

A compression scheme for radio data in high performance computing

K. Masui^{a,b,*}, M. Amiri^a, L. Connor^{c,d,e}, M. Deng^a, M. Fandino^a, C. Höfer^a, M. Halpern^a, D. Hanna^f, A.D. Hincks^a, G. Hinshaw^a, J.M. Parra^f, L.B. Newburgh^d, J.R. Shaw^{a,c}, K. Vanderlinde^{e,d} *Astronomy and Computing* 12 (2015) 181–190

Information

#6

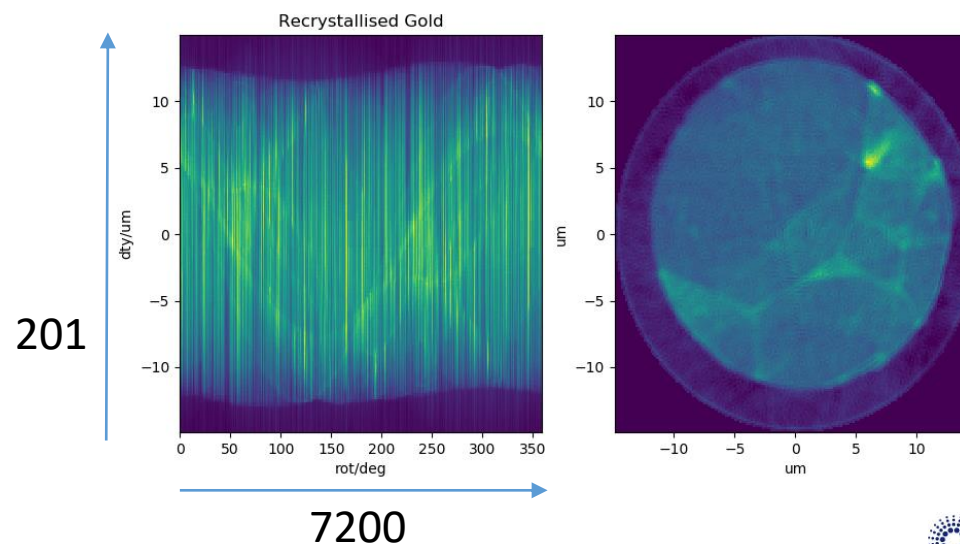
To ensure highest stability at full frame rates, DECTRIS strongly advises using *bslz4* compression.



500
fps

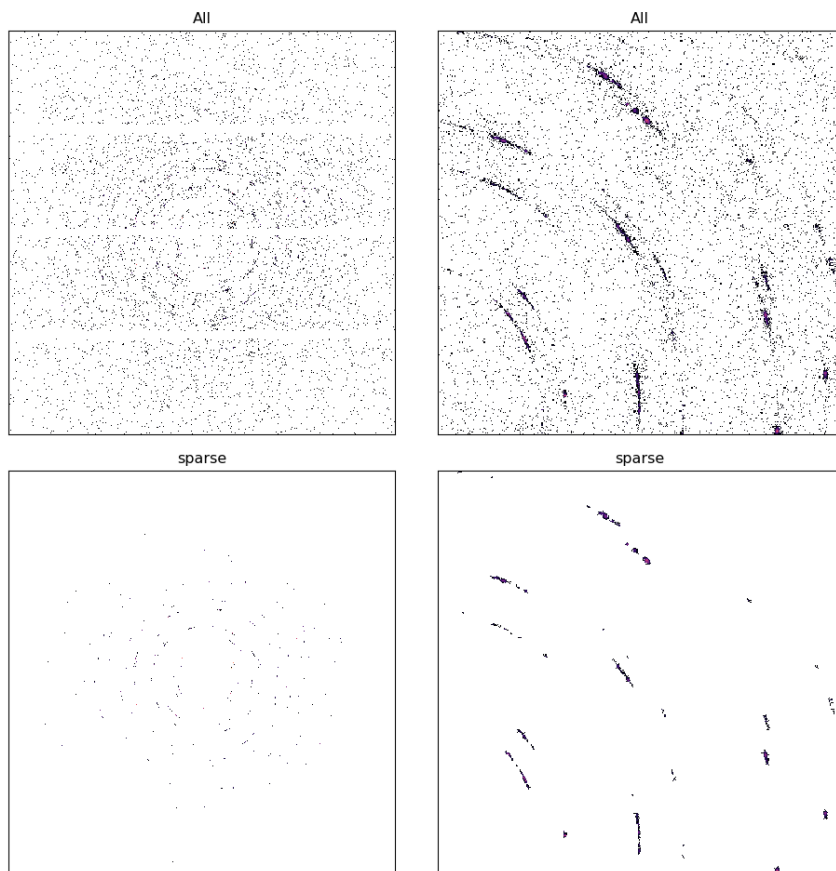
Data rate
(raw uncompressed)
... 4 GB/s
... 14 TB / hour
... 345 TB / day

