Exascale Computing for Radio Astronomy: Mapping Pulsar Search on DataFlow

Kees van Berkel and Frank Boerman

MPSoC 2019, July 8-12

Hakone, Kanagawa, Japan

Pulsar

- A highly magnetized rotating neutron star (or white dwarf).
- Lots of open research questions.
- Radiation beam may hit earth once per rotation (lighthouse).
- Pulsars can be studied on earth by analyzing the beam.

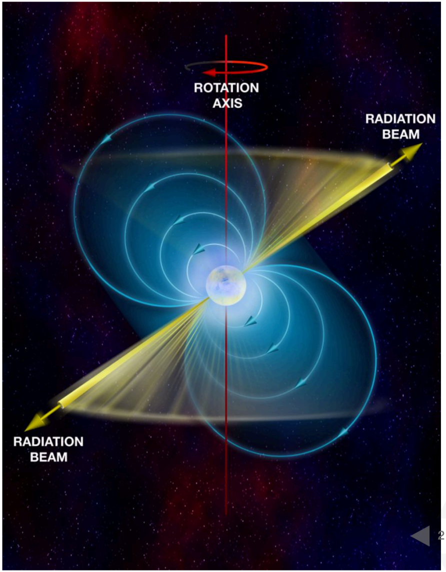

But first we need to pulsar search!

Pulsar population

Pulsar rotation period $P(s)$:

- milli second (!) vs second.

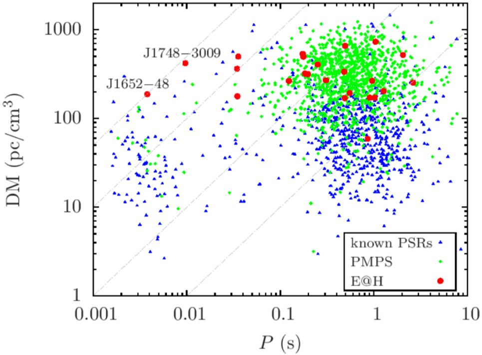

Dispersion:

- The interstellar medium, contains ionized gas.
- Group velocity $v_g = \mu(f)c$: longer waves propagate slower than shorter waves.
- After many light years, a pulse smears, typically >> P. In the gallactic plane: lots of dispersion (the green population).
- The dispersed pulse is deeply burried in noise.

Pulsar dispersion

B1256-60, $P = 128ms$:

- x-axis = pulsar phase $[0..1]$.
- In channel 95:

  pulse arrives at *phase* $= 0.12$.
- In channel 91: *phase* $= 0$
- In channel 59: *phase* $= 0.12$,

  of the *next* pulse.

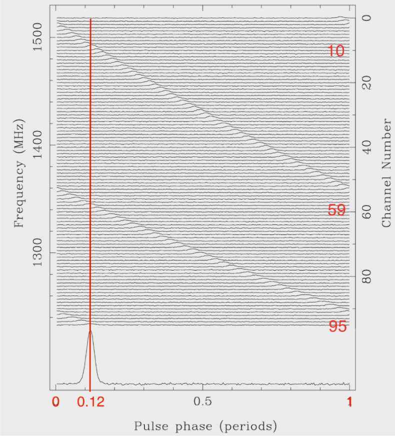

Pulse is smeared over many P.

$\Delta t = (f_1^{-2} - f_2^{-2}) \times DM \times 4.15 \times 10^6$ ms,
with f [MHz] and *Dispersion Measure* DM = 295 $cm^{-3}pc$.

Pulsar dispersion

Dispersion can be described as a *phase-only* filter [Lorimer, 2005]

$$V(f_0 + f) = V_{int}(f_0 + f) \times H(f_0 + f) \ ,$$

where $V(f)$ and $V_{int}(f)$ are the observed and emited signals around a center f_0 within a Δf, and the filter transfer function $H(f)$

$$H(f_0 + f) = \exp\left[+i \frac{2\pi D f^2}{(f+f_0)f_0^2} DM\right] \ ,$$

where:

- D is a dispersion constant, related to the plasma frequency.
- DM, Dispersion Measure, is the integrated column density of free electrons between an observer and a pulsar.

