Non-imaging Processing

Aris Karastergiou, Ben Stappers, Michael Kramer

Outline

- What is non-imaging processing?
- Where are the problems?
- Connecting the science to the engineering
- What are the processing challenges?
- What is the present state-of-the art?
- What are the potential solutions?
- What tradeoffs need to still be considered?
- A tool for specifications

What is non-imaging processing

- A working definition is anything that processes data that comes out of the central beamformer.
- It is therefore predominantly associated with the processing of data in the time domain.
- The time resolution required typically far exceeds what the correlator can deliver.
- In general it relies on more than one beam being available at one time.

Defining the problem.

- One of the key aims so far has been to better describe the requirements.
- The parameter space is so large that it is hard to choose an optimal set.
- So far this has predominantly been from the point of view of pulsar and fast transient processing.
- Strongly influenced also by design choices, in particular:
  - Types and numbers of receptors
  - Size of the core

Science Goals: Gravitational Waves

- Required: additional MSPs widely distributed on the sky
- Implies: all-sky survey for MSPs
- Frequency: 400 MHz AAs and/or 800 MHz dishes, balance spectra and propagation effects
- Maximum DM trials: Depends on if in plane or not.
- Survey speed: but cannot trade (A/T) for FoV

- Required: Higher precision timing of known & future MSPs
- Implies: pointed observations toward "timers"
- Sensitivity: Maximum possible
- Field of View: Larger + Beams gives greater efficiency.
- Frequency: 0.8-3 GHz

Science Goals: Tests of GR

- Required: additional ultra-relativistic binaries
- Implies: low- and moderate Galactic latitude survey for binaries
- Survey speed: but cannot trade (A/T) for FoV
- Frequency: 1.4 GHz (and ~ 2 GHz?)
- maximum DM trials: Depends on BW/in Gal. plane.

- Required: higher precision timing of known and future ultra-relativistic binaries
- Implies: pointed observations toward "timers"
- Sensitivity: Maximum possible required
- Frequency range: 1.4 GHz and 2 GHz

Begin with use cases

- Use cases provide one way to narrow down specifications and this is what we have done.
- By considering specific scenarios we have been able to describe an end to end description of the requirements.
- This has been done for both searching with the sparse AAs and pulsar timing.

Processing Challenges

- Beamforming – sparse Array
- Dedispersion
  - Incoherent for pulsar surveys
  - Coherent for high precision timing
- Folding (only for known pulsars)

Processing Challenges

- Beamforming – sparse Array
- Dedispersion
  - Incoherent for pulsar surveys
  - Coherent for high precision timing
- Searching: i.e. Fourier Transforming
- Acceleration processing
- Candidate identification

Pulsar Search Implications

- Beams - > few thousand
  – Depends on A/T, survey time, core size
- Dedispersion -> few thousand
  – Depends on frequency / Area of sky
- Searching: i.e. Fourier Transforming
  – Requires accumulation step of data!
- Acceleration processing – few hundred
  – Frequency domain vs time domain
- Candidate identification
  – AI, Neural Nets or other pattern recognition.

State of the Art Searching

- Parkes Multibeam
  - 13 beam analog until 2009
  - 13 beam digital since 2009
- PALFA – Arecibo
  - 7 beam digital (wide-band) since 2010(ish)
- WSRT interferometer
  - Up to a few thousand beams (formed offline)
- LOFAR
  - Few hundred beams simultaneous
- In all cases no real time data reduction at full resolution; Real-time systems only developed for single-pulse searches, i.e. no folding
- So far apart from channelisation, all done in CPU/GPU processors.

Pulsar Timing Implications

- Simultaneously process 10-100 beams
- Directly related to number of pulsars found (20,000/2000), FoV and A/T
- Estimates of few days to get through all sources
- Need different approach for high precision and "normal".
- Coherent dedispersion needed for high precision
- BW ~ 1GHz, Tres ~ 0.2 – 100 μs, Tint > 120s

State of the Art Timing

All coherent dedispersion based

- CASPSR
  - ~ 1 GHz bandwidth/Direct sampled/FFT/GPU
- GUPPI
  - ~ 1 GHz bandwidth/PFF/GPU
- ROACH – Jodrell & Effelsberg
  - 512 MHz BW/PPF/CPUs -- x2 to get 1 GHz.

