CENTRE OF EXCELLENCE FOR HPC
ASTROPHYSICAL APPLICATIONS

Enhancing the Energy Efficiency of SPACE Applications

João Barbosa
(IT4I@VSB, Czech Republic)

Co-funded by
the European Union

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.

EuroHPC
Joint Undertaking

Why Energy Efficiency, Why Now?

- Energy = growing bottleneck in HPC scalability
- 20 MW system → 1 MW saved at 5% gain
- Static tuning = low-effort, high-impact solution
- Applicable without modifying scientific code

“If one considers that a machine consumes 20MW, energy savings of 5–10% are in orders in megawatts…”

Enhancing the Energy Efficiency of SPACE Apps

SPACE CoE: optimizing HPC energy efficiency

Target: astrophysics applications on modern architectures

Strategy: static frequency tuning using MERIC

Platforms analyzed: NVidia A100, Intel Sapphire Rapids, NVidia Grace

Tools for energy efficiency evaluation

H2020 **READEX** (2015-2018): Complex parallel application has different requirements during execution, so it gives a possibility to be dynamically tuned for energy savings without performance penalty.

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

Example of power consumption timeline of a single node and its components (Gadget)

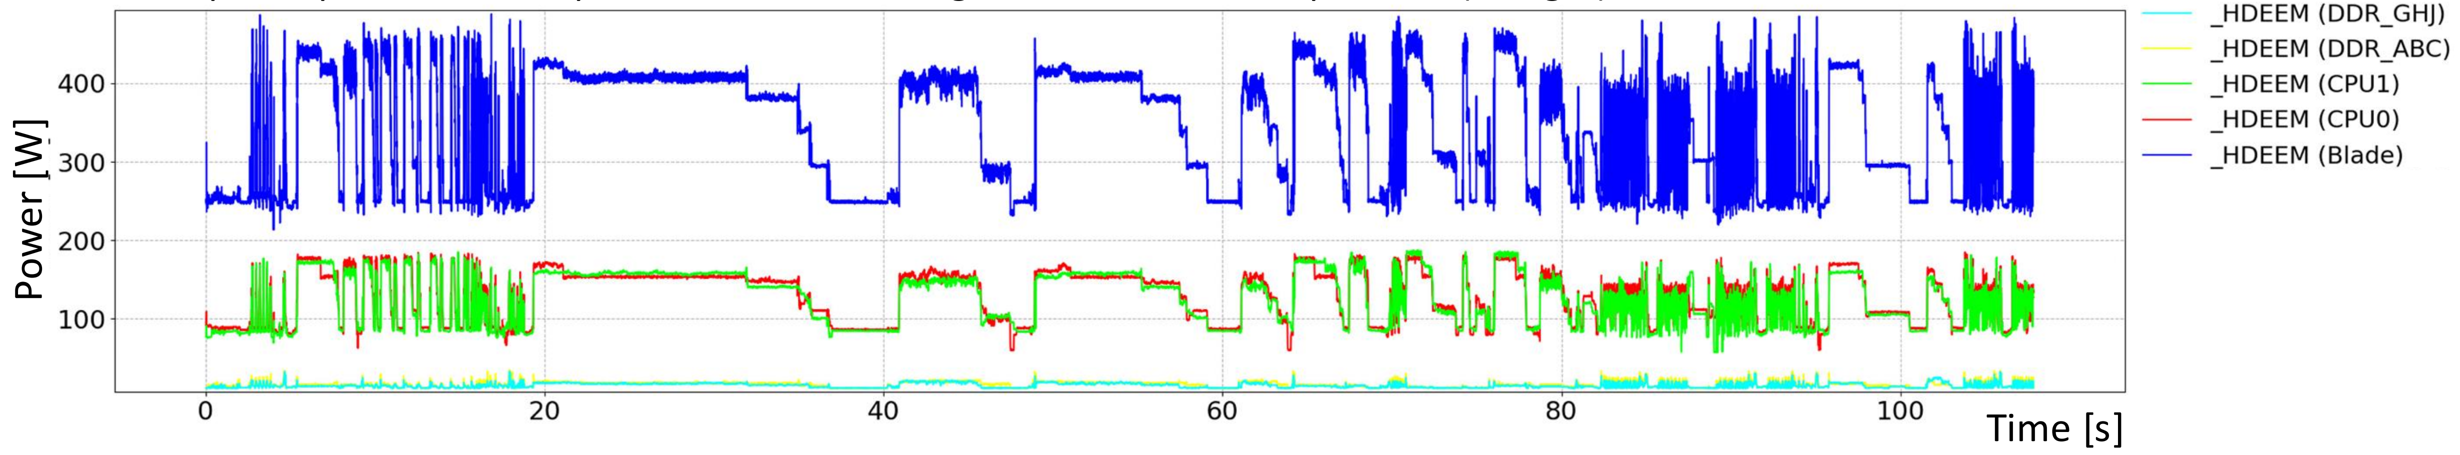

MERIC runtime system provides dynamic application tuning

- lightweight & easy to install & easy to use
- C/C++ API and Fortran module
- MPI, OpenMP and CUDA parallelization
- performance and power-aware
- support for a wide range of architectures and power monitoring systems

Our Approach: Static Tuning + MERIC

Hardware	Power Monitoring	Node Power Baseline	Tuned Power Knob
Nvidia A100-SXM4 GPU - 400 W TDP	NVML	600 W	SMF (1.41 - 0.21 GHz)
Intel Xeon Max 9468 CPU (Sapphire Rapids) + HBM - 48 cores - 2.1 GHz nominal frequency - 3.5 GHz turbo frequency (2.6 GHz with all cores utilized) - 350 W TDP	Intel RAPL	190 W	CF (2.6 - 0.8 GHz), UCF (2.5 - 0.8 GHz)
Nvidia Grace CPU superchip - 2x 72 cores (2x Grace CPUs) - 3.1 GHz nominal frequency - 500 W TDP	PDU	-	CF (3.3 - 0.9 GHz)

Table 4: Hardware architectures, available monitoring systems, and tuned power knobs

- Static tuning = set frequency once at job start
- Metrics tracked: runtime, energy use, FLOPs/Watt
- MERIC: open-source runtime, non-intrusive, SLURM-ready
- Easily applied across nodes & jobs