Pulsar search ≈ dedispersion

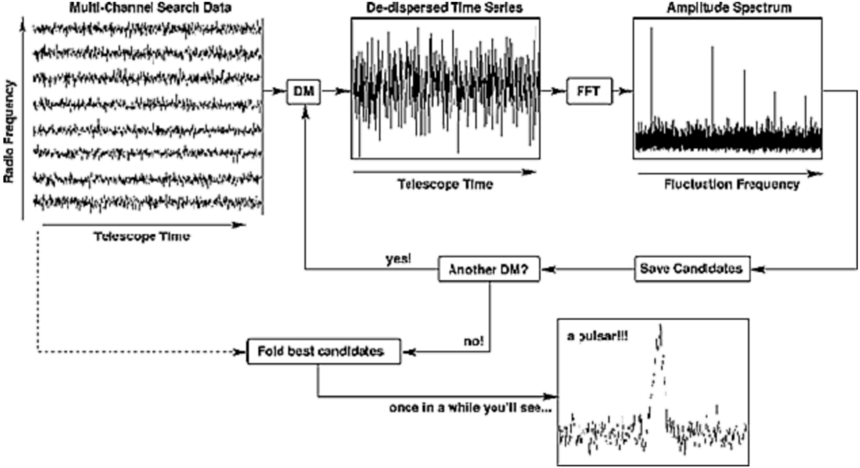

Todays trend is towards more advanced algorithms.

Pulsar de-dispersion

Coherent de-dispersion now is simple *in principle*:

$$V_{int}(f_0 + f) = V(f_0 + f) \times H^{-1}(f_0 + f)$$

The problem is that we do not know H.

So we try many H^{-1} (many DM). For SKA1-Mid, 2023 [Levin'17]:

- 6000 trials for DM
- $\times$ 16kHz baseband sample rate (Re, Im).
- $\times$ 4 polarizations $\times$ 4096 channels
- $\times$ 1500 simultaneous beams $\times$ 24/7

The effective de-dispersion sample rate = 2.4×10^{15} Hz .

Dedispersion: dataflow graph

```
G (Graph)        : build    : no errors
```

Out[17]:

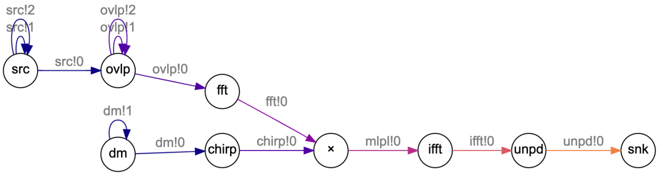

Coherent dedispersion in frequency domain, "overlap-save":

- `ovlp` makes blocks of $N = 2^{16}$ samples, overlap of $M = 2^{13}$.
- `chirp` produces transfer function $H(DM)$ in freq. domain.
- $\times$ in frequency domain = convolution in time domain.
- `unpd` removes overlap & tests if candidate pulsar.

Dedispersion: simulation

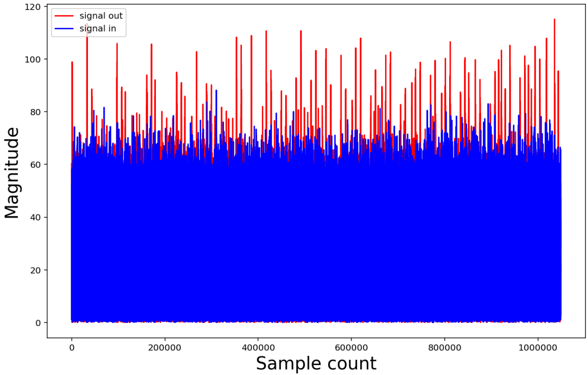

For $DM = 100$ only: pulsar! (1 channel: only very luminous ones.)

Strong aliasing effect as pulsar period $\approx$ FFT size N. Taper needed.

Dedispersion: $5 \times DM$ in parallel

```
        G2 (Graph)       : build    : no errors
```

Out[27]:

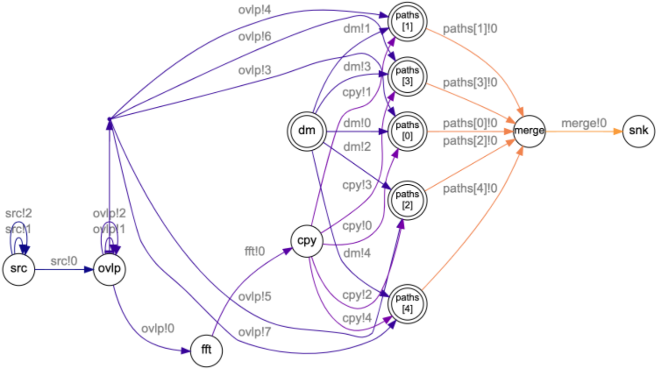

Work in progress

Dedispersion: 5 DM values in ||

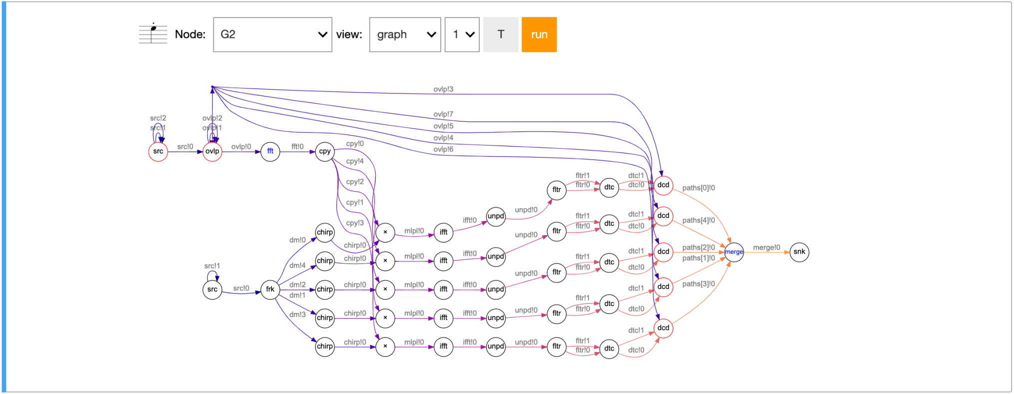

Input stream $1\times$ FFT-ed, and $5\times$ IFFT-ed with different DM values.

Only candidate pulses are merged towards output (= simplification).

Dedispersion $5\times$: simulation, $DM = 100$

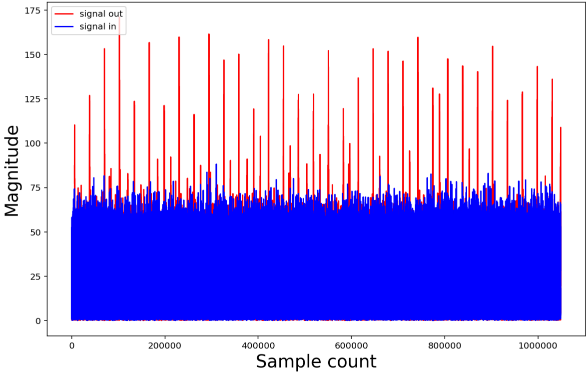

Only for $DM = 100$ output passed: pulsar detected!

Other dedispersed blocks: no pulsar candidate detected, no output.

Dedispersion $5\times$: simulation, $DM = 100$

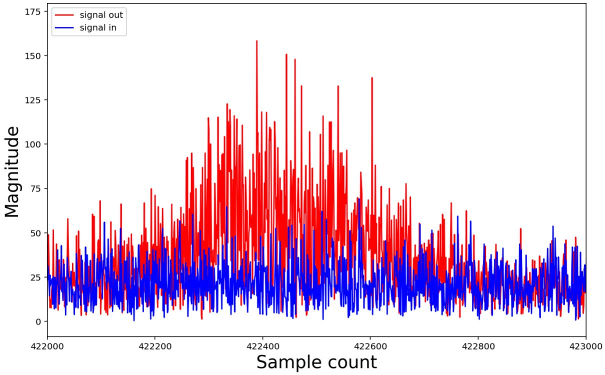

Zoomed-in: a dedispersed pulse, rising above the noise.

Pulsar search by LOFAR**

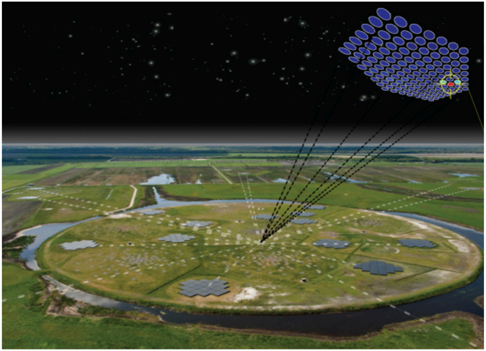

Dedispersion on DRAGNET [Bassa et al, 2016]:

- 23 nodes × 4 Titan X GPU,

1 Titan X: 6.14 TFlops,

- dedispersion: 0.24 TFlops,
- only 4% of max throughput.
- ... ?

