

Towards a Streaming Algorithm for AGIPD Calibration and Photon Counting


Felix Brausse
Postdoc @ MID

Data-Analysis Satellite Meeting
Hamburg, Jan. 25th, 2022




Case Study: XPCS Experiments

- EuXFEL Instruments cover a vast range of experiments with widely varying conditions
 - For certain operation conditions, we may need more tailored tools
- Take XPCS: X-Ray Photon Correlation Spectroscopy
 - Based on **Speckle**
 - Speckle patterns change **stochastically** with dynamics
 - Autocorrelation reveals the driving dynamics
- Characteristic conditions:
 - Very often *very* low count rates, $\sim 10^{-2}$ ph/pix/pulse; → many runs for good statistics
 - Analysis requires **temporal autocorrelation** function





Laser-pointer speckle



Two-time correlation function

Missed Opportunities in the Current Process

- A) The Calibration Pipeline for AGIPD is very good, but it
 - i) General Purpose, based on batch processing and calibration management
 - ii) Struggles with some quirks of the detector, like flickering baselines
 - iii) Resource intense
 - iv) Stores **dense** *proc* data (compression ratio: 1) when the number of zeros $\sim (1-\lambda)$
- B) We currently cannot get temporal correlations live





Ad 1) Wobbling Baselines

What does AGIPD Raw Data look like at Low Intensities?

- Single-photon resolution up to $k \sim 4$ or so
- “Every statistical model is wrong”. Here is how I get it wrong.
 - Convolution: Poisson * Gaussian
 - → Everything from exponential family
 - Gain a , offset b , count rate λ , (fixed) width σ

$$F(x, a, b, \lambda, \sigma) = \sum_k N * P(\lambda, k) * \exp \left\{ - \left(\frac{x - (ak + b)}{a \sigma} \right)^2 \right\}$$

$$P(\lambda, k) = \frac{\lambda^k}{k!} e^{-\lambda}$$



0 1 2 photons

How Do We Fit the Model (without Histograms)?

- The Maximum-Likelihood principle tells us how; Maximize the likelihood function: $L(\boldsymbol{\theta}, \mathbf{x}) = \prod_i^n F(x_i, a, b, \lambda, \sigma)$
- Requires gradients and many iterations.
- Even better: Expectation Maximization
 - Idea: There must be a latent variable; here: photon count k
- We can decompose $F(x_i, a, b, \lambda, \sigma) = \sum_k f_k(\boldsymbol{\theta}, x_i)$ to get the likelihood function $L(\boldsymbol{\theta}, \mathbf{x}, \mathbf{k}) = \prod_i^n \prod_k f_k(\boldsymbol{\theta}, x_i)$

How Do We Fit the Model (without Histograms)?

Now we can optimize each component individually; using the weights


$$T_{i,k} = \frac{f_k(\boldsymbol{\theta}, x_i)}{\sum_k f_k(\boldsymbol{\theta}, x_i)} \quad \text{we can use the Maximum-Likelihood update for a Gaussian: the mean } \mu_k^{(t+1)} = \frac{\sum_i T_{i,k} x_i}{\sum_i T_{i,k}}$$

(this looks like a *softmax*!)

(this is the *k*-means algorithm!)

Simplify even further: set $\max_k \{T_{i,k}\}$ to 1, the rest to 0



→ “The-Winner-Takes-It-All” Flavor



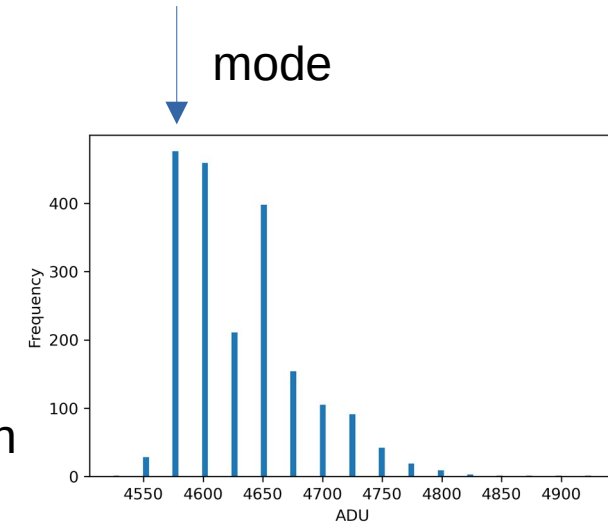
Problem: We Need a Starting Guess

Solution: Look at the Mode (Most Frequent Element)