Static tuning of HPC applications

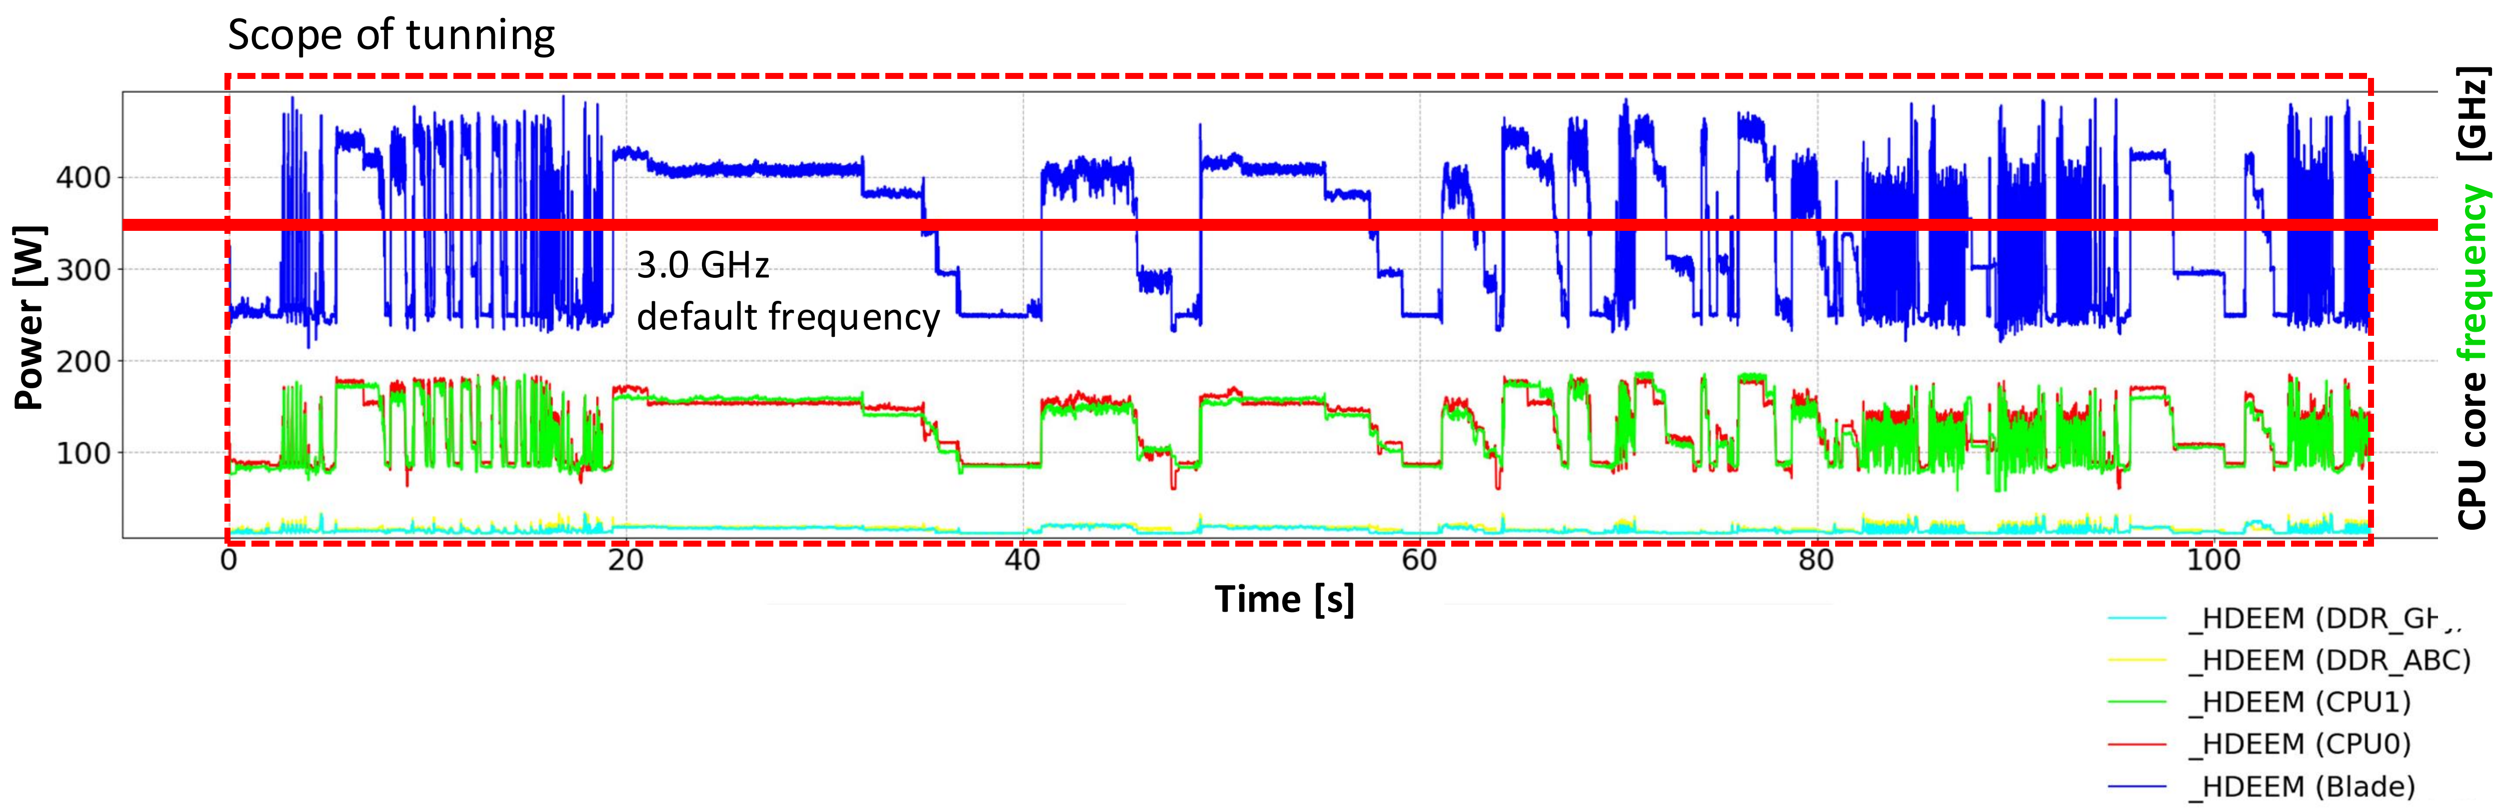

Static tuning of HPC applications

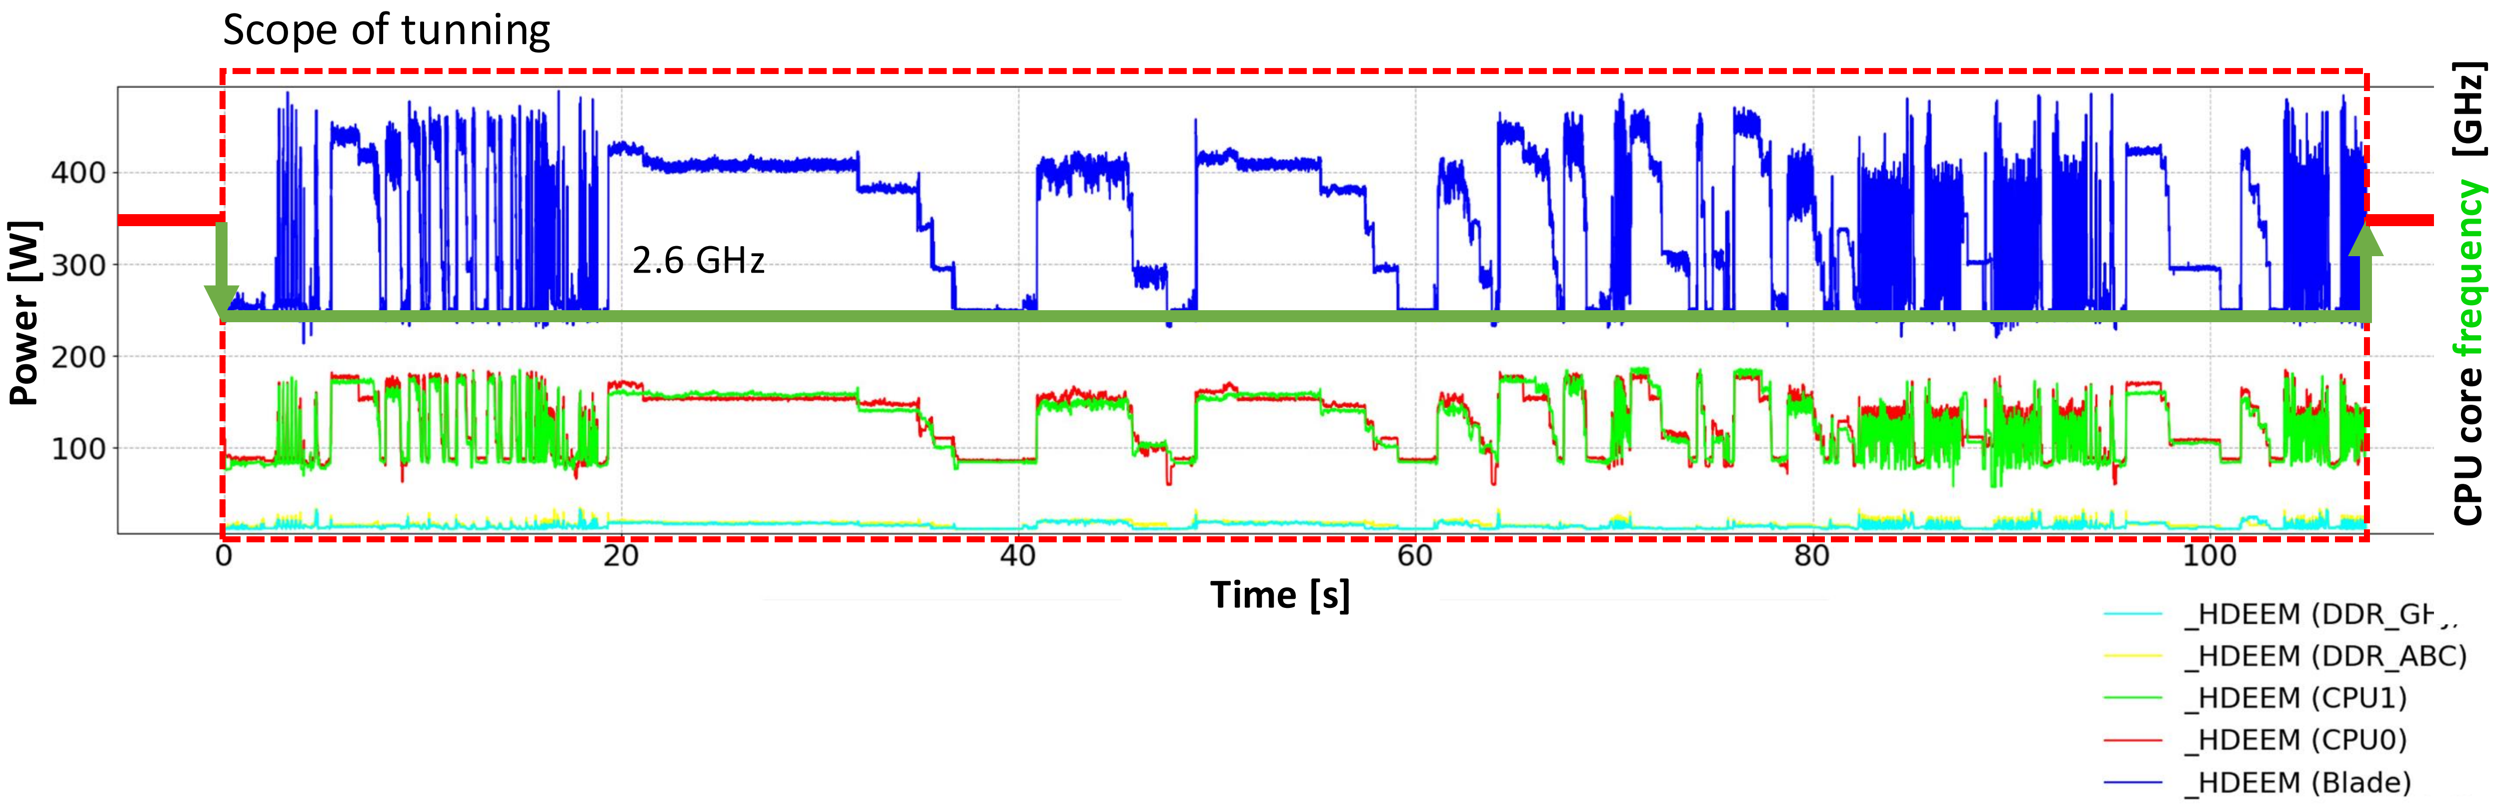

Example of static tuning on Gadget

uncore [GHz] / core [GHz]	1.2	1.4	1.6	1.8	2.0	2.2	2.4
1.3	106.64	95.07	89.08	79.95	74.68	71.86	70.56
1.5	90.81	77.83	72.3	64.95	59.81	58.26	55.46
1.7	79.73	68.37	59.76	52.62	46.43	44.1	42.27
1.9	78.24	60.6	50.98	42.75	38.39	36.25	33.26
2.1	71.83	52.24	45.23	36.16	31.4	27.74	23.78
2.3	64.06	52.49	39.01	30.3	25.97	22.34	19.35
2.5	68.28	46.39	37.38	29.44	24.35	17.48	16.44
2.6	69.78	47.34	36.22	24.38	20.38	17.05	13.32
2.7	67.6	42.46	34.08	24.44	17.59	14.37	10.33
2.8	65.64	45.95	30.55	24.79	16.31	13.2	7.78
2.9	63.26	50.33	31.94	22.67	13.99	10.46	7.53
3	63.02	46.02	27.63	21.36	13.37	8.03	2.73
3.1	59.03	45.2	27.36	22.01	12.54	6.57	1.49
3.2	56.33	43.32	28.45	19.56	13.67	5.35	4.67
3.3	56.62	42.06	34.96	16.65	13.27	4	0.57

Runtime extension [%]

	1.2	1.4	1.6	1.8	2.0	2.2	2.4
1.3	8.36	5.04	4.44	2.32	3.94	7.13	12.08
1.5	3.14	−1.03	−1.99	−3.04	−1.96	1.89	5.32
1.7	1.58	−1.79	−4.29	−5.89	−5.69	−2.82	0.92
1.9	3.33	−3.84	−7.02	−9.81	−8.42	−5.59	−2.62
2.1	2.43	−5.97	−8	−11.15	−10.56	−8.8	−6.5
