AutoCFD HPC

Herbert Owen - Barcelona Supercomputing Center
Benet Eiximeno - Nvidia

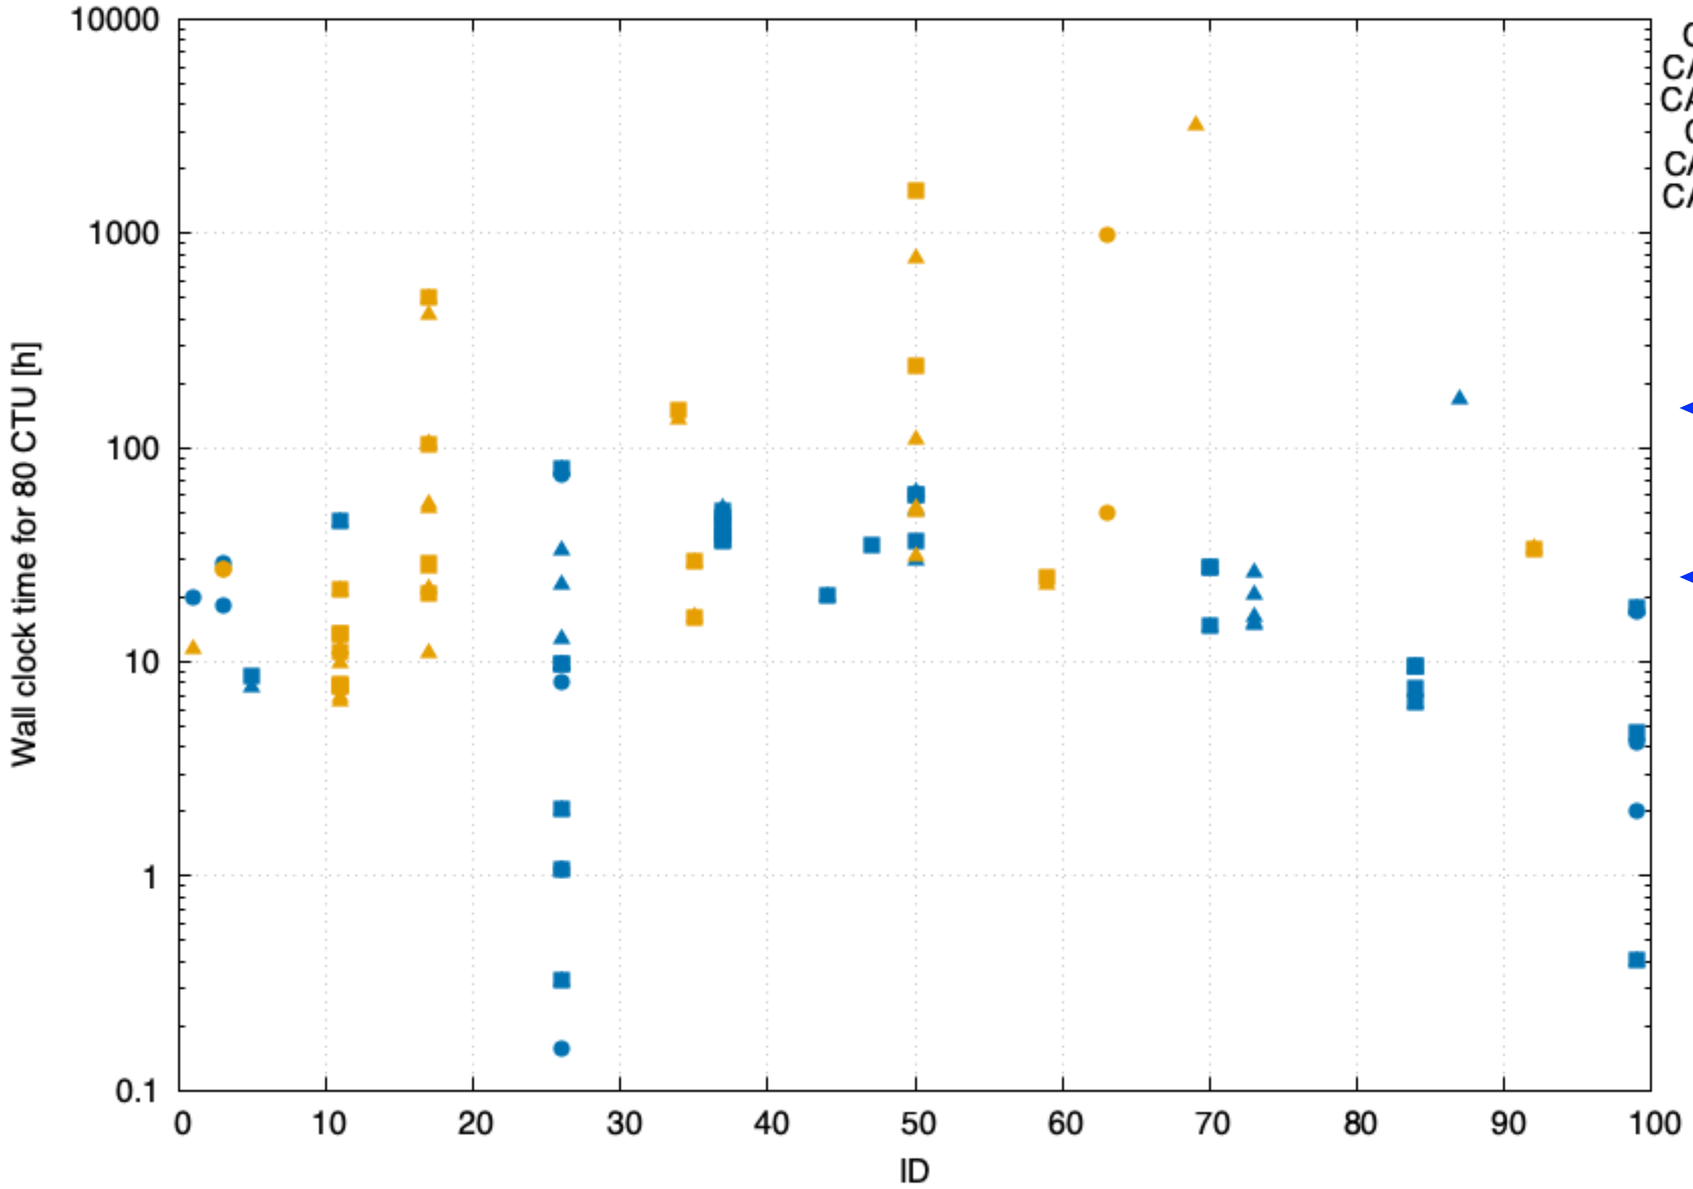
Motivation

How to somehow analyse the cost of different runs

- codes and numerical methods
- mesh sizes & time step sizes
- total physical simulation time
- performed on different hardware

First solution extrapolate to a run for 80 CTU (flow throughs)

