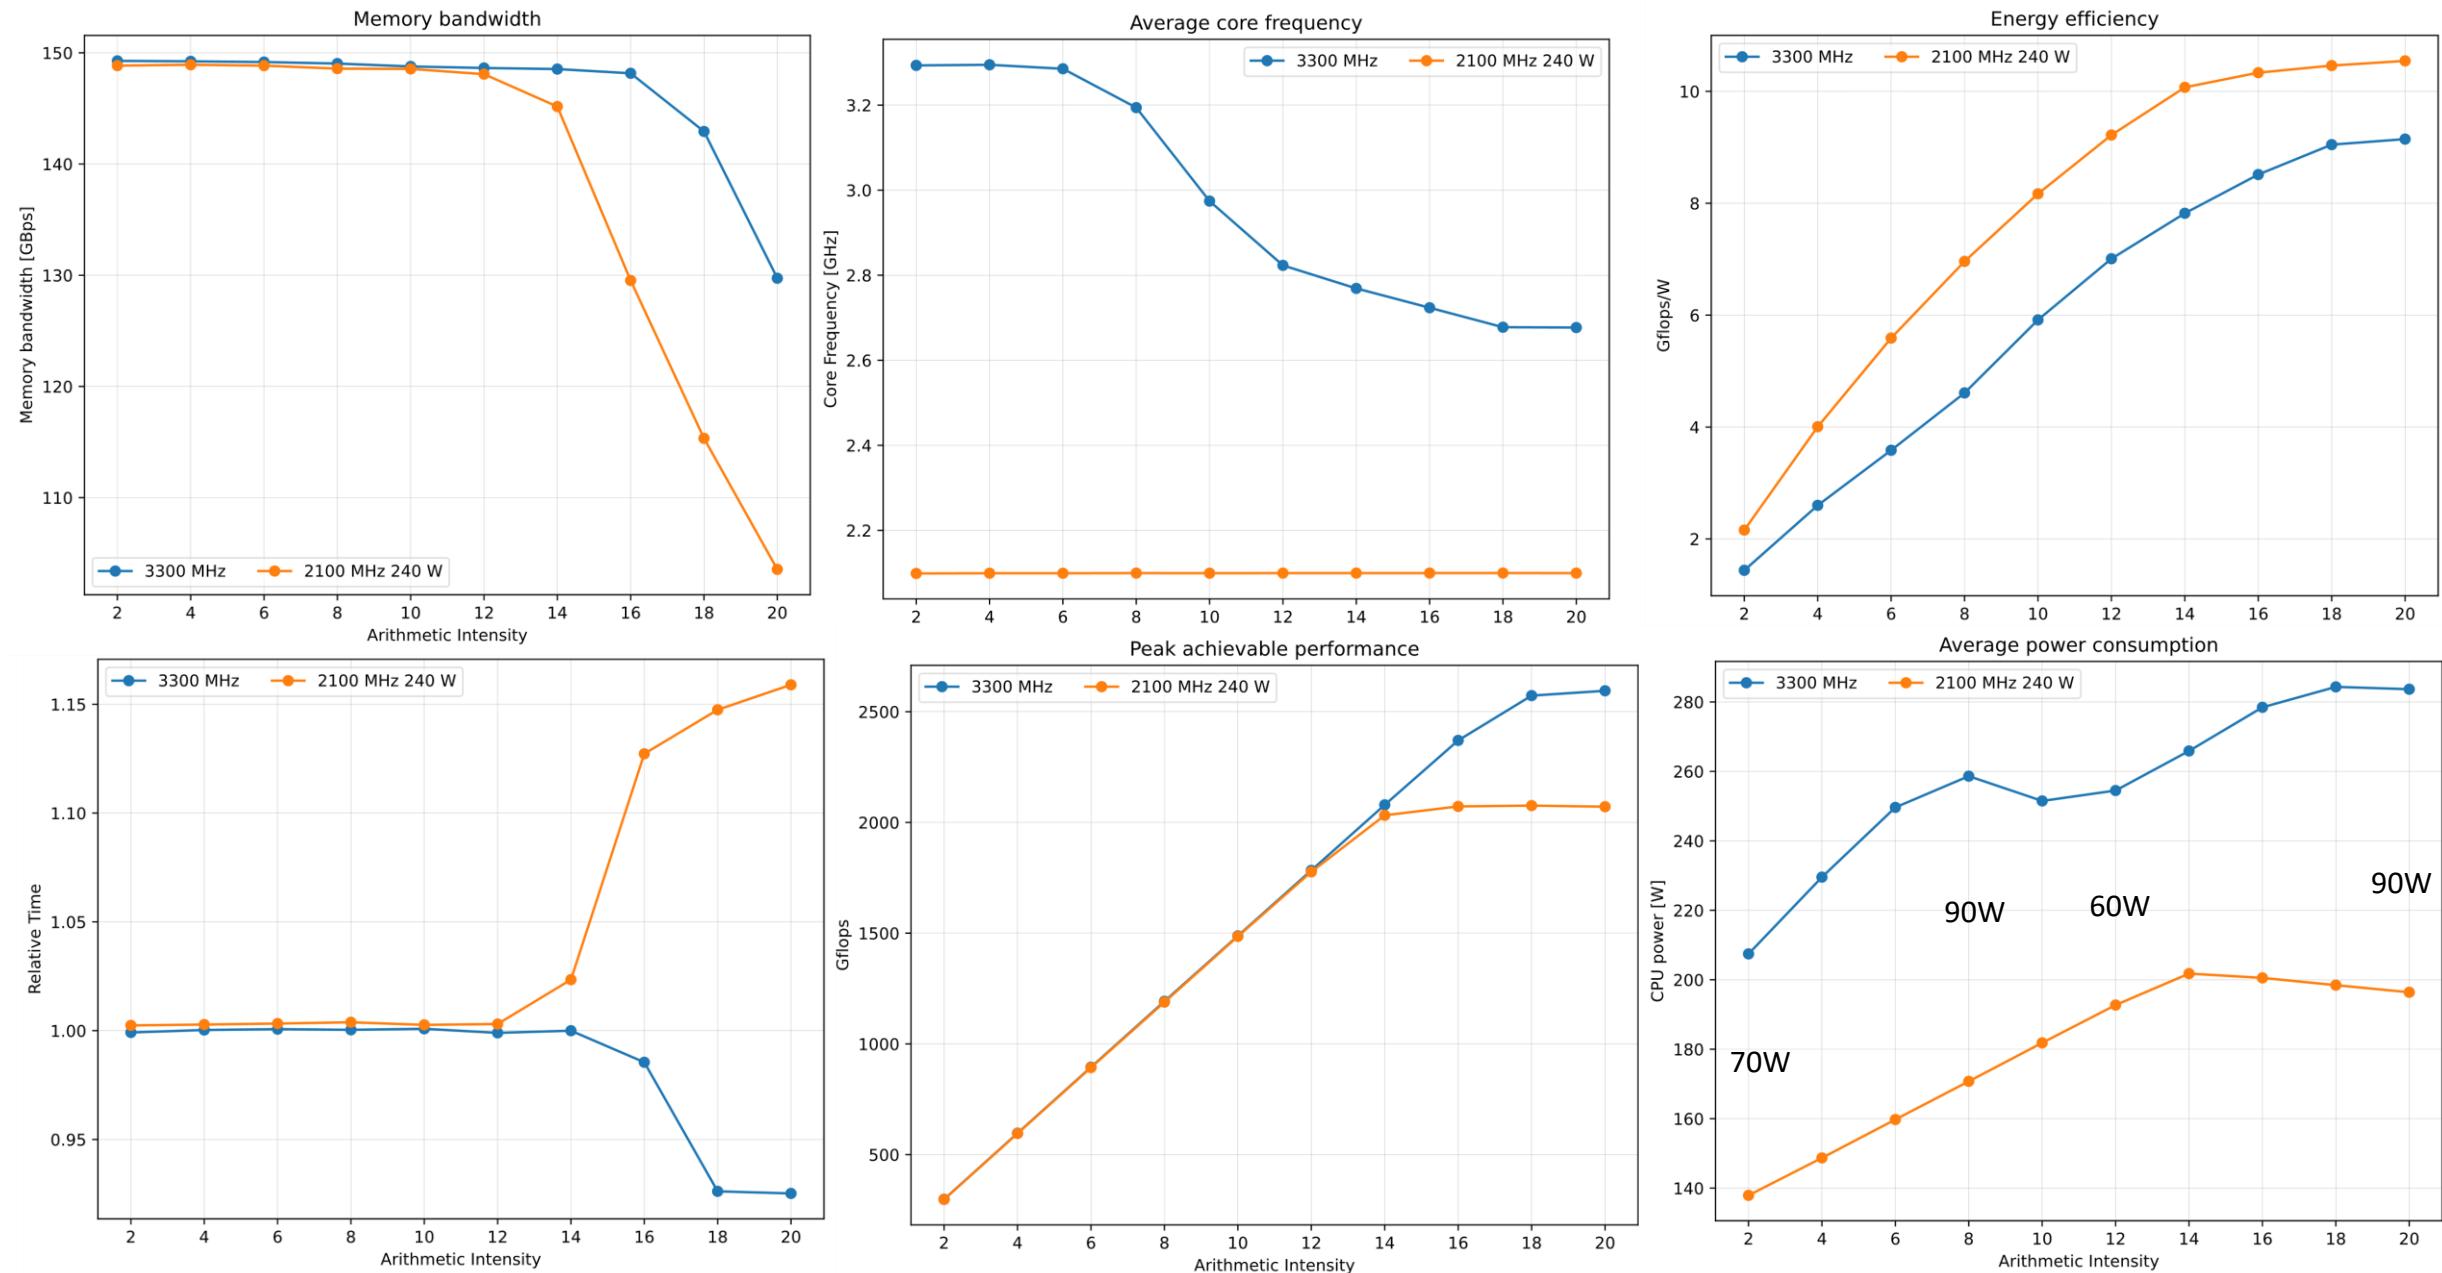


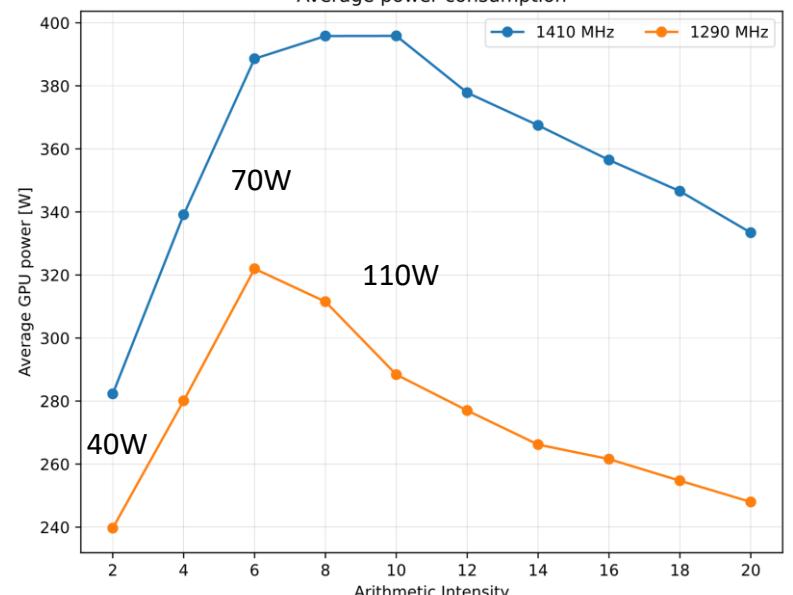
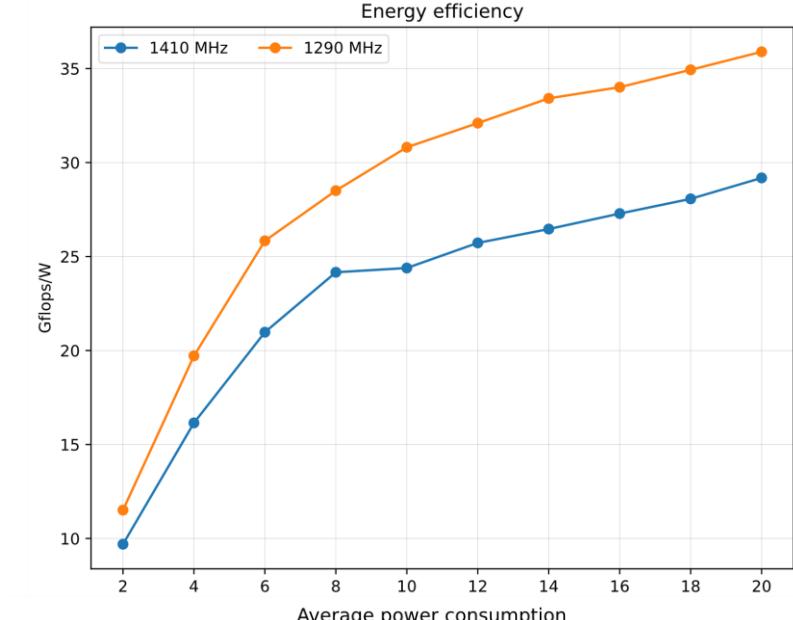
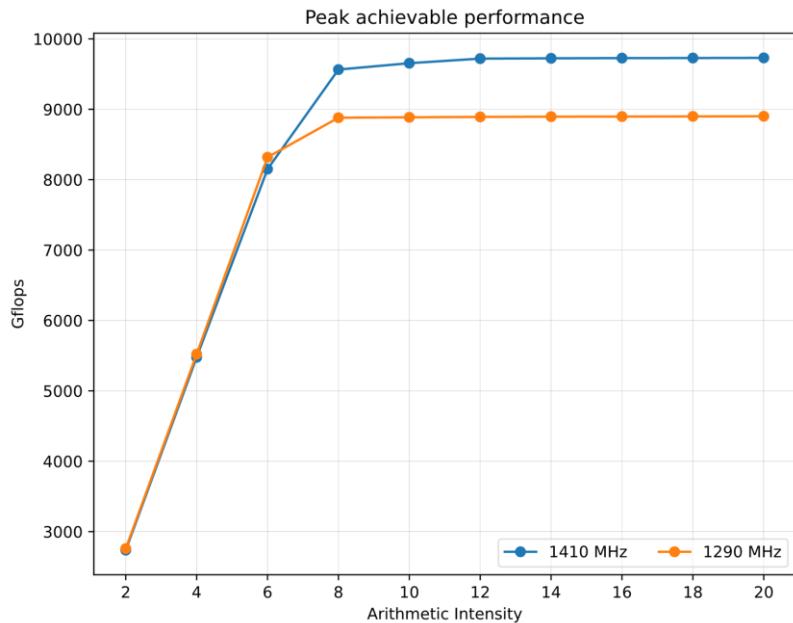
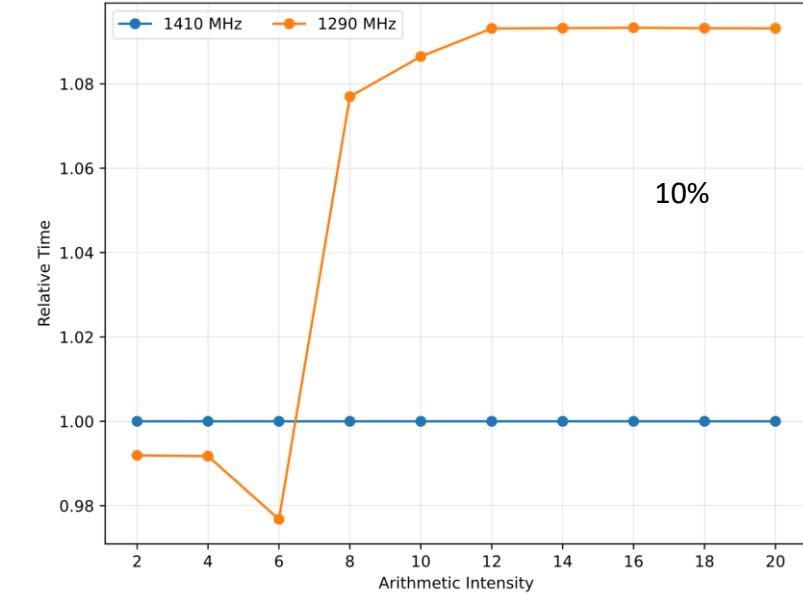
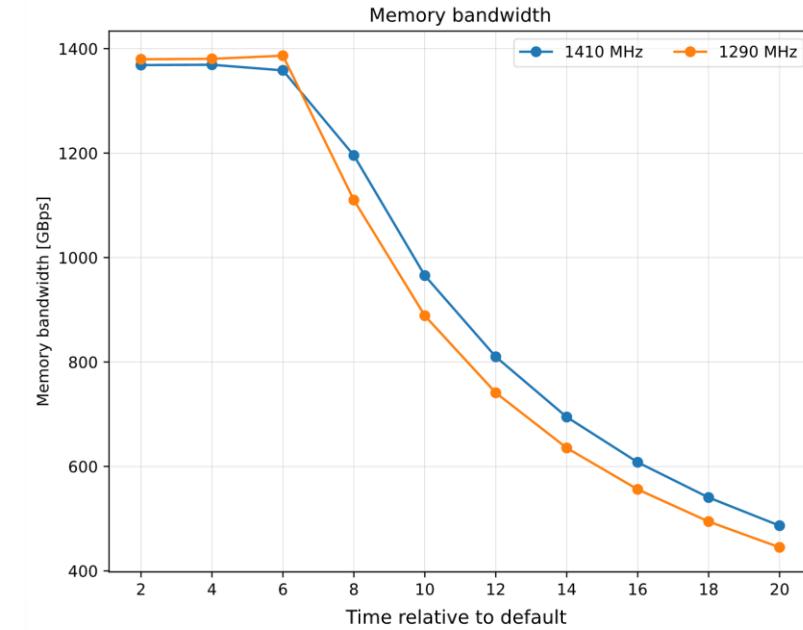


IT4Innovations national supercomputing center Karolina energy saving measures