2.3	1.66	−2.74	−8.55	−11.82	−10.85	−9.29	−7.07
2.5	8.74	−1.75	−5.89	−10.85	−9.56	−8.88	−5.21
2.6	11.55	0.27	−4.26	−9.88	−9.01	−8.1	−6.31
2.7	12.16	−1.02	−4.29	−8.45	−9.43	−7.99	−6.66
2.8	12.99	2.7	−4.8	−6.45	−8.77	−7.43	−7.43
2.9	13.87	7.94	−1.67	−5.83	−8.07	−7.52	−5.47
3	16.57	7.73	−1.85	−4.14	−6.68	−6.94	−6.82
3.1	17.16	10.53	0.39	−1.34	−4.83	−5.67	−5.43
3.2	19	12.09	3.72	−0.38	−1.55	−3.88	−1.22
3.3	22.21	14.16	10.62	0.2	0.69	−0.96	−2.36

Energy savings [%]

Example of static tuning on ChaNGa

uncore [GHz] / core [GHz]	1.4	1.6	1.8	2	2.2	2.4
1.3	15.27	13.27	13.25	13.53	12.44	12.78
1.5	8.39	8.64	7.24	9.82	9.31	8.78
1.7	4.52	4.82	4.15	5.86	4.12	3.4
1.9	4.46	4.12	3.81	1.52	3.11	2.35
2.1	3.4	3.53	3.24	1.58	1.98	1.7
2.3	2.79	1.22	1.4	1.09	1.12	1.16
2.5	1.37	1.26	1.01	0.31	0.69	0.8
2.6	1.8	1.47	1.16	−0.11	0.2	−0.14
2.7	1.51	0.55	0.49	−0.41	1.29	0.34
2.8	1.98	1.05	0.9	0.04	0.47	−0.33
2.9	1.51	0.7	1.19	0.1	0.95	−0.46
3	0.99	−0.1	−0.43	−1.14	−0.29	−1.37
3.1	0.41	−0.21	−0.43	0.92	−0.87	−0.32
3.2	0.7	−0.09	0.99	−0.24	−1.14	−2.26
3.3	0.11	−0.55	−0.25	−0.43	−0.33	−1.18

Runtime extension [%]

	1.4	1.6	1.8	2	2.2	2.4
1.3	−34.22	−33.49	−32.13	−29.62	−26.95	−23.38
