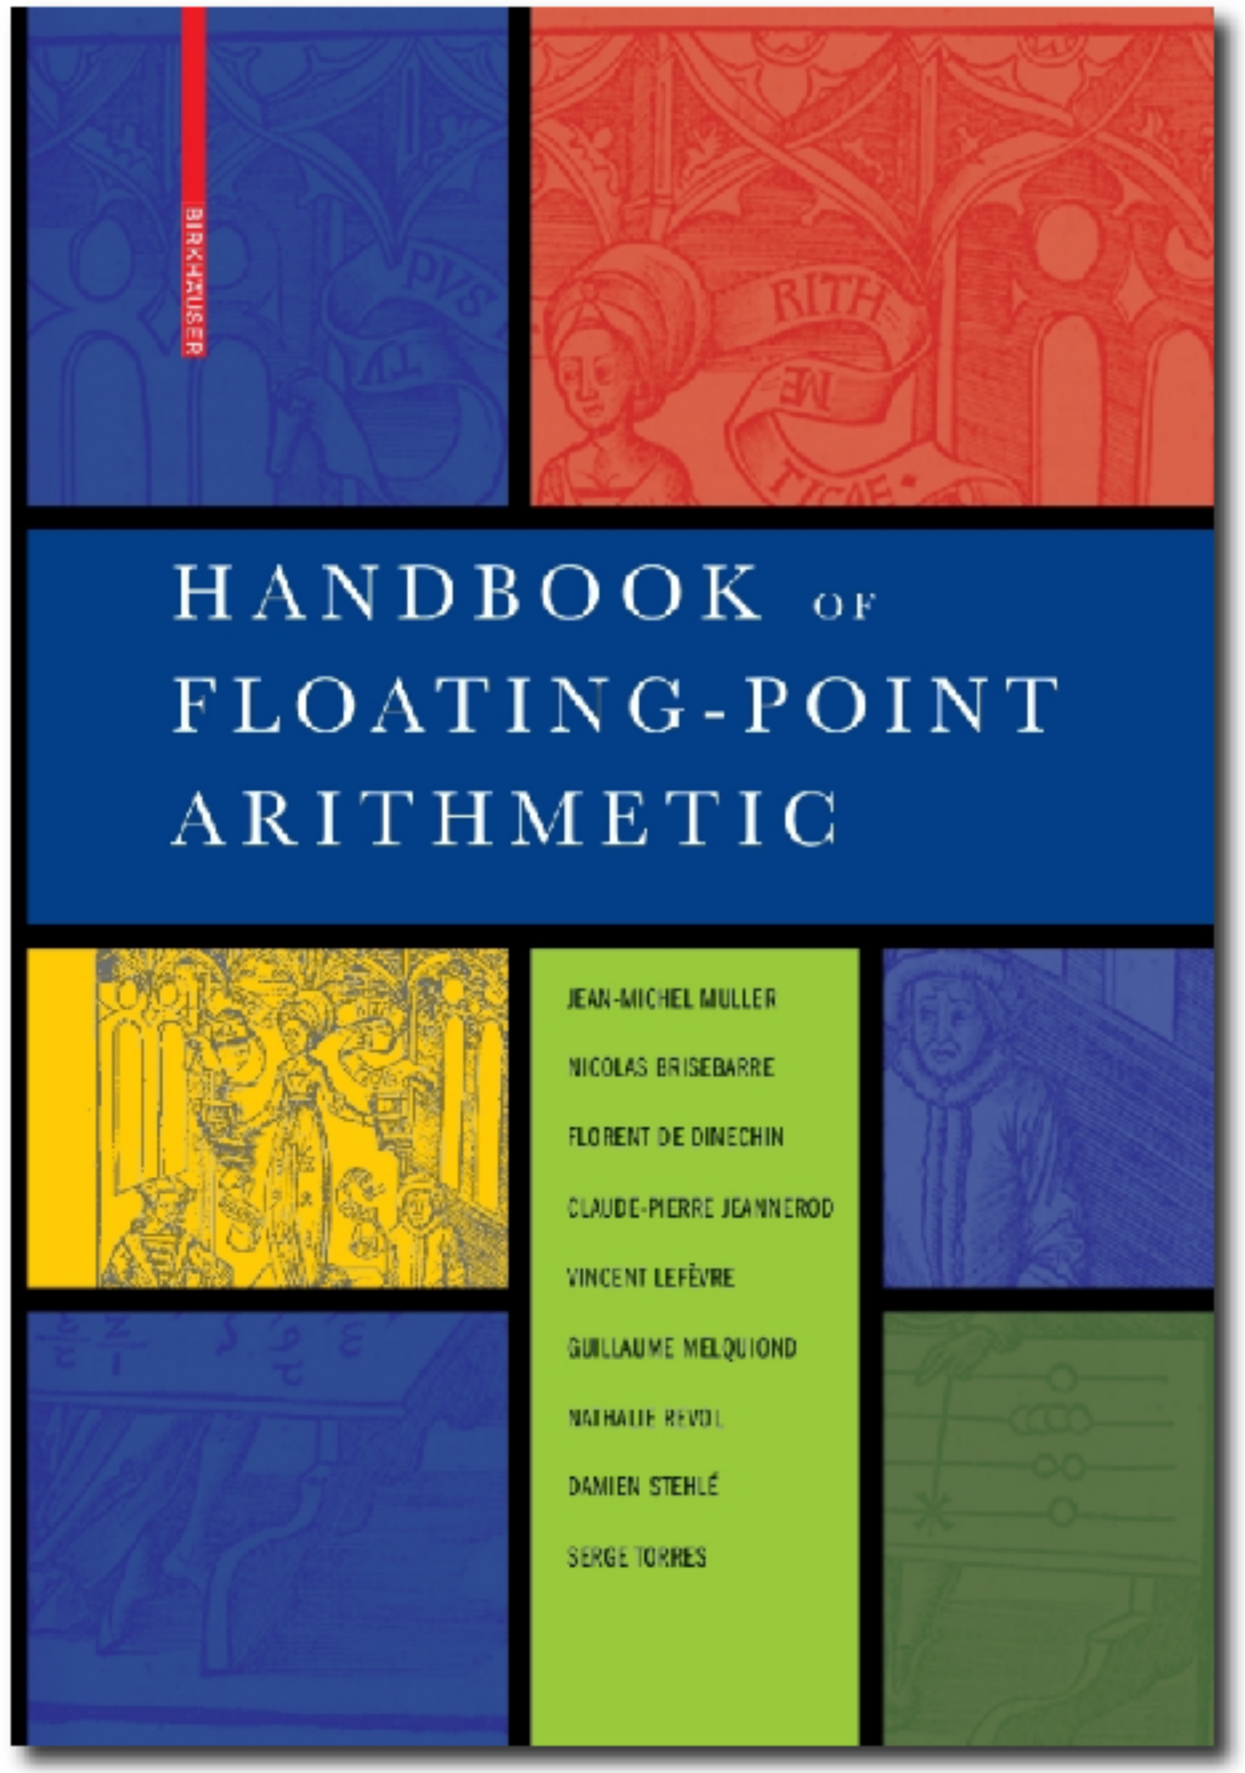

Why Number Precision?

		TFLOP/s		
Model	GB/s	Single	Double	Ratio
AMD Radeon R9 Nano	512	8.19	0.51	16
AMD FirePro W9100	320	5.24	2.62	2
Intel Xeon E5-2699 v4	77	1.55	0.77	2
Intel Xeon Phi 7120A	352	2.42	1.21	2
NVIDIA Tesa K40c	288	4.29	1.43	3
NVIDIA Tesa M40	288	7.00	0.21	32

Potential Speedups

- If a code region is limited by…
  - FLOPs = **2× to 32×**
  - Memory bandwidth = **2x**
  - Disk I/O = **2x**
  - Latency (memory, disk, network, …) = **1x**

The Status Quo

- Extensive research in bars indicates that, if given the choice between a single and a double measure, the double wins every time.

- CFD codes are no exception.

Do We Need Double Precision?

- Very little research in the CFD space.
- Results mostly limited to **steady state** computations where double precision **does appear to be necessary**.

Methodology

- Rerun several of our **previous published** test cases using single precision arithmetic.
- Compare the results and assess the performance.

Experiments

- Using PyFR we have evaluated several **unsteady viscous test cases**.
  - Taylor–Green vortices.
  - Flow over a circular cylinder.
  - Flow over a NACA 0021.

3D Taylor–Green Vortex

- Standard test case for DG.