Rooflines!

**Input sample stream + algorithm parameters for our simulations provided by LOFAR.

Operational Intensity

Arithmetic intensity I_A = amount of compute / unit problem size:

$$I_A = \frac{\text{number\_of\_operations}}{\text{size\_of (input + output) [bytes]}}$$

Operational intensity I_O = amount of compute / unit DRAM traffic:

$$I_O = \frac{\text{number\_of\_operations}}{\text{amount\_of\_DRAM\_traffic (input + output) [bytes]}}$$

$I_O = I_A$ only if the entire problem fits in on-chip memory.

In practice $I_O << I_A$.

I_O depends on algorithm choices and on available on-chip memory.

FFT on a GPU

$I_O \approx 1$, because [Govindaraju 2008, MPSoC'2016]:

- per scalar core 10s of threads needed

  to hide register read-after-write latencies.
- per thread, up to 128 registers

  $\Rightarrow$ most on-chip memory spent on registers, almost all idle.
- for 100s of scalar cores, many 1000s of threads needed

  $\Rightarrow$ 1 thread / radix-8 butterfly.
- full FFT block write+read every 3 FFT stages, out of $\log_2(N)$.

  $\Rightarrow I_O = \frac{I_A}{\log_8(N)} = \frac{5/9 \times \log_2(N)}{\log_8(N)} \approx 2$.

FFT on a FPGA

Pipelined FFT data read once, write once ⇒ $I_O = I_A$.

Furthermore, for dedispersion, I_A >>, e.g. by 64 DM paths in ||,

or by unfolding** the IFFT 64x:

FFT 1x 64x, parallel 64-unfolded

FFT	1x	64x, parallel	64-unfolded**
throughput	1	64	64
costs (dsp)	1	64	64
costs (mem)	1	64	1
latency	1	1	1/64

**For $N = 2^{16}$-point FFT on Xilinx V6 [Garrido 2014]:
64 complex inputs each clock cycle (at f_{clock} of 115 MHz):

Dedispersion rooflines

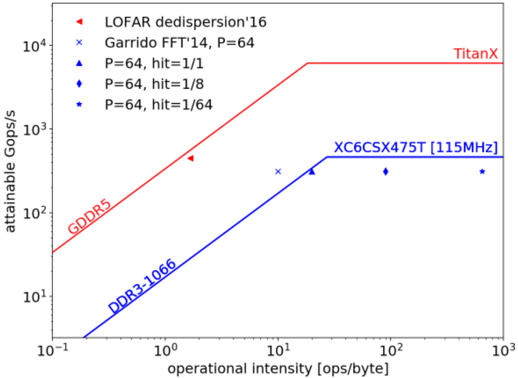

I_O for dataflow dedispersion depends on pulsar-candidate hit-rates.

Garrido paper did not discuss off-chip I/O.

Mapping Pulsar Search on DataFlow

Lifting a highly dispersed pulsar signal above the noise:

- is extremely computationally intensive (exascale for SKA);
- has high arithmetic intensity and lots of data parallelism.

The dataflow programming model + analysis + transformations:

- supports quantitative exploration of various forms of parallelism.

FPGA relative to GPU offers, assuming comparable rooflines:

- superior operational intensity, 10-100 $\times$, and hence
- $> 10\times$ pulsars/year and $> 10\times$ pulsars/MWyear.

Next: Mapping Pulsar Search on DataFlow on FPGA.

References

- Lorimer and Kramer, Handbook of Pulsar Astronomy, Cambridge University Press, 2005.
- Levin et al, *Pulsar Searches with the SKA*, Proceedings IAU Symposium No. 337, 2017.
- Bassa et al, *Enabling pulsar and fast transient searches using coherent dedispersion*, Astr. and Comp., Vol 18. pp 40-46, 2017.
- Govindaraju et al, High Performance Discrete Fourier Transforms on Graphics Processors, Proc. of the 2008 ACM/IEEE conference on Supercomputing, article #2.
- Garrido et al, Challenging the Limits of FFT Performance on FPGAs, Proc. ISIC 2014, pp 172-175.