Wall clock for 80 CTU



1 week

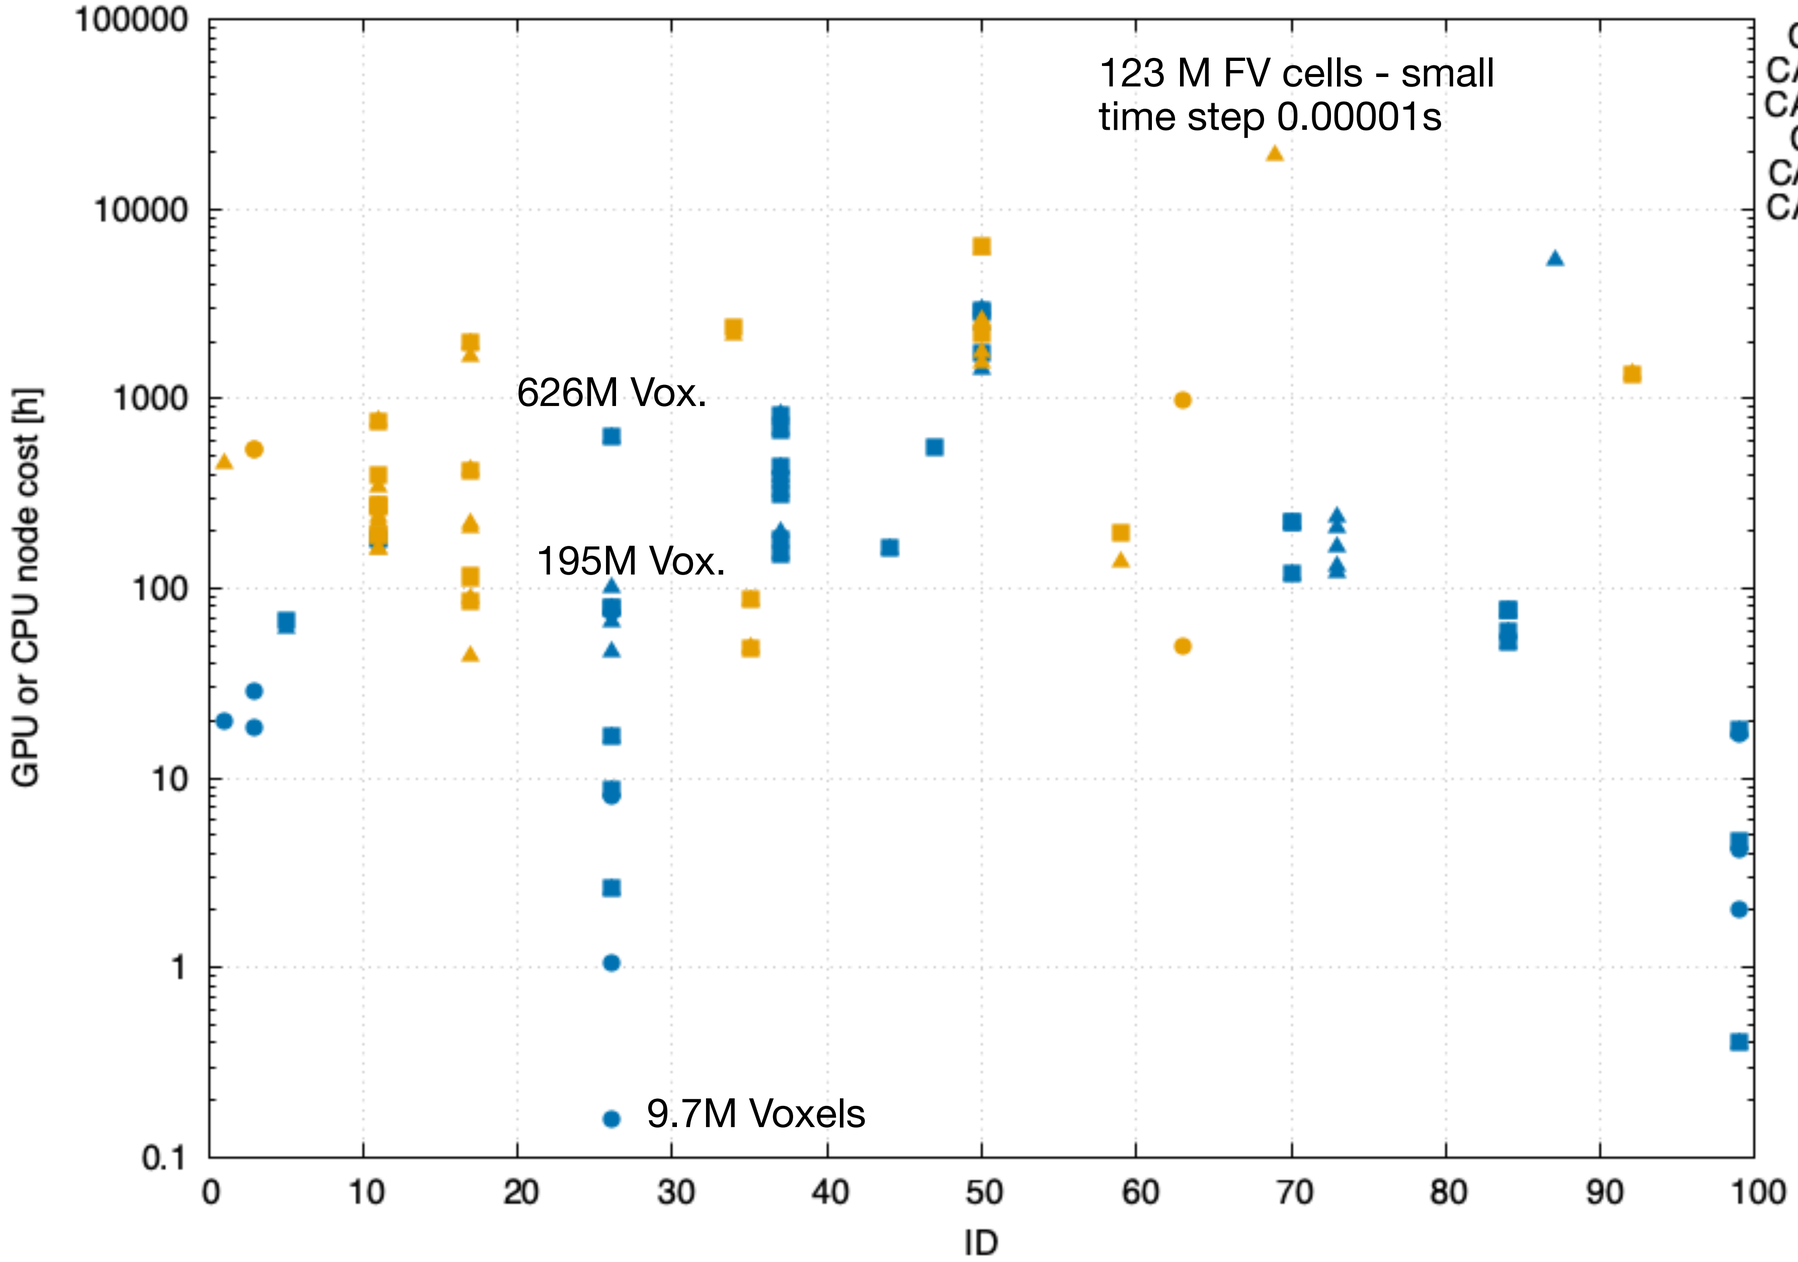
1 day

Gives idea of speed.

