


Flagship code SPECFEM3D



Tools status (OK)	Tracing (YES)	Analysing (YES)	POP metrics (YES)	Other metrics (NO)
Scorep (MPI+CUDA)	YES	YES	YES	NO

- **Tools status:** the interaction of POP tools with your code in the audit configuration (MPI+CUDA)
- **Tracing:** if traces from a given tool are correctly produced
- **Analysing:** if POP experts are analyzing obtained traces
- **POP metrics:** if you already produced POP metrics
- **Other metrics:** if you already produced other metrics (application, TALP....)
- **Notes:** please, insert here any issue or further comments



v2: Efficiency model (strong scaling)




Problem size		1024x1024	1024x1024	1024x1024	1024x1024	1024x1024	1024x1024
MPI GPU ranks		8	32	64	128	256	512
v1	Wall time [s]	2001.806		255.948	135.478	73.846	61.400
GPU	Global scaling efficiency	0.995		0.973	0.919	0.843	0.507
	- Computation time scaling	1.000		0.987	0.972	0.950	0.937
	- Parallel efficiency	0.995		0.986	0.945	0.887	0.541
	- - Load balance efficiency	1.000		0.998	0.996	0.994	0.989
	- - Orchestration efficiency	0.995		0.988	0.948	0.892	0.547
v2	Wall time [s]	1996.055	505.209	259.019	133.395	77.756	46.871
GPU	Global scaling efficiency	0.997	0.985	0.960	0.932	0.800	0.664
	- Computation time scaling	1.000	0.997	0.989	0.976	0.957	0.952
	- Parallel efficiency	0.997	0.988	0.971	0.956	0.836	0.697
	- - Load balance efficiency	1.000	0.999	0.998	0.991	0.990	0.992
	- - Orchestration efficiency	0.997	0.989	0.973	0.964	0.845	0.703


- Orchestration efficiency
 - MPI communication & CUDA management
 - aka Communication efficiency

- Similar GPU strong scaling efficiency to 256 GPUs (excellent load balance)
- Considerably improved global scaling & parallel efficiencies for 512 GPUs
 - due to much improved orchestration efficiency using GPU-Aware MPI


Key:



Strong scaling (1024x1024 total)



v2: Strong scaling (1024x1024 t)



v2: Summary of observations



- *iterate_time* (solver) used as focus of analysis
- Weak scaling performance unchanged (in regime of measurement)
- Strong scaling performance improved considerably for 512 GPUs
 - otherwise similar to previous version
 - much more efficient orchestration of GPU devices
 - via exploitation of GPU-Aware MPI
 - CUDA kernel execution time only slightly faster (1.5%)