- All capable of real time but no multi-beaming.

Potential hardware solutions

- Timing requirements, even though include computationally expensive CoDeDi, are small perturbation on searching, as will be performed on many less beams.
- However, have specific demands on nature of data.
- More of the same
  - Problems with the amount of data to store (linked to time for survey/beams etc..)
- New generation computers, e.g. BG
  - Problems with accumulating data
- Fastest possible processing is desirable so need a streaming solution.
- GPUs, FPGAs, ASICS, Hybrids

Software: Pulsar Timing

- Data acquisition, real-time coherent dedispersion and folding codes mature
- Many groups using same software suites.
- Optimised and running on CPUs and GPUs
- Development work ongoing for CoDeDi and folding to also be done in fpga-like hardware.
- Porting to a multi-beam (10s) environment will not be a show-stopper.

Software: Searching

- Community converged to 1 or 2 codes
- Almost certainly need to be reconfigured in a large multi-beam environment
- Not completely optimised for real time use, with exception of dedispersion (Artemis, Swinburne)
- Only partially ported to GPU-like architectures so far.
- Acceleration searching is the big issue as dominates the required resources.

Data Products: Timing

- Main data products will be calibrated and accurately time-stamped folded pulse profiles.
- Of course these need to come with all appropriate meta-data
- The exact data volume will depend on the required time/phase/freq resolution required for the science goals.
- Reduction in data size is usually more than three orders of magnitude over raw.

Data Products: Searching

- Will depend strongly on whether done in real time or not.
- If real time, may only be so-called candidate plots of a few megabytes in size
- If not real-time, intermediate data products will be very large, and of the order of petabyte storage capacity will be required.
- Likely to be somewhere in between with partial data products stored.
- Further investigation is required in coming year.