But not of COST.

Could have been run on 10 or 100 GPUs

Motivation - Significantly different costs - CPU node or GPU hour for 80 CTU



Misleading due to:

1) different dt
2) mesh size
3) wide variation in compute power between different hardware - Introducing the hardware cost is the best solution we found

Hardware cost – detail on CPU prices

Google Cloud Compute Engine | Resource-based CUD, 3 years | USD

CPU	Machine type	Physical cores	USD/node hour	USD/core hour
Intel	c4-standard-288-metal	144	6.4045	0.04448
Intel	n4-standard-80	40	1.6326	0.04082
Intel	c3-standard-192-metal	96	4.3548	0.04536
Intel	n2-standard-128	64	2.7972	0.04371
AMD	c4d-standard-384-metal	192	8.2455	0.04295
AMD	n4d-standard-96	48	1.8298	0.03812
AMD	c3d-standard-360	180	7.3537	0.04085
AMD	n2d-standard-224	112	4.2587	0.03802

Source: cloud.google.com/products/compute/pricing/general-purpose

Thanks Ruggero Poletto!

Average values are used for all CPUs in this work according to their number of cores

ID 1 & 11 reported 152 cores per node for AMD EPYC 9654 but it is 192
ID 17 reported 52 cores per node for Intel Xeon E5-2640-4v but it is 20

Hardware cost – detail on GPU prices

Selected reference rates | USD per GPU-hour, including allocated host resources

GPU	USD/GPU hour	8-GPU VM USD/hour	Basis
B200	7.0884	56.7072	CUD 3y
H200	4.6511	37.2084	CUD 3y
H100	5.0813	40.6501	CUD 3y
A100[1]	2.4000	—	DWS
RTX PRO 6000 Blackwell	1.9795	15.8356	CUD 3y
RTX 6000 Ada 48GB	1.6275	—	Estimate
L40S 48GB	1.6275	—	Estimate
RTX 5090 32GB	1.0117	—	Estimate
AMD Instinct MI250X[2]	1.0000	—	Estimate

[1] A100 memory was unspecified: use $2.40/GPU-hour as the explicit fallback.
[2] MI250X: rate per GCD; 4 modules = 8 GCDs/node ($8/node-hour).

Quoted rates: Google Cloud

B200, H200, H100, RTX PRO 6000:
Resource-based CUD, 3 years.
A100: DWS Flex-start fallback.

Source: https://cloud.google.com/products/compute/pricing/accelerator-optimized#accelerator-optimized-machine-type-family

Some observations on the use of dollar cost

We should go into finer detail on the hardware: Is A100 40GB or 80 GB, what about the different bandwidths? That affects cost and speed.

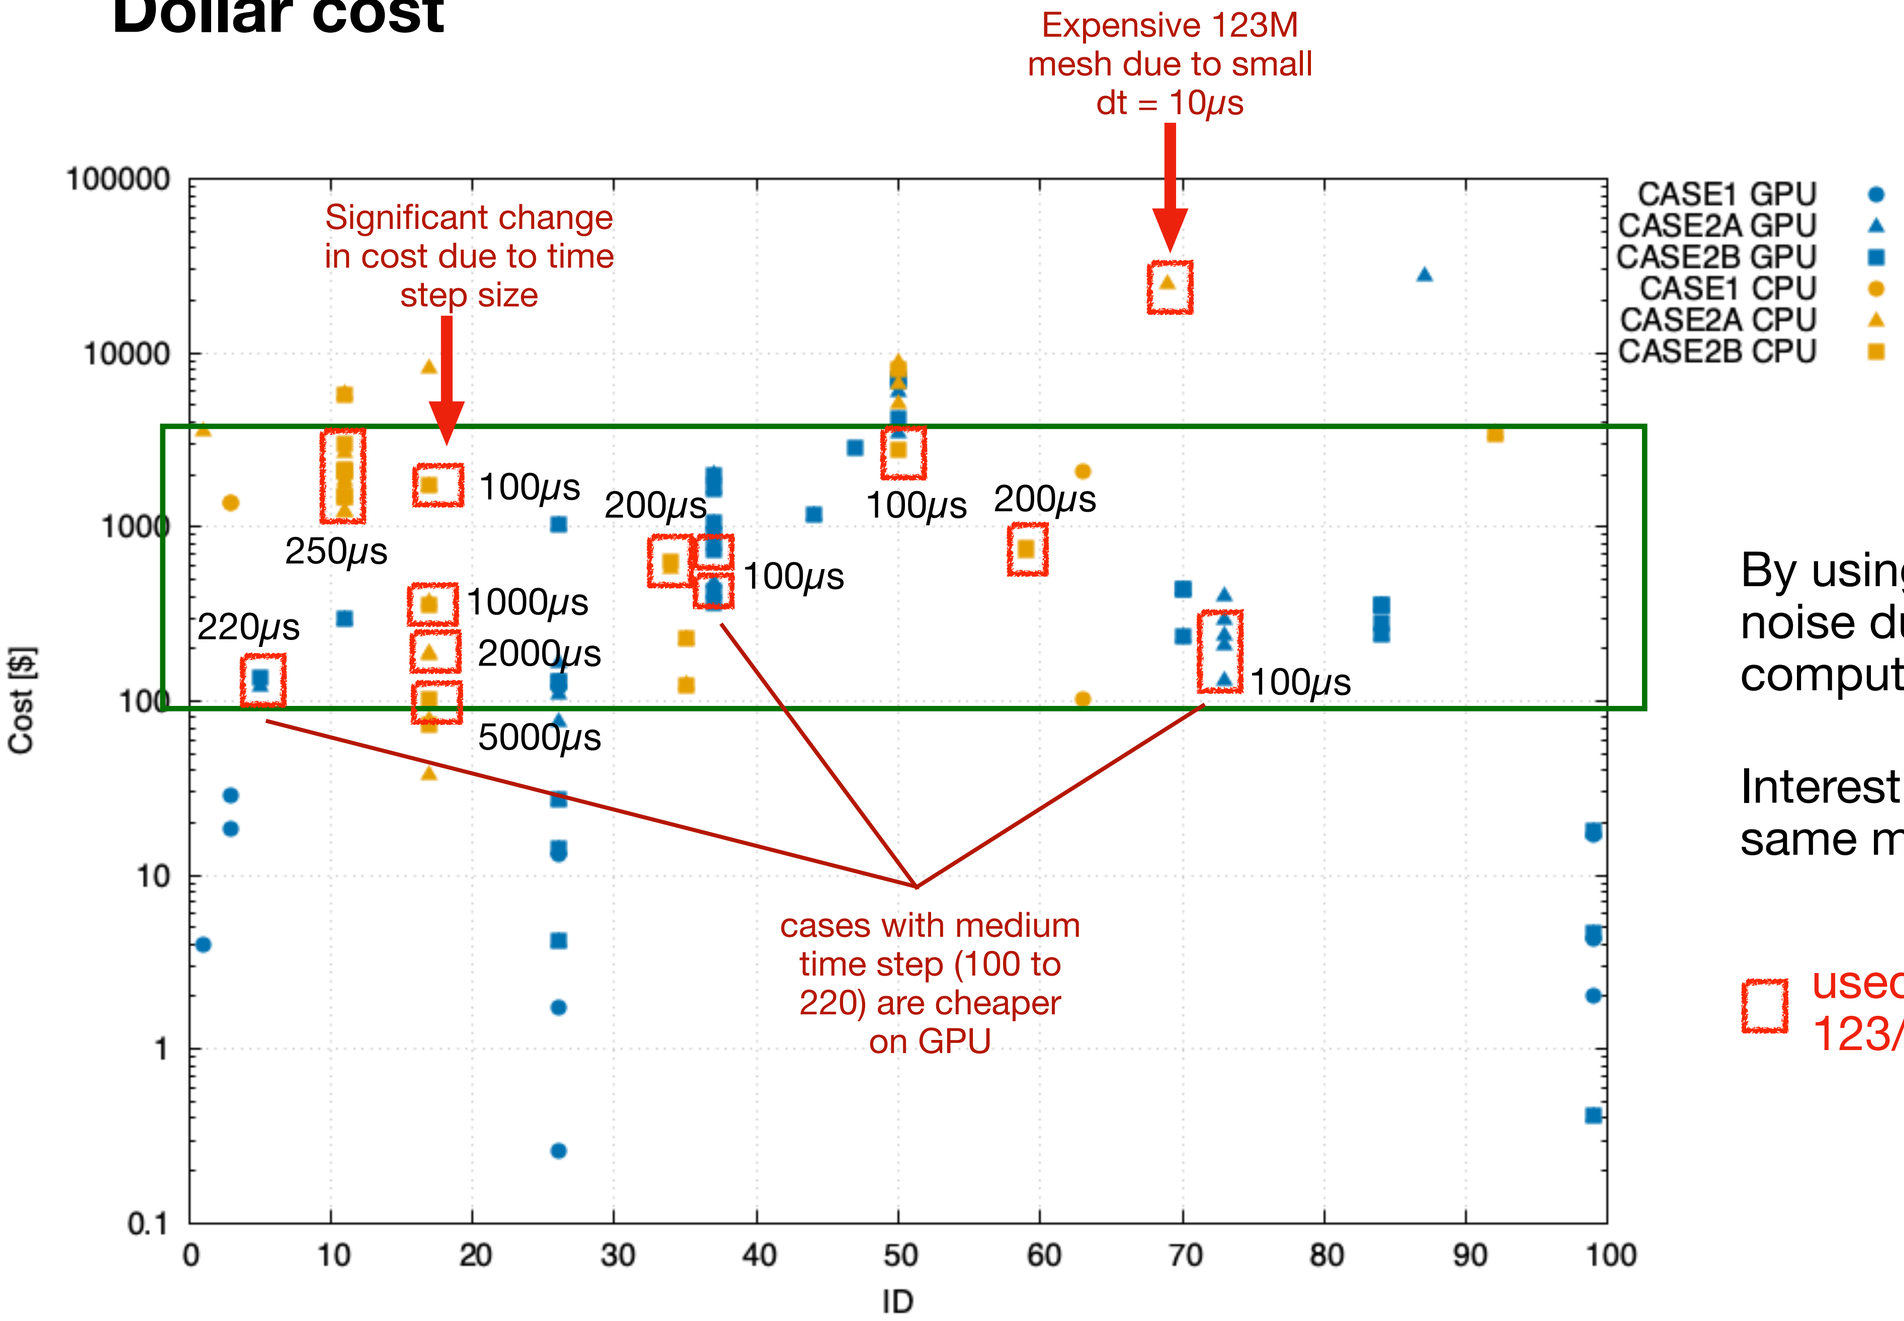
We have not taken into account the cost of the different networks