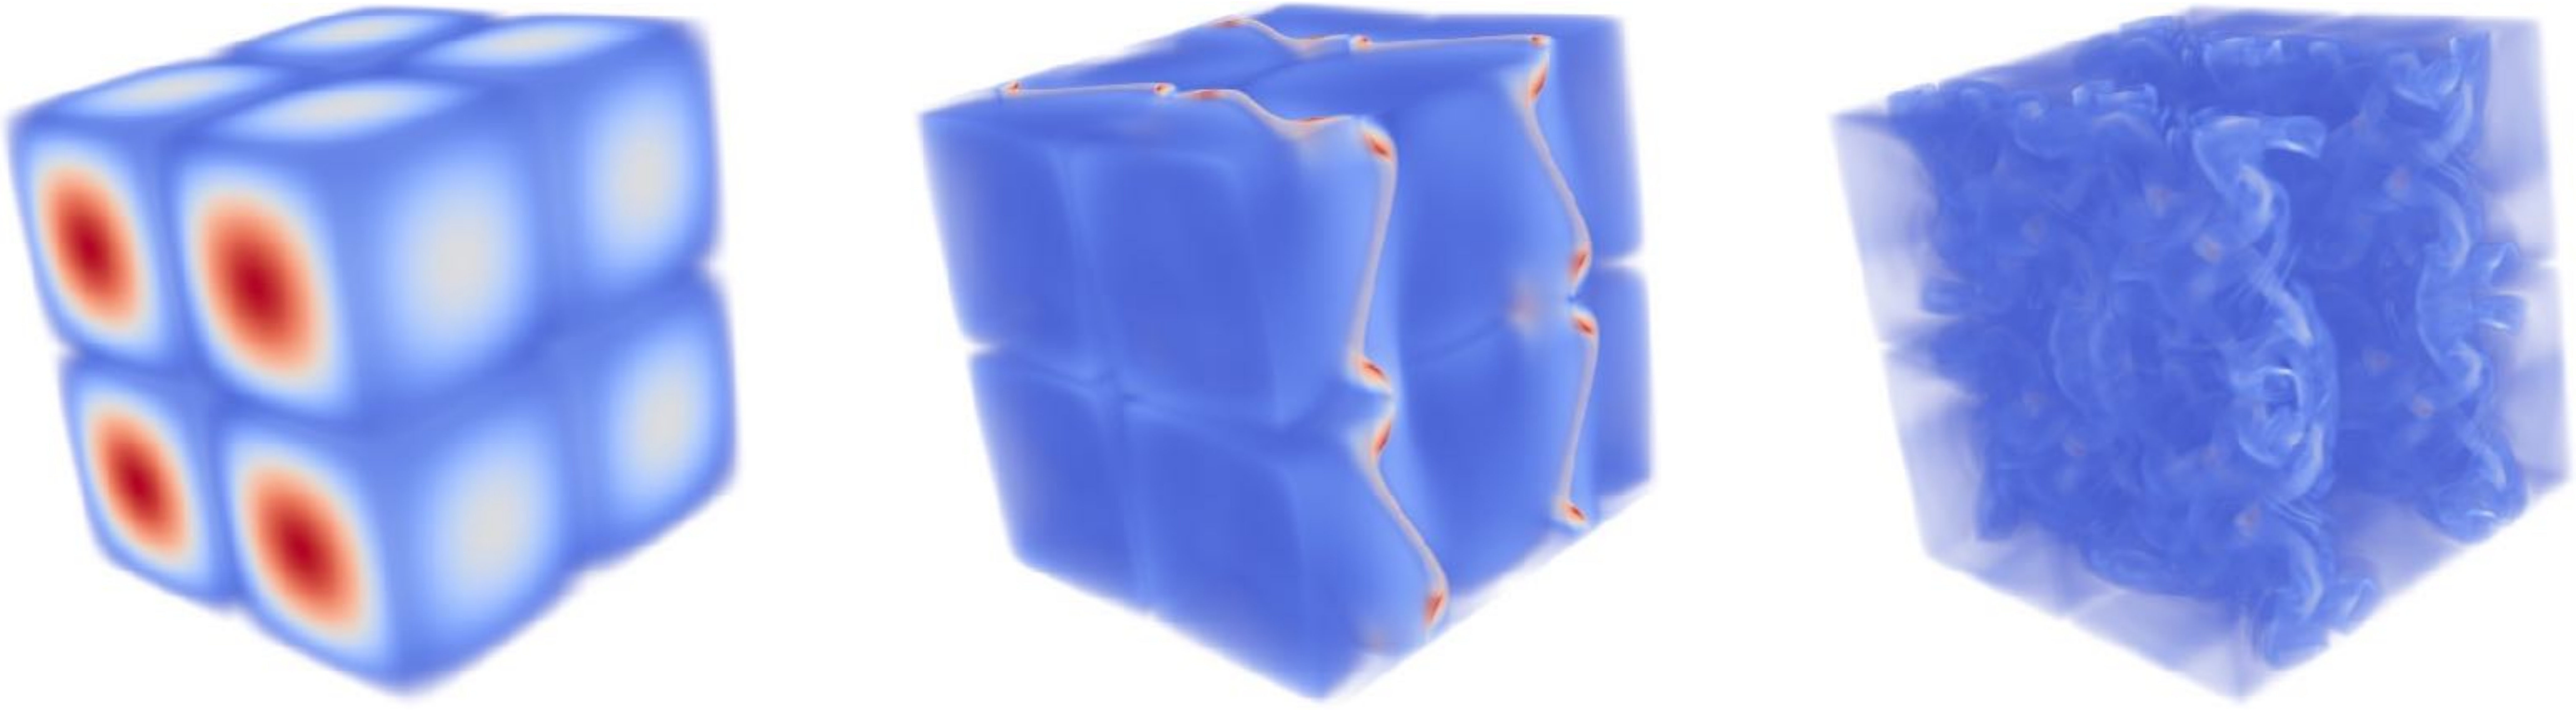

3D Taylor–Green Vortex

- Four **structured grids** with roughly **constant DOF count**.

			Memory / GiB	
Order	N_E	$\sum N_u$	Single	Double
$\wp = 2$	86^3	258^3	6.4	12.2
$\wp = 3$	64^3	256^3	5.4	10.3
$\wp = 4$	52^3	260^3	5.1	9.8
$\wp = 5$	43^3	258^3	4.6	9.0

3D Taylor–Green Vortex

- Consider **kinetic energy decay rate**.
- Compare with **van Rees et al.**
- No difference between single and double.