Real-time processing: a framework

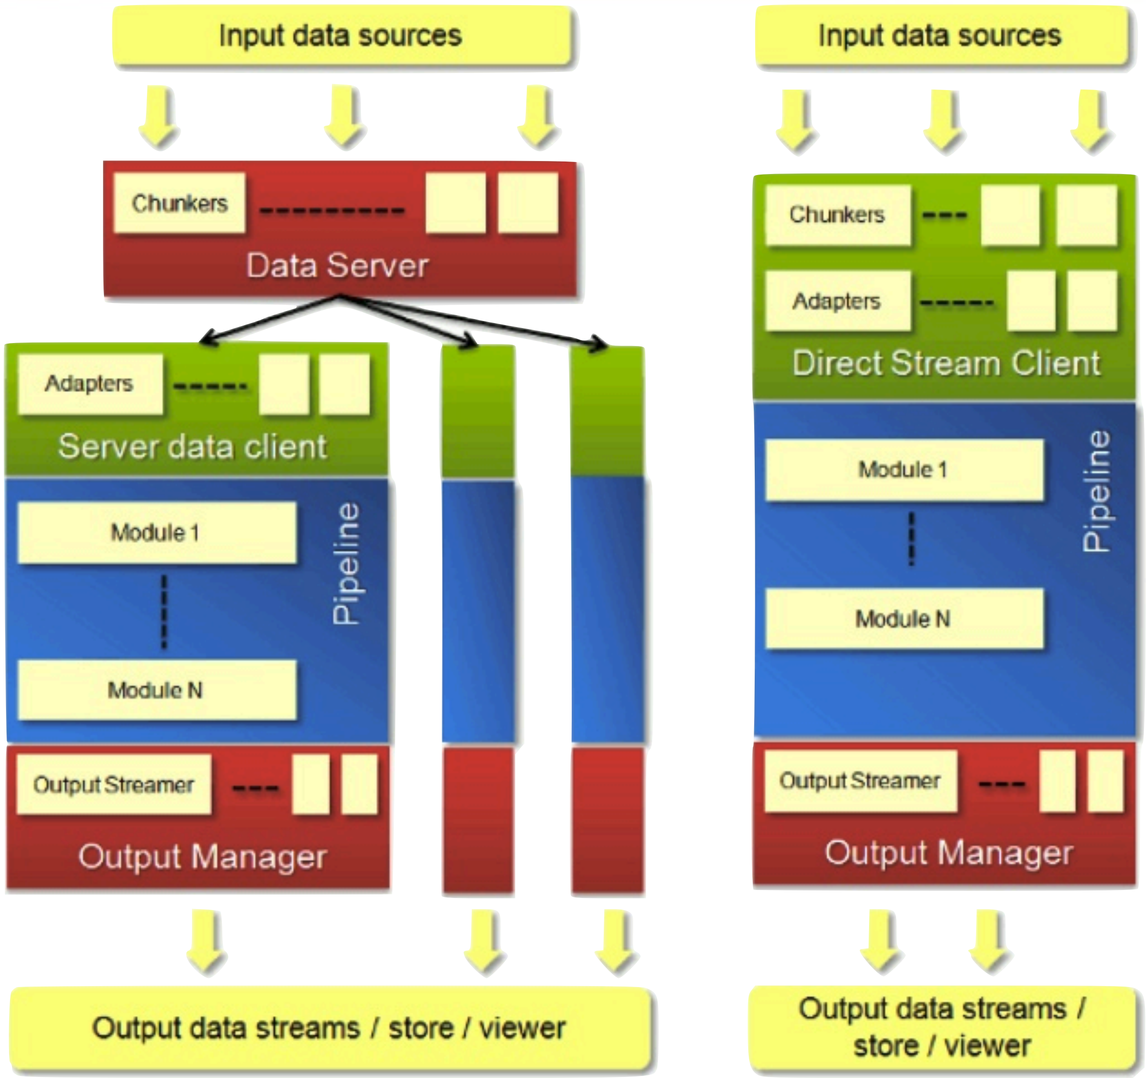

PELICAN

The Pelican framework is designed to allow flexible deployment of hardware resources to cope with data rates and data processing requirements of an incoming data stream. Pelican maintains three layers with well defined interfaces between them, namely data acquisition, data processing, and data output.