Are the computations being done in FP32 or FP64?

Proposal for next edition:

Do we want to go into finer detail? At least for some of the participants? They should provide the information on the hardware they will use in advance so we can better classify and price it.

Dollar cost



By using dollar cost we reduce noise due to different computational power

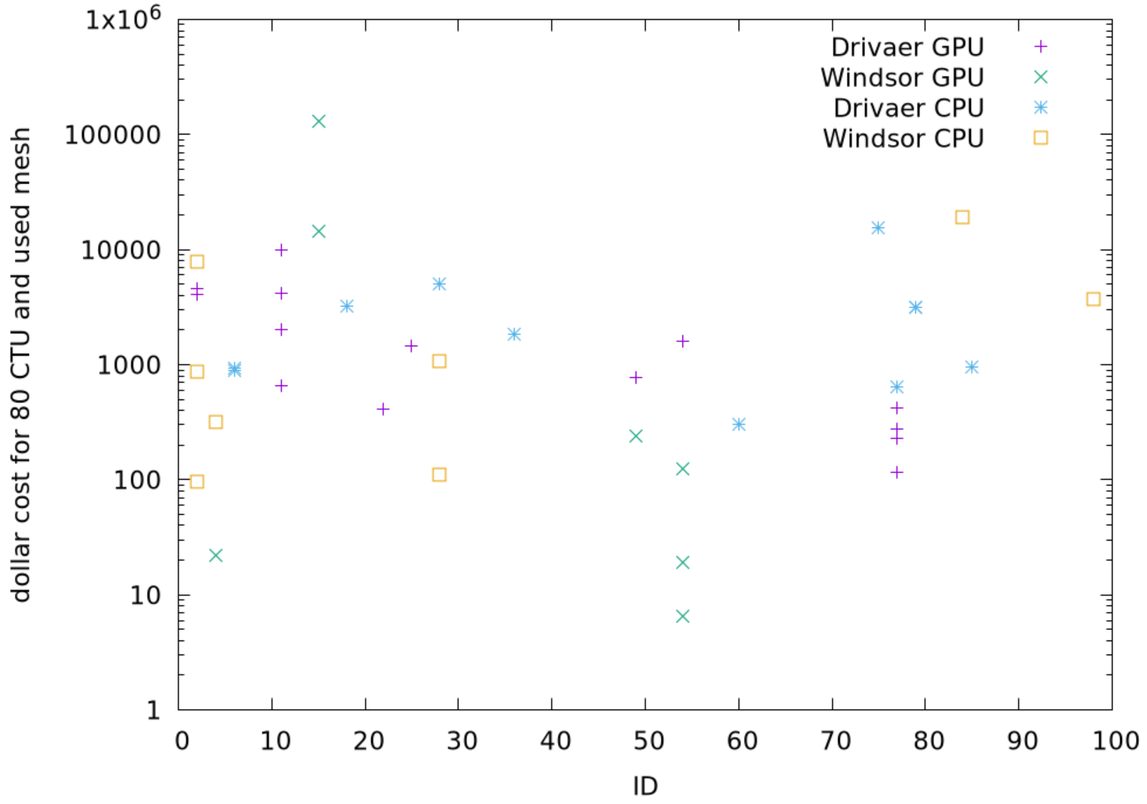
Interesting for comparing using same mesh