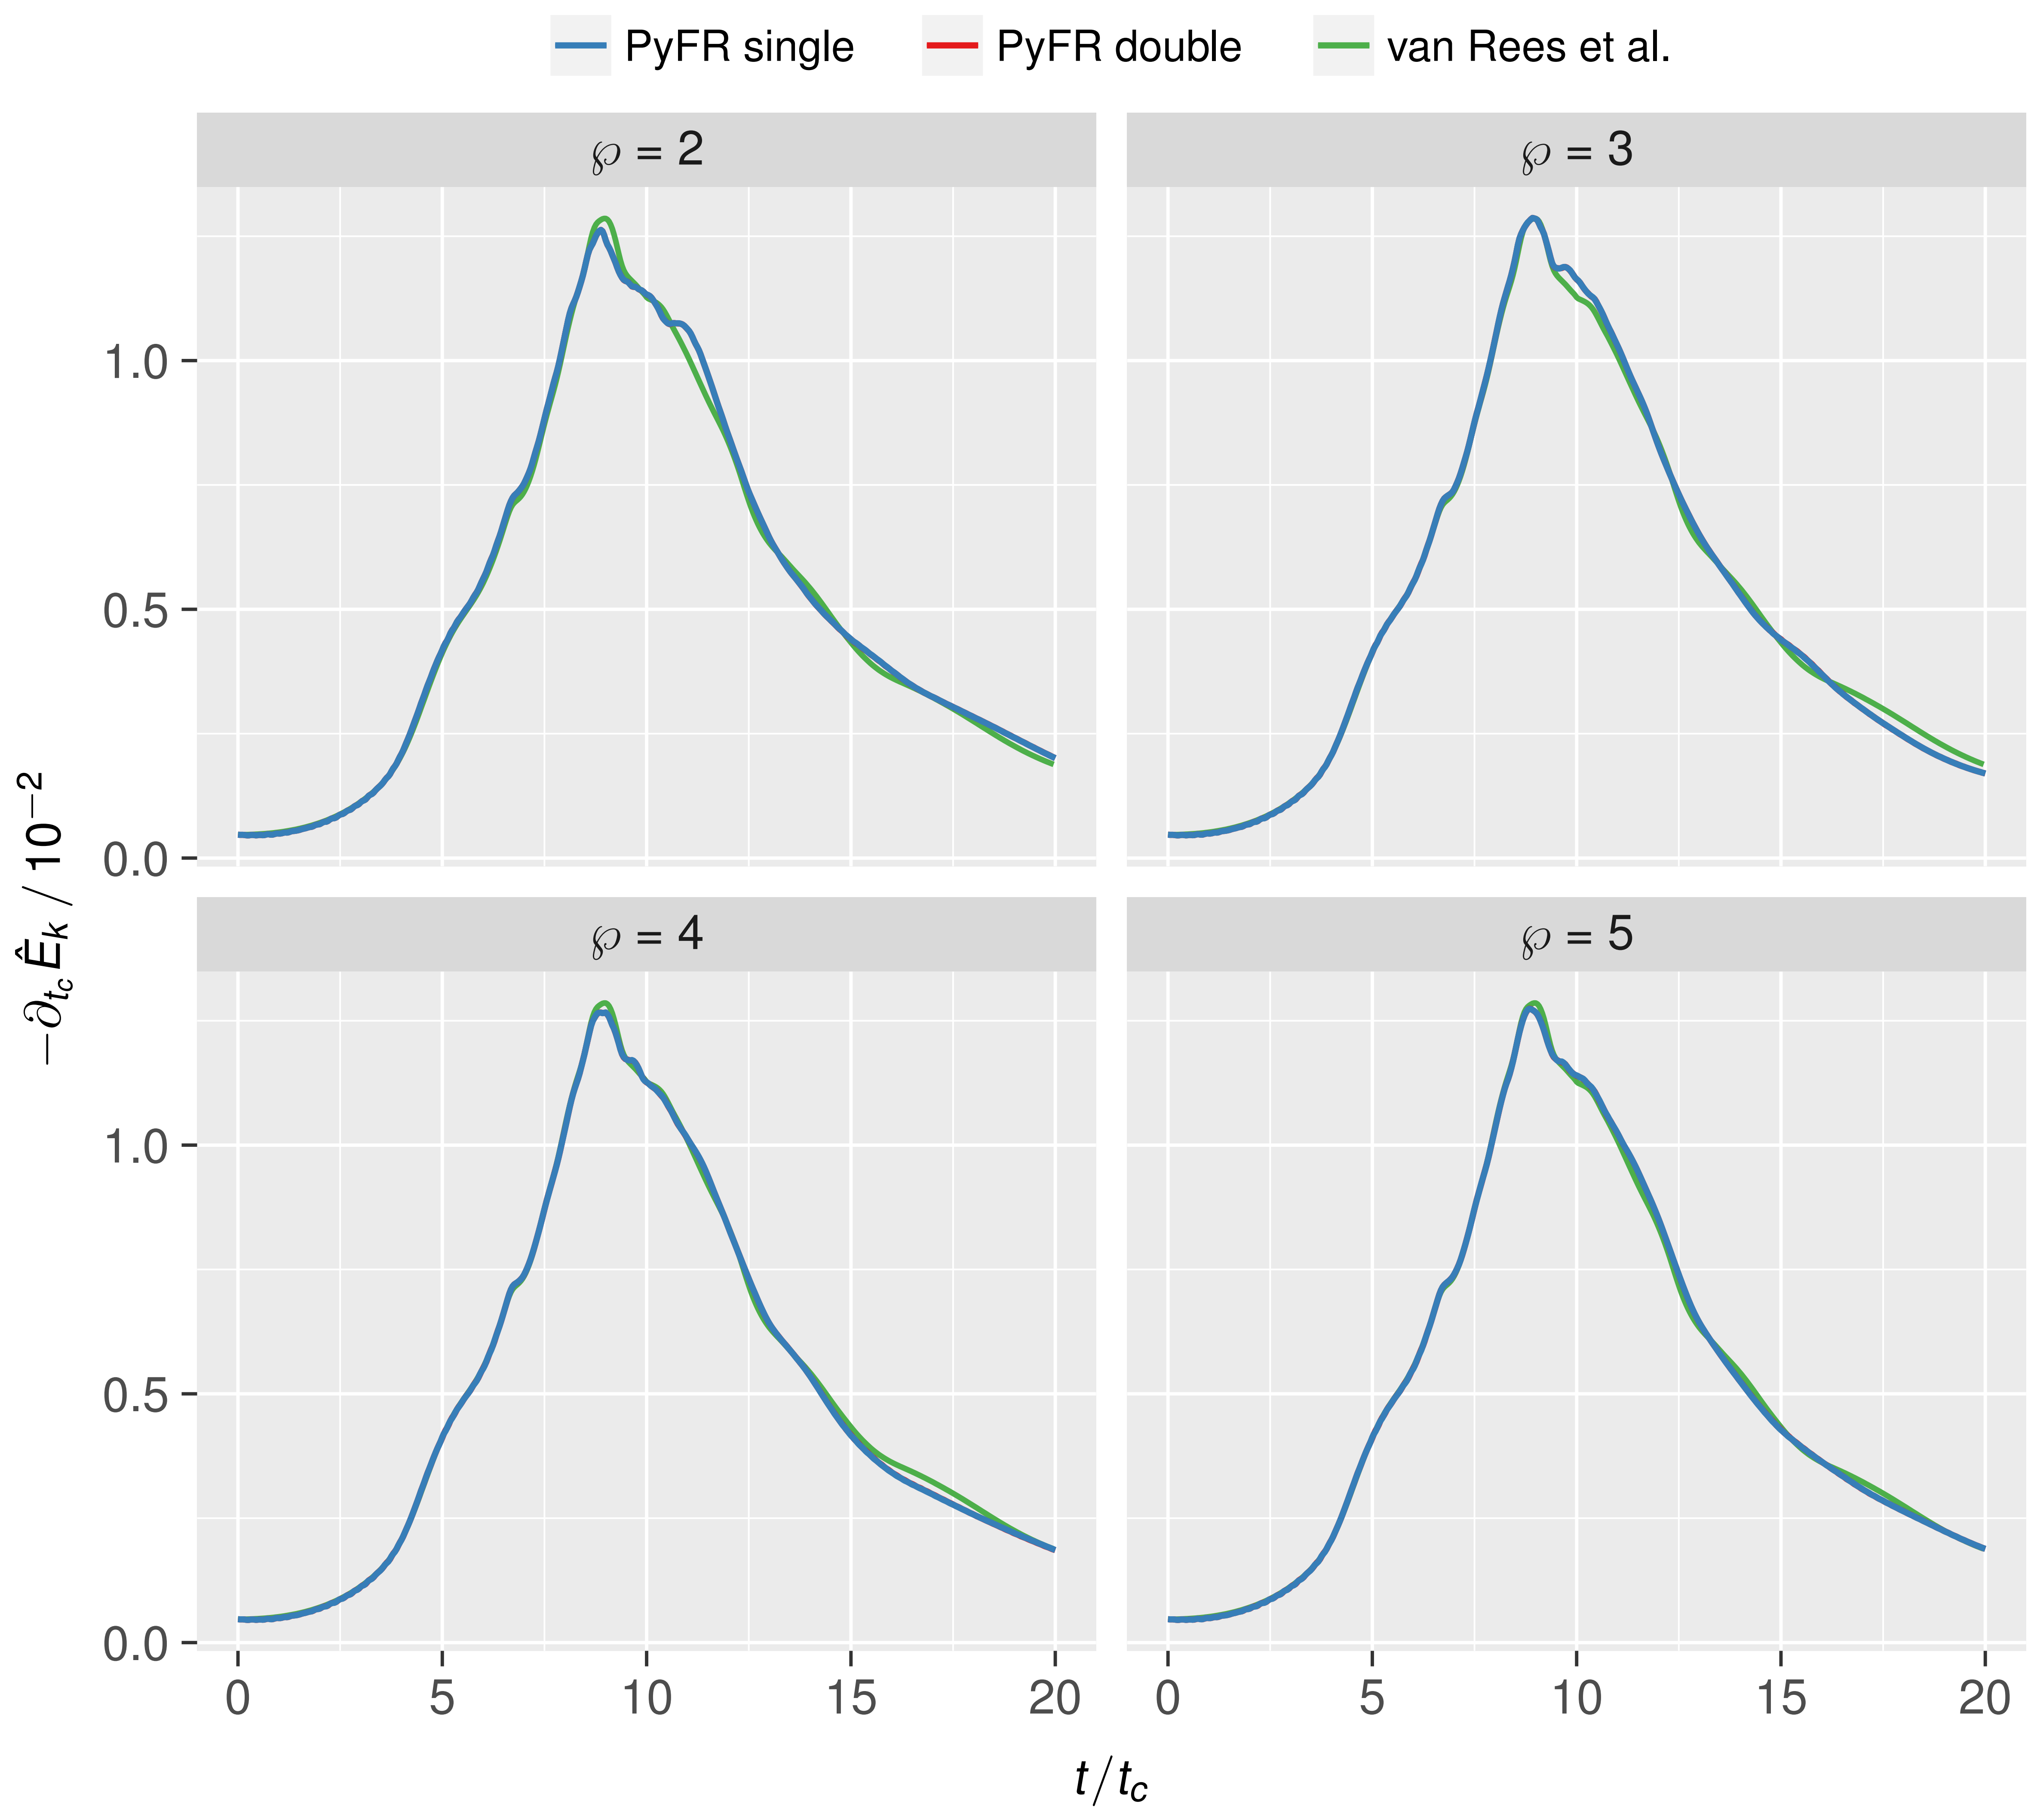

3D Taylor–Green Vortex

- Performance on a **two NVIDIA K40c's** with GiMMiK.

		$t_w / \sum N_u$ / 10^{-9}s		GFLOP / s		
Order	GFLOP / stage	Single	Double	Single	Double	Speedup
$\wp = 2$	1.84×10^1	4.8	8.9	222.1	120.5	1.84
$\wp = 3$	1.82×10^1	4.2	7.9	252.3	134.6	1.88
$\wp = 4$	1.92×10^1	4.4	8.6	255.9	129.7	1.97
$\wp = 5$	1.96×10^1	4.5	13.1	250.8	87.0	2.88