1 day XRDCT : [30] [201] [7200] [2162] [2068]



Diffraction data : Average bits per pixel is very low

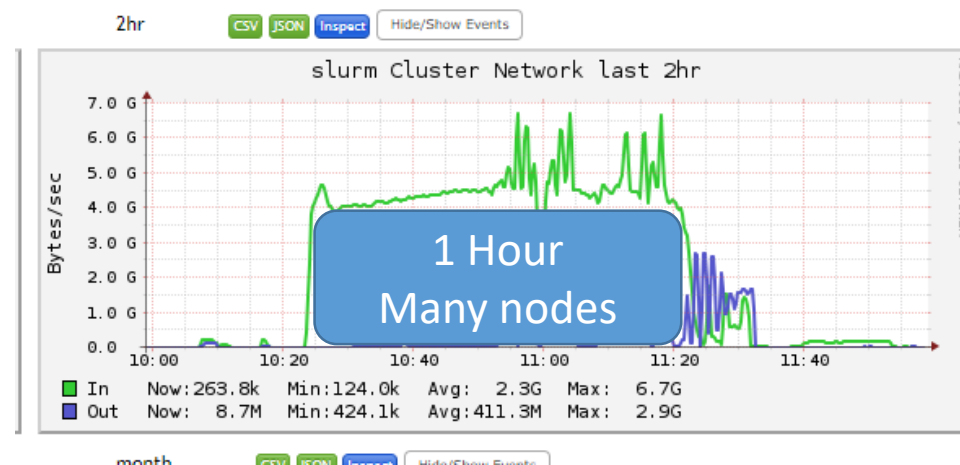
Peak $\sim 10,000$ ph/px
Background < 0.1 ph/px



Recent example:

Dectris bslz4 in hdf5 containers
10 - 45X compression

~ 500 TB raw \rightarrow 13 TB disk \rightarrow 62 GB sparse



CPU data reduction from BSLZ4

- “Typically” 1 GB/s/core for LZ4 decompression*
- 1 core is 4X slower than the detector
- Stream benchmark gives >10 GB/s single core (50 GB/s all cores)
- GPFS supplies hdf5 data at 2 GB/s
 - Network storage can offer 20 – 80 GB/s uncompressed
 - Needs 20-80 cores
- GPU’s are useful for image processing anyway (background estimate, etc)

(* ignoring: LZ4 from Intel IPP, threading per frame, reductions inside the decompressor)

GPU decompression : nvcomp offers LZ4

- BSLZ4 HDF5 plugin uses an internal block size of 8 kB
- ESRF Lima uses 1 frame hdf5 chunk-size (2162x2068 pixels)
 - u16 : 8.5 MB = 1094.5859375 blocks
 - u32 : 17.1 MB = 2189.171875 blocks
- Scan through chunks on CPU to locate internal blocks
 - 58 μ s / frame h5py for IO
 - 76 μ s / frame with scanning (overlap with IO?)
 - ... ROI decompression

Tesla V100-SXM2-32GB, 1000 frames, 16 bit

```
==1769648== Profiling application: python3 -O test_itercuda.py /tmp_14_days/eiger_0000.h5 /entry_0000/ESRF-ID11/eiger/data
==1769648== Profiling result:
```

Type	Time (%)	Time	Calls	Avg	Min	Max	Name
GPU activities:	41.37%	74.668ms	1000	74.668us	67.931us	117.72us	h5lz4dc
	35.20%	63.543ms	1000	63.543us	62.844us	64.507us	shuf_8192_16
	13.57%	24.486ms	1000	24.485us	23.772us	25.628us	reduce_kernel_stage1
	6.93%	12.505ms	2000	6.2520us	1.2470us	12.704us	[CUDA memcpy HtoD]
	1.34%	2.4179ms	1000	2.4170us	2.1400us	5.4040us	simple_shuffle_end
	0.91%	1.6496ms	1000	1.6490us	1.4360us	2.0440us	reduce_kernel_stage2
	0.69%	1.2366ms	1000	1.2360us	1.2150us	1.3120us	[CUDA memcpy DtoH]
API calls:	44.70%	254.97ms	1	254.97ms	254.97ms	254.97ms	cuCtxCreate
	15.60%	88.971ms	3003	29.627us	9.1110us	291.51us	cuMemFree
	10.42%	59.418ms	1	59.418ms	59.418ms	59.418ms	cuCtxDetach
	9.88%	56.362ms	5000	11.272us	9.2210us	43.504us	cuLaunchKernel
	7.55%	43.067ms	2000	21.533us	13.839us	154.85us	cuMemcpyHtoD
	5.68%	32.404ms	3003	10.790us	5.3960us	4.9651ms	cuMemAlloc
	4.88%	27.857ms	1000	27.856us	25.767us	45.389us	cuMemcpyDtoH
	0.72%	4.1128ms	5000	822ns	617ns	5.7850us	cuFuncSetBlockShape
	0.26%	1.4547ms	7	207.81us	172.81us	297.96us	cuModuleLoadDataEx
	0.25%	1.4231ms	3011	472ns	336ns	6.2590us	cuCtxGetDevice
	0.04%	233.70us	7	33.386us	17.159us	106.24us	cuModuleUnload
	0.00%	17.737us	22	806ns	295ns	2.6830us	cuDeviceComputeCapability
	0.00%	16.996us	31	548ns	283ns	1.8260us	cuDeviceGetAttribute
	0.00%	10.711us	10	1.0710us	554ns	1.5800us	cuModuleGetFunction
	0.00%	8.9610us	2	4.4800us	3.8030us	5.1580us	cuDeviceGetPCIBusId
	0.00%	2.2600us	1	2.2600us	2.2600us	2.2600us	cuCtxPopCurrent
	0.00%	2.1060us	3	702ns	569ns	836ns	cuDeviceGet
	0.00%	2.0090us	3	669ns	394ns	953ns	cuDeviceGetCount
	0.00%	1.2480us	1	1.2480us	1.2480us	1.2480us	cuCtxPushCurrent