1.5	−35.26	−33.88	−33.12	−28.94	−26.63	−23.78
1.7	−35.76	−34.11	−32.86	−30.14	−27.82	−25.15
1.9	−34.2	−32.82	−31.43	−30.26	−26.49	−24.66
2.1	−31.32	−30.93	−29.53	−28.09	−25.59	−22.7
2.3	−29.79	−29.32	−27.84	−25.34	−22.3	−19.62
2.5	−27.84	−26.46	−24.81	−22.87	−19.76	−16.47
2.6	−26.18	−25.13	−23.17	−21.81	−18.9	−16.07
2.7	−27.64	−26.53	−25	−22.82	−19.46	−16.9
2.8	−25.64	−24.33	−22.86	−21.24	−17.74	−15.11
2.9	−23.35	−22.06	−20.08	−18.67	−14.79	−12.55
3	−20.82	−19.92	−18.84	−16.5	−13.34	−10.42
3.1	−18.08	−16.88	−14.02	−12.11	−10.83	−6.68
3.2	−15.61	−14.17	−12.05	−10.18	−7.93	−6.39
3.3	−12.51	−10.9	−10.86	−7.94	−5.95	−4.4

Energy savings [%]

Hardware Platforms Tested

- **A100 GPU** (Karolina): flagship accelerator node
- **Sapphire Rapids CPU**: tested with DDR & HBM
- **Grace CPU**: ARM-based, power-efficient designAll nodes profiled using node-wide energy metrics

Hardware	Node hardware	System name
Nvidia A100	2× AMD 7763 CPU, 16× DDR4, 8× A100 GPU (with 40 GB HBM2 per GPU)	IT4I Karolina
Intel Sapphire Rapids	2× CPU, 2× 64GB HBM, 2× 128 GB (DDR5)	IT4I CS
Nvidia Grace CPU	1× superchip (2× 72 CPU cores) with 2× 120 GB LPDDR5X	IT4I CS