- For $\lambda < 1$ the mode is always 0.
- For n data points we have to reserve n memories + n counters *worst case*
- Special case of 1 memory: the Boyer-Moore Majority-Vote Algorithm



floor
division



- Extend majority vote to k memories
- SpaceSaving(k) works really well for $k = 2$

Algorithm 3: SPACESAVING(k)

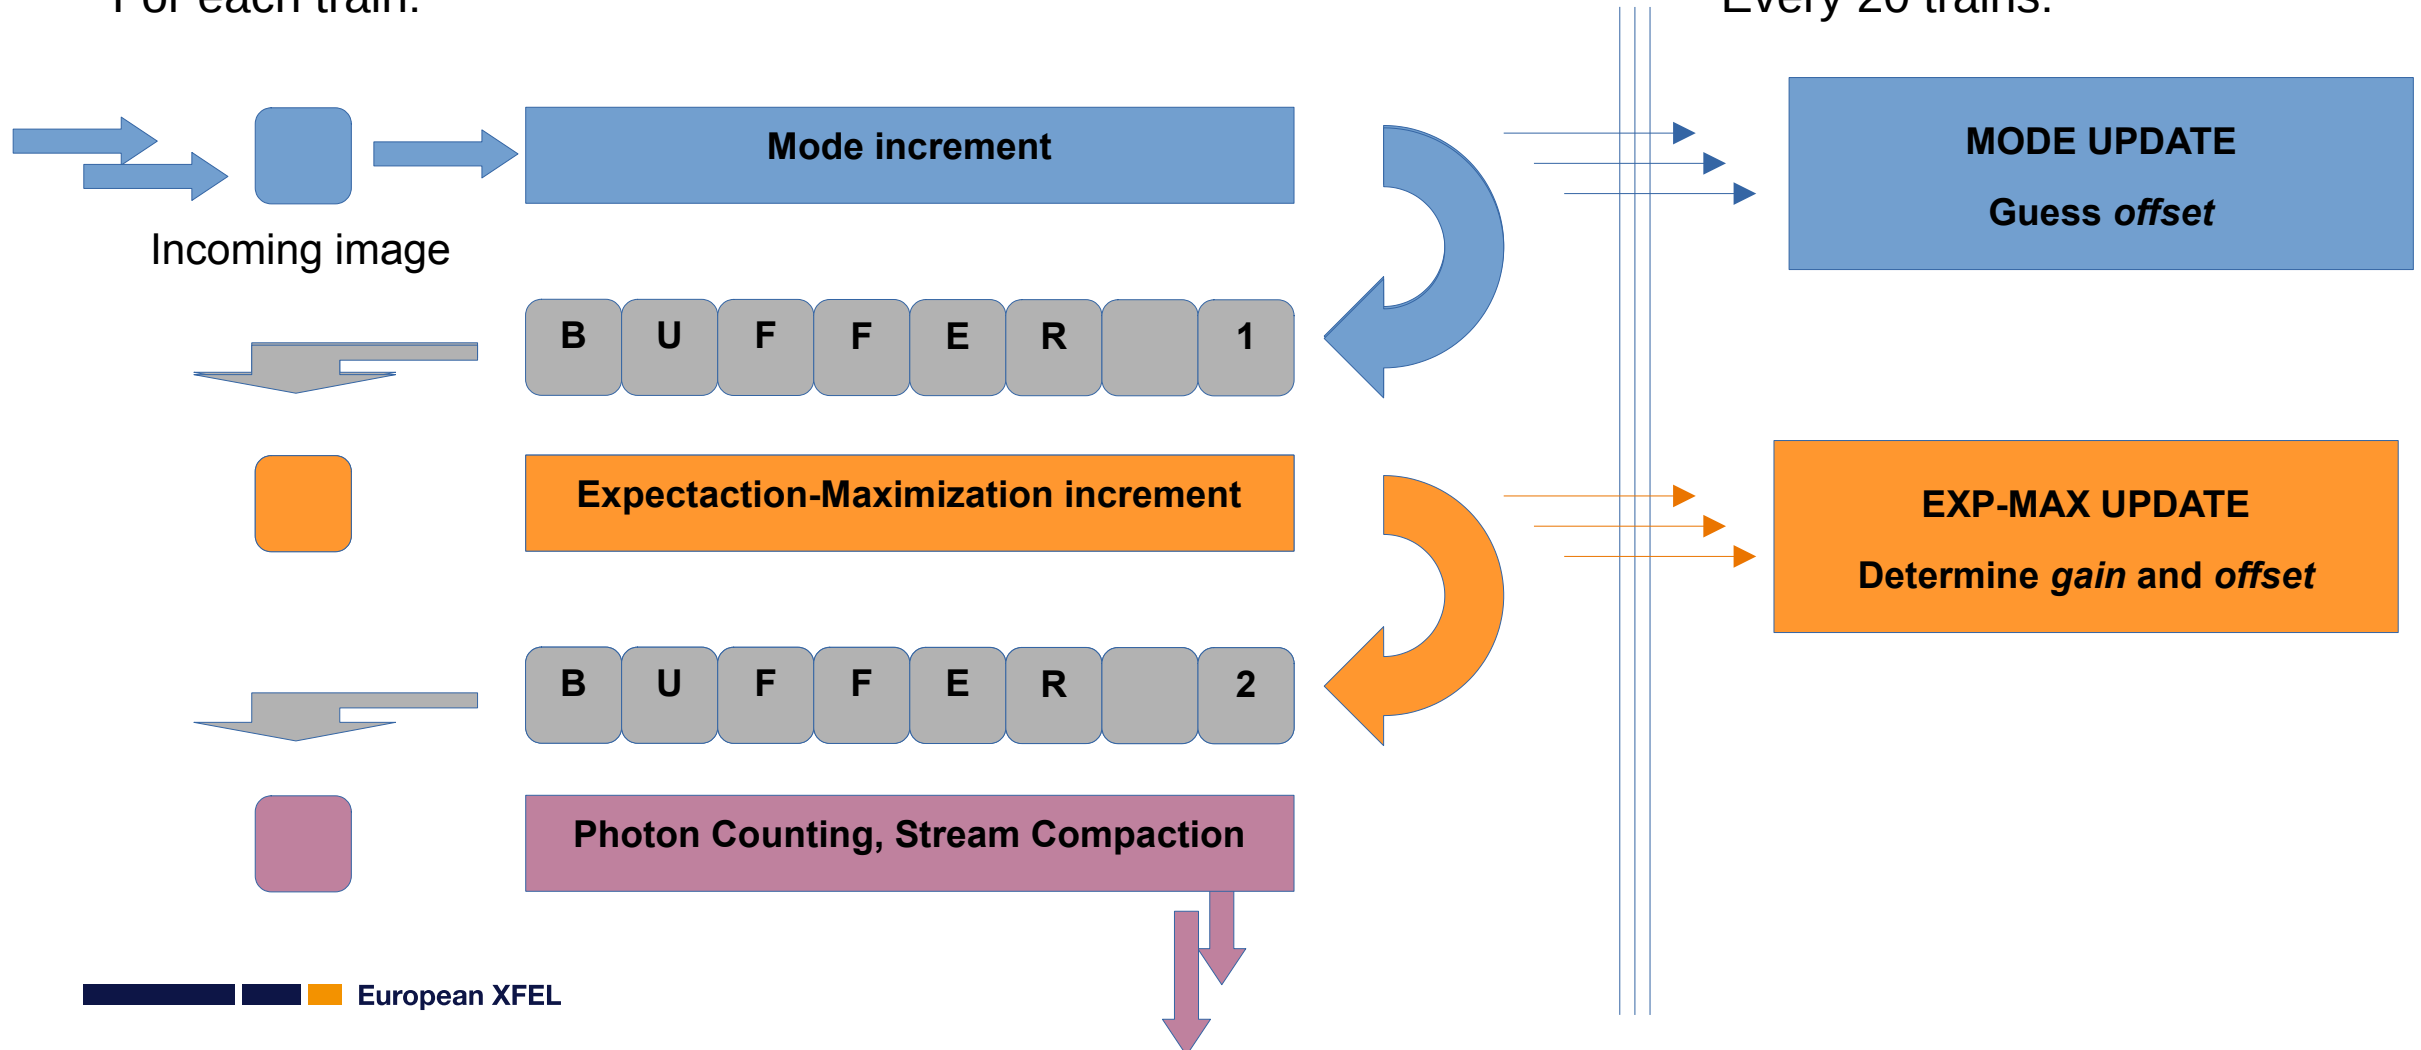
```

 $n \leftarrow 0$ ;
 $T \leftarrow \emptyset$ ;
foreach  $i$  do
     $n \leftarrow n + 1$ ;
    if  $i \in T$  then  $c_i \leftarrow c_i + 1$ ;
    else if  $|T| < k$  then
         $T \leftarrow T \cup \{i\}$ ;
         $c_i \leftarrow 1$ ;
    else
         $j \leftarrow \arg \min_{j \in T} c_j$ ;
         $c_i \leftarrow c_j + 1$ ;
         $T \leftarrow T \cup \{i\} \setminus \{j\}$ ;