- Scalability
- Flexibility
- Minimum overhead

Real-time processing: an example

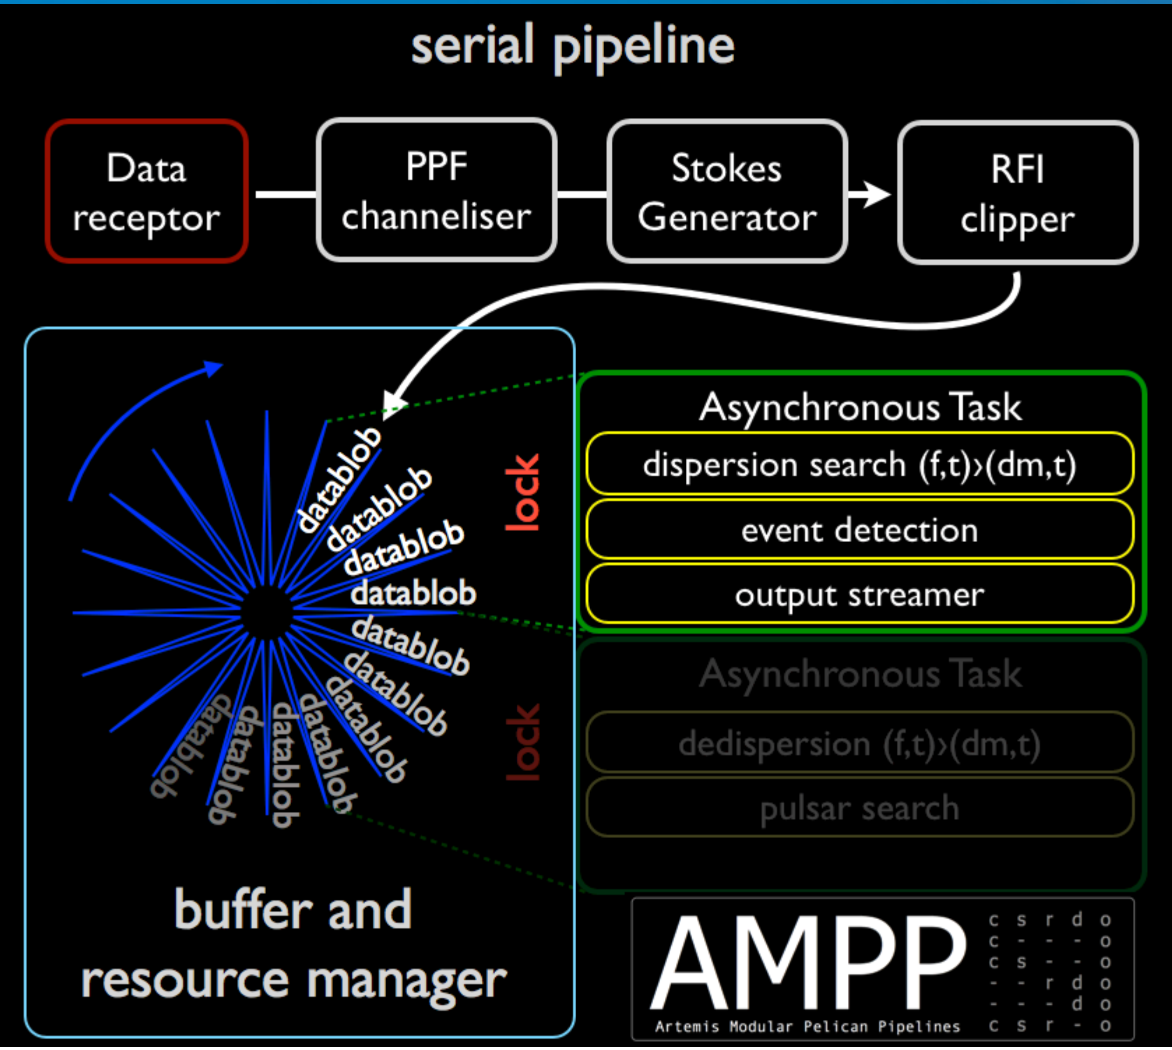

ARTEMIS

Real-time processing for individual radio pulse searches has been developed within the PELICAN framework and successfully deployed on international LOFAR stations for testing. Results for cost, power and performance have been presented in the Central Signal Processing CoDR.