Scientific Codes & Benchmarks

7 codes: Pluto, OpenGADGET, iPIC3D, RAMSES, BHAC, FIL, ChaNGa

Representative of MHD, N-body, particle & relativistic codes

Benchmarks: real input cases, 1–5 minute runs

Full-node usage for both GPU and CPU tests

Hardware Platform	**Nvidia A100 GPU**	**Sapphire Rapids w. DDR and HBM**	**Nvidia Grace CPU**
Where	IT4I Karolina production system	IT4I complementary system	IT4I complementary system
Key features	Evaluation of accelerated platform	evaluation of modern x86 CPU and effect of DDR and HBM memory on energy efficiency	Evaluation of energy efficiency of modern ARM-based CPU to have comparison with x86
Tuning	static tuning of GPU streaming multiprocessor frequency	static tuning of CPU core frequency and CPU uncore frequency	static tuning of CPU core frequency
OpenGadget	yes	yes	yes
Pluto	yes	yes	yes
ChaNGa	yes	yes	yes
iPiC3D	yes	yes	yes
RAMSES	no	yes	yes
FIL	no	yes	yes
BHAC	no	yes	yes

A100 GPU – Modest but Measurable Gains

Set GPU frequency [MHz]	Runtime [s]	Runtime extension	Average power consumption of the GPU [W]	Average power consumption of the CPU [W]	CPUs + GPUs energy consumption [kJ]	Energy savings for GPUs + CPUs	Node energy consumption [kJ]	Compute node energy savings
1410	187.4	100.0%	185.8	118.4	322.9	0.0%	435.3	0.0%
1350	189.9	101.3%	172.7	119.0	307.5	4.8%	421.5	3.2%
1290	192.1	102.5%	160.0	119.6	291.8	9.6%	407.1	6.5%
1230	197.3	105.3%	149.0	119.7	282.4	12.5%	400.8	7.9%
1170	204.3	109.0%	138.9	120.1	276.1	14.5%	398.7	8.4%
1110	211.7	113.0%	129.5	119.4	270.0	16.4%	397.0	8.8%
1050	217.8	116.2%	122.1	119.8	264.8	18.0%	395.5	9.2%
990	225.7	120.4%	118.6	119.5	268.1	17.0%	403.5	7.3%
930	233.8	124.8%	115.9	119.4	272.8	15.5%	413.1	5.1%
870	235.1	125.5%	115.0	119.4	272.6	15.6%	413.7	5.0%
810	250.0	133.4%	111.1	119.1	281.7	12.7%	431.7	0.8%
750	265.3	141.6%	107.6	119.3	291.6	9.7%	450.8	-3.5%
690	280.5	149.7%	104.4	119.2	301.1	6.8%	469.4	-7.8%

- Pluto: 9% savings at 13% runtime overhead
- OpenGADGET: 7%, iPIC3D: 3–5%
- Power draw far below TDP (e.g., Pluto ~186W)
- Tuning worthwhile even with built-in DVFS

A100 GPU – Modest but Measurable Gains

Set GPU frequency [MHz]	Runtime [s]	Runtime extension	Average power consumption of the GPU [W]	Average power consumption of the CPU [W]	CPUs + GPUs energy consumption [kJ]	Energy savings for GPUs + CPUs	Node energy consumption [kJ]	Compute node energy savings
1410	187.4	100.0%	185.8	118.4	322.9	0.0%	435.3	0.0%
1350	[illegible]	101.3%	172.7	[illegible]	307.5	4.8%	421.5	3.2%
1290	192.1	102.5%	160.0	119.6	291.8	9.6%	407.1	6.5%
1230	197.3	105.3%	149.0	119.7	282.4	12.5%	400.8	7.9%
1170	204.3	109.0%	138.9	120.1	276.1	14.5%	398.7	8.4%
1110	211.7	113.0%	129.5	119.4	270.0	16.4%	397.0	8.8%
1050	217.8	116.2%	122.1	119.8	264.8	18.0%	395.5	9.2%
990	225.7	120.4%	118.6	119.5	268.1	17.0%	403.5	7.3%
930	233.8	124.8%	115.9	119.4	272.8	15.5%	413.1	5.1%
870	235.1	125.5%	115.0	119.4	272.6	15.6%	413.7	5.0%
810	250.0	133.4%	111.1	119.1	281.7	12.7%	431.7	0.8%
750	265.3	141.6%	107.6	119.3	291.6	9.7%	450.8	-3.5%
690	280.5	149.7%	104.4	119.2	301.1	6.8%	469.4	-7.8%

- Pluto: 9% savings at 13% runtime overhead
- OpenGADGET: 7%, iPIC3D: 3–5%
- Power draw far below TDP (e.g., Pluto ~186W)
- Tuning worthwhile even with built-in DVFS

Sapphire Rapids – Tunable & Versatile

- DDR config: 6–14% energy savings with minimal slowdown
- HBM config: better perf/W, smaller tuning window
- ChaNGa: 20% energy savings with 7% longer runtime
- CPU core & uncore tuning is impactfu

CF/UnCF [GHz]	0	2,5	2,4	2,2	2	1,8	1,6	1,4	1,2	1	0,8
0	100,0%	101,1%	101,1%	101,3%	101,6%	101,5%	101,1%	101,3%	102,2%	103,1%	105,7%
2,6	100,1%	101,5%	101,3%	103,2%	101,5%	102,7%	103,7%	105,4%	107,3%	107,8%	107,7%
2,5	104,6%	103,0%	104,9%	105,5%	104,8%	103,0%	105,4%	109,2%	110,3%	109,4%	107,4%
2,3	112,4%	110,0%	110,1%	109,5%	108,8%	107,7%	108,5%	111,2%	113,0%	112,2%	107,6%
2,1	123,2%	117,5%	116,9%	113,9%	114,0%	113,3%	113,9%	116,5%	120,2%	120,8%	116,0%
1,9	117,7%	124,7%	123,8%	122,1%	124,1%	125,9%	124,6%	124,4%	125,5%	126,4%	122,4%
1,7	132,0%	134,3%	134,0%	133,8%	135,7%	136,7%	134,7%	133,1%	133,7%	134,4%	133,5%
1,5	144,6%	147,8%	144,4%	146,8%	148,5%	150,2%	147,6%	145,4%	145,1%	145,3%	145,2%
1,3	176,3%	165,3%	162,2%	167,0%	165,1%	165,2%	163,3%	163,9%	165,0%	164,3%	166,2%
1,1	187,1%	184,4%	183,1%	188,9%	183,0%	182,6%	181,4%	187,1%	188,3%	186,0%	187,3%