```


Putting it all together

For each train:



Let Us Look At Some Code

Iteration of the SpaceSaving(2)

```

1 import cupy as cp
2
3 def iterMajority(im, args):
4     val1, val2, cnt1, cnt2 = args
5
6     im = im // 25
7
8     tst1 = im == val1
9     tst2 = im == val2
10
11     cnt1 = cp.where(tst1, cnt1+1, cnt1)
12     cnt2 = cp.where(tst2, cnt2+1, cnt2)
13
14     argmin = cnt1 < cnt2
15     updMsk = ~(tst1 | tst2)
16
17     val1 = cp.where(updMsk & argmin, im, val1)
18     cnt1 = cp.where(updMsk & argmin, cnt1+1, cnt1)
19
20     val2 = cp.where(updMsk & ~argmin, im, val2)
21     cnt2 = cp.where(updMsk & ~argmin, cnt2+1, cnt2)
22
23     return val1, val2, cnt1, cnt2
24
25

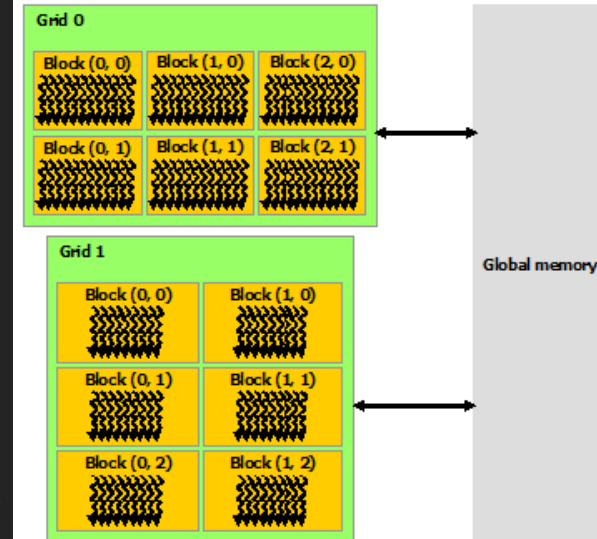
```

```

35 extern "C" __global__
36 void iterMajorBits(unsigned short* im, unsigned int* val, unsigned int* cnt) {
37
38     int tid = blockDim.x * blockIdx.x + threadIdx.x;
39     char cs = 0;
40
41     short imCast = im[tid] >> 4;
42     unsigned int tmpCnt = cnt[tid], tmpVal = val[tid];
43
44     short c1 = tmpCnt >> 16;
45     short c2 = tmpCnt & 0xFFFF;
46     short v1 = tmpVal >> 16;
47     short v2 = tmpVal & 0xFFFF;
48
49     cs |= (v1 == imCast) << 0;
50
51     c1 += 1;
52     cnt[tid] = (tmpCnt & 0x0000FFFF) | (c1 << 16);
53     break;
54
55     case 2:
56         c2 += 1;
57         cnt[tid] = (tmpCnt & 0xFFFF0000) | c2;
58         break;
59
60     case 4:
61         c1 += 1;
62         cnt[tid] = (tmpCnt & 0x0000FFFF) | (c1 << 16);
63         val[tid] = (tmpVal & 0x0000FFFF) | (v1 << 16);
64         break;
65
66     case 0:
67         c2 += 1;
68         cnt[tid] = (tmpCnt & 0xFFFF0000) | c2;
69         val[tid] = (tmpVal & 0xFFFF0000) | v2;
70         break;
71
72     }
73 }

```

~20x faster



```

itMajBits_kernel( (np.prod(img.shape) >> 10,), # no. of blocks
                  (2**10,), # no. of threads
                  (img, val, cnt) ) # arguments

```

Let's Look at It in Action

■ (video)