Flow Over a Cylinder

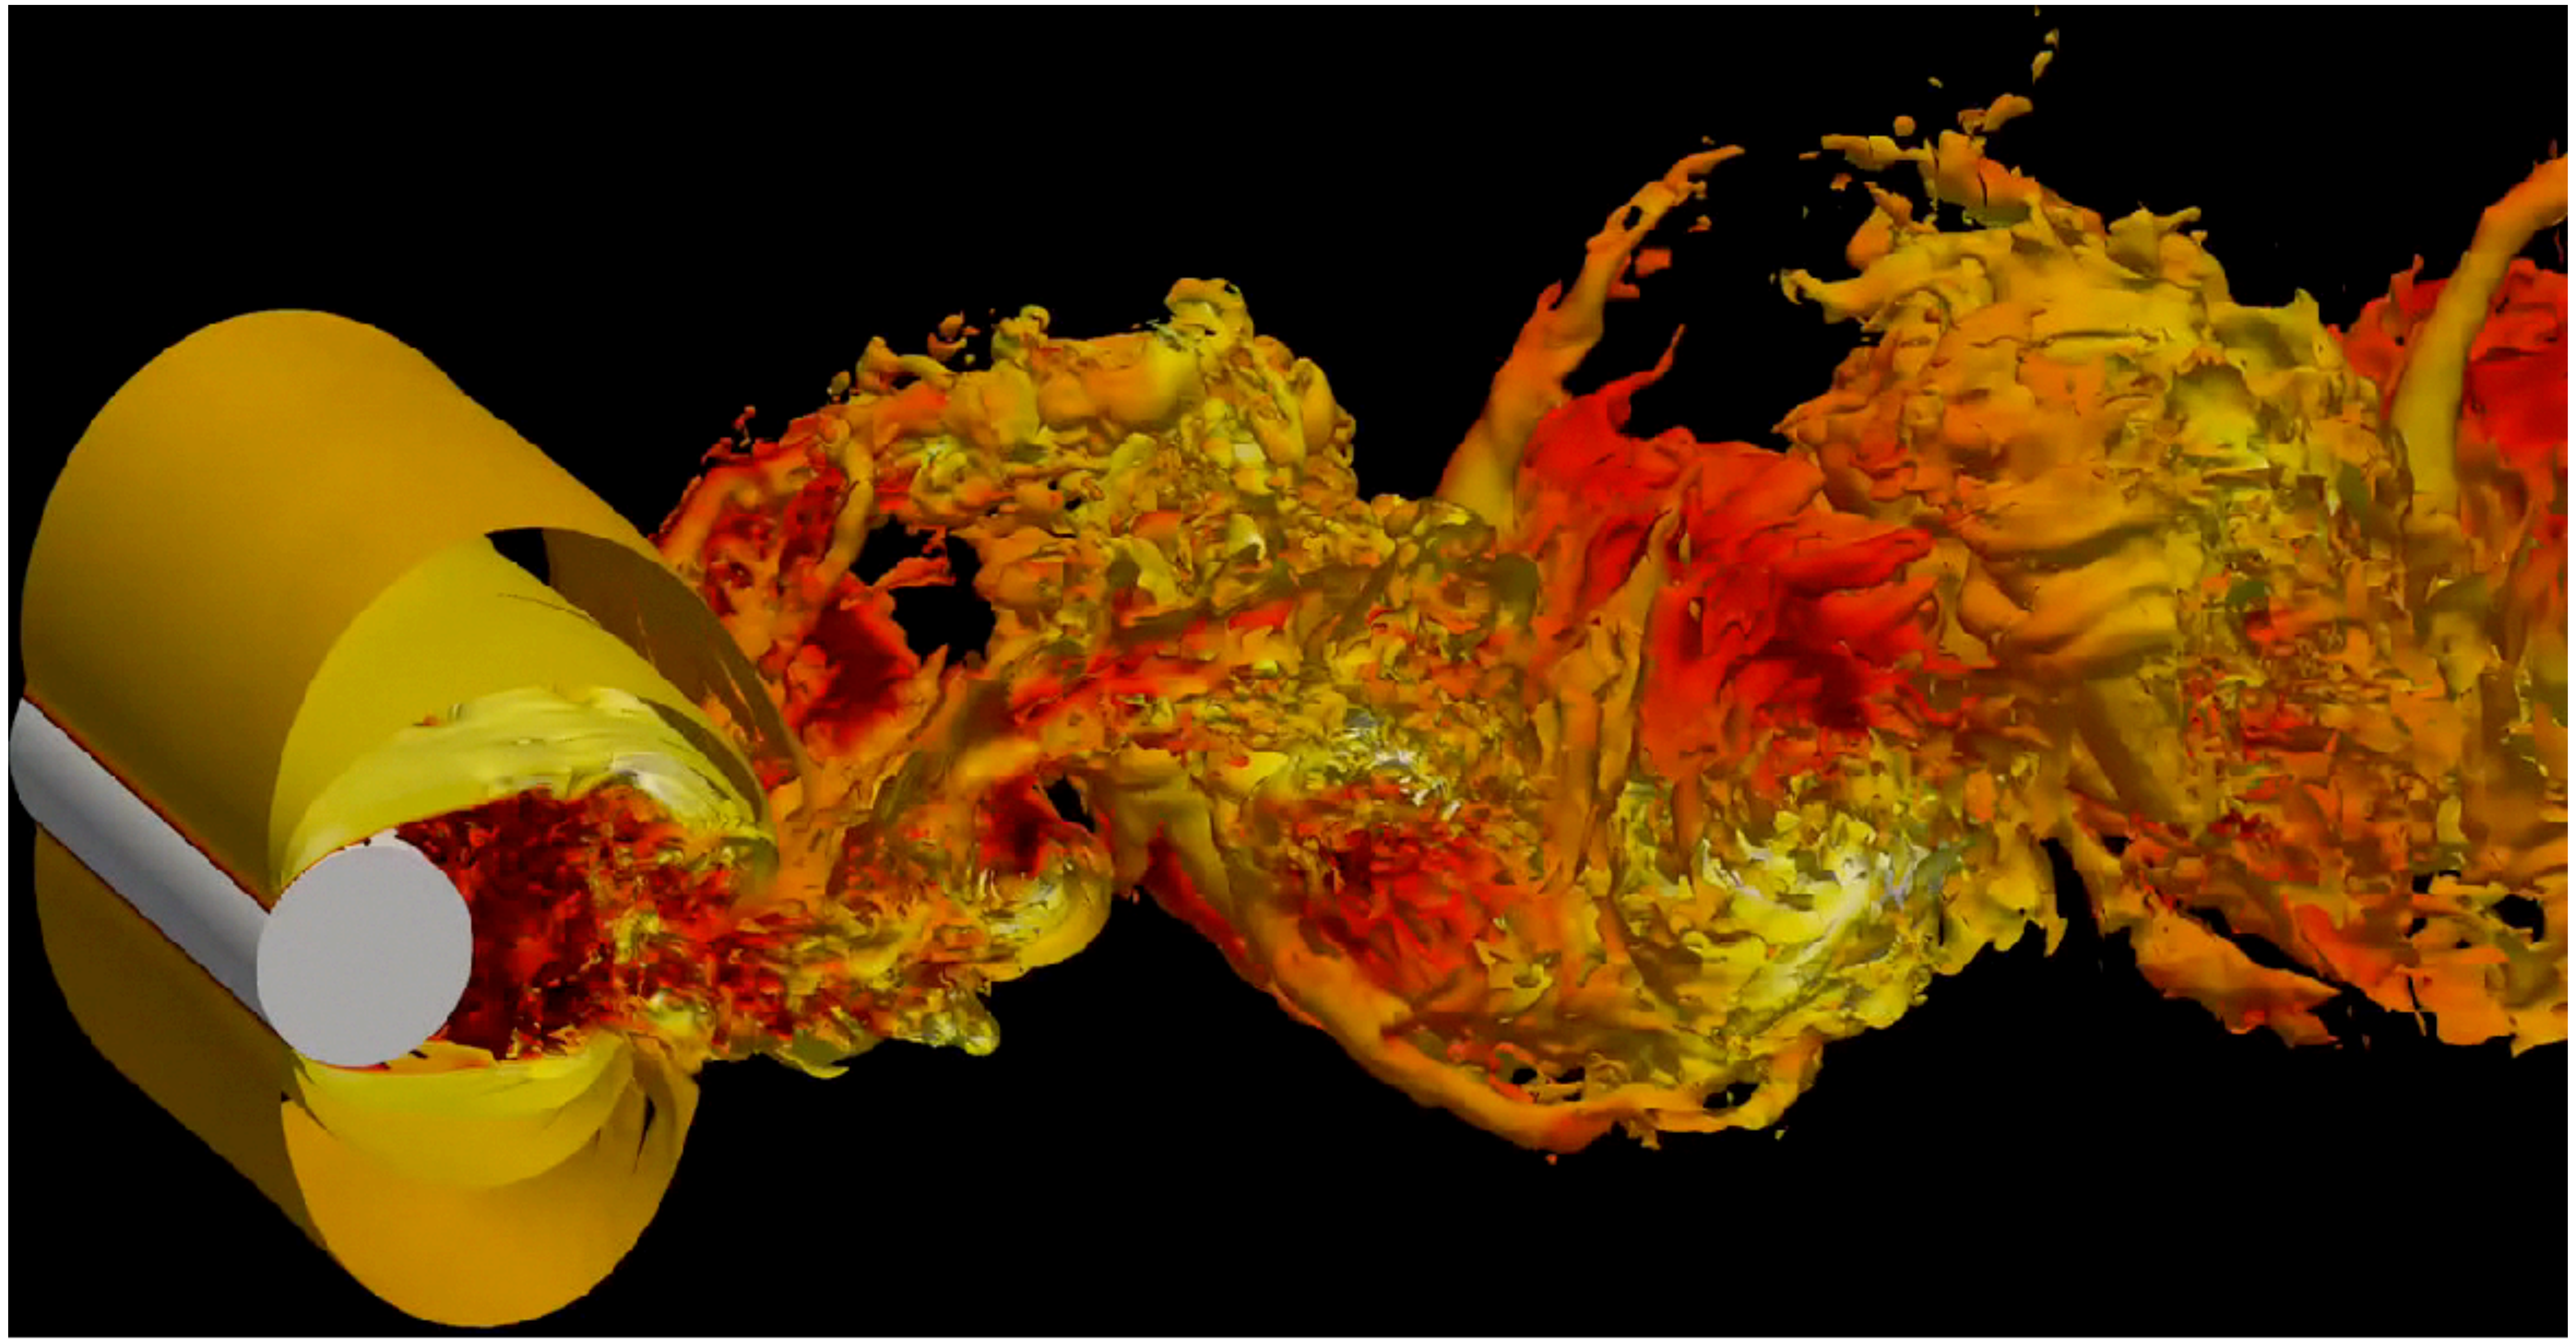

Flow Over a Cylinder

- Cylinder at **Re = 3900**, and **Ma = 0.2** with ***p* = 4**.
- **Mixed prism/tet grid** of span πD.