CF/UnCF [GHz]	0	2,5	2,4	2,2	2	1,8	1,6	1,4	1,2	1	0,8
0	100 %	102 %	100 %	97 %	91 %	89 %	89 %	90 %	89 %	86 %	84 %
2,6	100 %	101 %	101 %	100 %	96 %	92 %	91 %	90 %	88 %	85 %	84 %
2,5	104 %	103 %	104 %	101 %	97 %	92 %	90 %	89 %	87 %	85 %	82 %
2,3	111 %	110 %	108 %	103 %	98 %	93 %	90 %	88 %	87 %	84 %	80 %
2,1	121 %	112 %	112 %	105 %	100 %	95 %	92 %	90 %	89 %	87 %	83 %
1,9	82 %	113 %	117 %	110 %	105 %	101 %	97 %	93 %	91 %	89 %	85 %
1,7	88 %	115 %	123 %	116 %	111 %	107 %	102 %	97 %	94 %	92 %	89 %
1,5	94 %	122 %	130 %	125 %	119 %	114 %	108 %	103 %	99 %	97 %	94 %
1,3	109 %	133 %	142 %	138 %	129 %	122 %	116 %	112 %	109 %	105 %	104 %
1,1	113 %	145 %	156 %	152 %	139 %	132 %	125 %	124 %	120 %	115 %	113 %

Sapphire Rapids – Tunable & Versatile

- DDR config: 6–14% energy savings with minimal slowdown
- HBM config: better perf/W, smaller tuning window
- ChaNGa: 20% energy savings with 7% longer runtime
- CPU core & uncore tuning is impactfu

CF/UnCF [GHz]	0	2.5	2.4	2.2	2	1.8	1.6	1.4	1.2	1	0.8
0	100,0%	101,1%	101,1%	101,3%	101,6%	101,5%	101,1%	101,3%	102,2%	103,1%	105,7%
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
2,5	104,6%	103,0%	104,9%	105,5%	104,8%	103,0%	105,4%	109,2%	110,3%	109,4%	107,4%
2,3	112,4%	110,0%	110,1%	109,5%	108,8%	107,7%	108,5%	111,2%	113,0%	112,2%	107,6%
2,1	123,2%	117,5%	116,9%	113,9%	114,0%	113,3%	113,9%	116,5%	120,2%	120,8%	116,0%
1,9	117,7%	124,7%	123,8%	122,1%	124,1%	125,9%	124,6%	124,4%	125,5%	126,4%	122,4%
1,7	132,0%	134,3%	134,0%	133,8%	135,7%	136,7%	134,7%	133,1%	133,7%	134,4%	133,5%
1,5	144,6%	147,8%	144,4%	146,8%	148,5%	150,2%	147,6%	145,4%	145,1%	145,3%	145,2%
1,3	176,3%	165,3%	162,2%	167,0%	165,1%	165,2%	163,3%	163,9%	165,0%	164,3%	166,2%
1,1	187,1%	184,4%	183,1%	188,9%	183,0%	182,6%	181,4%	187,1%	188,3%	186,0%	187,3%

CF/UnCF [GHz]	0	2.5	2.4	2.2	2	1.8	1.6	1.4	1.2	1	0.8
0	100 %	102 %	100 %	97 %	91 %	89 %	89 %	90 %	89 %	86 %	84 %
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
2,5	104 %	103 %	104 %	101 %	97 %	92 %	90 %	89 %	87 %	85 %	82 %
2,3	111 %	110 %	108 %	103 %	98 %	93 %	90 %	88 %	87 %	84 %	80 %
2,1	121 %	112 %	112 %	105 %	100 %	95 %	92 %	90 %	89 %	87 %	83 %
1,9	82 %	113 %	117 %	110 %	105 %	101 %	97 %	93 %	91 %	89 %	85 %
1,7	88 %	115 %	123 %	116 %	111 %	107 %	102 %	97 %	94 %	92 %	89 %
1,5	94 %	122 %	130 %	125 %	119 %	114 %	108 %	103 %	99 %	97 %	94 %
1,3	109 %	133 %	142 %	138 %	129 %	122 %	116 %	112 %	109 %	105 %	104 %
1,1	113 %	145 %	156 %	152 %	139 %	132 %	125 %	124 %	120 %	115 %	113 %

Sapphire Rapids – Tunable & Versatile

- DDR config: 6–14% energy savings with minimal slowdown
- HBM config: better perf/W, smaller tuning window
- ChaNGa: 20% energy savings with 7% longer runtime
- CPU core & uncore tuning is impactfu

CF/UnCF [GHz]	0	2,5	2,4	2,2	2	1,8	1,6	1,4	1,2	1	0,8
0	100,0%	101,1%	101,1%	101,3%	101,6%	101,5%	101,1%	101,3%	102,2%	103,1%	105,7%
2,6	100,1%	101,5%	101,3%	103,2%	101,5%	102,7%	103,7%	105,4%	107,3%	107,8%	107,7%
2,5	104,6%	103,0%	104,9%	105,5%	104,8%	103,0%	105,4%	109,2%	110,3%	109,4%	107,4%
2,3	112,4%	110,0%	110,1%	109,5%	108,8%	107,7%	108,5%	111,2%	113,0%	112,2%	107,6%
2,1	123,2%	117,5%	116,9%	113,8%	114,0%	113,3%	113,8%	116,5%	120,2%	120,6%	116,0%
1,9	117,7%	124,7%	123,8%	122,1%	124,1%	125,9%	124,6%	124,4%	125,5%	126,4%	122,4%
1,7	132,0%	134,3%	134,0%	133,8%	135,7%	136,7%	134,7%	133,1%	133,7%	134,4%	133,5%
1,5	144,6%	147,8%	144,4%	146,8%	148,5%	150,2%	147,6%	145,4%	145,1%	145,3%	145,2%
1,3	176,3%	165,3%	162,2%	167,0%	165,1%	165,2%	163,3%	163,9%	165,0%	164,3%	166,2%
1,1	187,1%	184,4%	183,1%	188,9%	183,0%	182,6%	181,4%	187,1%	188,3%	186,0%	187,3%

