

Fusion: tokamak infrared cameras

Lossless, Phidelity: $\times 2.3$ over the archive's actual storage, -35% against FFV1 with best settings, -38% against JPEG XL at maximum effort.

Data

FAIR-MAST (UKAEA, CC BY-SA 4.0), MAST tokamak, shots 30390 to 30420: InSb camera 8×320 at 5 kHz and CMT camera 32×256 at 0.8 kHz, 14 bits; 15 sequences, 163 MB raw.

On PC — one core

Format	Ratio (raw / compressed)	Encoding (M samples/s)	Decoding (M samples/s)
Phidelity compact	4.31	12	12
Phidelity light	3.86	70	80
FFV1 (1 slice, large context, g 300)	2.81	33	25
JPEG XL e9	2.69	0.4	16
JPEG XL e7	2.60	2.0	16
MAT compressed (deflate)	2.66	9	215
JPEG-LS	2.55	55	62
HDF5 gzip + shuffle	2.09	40	134
Zarr blosc lz4 (storage FAIR-MAST)	1.90	—	—
NPZ, MAT, HDF5 defaults	1.00	—	—
Phidelity compact ± 2	9.73	13	15
Phidelity light ± 2	7.42	53	86

Good to know

- These are narrow high-rate windows, not full frames.
- JPEG XL and JPEG-LS code each sequence as one tall image (no video mode).

Conditions

- PC: AMD Ryzen 7 8700F, a single core; every file decoded and compared with the original (zero error lossless, at most k in $\pm k$).
- $\pm k$: guaranteed maximum error of k on every sample. Phidelity and JPEG-LS (NEAR) hold it natively; for JPEG XL and the reference format, the image is first rounded to a $2k + 1$ step, then coded losslessly.
- Codecs: libjxl 0.12 (efforts 3, 7 and 9), CharLS (JPEG-LS), OpenJPEG (JPEG 2000), FFV1 through ffmpeg, LERC, CCSDS 123.0-B-2 (encoder written for the benchmark, byte-identical to the NTNU verification model).
- Memory: peak private memory during the call only, measured on the largest item of the set.
- Embedded: instruction counts in emulation (QEMU), not board timings; cropped samples, sizes comparable with each other only. Compact and light through the 1.7 kit, fast through its core; in both cases with a caller-provided output buffer. Byte-identical streams on PC and ARM. Code on Cortex-M7: fast 73 KB, compact 92 KB, CharLS 205 KB, CCSDS 123 6 KB.