An FPGA-Based Back End for Real Time, Multi-Beam Transient Searches Over a Wide Dispersion Measure Range

Larry D'Addario[1], Nathan Clarke[2],
Robert Navarro[1], and Joseph Trinh[1]
[1]Jet Propulsion Laboratory, California Institute of Technology
[2]Curtin University, Perth, Australia

Copyright 2013 California Institute of Technology

v2.7 20130110

Introduction

- Objective: Survey the radio sky for short (<1 s) transient pulses.
- Requires a sensitive, wide-field-of-view telescope
  - Targeting ASKAP with phased array feeds
  - 36x 12m dishes, each with up to 36 beams, 30 deg^2 FoV at 0.7-1.8 GHz
- Requires removal of interstellar dispersion for many trial dispersion measures.
- Performance achieved in our implementation
  - Real-time incoherent de-dispersion for 442 DMs and 36 beams
  - Real-time automatic transient detection by examining all 14,400 de-dispersed time series
  - Integrating time (time resolution) < 1 ms
  - Detection latency < 35 ms (allows capture of raw voltage samples near each tentative detection for off-line analysis)
  - To our knowledge, this is the highest-performance transient search engine implemented to date.
- Context
  - CRAFT collaboration's proposal for a commensal transient survey was one of 10 science programs selected for ASKAP.
  - SKA: Our investigation of automated systems for detecting fast transients should inform design of SKA's high-time-resolution features.

Design Considerations

- ASKAP will have
  - up to 36 dual-polarization beams per antenna, producing ~30 deg^2 FoV
  - 304 MHz instantaneous bandwidth in the range 700-1800 MHz; 1 MHz channels.
- Antenna combining options
  - Full cross correlation and imaging each beam's FoV: too slow. ASKAP correlator's minimum integration time is 5 s.
  - Coherent beamforming – form multiple array beams from each set of corresponding antenna beams: high sensitivity and fine time resolution, but small total FoV since only a few narrow array beams can be formed.
  - Incoherent combining – sum power spectra across antennas: N/sqrt(N) = 6 times less sensitivity than coherent beamforming, but preserves entire FoV. For ASKAP, this provides the highest survey speed.
- De-dispersion
  - Incoherent de-dispersion is the only option
  - To cover DM = 10 to 3000 pc/cm^3 near 1 GHz, ~400 DMs must be searched
- Raw sample buffer capture: preserve voltage samples near tentative detection events to allow more detailed analysis
  - Requires real-time event recognition and short latency
- Time resolution: desire best possible; < 1 ms achieved.

CRAFT Back-End Implementation (Tardis-ASKAP)