CF/UnCF [GHz]	0	2,5	2,4	2,2	2	1,8	1,6	1,4	1,2	1	0,8
0	100 %	102 %	100 %	97 %	91 %	89 %	89 %	90 %	89 %	86 %	84 %
2,6	100 %	101 %	101 %	100 %	96 %	92 %	91 %	90 %	88 %	85 %	84 %
2,5	104 %	103 %	104 %	101 %	97 %	92 %	90 %	89 %	87 %	85 %	82 %
2,3	111 %	110 %	108 %	103 %	98 %	93 %	90 %	88 %	87 %	84 %	80 %
2,1	121 %	112 %	112 %	[illegible]	[illegible]	[illegible]	92 %	[illegible]	[illegible]	[illegible]	[illegible]
1,9	82 %	113 %	117 %	110 %	105 %	101 %	97 %	93 %	91 %	89 %	85 %
1,7	88 %	115 %	123 %	116 %	111 %	107 %	102 %	97 %	94 %	92 %	89 %
1,5	94 %	122 %	130 %	125 %	119 %	114 %	108 %	103 %	99 %	97 %	94 %
1,3	109 %	133 %	142 %	138 %	129 %	122 %	116 %	112 %	109 %	105 %	104 %
1,1	113 %	145 %	156 %	152 %	139 %	132 %	125 %	124 %	120 %	115 %	113 %

Grace CPU – Energy Efficiency Leader

- Outperformed all CPUs in MFLOPs/W efficiency
- 8–26% savings with <3% runtime increase
- Up to 30% savings under relaxed runtime constraints
- Ideal for energy-aware workloads

Set CPU frequency [GHz]	Runtime [s]	Runtime extension	Average power consumption of the node [W]	Node energy consumption [kJ]	Compute node energy savings	MFLOPS/W
0.00	69.0	100.0%	617.7	42.6	0.0%	218.1
3.30	71.0	102.9%	613.8	43.6	-2.3%	213.3
3.10	71.0	102.9%	528.8	37.5	11.9%	247.6
2.90	71.0	102.9%	481.4	34.2	19.8%	272.0
2.70	71.0	102.9%	447.0	31.7	25.5%	292.9
2.50	72.0	104.3%	417.1	30.0	29.5%	309.6
2.30	75.0	108.7%	388.0	29.1	31.7%	319.4
2.10	75.0	108.7%	363.7	27.3	36.0%	340.7
1.90	81.0	117.4%	345.1	28.0	34.4%	332.5
1.70	82.0	118.8%	337.3	27.7	35.1%	336.1
1.50	83.0	120.3%	328.0	27.2	36.1%	341.4
1.30	86.0	124.6%	320.4	27.6	35.3%	337.2
1.10	90.0	130.4%	314.1	28.3	33.7%	328.8
0.90	99.0	143.5%	300.6	29.8	30.2%	312.3

Grace CPU – Energy Efficiency Leader

- Outperformed all CPUs in MFLOPs/W efficiency
- 8–26% savings with <3% runtime increase
- Up to 30% savings under relaxed runtime constraints
- Ideal for energy-aware workloads

Set CPU frequency [GHz]	Runtime [s]	Runtime extension	Average power consumption of the node [W]	Node energy consumption [kJ]	Compute node energy savings	MFLOPS/W
0.00	69.0	100.0%	617.7	42.6	0.0%	218.1
3.30	71.0	102.9%	613.8	43.6	-2.3%	213.3
3.10	71.0	102.9%	528.8	37.5	11.9%	247.6
2.90	71.0	102.9%	481.4	34.2	19.6%	272.0
2.70	71.0	102.9%	447.0	31.7	25.5%	292.9
2.50	72.0	104.3%	417.1	30.0	29.5%	309.6
2.30	75.0	108.7%	388.0	29.1	31.7%	319.4
2.10	75.0	108.7%	363.7	27.3	36.0%	340.7
1.90	81.0	117.4%	345.1	28.0	34.4%	332.5
1.70	82.0	118.8%	337.3	27.7	35.1%	336.1
1.50	83.0	120.3%	328.0	27.2	36.1%	341.4
1.30	86.0	124.6%	320.4	27.6	35.3%	337.2
1.10	90.0	130.4%	314.1	28.3	33.7%	328.8
0.90	99.0	143.5%	300.6	29.8	30.2%	312.3

Grace CPU – Energy Efficiency Leader

- Outperformed all CPUs in MFLOPs/W efficiency
- 8–26% savings with <3% runtime increase
- Up to 30% savings under relaxed runtime constraints
- Ideal for energy-aware workloads

Set CPU frequency [GHz]	Runtime [s]	Runtime extension	Average power consumption of the node [W]	Node energy consumption [kJ]	Compute node energy savings	MFLOPS/W
0.00	69.0	100.0%	617.7	42.6	0.0%	218.1
3.30	71.0	102.9%	613.8	43.6	-2.3%	213.3
3.10	71.0	102.9%	528.8	37.5	11.9%	247.6
2.90	71.0	102.9%	481.4	34.2	19.8%	272.0
2.70	71.0	102.9%	447.0	31.7	25.5%	292.9
2.50	72.0	104.3%	417.1	30.0	29.5%	309.6
2.30	75.0	108.7%	388.0	29.1	31.7%	319.4
2.10	75.0	108.7%	363.7	27.3	36.0%	340.7
1.90	81.0	117.4%	345.1	28.0	34.4%	332.5
1.70	82.0	118.8%	337.3	27.7	35.1%	336.1
1.50	83.0	120.3%	328.0	27.2	36.1%	341.4
1.30	86.0	124.6%	320.4	27.6	35.3%	337.2
1.10	90.0	130.4%	314.1	28.3	33.7%	328.8
0.90	99.0	143.5%	300.6	29.8	30.2%	312.3

Key Takeaways Across Platforms

- Grace: best energy efficiency, scalable tuning
- Sapphire Rapids: tunable, flexible with memory options
- A100: efficient baseline, gains still achievable
- Static tuning: 10–20% savings, <3% runtime cost