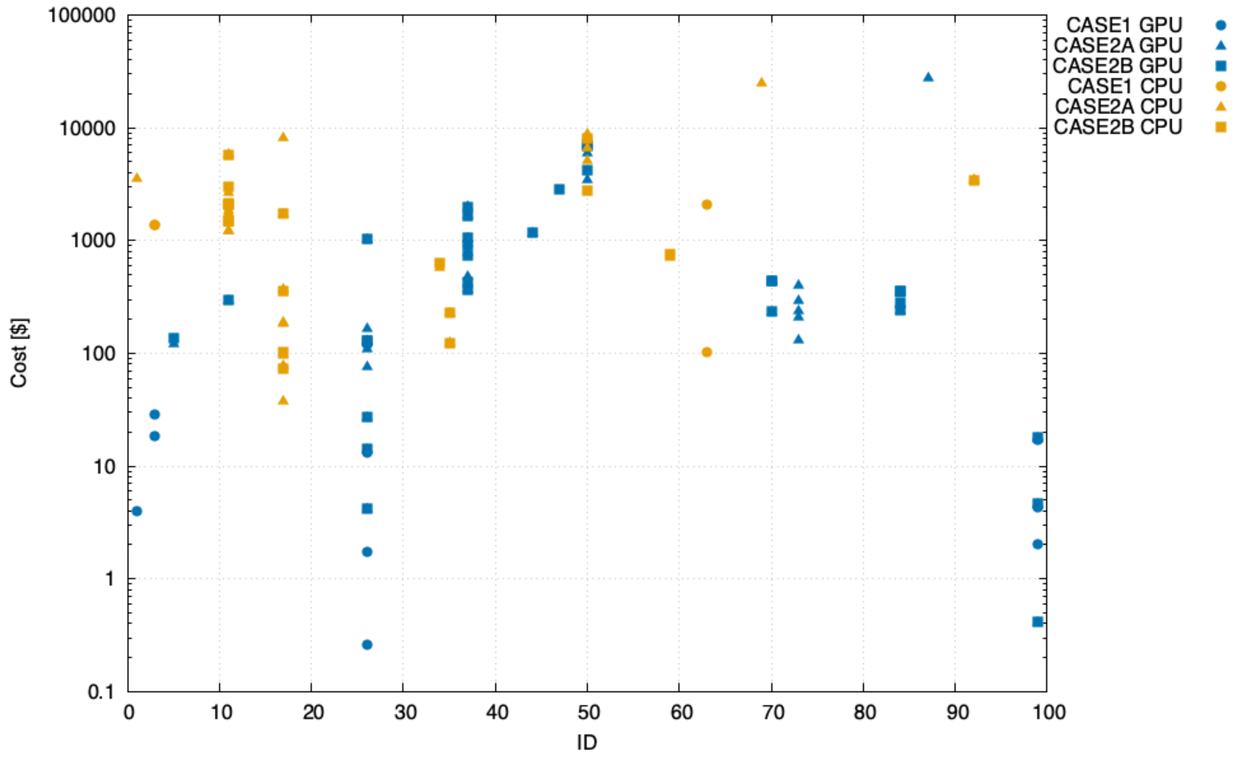
used to do mark runs with 123/126M Drivaer mesh

Dollar cost - comparison with 2024

Min and max values are lower but general trend is similar - high spread

Lower max value mainly due to used mesh size!!!!