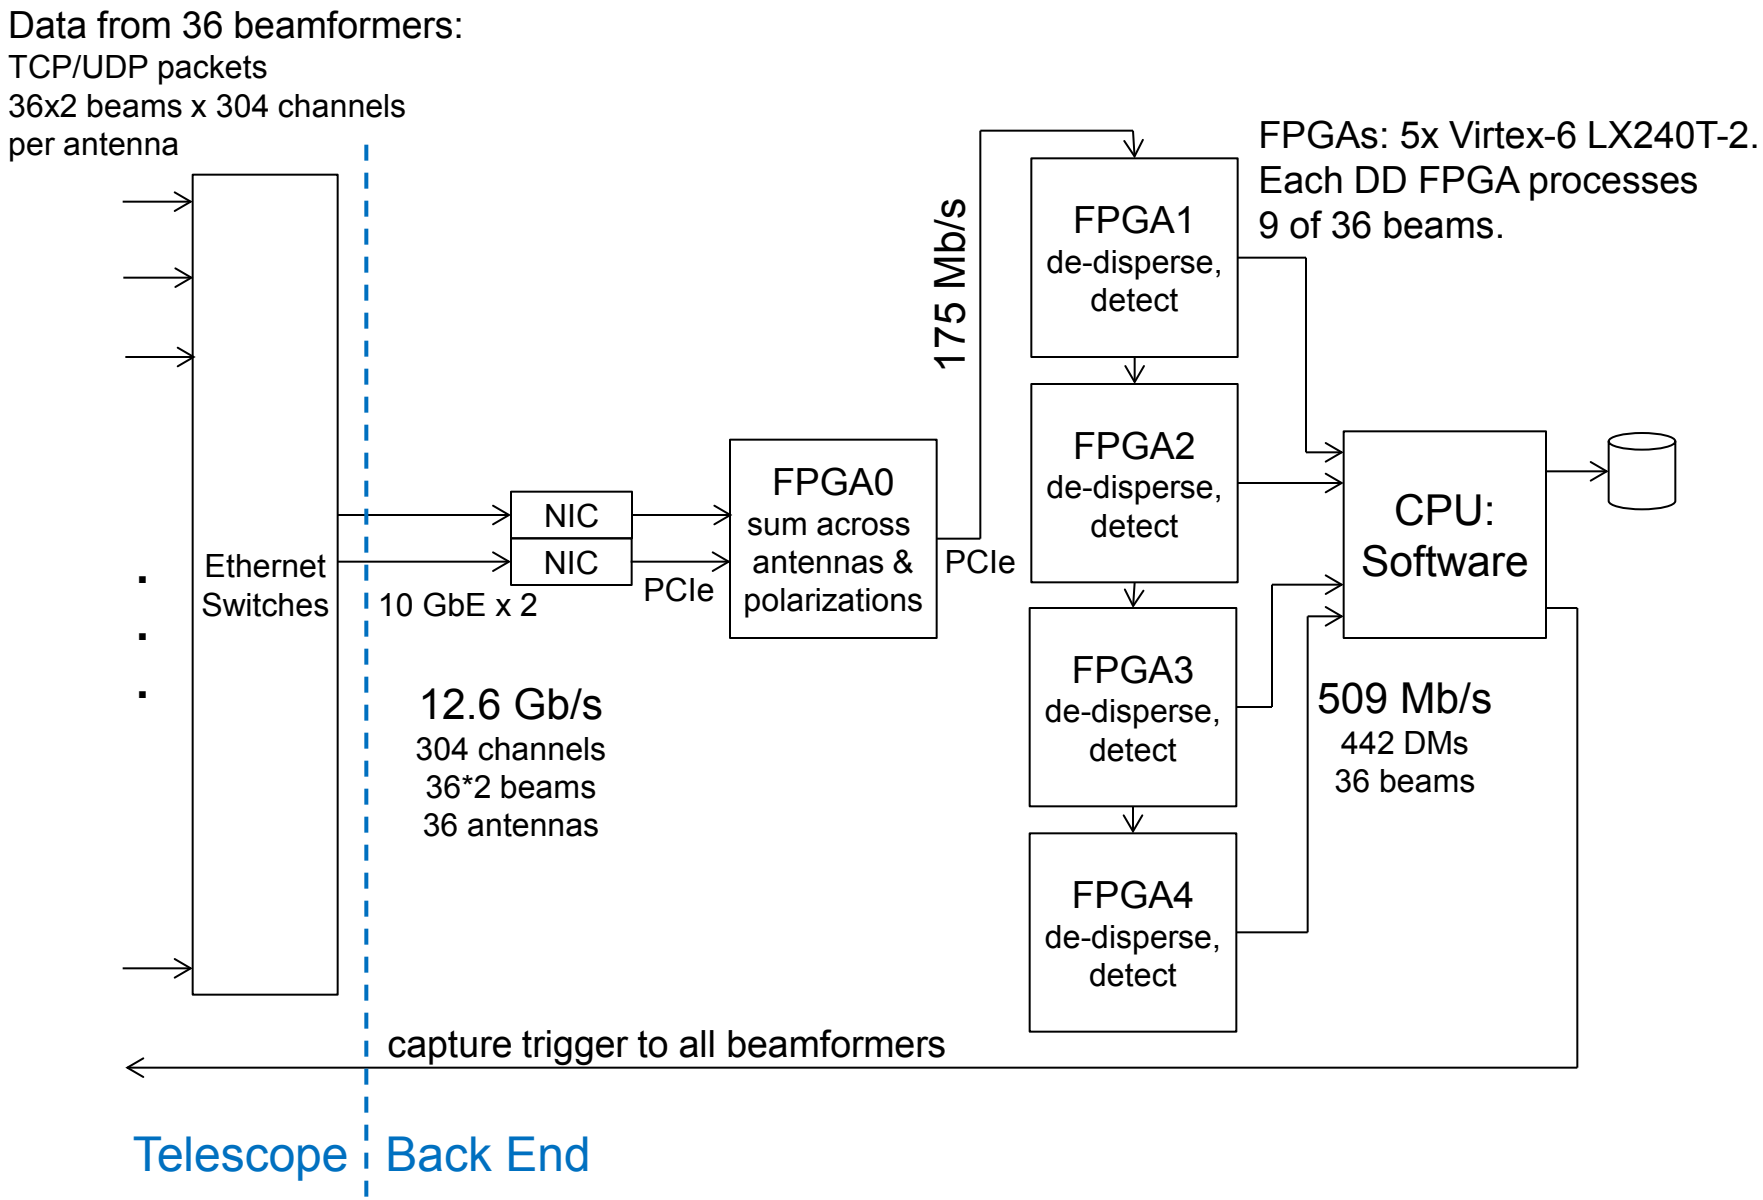

Hardware

- COTS PCIe card with plug in FPGA modules
  - Pico Computing EX-500 backplane and up to 6 M-501 FPGA modules.
  - Each module has one Virtex-6 LX240T-2 and 512 MB DDR3 memory.
  - One FPGA module for cross-antenna summing, four for de-dispersion.
  - Backplane has x16 Gen2 PCIe to host, 8GB/s bandwidth
- Data from beamformers received over dual 10GbE transceivers.

- Also Considered:
  - Casper Roach Board
  - ASKAP Redback Board.

EX-500 backplane with two M-501 FPGA modules installed.

Single-Dish Version (Tardis-SD)

- Large parts of the CRAFT back end were complete and had passed laboratory tests in late 2011 (including the de-dispersion engines).
- In early 2012 it became clear that delays in the ASKAP project and a decision to re-design its electronics to take advantage of technology advancements would not allow it to accommodate the CRAFT back end for several years.
- We therefore decided to create a second version, suitable for single-dish telescopes, that could be deployed sooner. A 34m antenna at the Goldstone DSN complex ("DSS13") is the primary target.
- Each DD FPGA's processing capacity was re-allocated to obtain finer time resolution:
  - ASKAP: 304 channels,1.0 ms integrations, 442 DMs, 9 beams/FPGA
  - Single Dish: 1024 channels, 0.1 ms integrations, 512 DMs, 1 beam/FPGA
- One Pico EX-500 motherboard can accept up to 6 FPGA boards, so we can process 6 beams, which are now independent of each other.
- DSS13 provides a dual-band receiver: 2.2-2.3 GHz and 8.2-8.6 GHz. Both bands can be observed simultaneously.
- We implemented two 1024-channel spectrometers in one ROACH1 board to support the back end.