LZ4 : 120 GB/s

~180 μ s/frame
11X detector

GeForce GTX TITAN X, 1000 frames, 16 bit

```
==49242== Profiling application: python -O test_itercuda.py /data/id11/nanoscope/blc12454/id11/WAu5um/WAu5um_DT3/scan0001/eiger_
==49242== Profiling result:
```

Time (%)	Time	Calls	Avg	Min	Max	Name
38.21%	189.53ms	1000	189.53us	188.00us	193.03us	shuf_8192_16
36.68%	181.92ms	1000	181.92us	169.60us	234.21us	h5lz4dc
13.19%	65.440ms	1000	65.440us	63.393us	68.129us	reduce_kernel_stage1
9.09%	45.078ms	2000	22.539us	1.1840us	52.289us	[CUDA memcpy HtoD]
1.30%	6.4596ms	1000	6.4590us	6.3360us	7.0720us	simple_shuffle_end
0.79%	3.9153ms	1000	3.9150us	3.6160us	7.1040us	reduce_kernel_stage2
0.54%	2.6655ms	1000	2.6650us	1.9200us	9.6960us	[CUDA memcpy DtoH]
0.20%	989.96us	1000	989ns	608ns	8.1290us	[CUDA memcpy DtoD]

LZ4 : 50 GB/s

~470 μ s/frame
4X detector

```
==49242== API calls:
```

Time (%)	Time	Calls	Avg	Min	Max	Name
60.28%	1.48397s	3004	494.00us	10.084us	4.1736ms	cuMemFree
17.38%	427.98ms	3004	142.47us	9.3940us	2.3956ms	cuMemAlloc
9.92%	244.25ms	1	244.25ms	244.25ms	244.25ms	cuCtxCreate
4.32%	106.42ms	2000	53.207us	10.887us	3.0975ms	cuMemcpyHtoD
2.81%	69.269ms	5000	13.853us	8.7170us	168.53us	cuLaunchKernel
2.39%	58.827ms	1	58.827ms	58.827ms	58.827ms	cuCtxDetach
1.32%	32.521ms	1000	32.521us	18.418us	162.30us	cuMemcpyDtoH
0.94%	23.032ms	1000	23.031us	17.821us	92.496us	cuMemcpyDtoD
0.30%	7.3015ms	7	1.0431ms	538.27us	1.6180ms	cuModuleLoadDataEx
0.17%	4.2039ms	5000	840ns	456ns	11.638us	cuFuncSetBlockShape
0.15%	3.6542ms	5012	729ns	351ns	357.79us	cuCtxGetDevice
0.01%	342.64us	7	48.949us	23.517us	140.86us	cuModuleUnload
0.00%	17.916us	10	1.7910us	498ns	3.4070us	cuModuleGetFunction
0.00%	15.065us	22	684ns	311ns	1.9420us	cuDeviceComputeCapability
0.00%	11.063us	25	442ns	337ns	1.1050us	cuDeviceGetAttribute
0.00%	4.3840us	2	2.1920us	1.9100us	2.4740us	cuDeviceGetCount

Status

- ✓ Parse bslz4 chunks to blocks
- ✓ Calls to nvcomp LZ4 CUDA decompressor from python
- ✓ Naïve bitshuffle CUDA kernels

Pull requests welcome:

<https://github.com/jonwright/bslz4decoders>

- Streams to overlap IO
- GPU image processing

Thanks!

- Bjoern Enders, Nick Wright (LBNL)
- @carsonswope (github)
- Jerome Kieffer (ESRF)