Real-time processing: multi-core

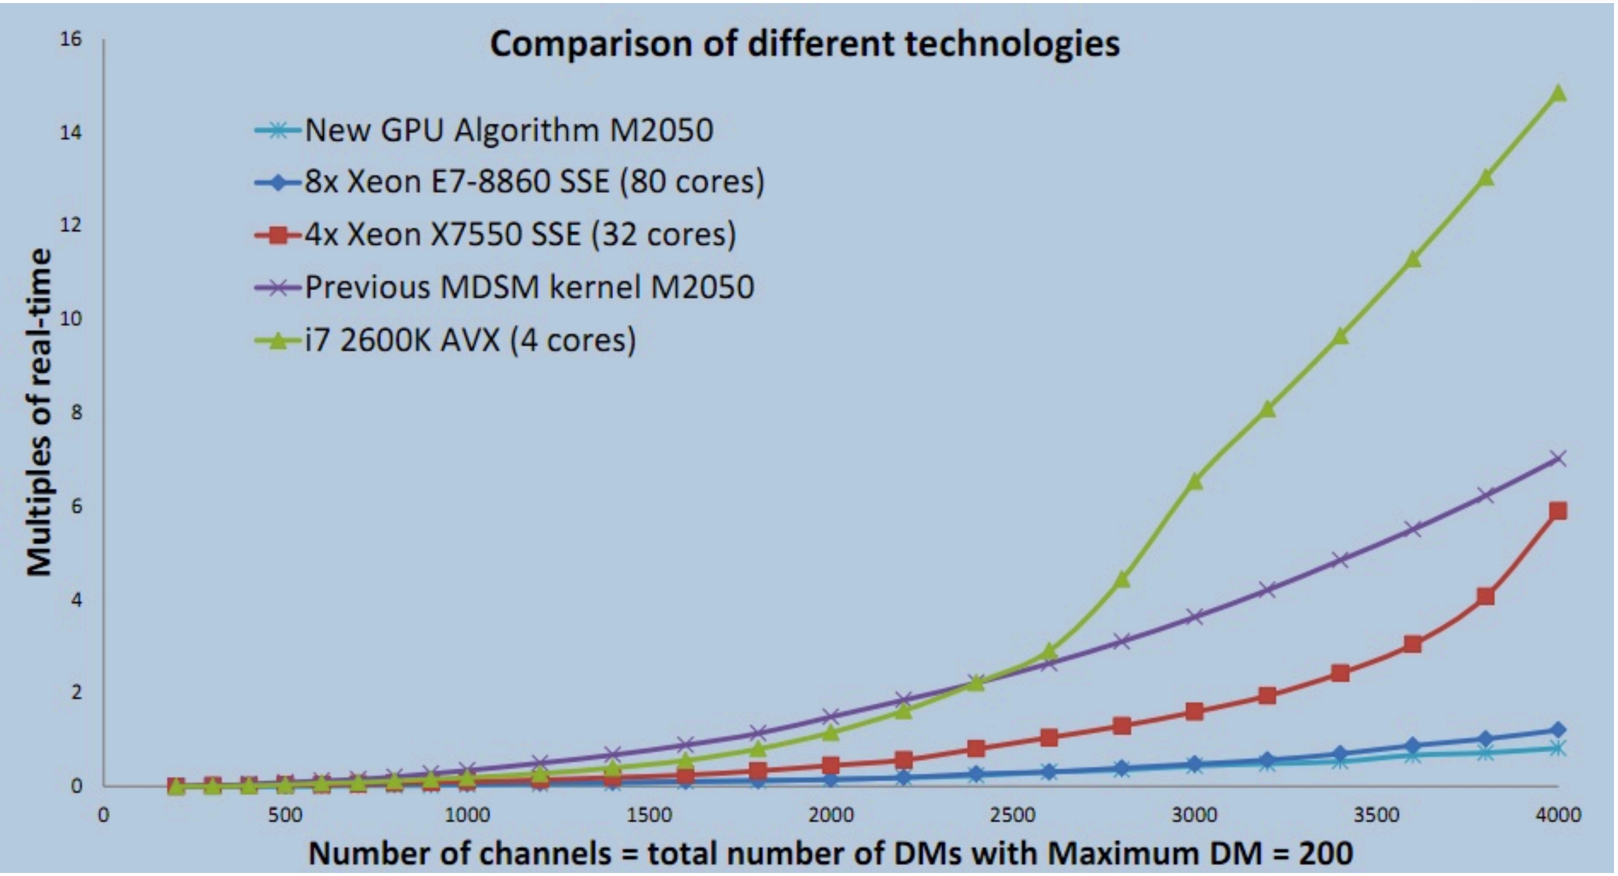

ARTEMIS

We are developing optimised code for CPU and GPU architectures and benchmarking the solutions against each other

Things still to consider: Tradeoffs

- Need to understand complex inter-dependencies:
  - beams vs bandwidth vs integration time vs acceleration vs gain
  - number of bits, many places in processing chain.
  - channel widths vs time resolution
  - timing accuracy
  - polarisation purity

A tool for defining the problem

- Complex systems which are highly interdependent.
- Software to compute dependencies between:
  - Dish/AA size, number of elements
  - Size of core, size of stations
  - Observing frequency / BW
  - Dispersion measures
  - Data rates / processing
- Work in progress: linkage between all, processing architectures.