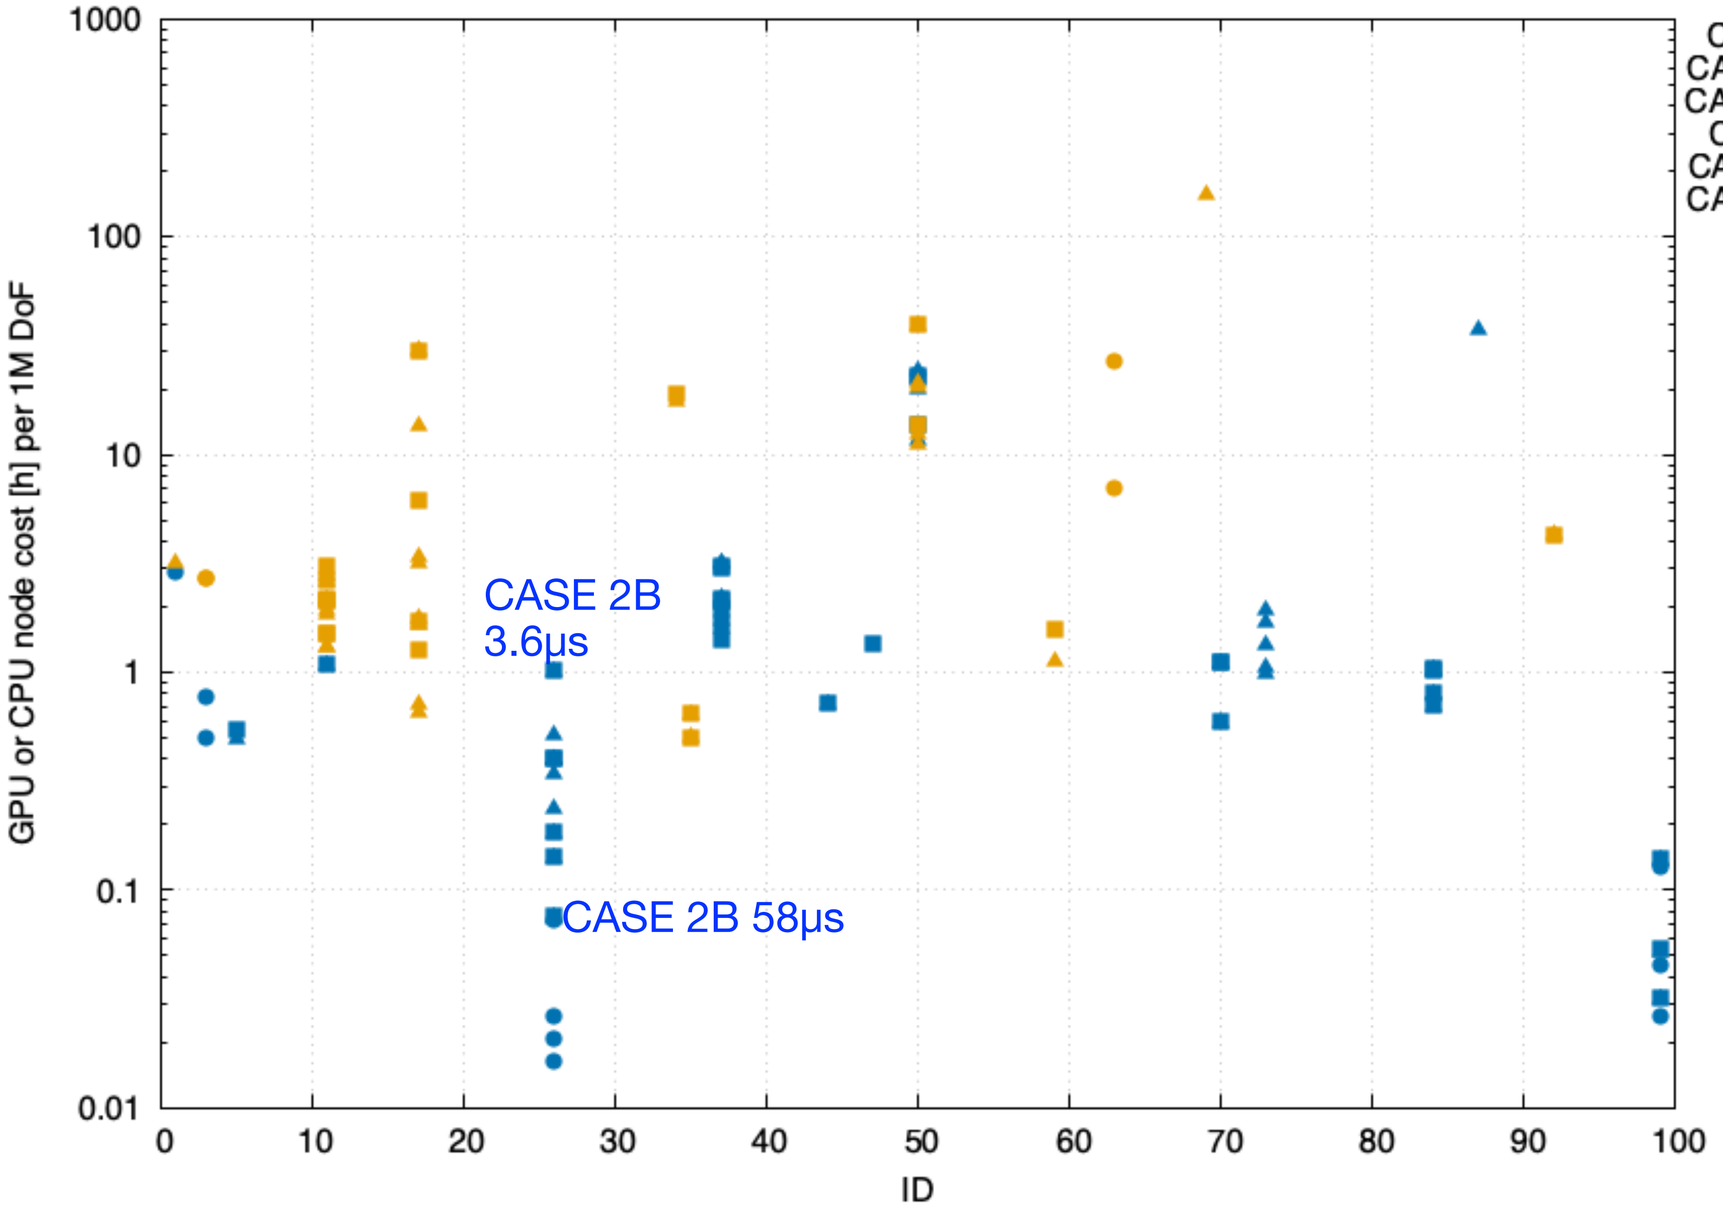




If we want to plot cost vs accuracy dollar cost is the best indicator - SRS TFG

when not taking accuracy into account other indicators are useful

Computational cost per 1M DoF (voxels, cells or nodes) - 80 CTU



Partially eliminates noise due to mesh - but finer meshes -> finer dt

As in previous images GPUs runs (blue) towards the bottom

Most of the codes between 1 and 10$ per 1M DoF

Computational time [ns] per cell and time step - COST

Wall time / average number of cells per GPU = Wall time * # of GPUs / # of cells (or voxels)