DSS13 Deployment Block Diagram

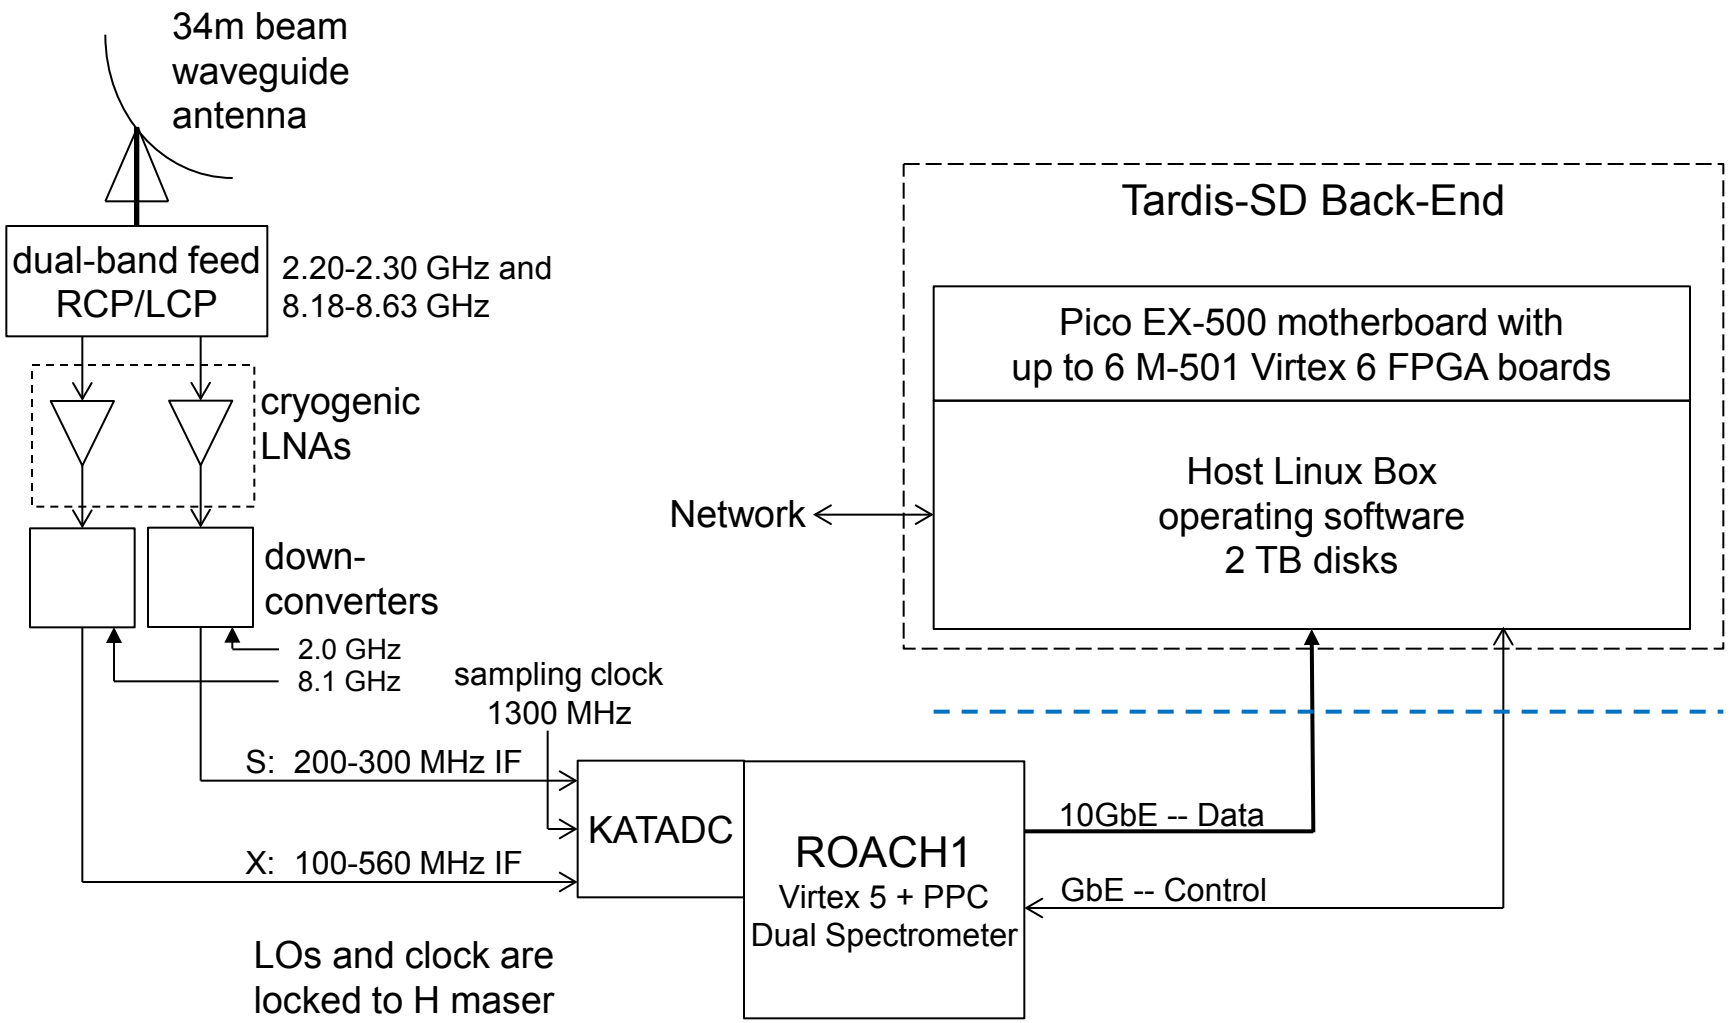

De-Dispersion Algorithm

A dynamic spectrum is an array of measurements of signal *energy* in frequency-time cells.

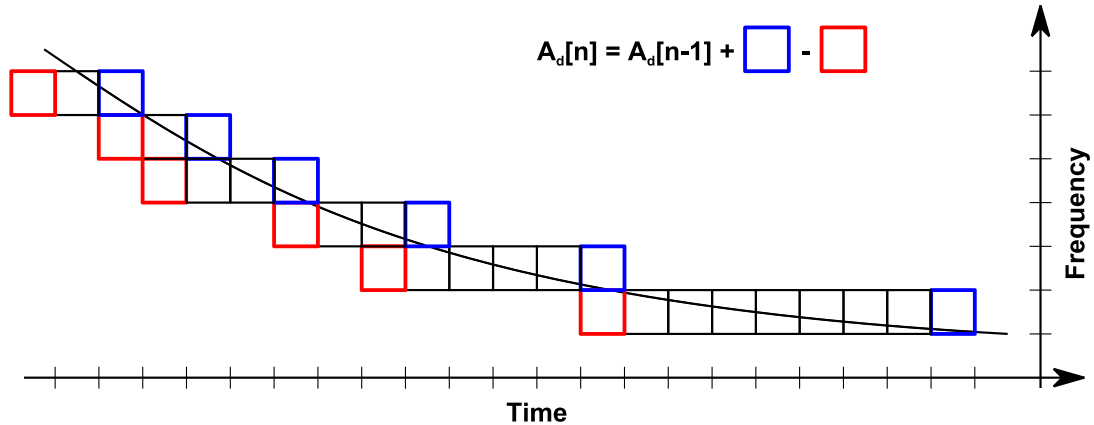

- All cells contain energy from the system noise. A dispersed pulse arriving at a given time and having a given DM deposits energy in *some* cells.
- One sample of the incoherently de-dispersed time series is obtained by summing the energies measured in cells containing a model pulse.
- Our algorithm includes a cell in the sum iff doing so increases the SNR. Unlike previous algorithms, this is closer to optimum because it
  - Does not use straight-line approximation for f-t profile (large fractional bandwidth).
  - Can include multiple samples per channel (high DMs).

  [For details see Clarke, Marquart, and Trott, *ApJS* 2013 (accepted).]
- We also distribute the trial DMs optimally across the search range so as to minimize the maximum loss of SNR at DMs between the search values (non linear, non exponential).
  - ASKAP: 700-1004 MHz, 442 DMs from 10 to 3000 pc/cm^3 → worst relative SNR of 0.746.
  - DSS13: 2.2-2.3 GHz, 512 DMs from 1 to 500 pc/cm^3 → worst relative SNR of 0.964.
- A fixed cell selection table for each trial DM is pre-computed and stored in a memory accessible to the FPGA. Each trial DM can have any value, and an arbitrary frequency-time profile can be used in place of the usual $\tau = D/f^2$.