	Pluto	**Open GADGET**	**iPIC3D**	**RAMSES**	**BHAC**	**FIL**	**ChaNGa**
Nvidia A100	**-6%** / 103% **-9%** / 113%	**-7%** / 102% **-7%** / 102%	**-3%** / 104% **-5%** / 111%	- - - -	- - - -	- - - -	- - - -
SPR w. DDR	**-9%** / 102% **-10%** / 106%	**-7%** / 102% **-7%** / 102%	**-7%** / 102% **-9%** / 108%	**-6%** / 102% **-7%** / 104%	**-10%** /103% **-14%** /110%	**-6%** / 102% **-7%** / 103%	**-14%**/ 103% **-20%** / 107%
SPR w. HBM	**-4%** / 101% **-6%** / 105%	**-9%** / 94% **-11%** / 98%	**-7%** / 101% **-7%** / 101%	**-7%** / 102% **-8%** / 104%	**-4%** / 99% **-4%** / 99%	**-4%** / 100% **-6%** / 104%	**-12%**/ 102% **-13%**/ 103%
Grace CPU	**-22%** /101% **-35%** /122%	**-13%** /103% **-33%** /128%	**-9%** / 103% **-29%** /126%	**-19%** /101% **-28%** /137%	**-26%** /103% **-36%** /109%	**-8%** / 102% **-20%** /117%	**-16%** / 102% **-30%** / 110%
Cascade Lake D2.2*	**-6%** /102% **-12%** /126%	**-9%** /103% **-12%** /130%	**-6%** / 101% **-13%** /115%	**-7%** /102% **-11%** /123%	**-5%** /102% **-11%** /118%	**-3%** / 102% **-13%** /127%	**30%** /102% **36%** / 110%

Key Takeaways Across Platforms

	Pluto
Nvidia A100	**-6%** / 103% **-9%** / 113%
[illegible]	[illegible]

Key Takeaways Across Platforms

- Grace: best energy efficiency, scalable tuning
- Sapphire Rapids: tunable, flexible with memory options
- A100: efficient baseline, gains still achievable
- Static tuning: 10–20% savings, <3% runtime cost

	Pluto	Open GADGET	iPIC3D	RAMSES	BHAC	FIL	ChaNGa
Nvidia A100	**-6%** / 103% **-9%** / 113%	**-7%** / 102% **-7%** / 102%	**-3%** / 104% **-5%** / 111%	- - - -	- - - -	- - - -	- - - -
SPR w. DDR	**-9%** / 102% **-10%** / 106%	**-7%** / 102% **-7%** / 102%	**-7%** / 102% **-9%** / 108%	**-6%** / 102% **-7%** / 104%	**-10%** /103% **-14%** /110%	**-6%** / 102% **-7%** / 103%	**-14%**/ 103% **-20%** / 107%
SPR w. HBM	**-4%** / 101% **-6%** / 105%	**-9%** / 94% **-11%** / 98%	**-7%** / 101% **-7%** / 101%	**-7%** / 102% **-8%** / 104%	**-4%** / 99% **-4%** / 99%	**-4%** / 100% **-6%** / 104%	**-12%**/ 102% **-13%**/ 103%
Grace CPU	**-22%** /101% **-35%** /122%	**-13%** /103% **-33%** /128%	**-9%** / 103% **-29%** /126%	**-19%** /101% **-28%** /137%	**-26%** /103% **-36%** /109%	**-8%** / 102% **-20%** /117%	**-16%** / 102% **-30%** / 110%
Cascade Lake D2.2*	**-6%** /102% **-12%** /126%	**-9%** /103% **-12%** /130%	**-6%** / 101% **-13%** /115%	**-7%** /102% **-11%** /123%	**-5%** /102% **-11%** /118%	**-3%** / 102% **-13%** /127%	**30%** /102% **36%** / 110%

Key Takeaways Across Platforms

- Grace: best energy efficiency, scalable tuning
- Sapphire Rapids: tunable, flexible with memory options
- A100: efficient baseline, gains still achievable
- Static tuning: 10–20% savings, <3% runtime cost

	Pluto	**Open GADGET**	**iPIC3D**	**RAMSES**	**BHAC**	**FIL**	**ChaNGa**
Nvidia A100	**-6%** / 103% **-9%** / 113%	**-7%** / 102% **-7%** / 102%	**-3%** / 104% **-5%** / 111%	- - - -	- - - -	- - - -	- - - -
SPR w. DDR	**-9%** / 102% **-10%** / 106%	**-7%** / 102% **-7%** / 102%	**-7%** / 102% **-9%** / 108%	**-6%** / 102% **-7%** / 104%	**-10%** /103% **-14%** /110%	**-6%** / 102% **-7%** / 103%	**-14%**/ 103% **-20%** / 107%
SPR w. HBM	**-4%** / 101% **-6%** / 105%	**-9%** / 94% **-11%** / 98%	**-7%** / 101% **-7%** / 101%	**-7%** / 102% **-8%** / 104%	**-4%** / 99% **-4%** / 99%	**-4%** / 100% **-6%** / 104%	**-12%**/ 102% **-13%**/ 103%
Grace CPU	**-22%** /101% **-35%** /122%	**-13%** /103% **-33%** /128%	**-9%** / 103% **-29%** /126%	**-19%** /101% **-28%** /137%	**-26%** /103% **-36%** /109%	**-8%** / 102% **-20%** /117%	**-16%** / 102% **-30%** / 110%
Cascade Lake D2.2*	**-6%** /102% **-12%** /126%	**-9%** /103% **-12%** /130%	**-6%** / 101% **-13%** /115%	**-7%** /102% **-11%** /123%	**-5%** /102% **-11%** /118%	**-3%** / 102% **-13%** /127%	**30%** /102% **36%** / 110%

Practical Relevance for HPC Operations

- 5% savings = 1 MW on a 20 MW system
- Job-level tuning via MERIC + SLURM = no code change
- Immediately deployable across architectures

“…can be applied job-wide using a job scheduler.”

Final Thoughts and What's Next

- Static tuning = ready-to-use, impactful optimization method
- Works on diverse codes and hardware platforms
- Future: dynamic tuning with real-time adjustments
- MERIC supports both static and dynamic strategies

Acknowledgement & Disclaimer

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain. Neither the European Union nor the granting authority can be held responsible for them

Co-funded by the European Union

EuroHPC Joint Undertaking

SPACE

CENTRE OF EXCELLENCE FOR HPC ASTROPHYSICAL APPLICATIONS

Thank you…

João Barbosa
(IT4i@VSB) joao.barbosa@vsb.cz

Co-funded by the European Union

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.

EuroHPC Joint Undertaking