It is the inverse of MIUPS

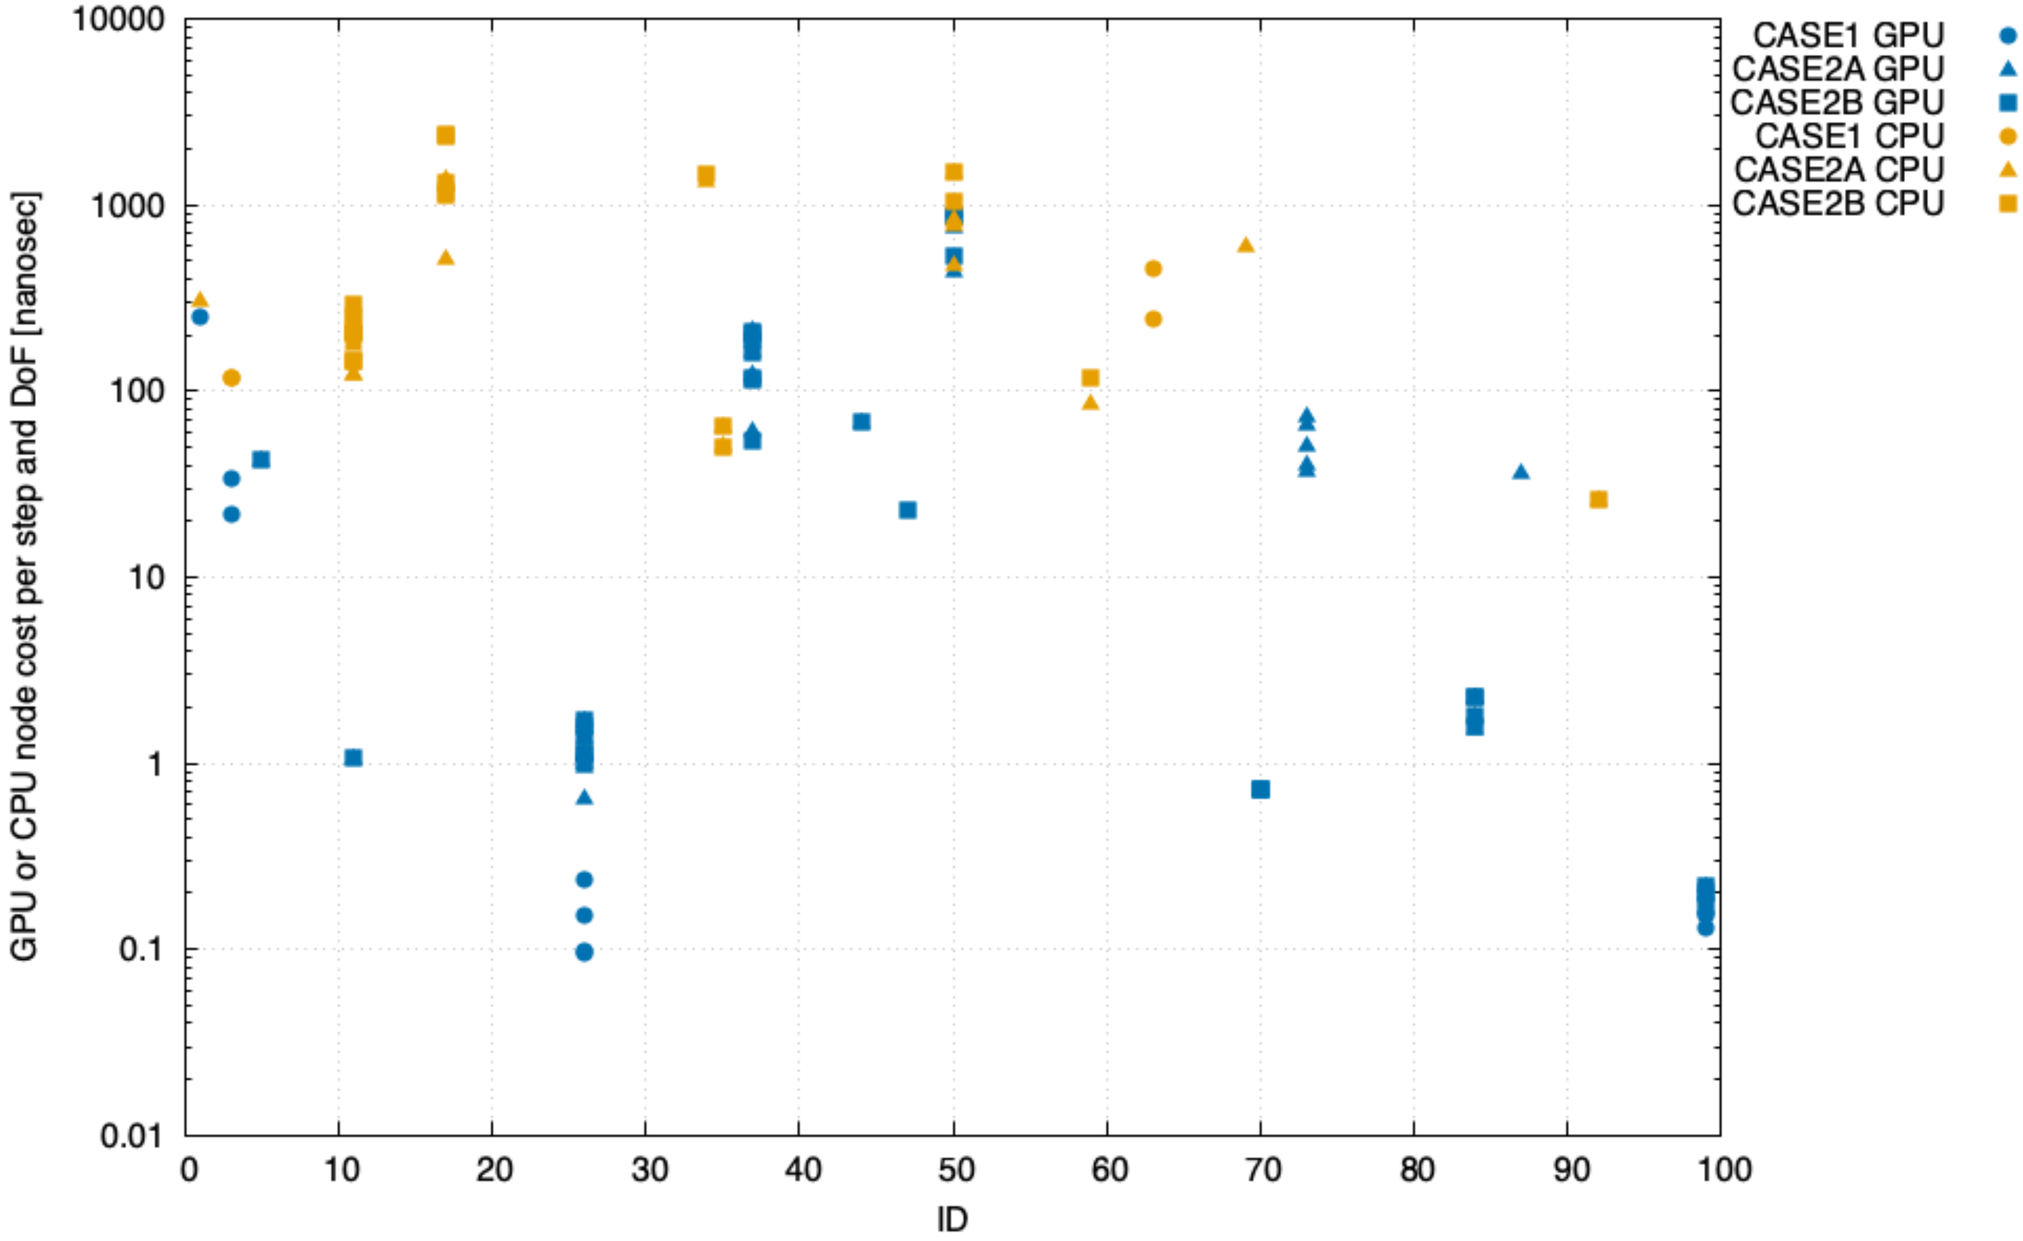
Advantages

- Allows to compare runs with different meshes and DOF per GPU (if within scalability range)
- Interesting from the HPC point of view - but little interest to plot against accuracy