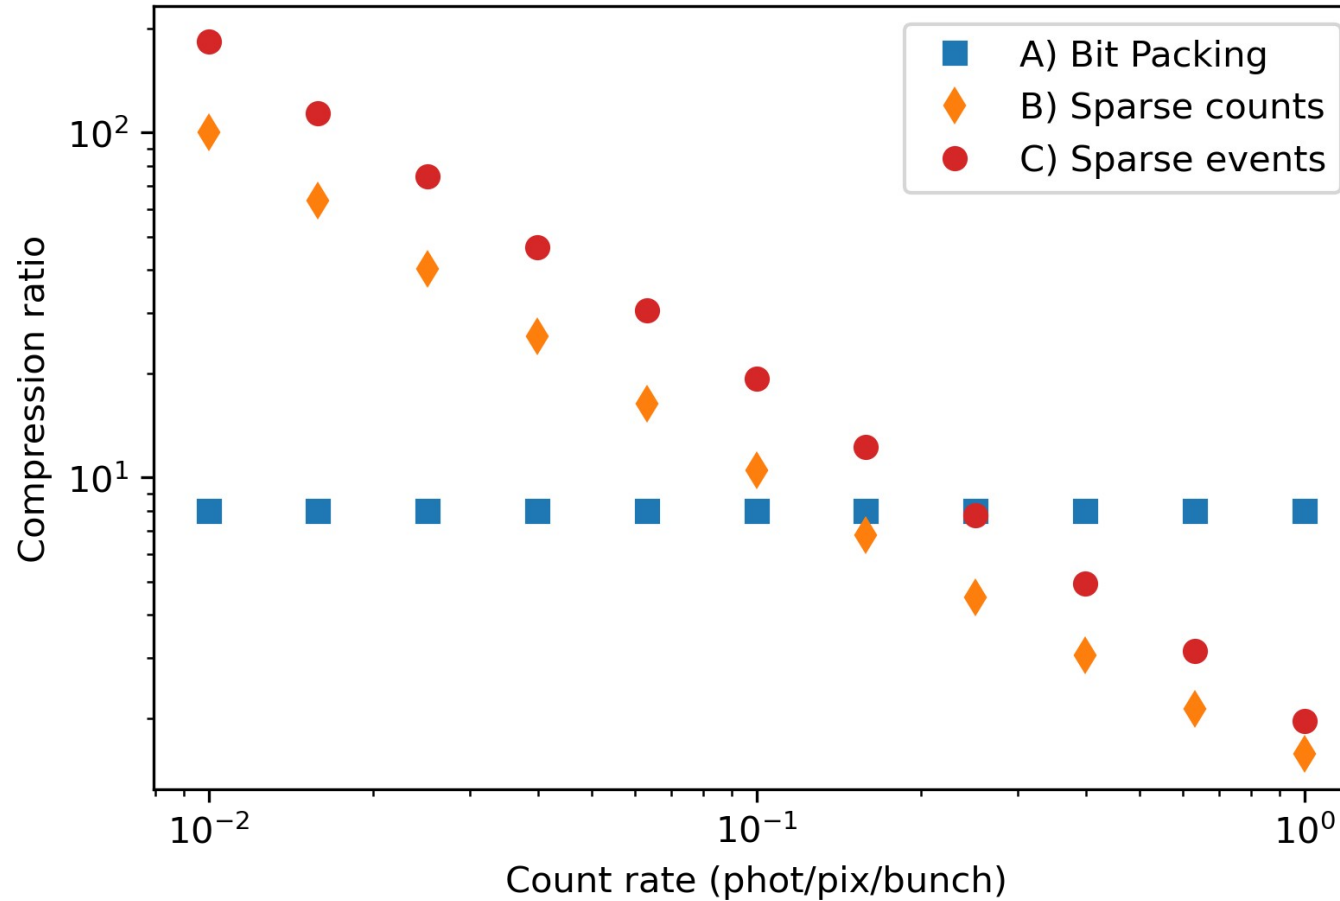
■ Overall performance: Mode and Exp-Max ~300Hz

■ Memory bandwidth → ~100 Hz

■ Reading from storage → 30 Hz max

Ad 2) Storage Requirements

AGIPD Data Compression



- Let's skip algorithmic compression
- Raw data: 2x 16 bit / pix (intensity + gain)
- Proc data: 32 bit float
- Strategy A: Bitpack into 8 x 4 bit (ratio 8)
- Strategy B: Sparse pixels (index + count)
- Strategy C: Sparse events (16bit index only)**
- B and C can be implemented efficiently through **prefix sums**
- B and C offer savings in sparse algorithms

Ad 3) Online Analysis

Results for Silica Particles

■ Proper analysis: Whole AGIPD geometry, define circles of like q



■ Calculate Two-Time Correlation Function from Outer Product

$$C_{i,j} = A_i * A_j$$

■ Super efficient with `einsum`, even better with cuTensor

■ BUT also extremely fast on multi-core CPU (OpenMP)

■ (video)



Summary

With the prototype of a streaming data-processing pipeline, three issues were addressed

- Shifting and flickering baselines in the AGIPD calibration
- Suppression of zeros
- On-line calculation of correlation
- 10 Hz (live) performance is easily within reach
- Depending on features >100Hz operation possible

GPUs can do a wonderful job for us!