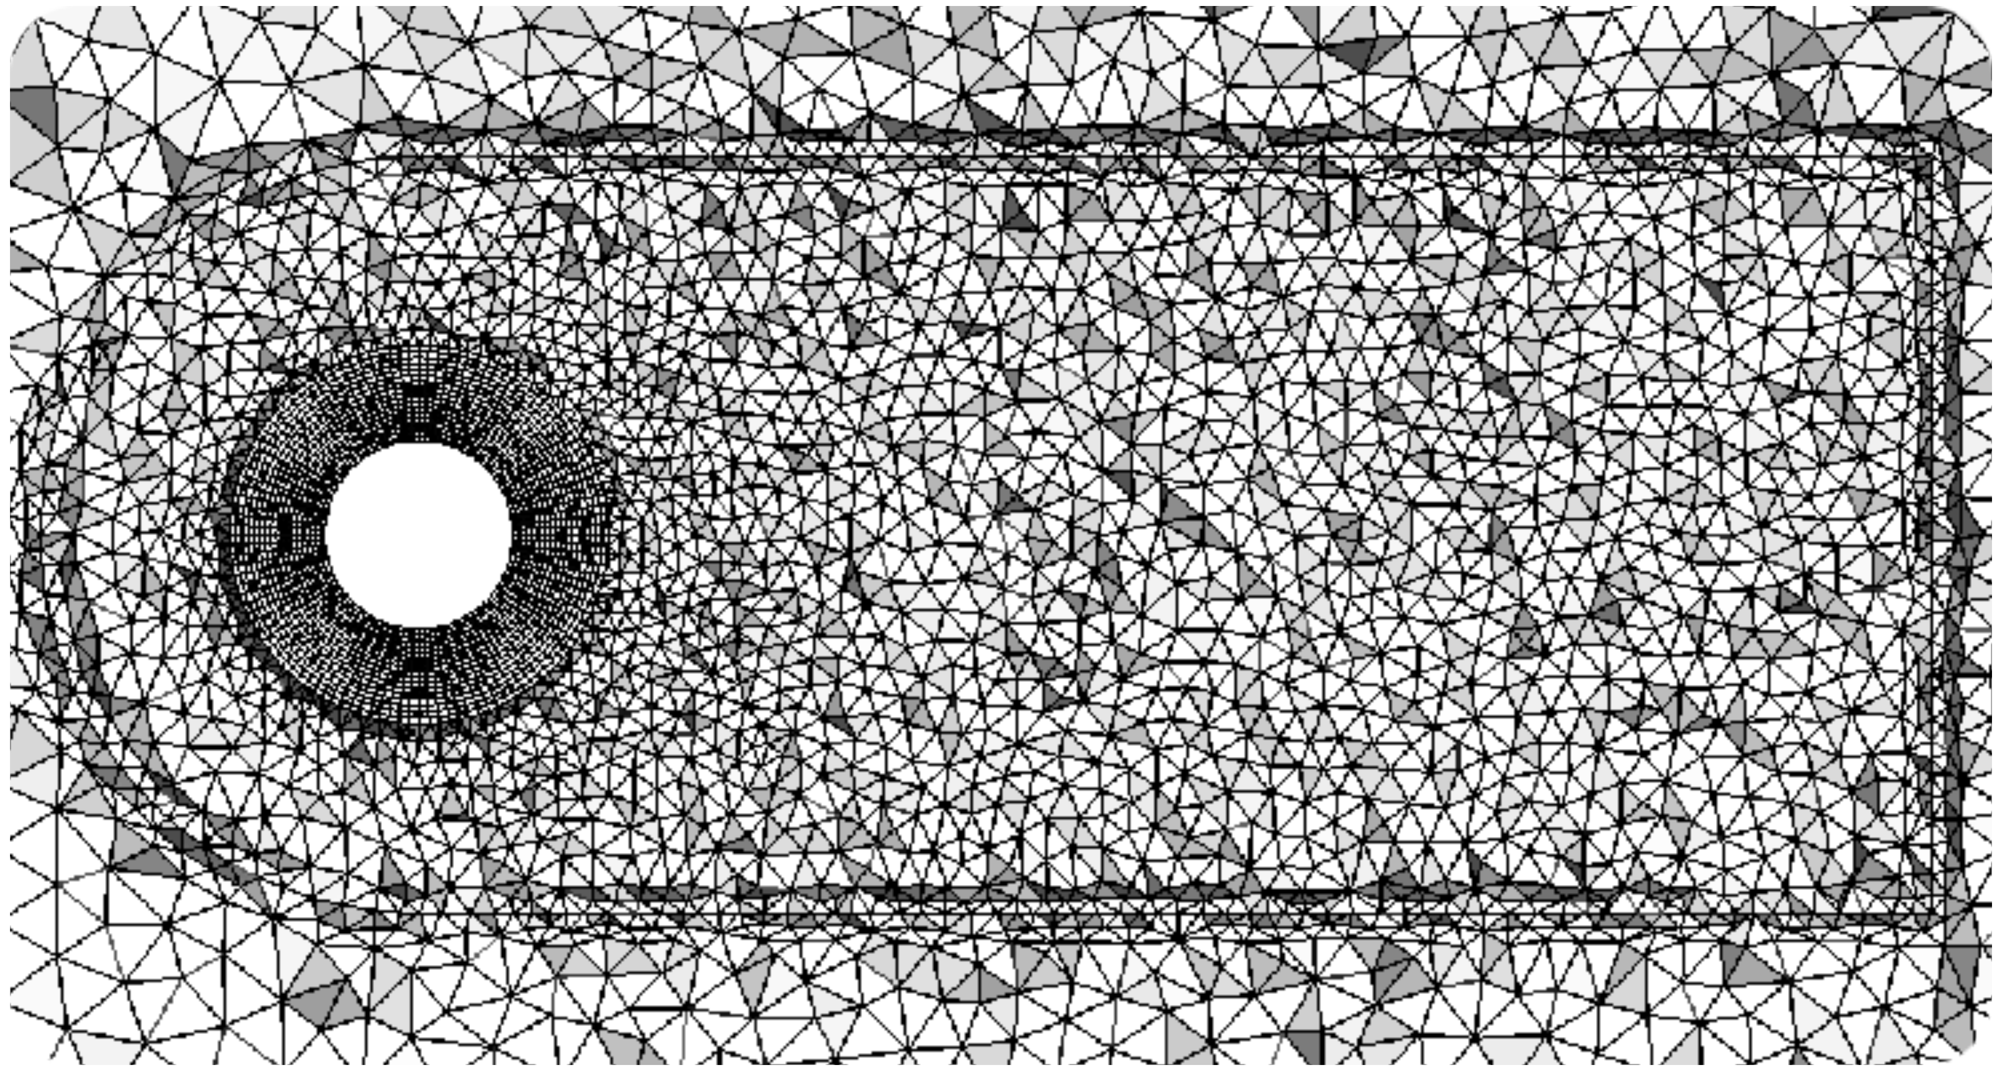

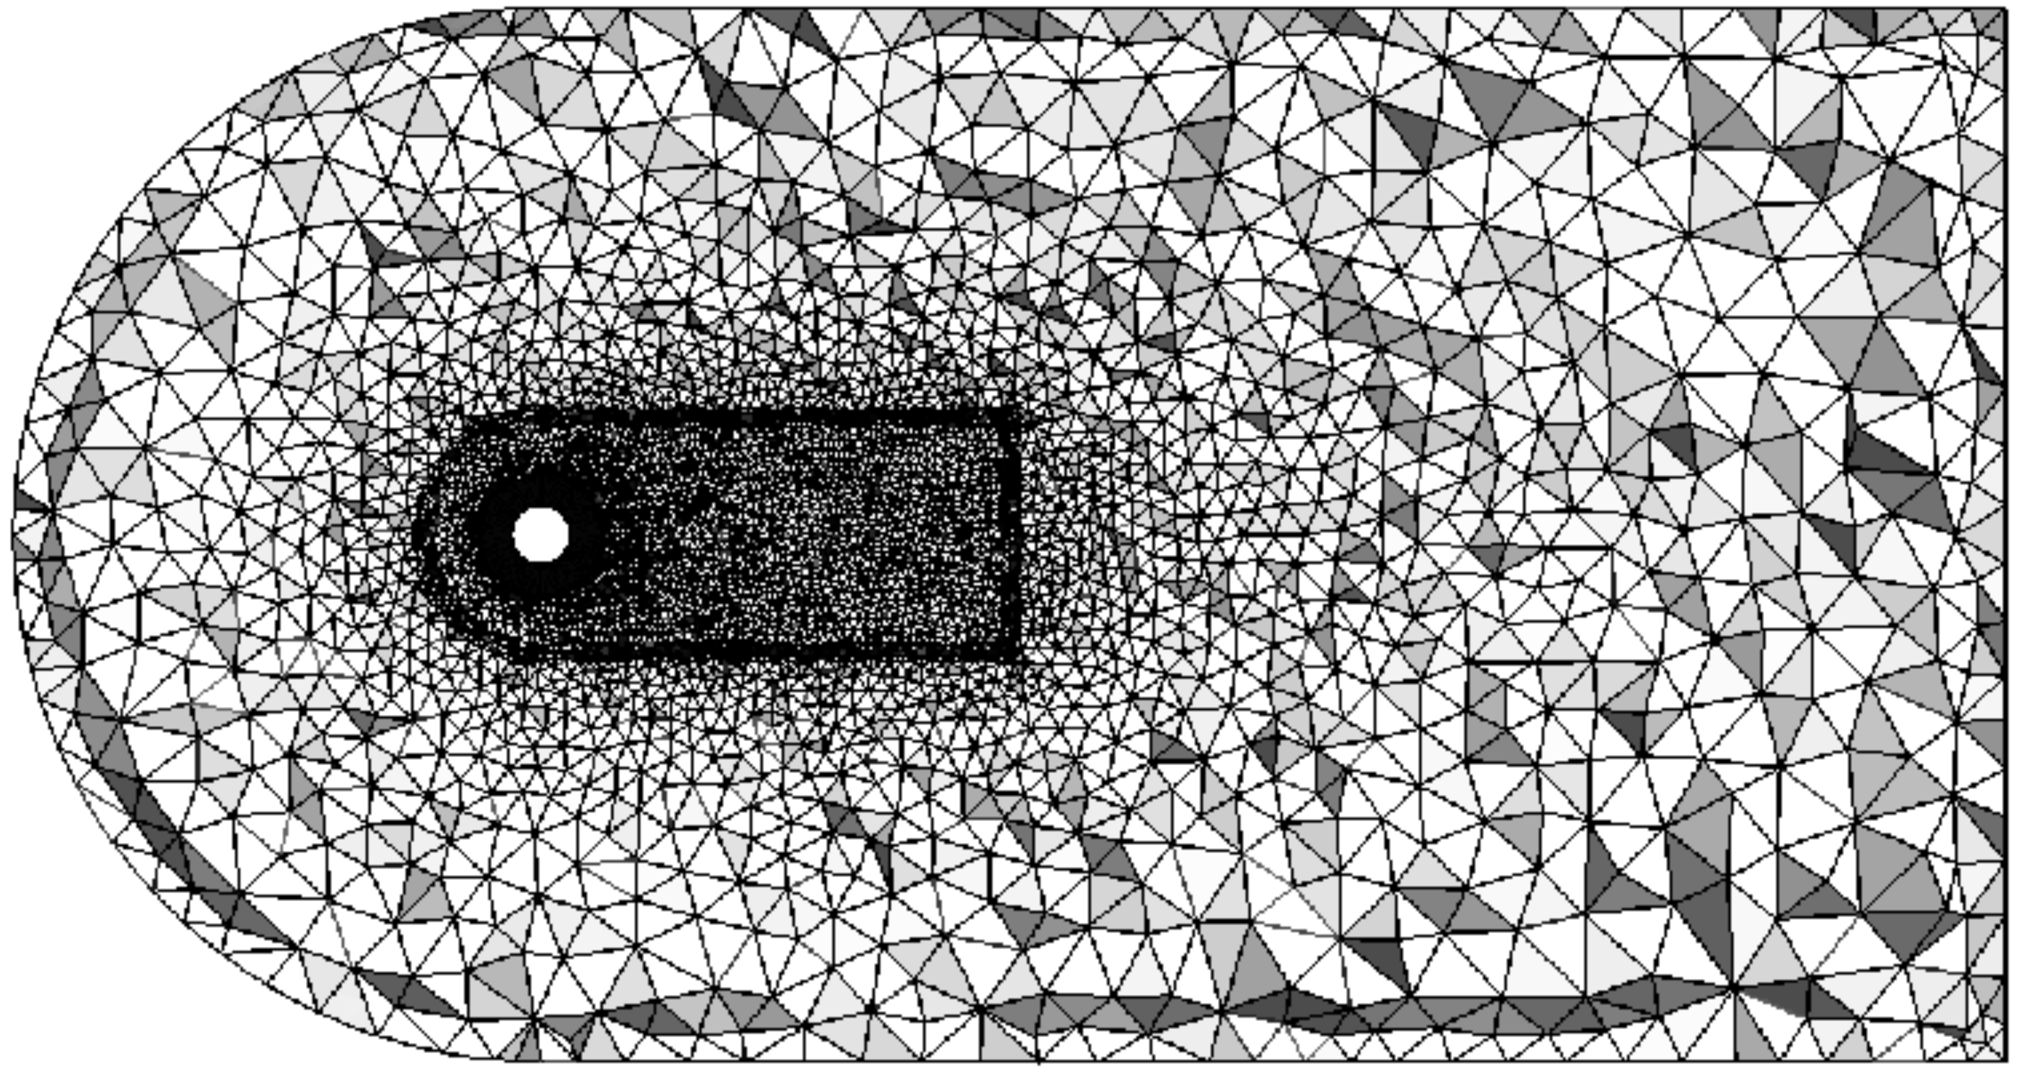

Flow Over a Cylinder

- Pressure coefficient on the surface.
- Compare with **Lehmkuhl et al**.