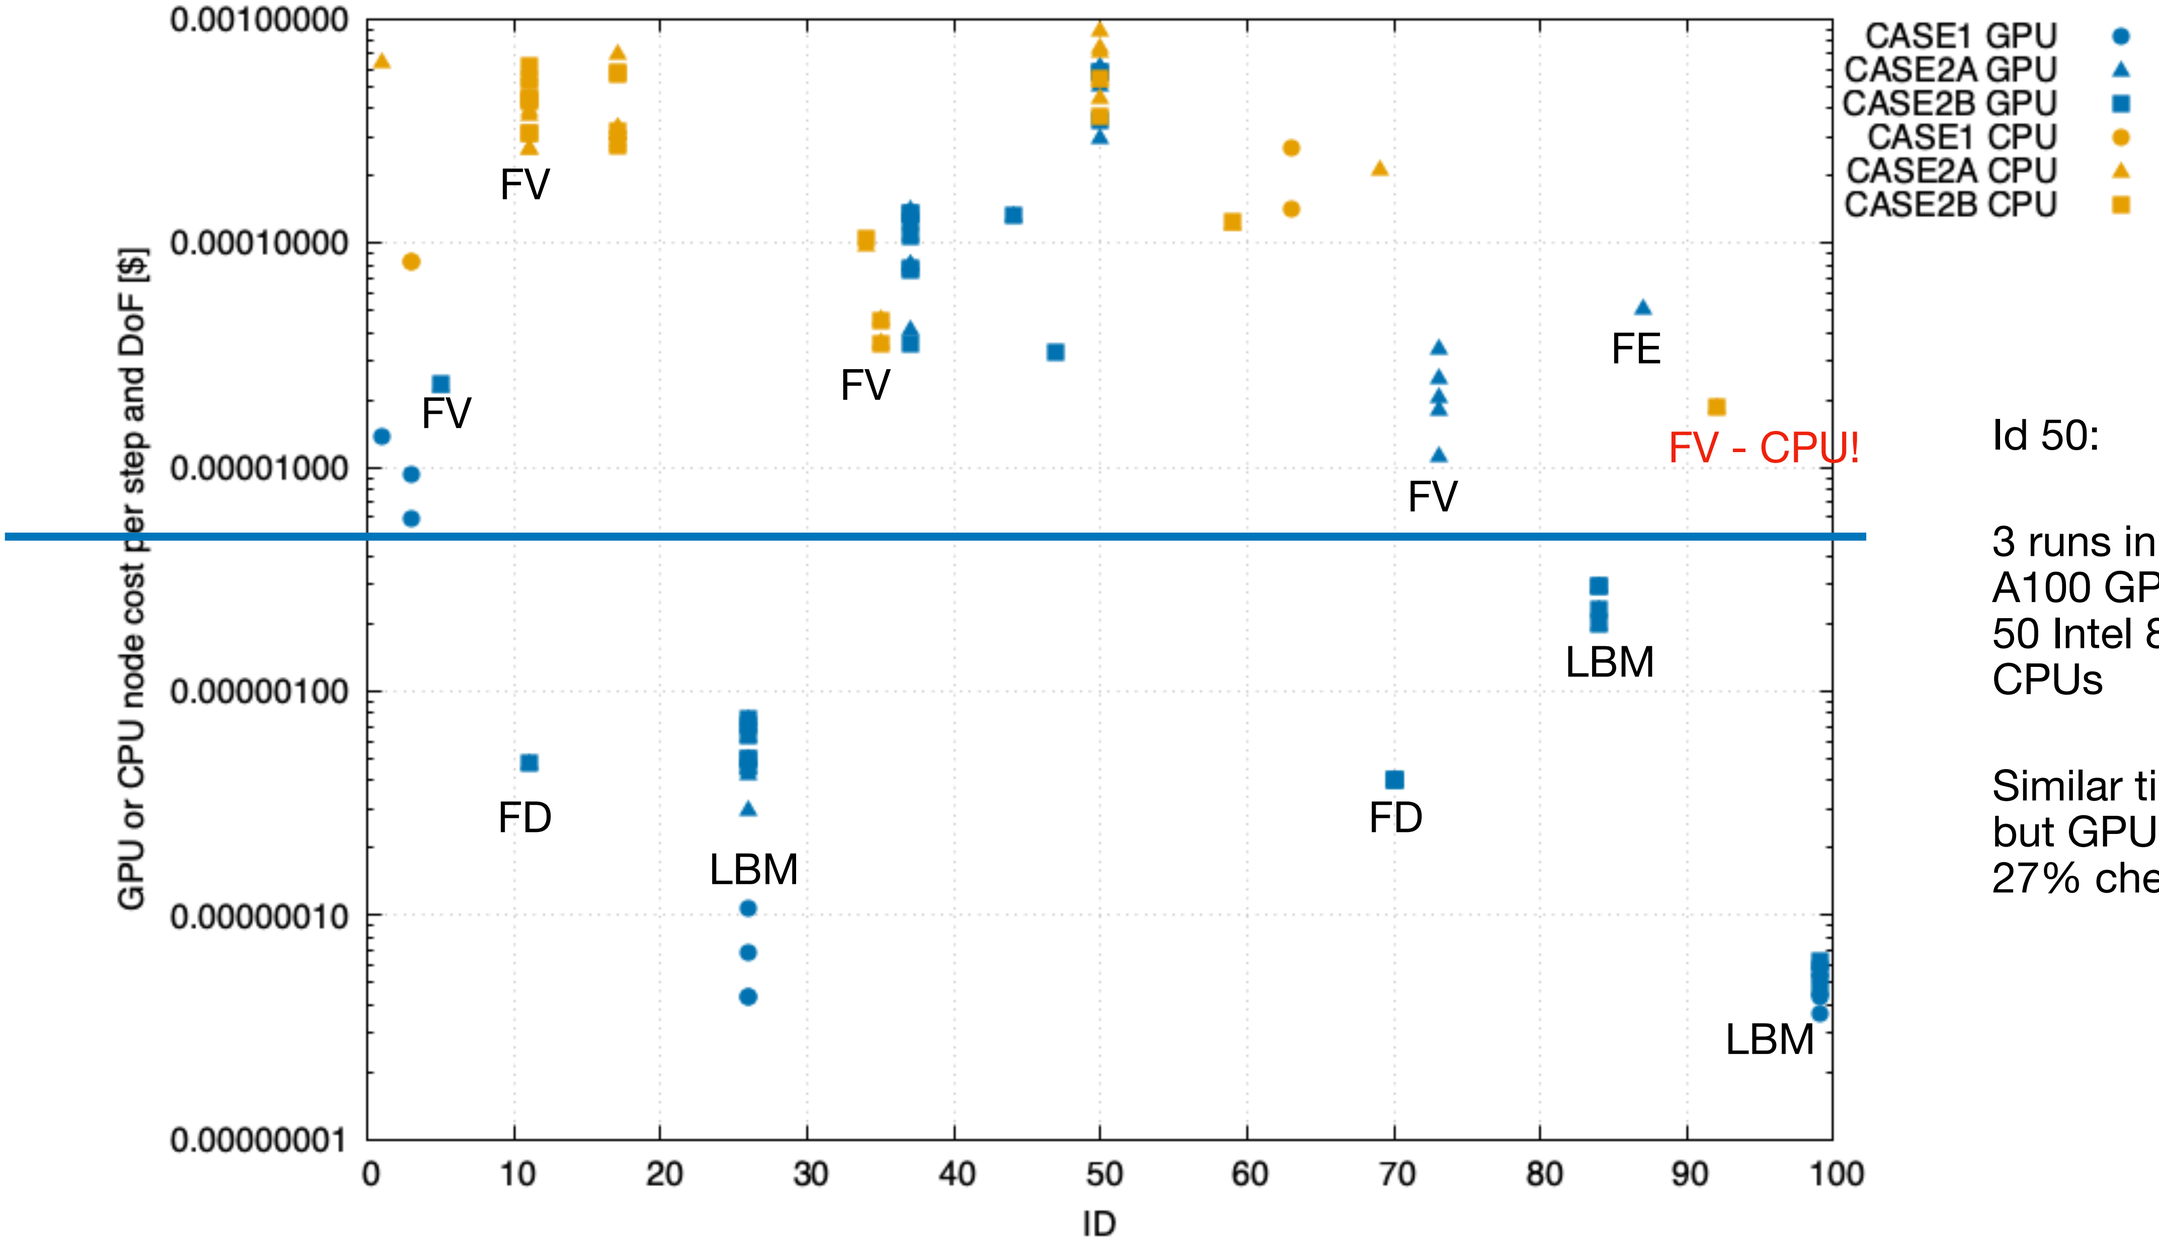
Disadvantages

- Makes little sense if different strategies are compared (explicit vs implicit). Explicit will be faster but at the cost of significantly smaller time steps
- Implicit will depend on tolerances used (LA solver and internal iterations)

Computational cost per DoF and time step



Dollar cost per DoF and time step - Observations



Id 50:

3 runs in 48 A100 GPU vs 50 Intel 8368 CPUs

Similar times but GPU 27% cheaper

Conclusions and Ideas for Future Work

We have somehow improved the cost we are using to compare between different hardware

Now it is based on publicly available data - links included in the presentation.

The analysed data consisted of 91 GPU results and 47 CPU results showing that GPU hardware is well established.

We have analysed the results using similar indicators to those used in 2024.

Accuracy metric gives another dimension in the interpretation of the results

AI methods and submissions are growing, they should also be compared to numerical simulations

How GPUs are used in each case? Is there room for online training of AI models?

Thanks!