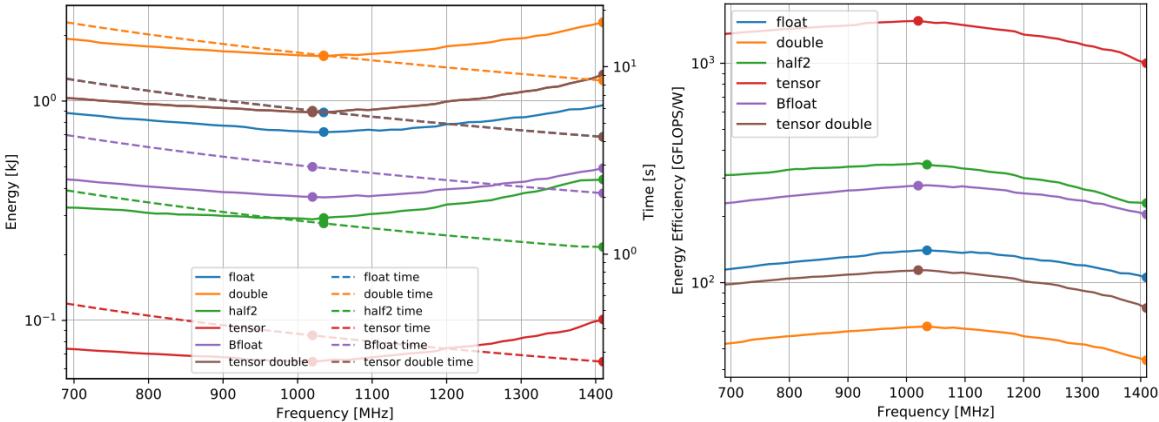
AMD 7h12 AVX2 deployed configuration



A100 deployed configuration



Spetko, Vysocky, Jansik, Riha "DGX-A100 Face to Face DGX-2—Performance, Power and Thermal Behavior Evaluation", Energies 2021



	Frequency [MHz]	Time [s]	Time Difference	Energy [J]	Energy Savings	Performance [TFLOPS]	Energy Efficiency [GFLOPS/W]
double	1410	8.43		2285		9.71	35.86
	1035	11.49	136.19%	1601	29.91%	7.13	51.16
float	1410	4.23		958		19.37	85.53
	1035	5.76	136.13%	721	24.76%	14.23	113.68
Bfloat	1410	2.11		494		38.75	165.74
	1035	2.88	136.19%	364	26.40%	28.46	225.17
half2	1380 *	1.09		439		75.04	186.69
	1020	1.48	135.65%	289	34.15%	55.32	283.53
tensor half	1410	0.27		101		307.02	810.80
	1020	0.37	138.18%	65	35.86%	222.18	1264.15
tensor double	1410	4.21		1321		19.44	62.02
	1020	5.82	138.19%	887	32.82%	14.07	92.32

Figure 9. DGX-A100 frequency scaling of Mandelbrot benchmark. The plot in the top left corner shows consumed energy and run-time of workload. The plot in the top right corner shows energy efficiency. The plots are in logarithmic scale. The table at the bottom of the figure shows the time difference, energy savings, and energy efficiency of the Mandelbrot benchmark at the base and at the optimal frequency. The marker in the plot highlights the two frequencies from the table.

* During half2 test the maximum frequency 1410 MHz was not reached due to power throttling.