


Flagship code FALL3D



Tools status (OK/KO)	Tracing (YES/NO)	Analysing (YES/NO)	POP metrics (YES/NO)	Other metrics (YES/NO)
OK	Yes	Yes	Yes	Yes Iterations/second

- Tools status: Extrae used. No critical issues reported
 - Tracing: Traces completed for 2 to 32 GPU configurations
 - Analysing: In progress
 - POP metrics: Available (see next slide)
- Notes: Intermediate meeting on 13/06. Final report should be available above 20/06





POP metrics


Procs	2	4	8	16	32
Host					
Global Scalability Efficiency	22.01	14.97	9.56	5.33	3.38
Parallel efficiency	22.01	30.07	39.08	45.40	46.91
- MPI Parallel efficiency	99.38	97.83	93.25	87.63	77.91
- Communication efficiency	99.45	98.44	95.96	90.86	93.86
- Load Balance	99.93	99.38	97.18	96.45	83.01
- CUDA Runtime Parallel efficiency	22.15	30.74	41.91	51.81	60.21
Computational Scalability	100.00	49.77	24.47	11.74	7.20
- Instruction Scalability	100.00	79.28	51.60	33.62	21.70
- IPC Scalability	100.00	97.40	98.70	105.19	110.39
- Frequency Scalability	100.00	64.45	48.04	33.19	30.07
Device					
Parallel efficiency	76.51	66.85	54.25	43.59	36.52
- Communication efficiency	99.63	99.33	98.90	98.65	99.04
- Scheduling efficiency	76.79	67.30	54.85	44.19	36.87
CPU/GPU					
Scheduling Runtime Overlap	98.94	98.04	96.76	96.00	97.08
Hybrid GPU					
CUDA Device LB	99.86	99.87	98.03	96.20	81.62
CUDA Host LB	99.60	99.61	97.85	96.07	81.85






Other metrics


Performance Comparison (Iterations per Second)



Performance Comparison (Iterations per Second)



Parallel Efficiency (wall time) Comparison



Strong Scalability (wall time) Comparison