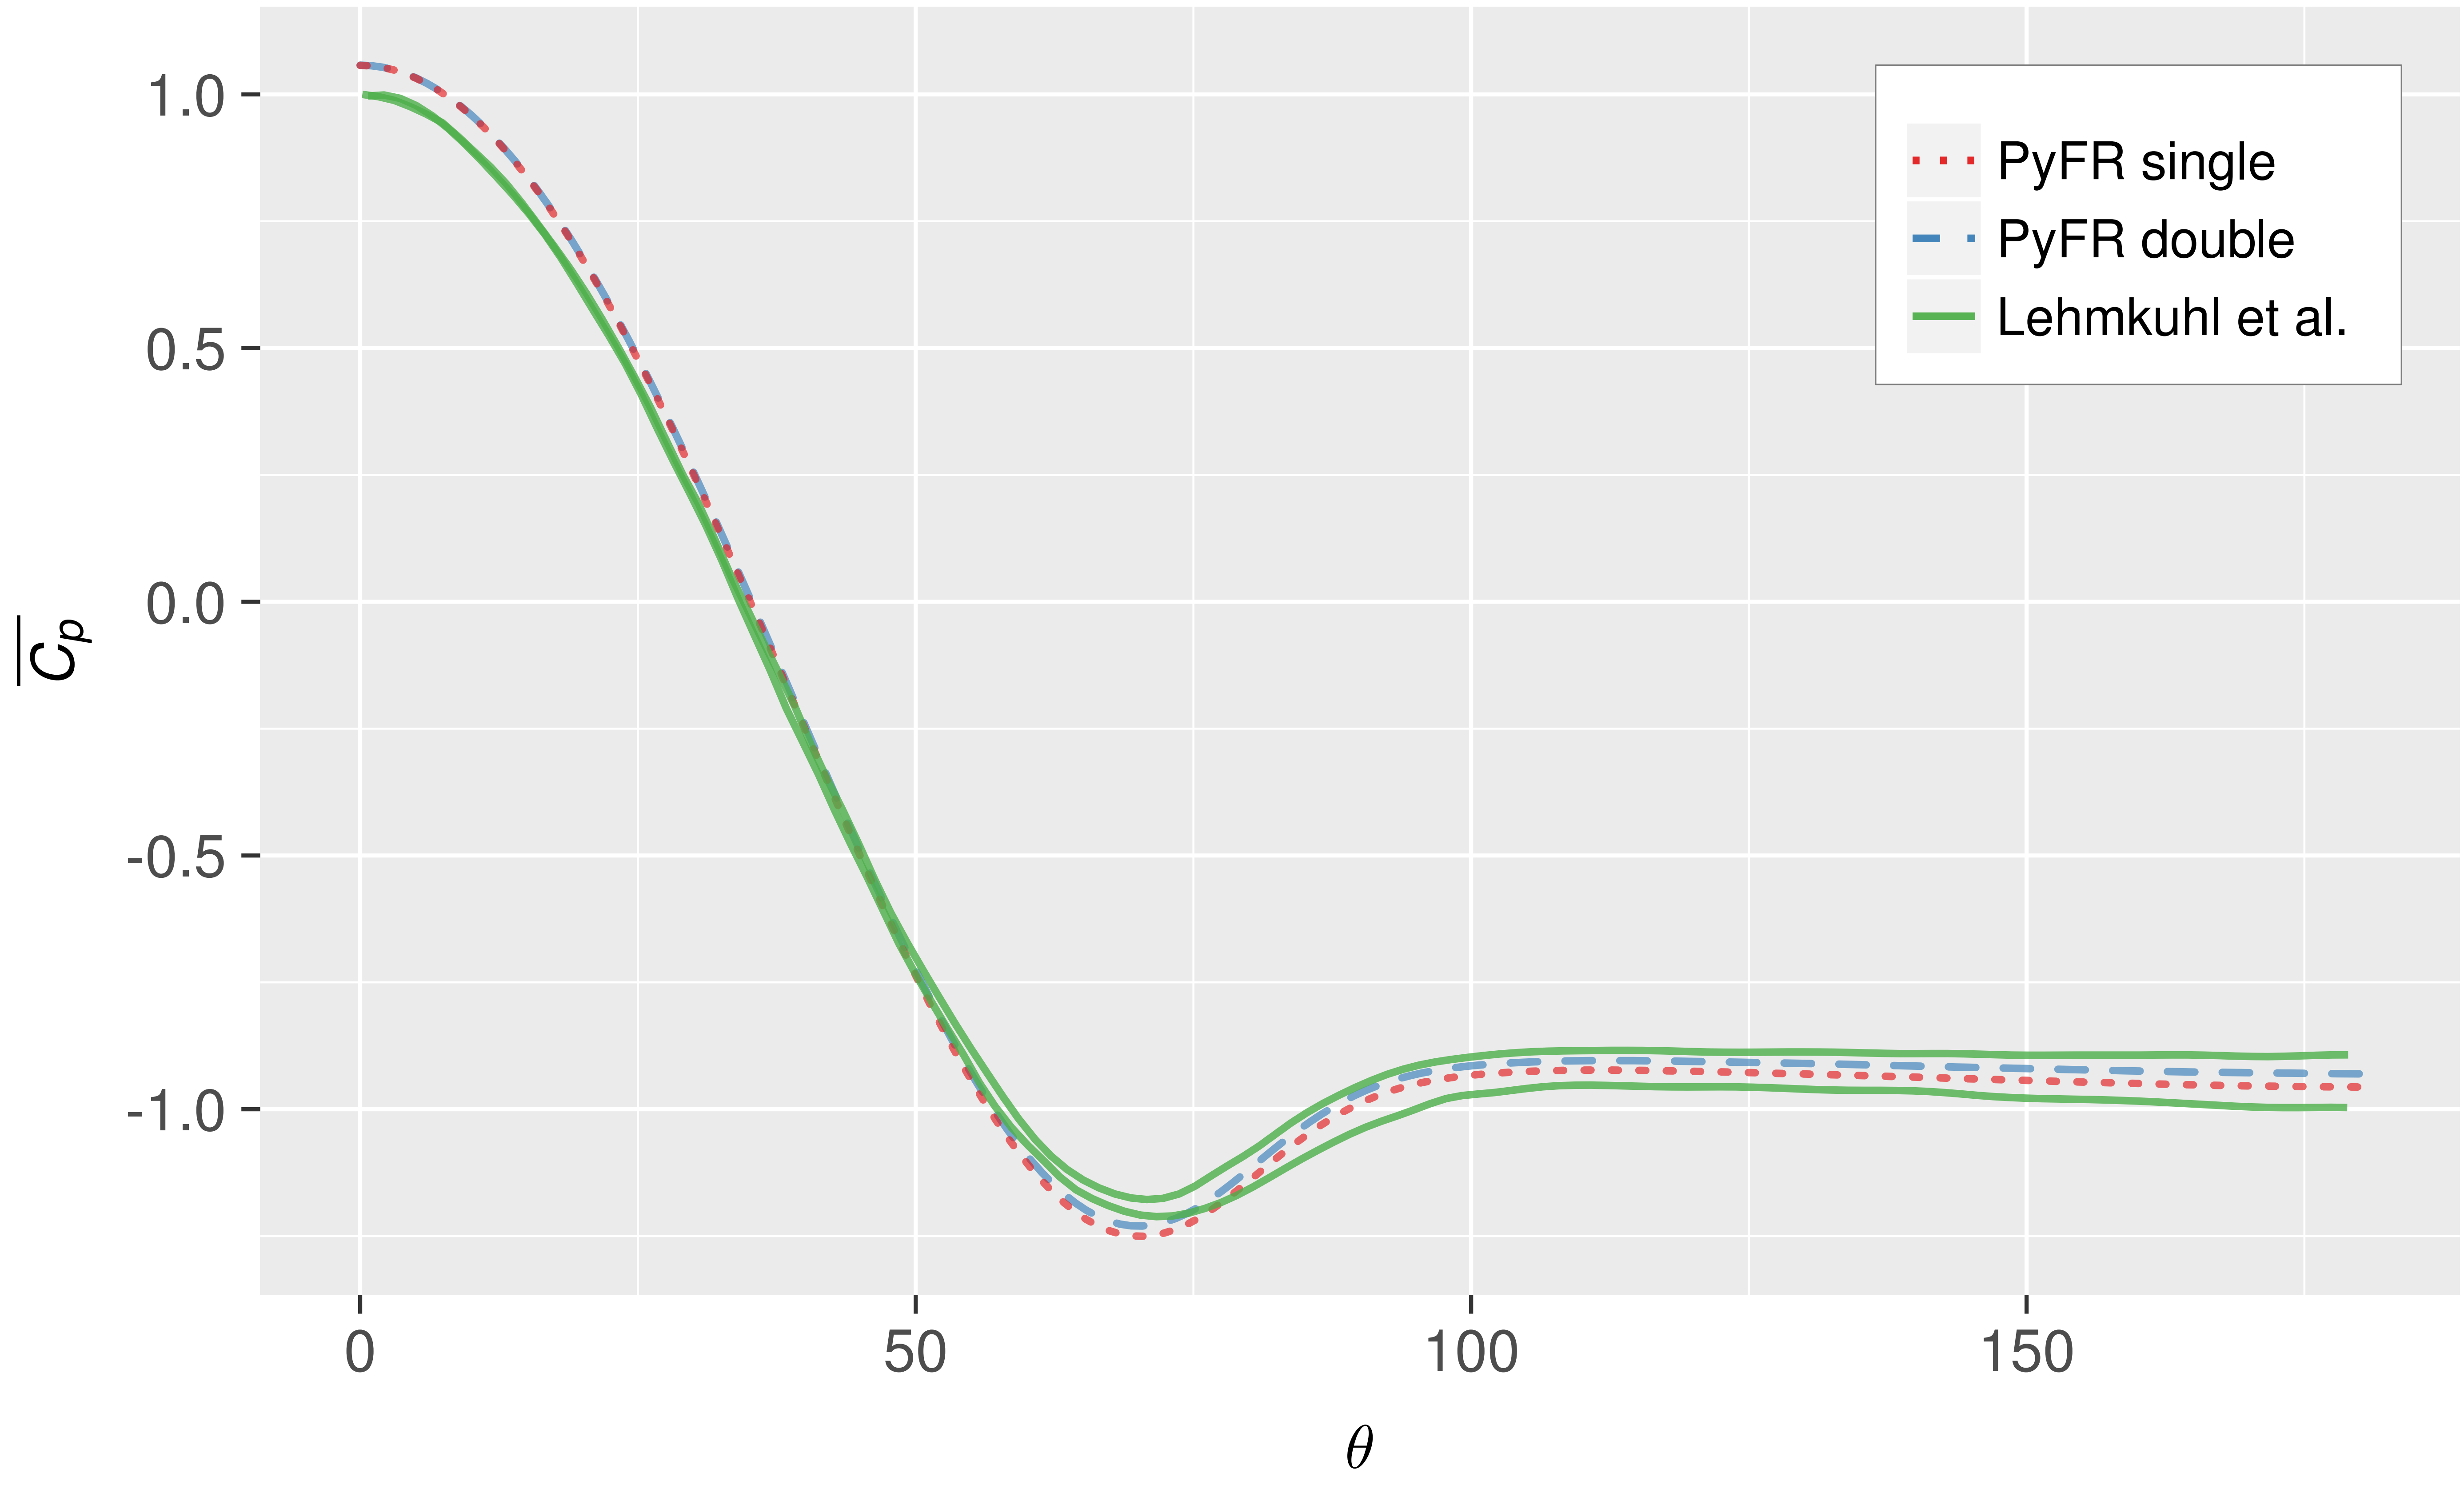

Flow Over a Cylinder

- Performance on a single **NVIDIA K40c** with GiMMiK.
  - Tet operator matrices are **small** and prisms **sparse**.
  - Overall speedup of **~1.6**.
- Simulation results in **heavy indirection**; thus experiences **less of an improvement** from single precision.

NACA 0021

- Flow over a NACA 0021 at 60 degree AoA.
- **Re = 270,000** and **Ma = 0.1**.
- Compare with experimental results of **Swalwell**.

NACA 0021

- **206,528** hexahedral elements.
- Span is four times the chord.
- Fourth order solution polynomials with **full anti-aliasing**.

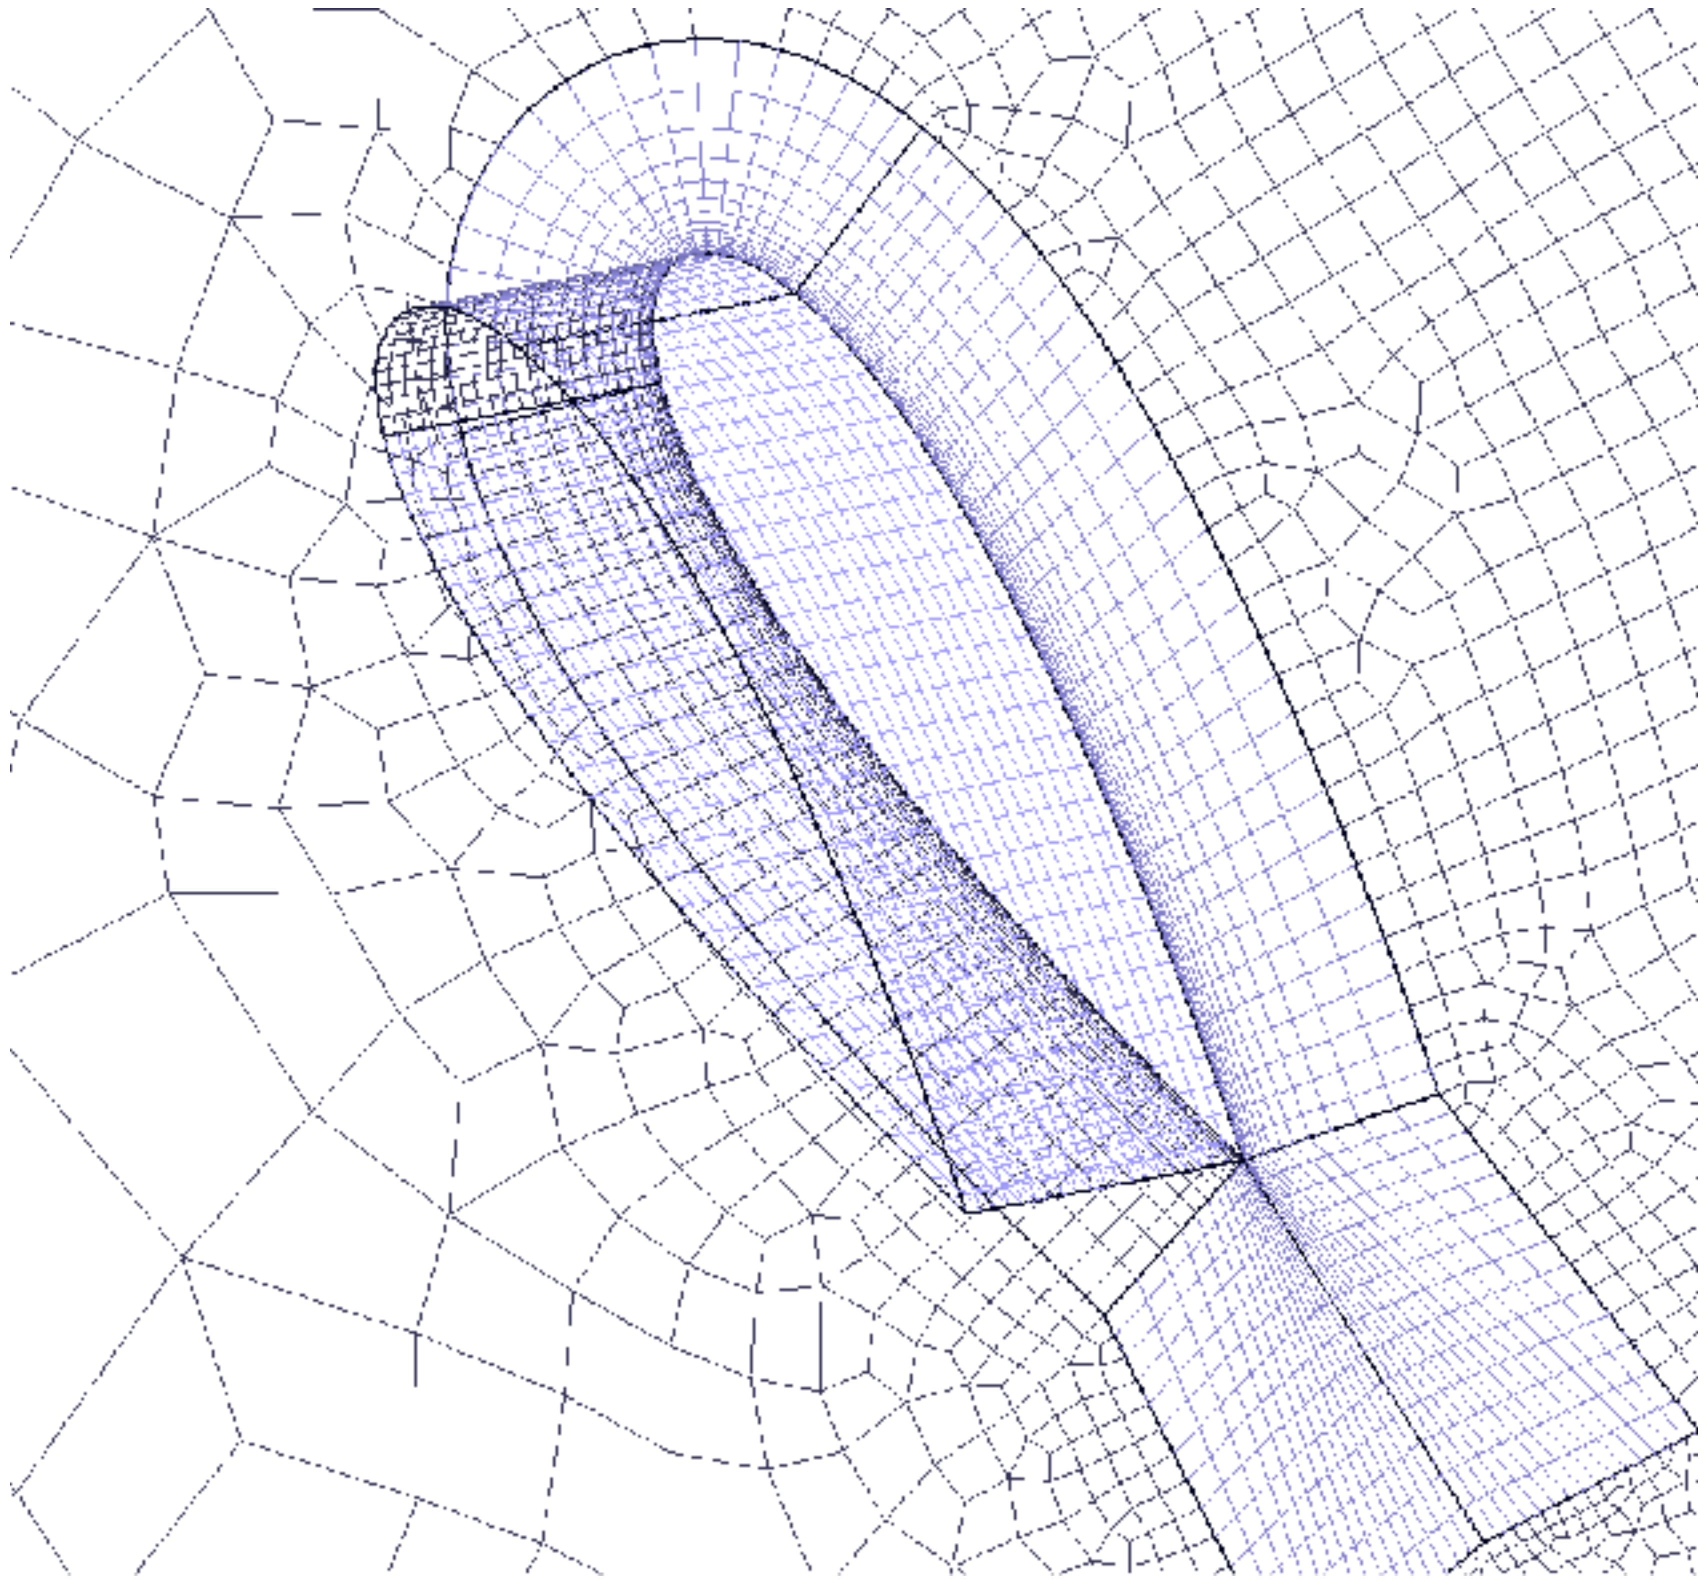

NACA 0021

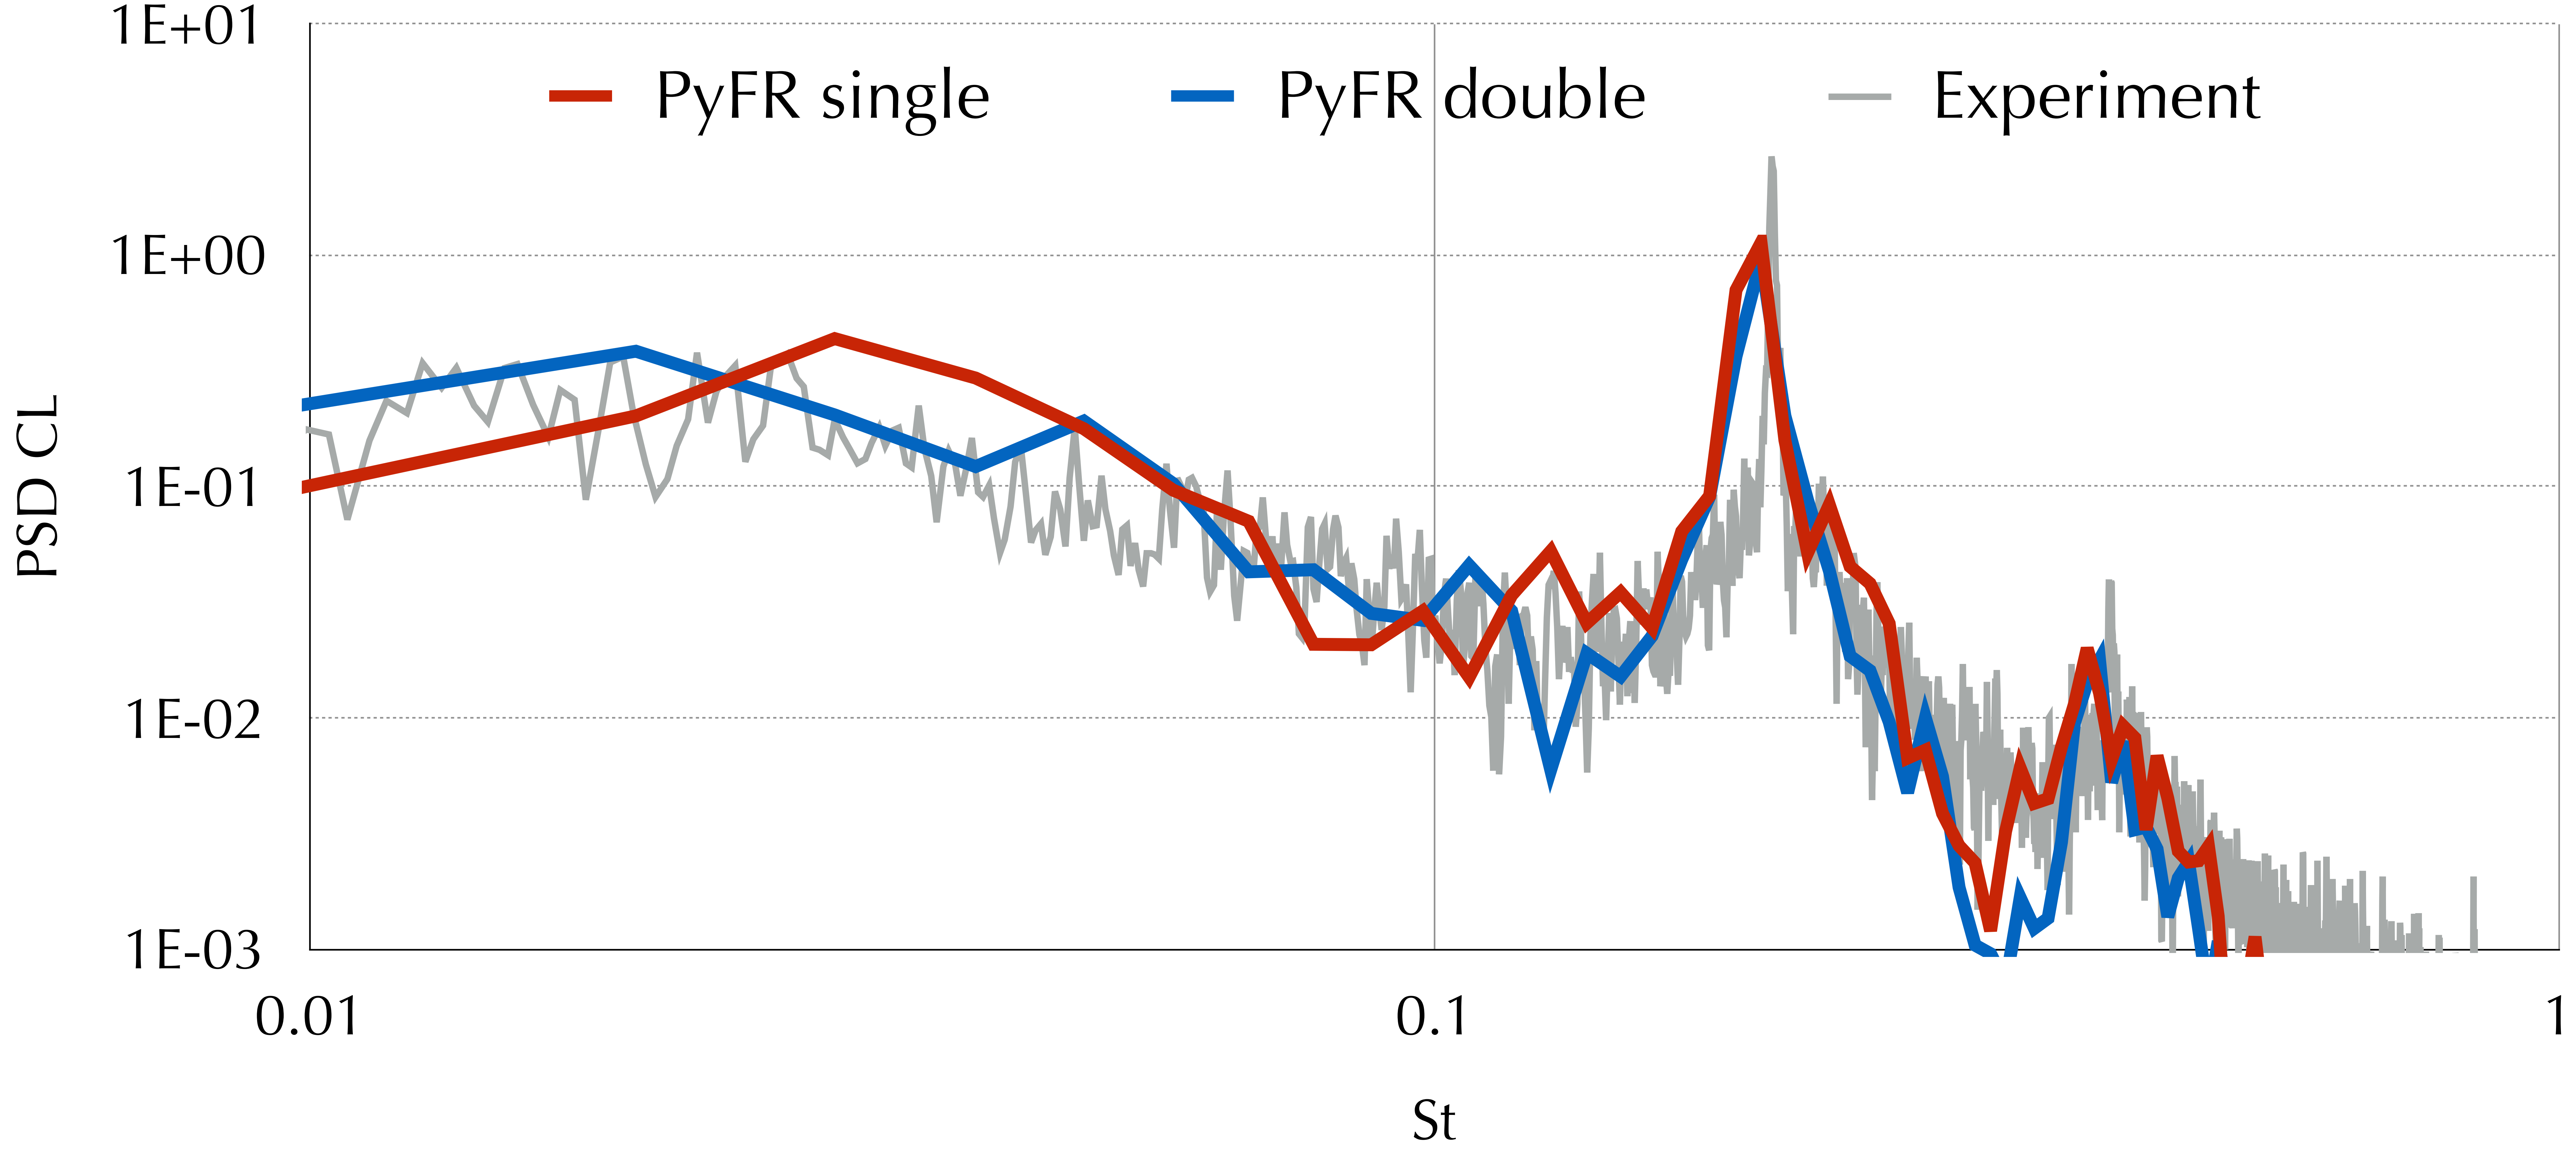

PyFR single
PyFR double
Experiment
1E+01
1E+00
1E-01
1E-02
1E-03
PSD CL
0.01
0.1
1
St

NACA 0021

- Performance on **16 NVIDIA K80's** (32 GPUs).
- All operators are **dense**.
- Near the limit of strong scaling.
- Overall speedup of **~1.8**.

Remarks and Closing Thoughts

For LES single precision is sufficient.

Acknowledgements

- Air Force Office of Scientific Research for their support under grant FA9550-14-1-0186.
- Margot Gerritsen for access to XStream.
- NVIDIA for hardware donations.