Improving Memory Bandwidth for De-dispersion

- The dynamic spectrum is held in a large buffer for processing by the FPGA's de-dispersion engine.
- Performance (minimum integration time, number of beams, number of DMs that can be handled in real time) is limited mainly by the bandwidth for reading the buffer into the FPGA.

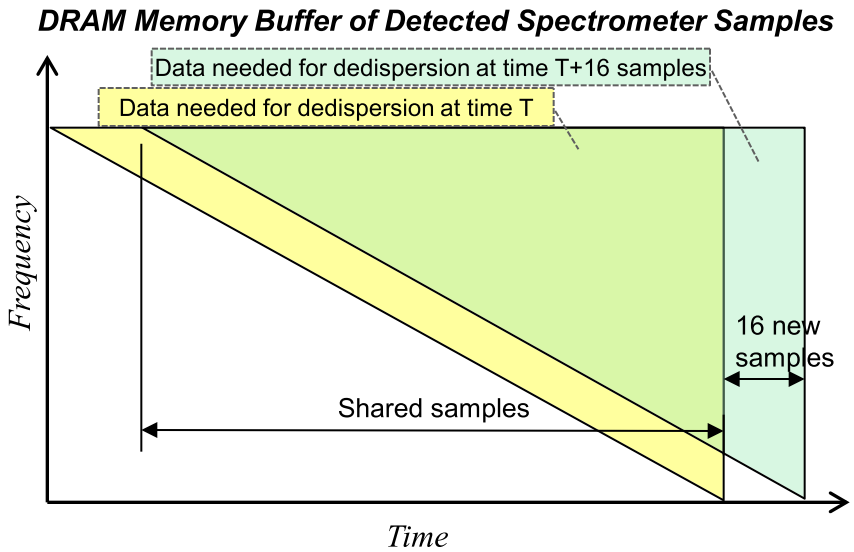

- Our implementation minimizes the required bandwidth by using the fact that successive de-dispersed time series samples use many of the same dynamic spectrum samples.
- It does this by computing de-dispersed samples in *groups,* reading all samples needed for each group together.
  - ASKAP version: group size = 16 samples (16 ms)
  - Single dish version: group size = 64 samples (6.4 ms)
- This procedure increases the latency of pulse detections in proportion to the group size, but we consider this acceptable as long as the latency is still much smaller than the length of the available capture buffer.
  - ASKAP: expecting capture buffers of ~10 s
  - DSS13: available spectrometer memory allows capture buffer of 0.77s at 8.4 GHz and 3.1 s at 2.2 GHz.

Transient Detection Algorithm

- A tentative detection is recorded for DM d whenever a threshold is exceeded by the current sample of the de-dispersed time series:

$$x_d > \mu_d + s\,\sigma_d$$

  where μ_d is the running mean and σ_d is the running standard deviation of that DM's de-dispersed time series; and s is a user-settable parameter.
- The mean and standard deviation are computed automatically using IIR filters, separately for each DM:

$$\mu_d[n] = \mu_d[n-1] + \frac{x_d[n] - \mu_d[n-1]}{M}$$

$$\sigma_d[n]^2 = \frac{(S-1)\sigma_d[n-1]^2 + (x_d[n] - \mu_d[n])(x_d[n] - \mu_d[n-1])}{S}$$

  where M and S are user-settable time constants, in samples. (M, S, and s are constant across all DMs.)
- Detection decisions are made as each group of de-dispersed samples is computed (16 samples/group in the ASKAP version and 64 in the single-dish version). A detection is flagged if any of the samples in that group exceeds the threshold.

Test Results from Goldstone 34m (1 of 2)

Pulsar J0332+5434, 2.2-2.3 GHz, 0.1 ms integrations, DM search range 1-500 pc/cm^3.
947s observation, 148,320 groups, 118 groups had TDs. Period=0.7145 s => 1325 pulses (9%).
Measured DM = 27.03±0.9 pc/cm^3; published DM = 26.83 pc/cm^3.

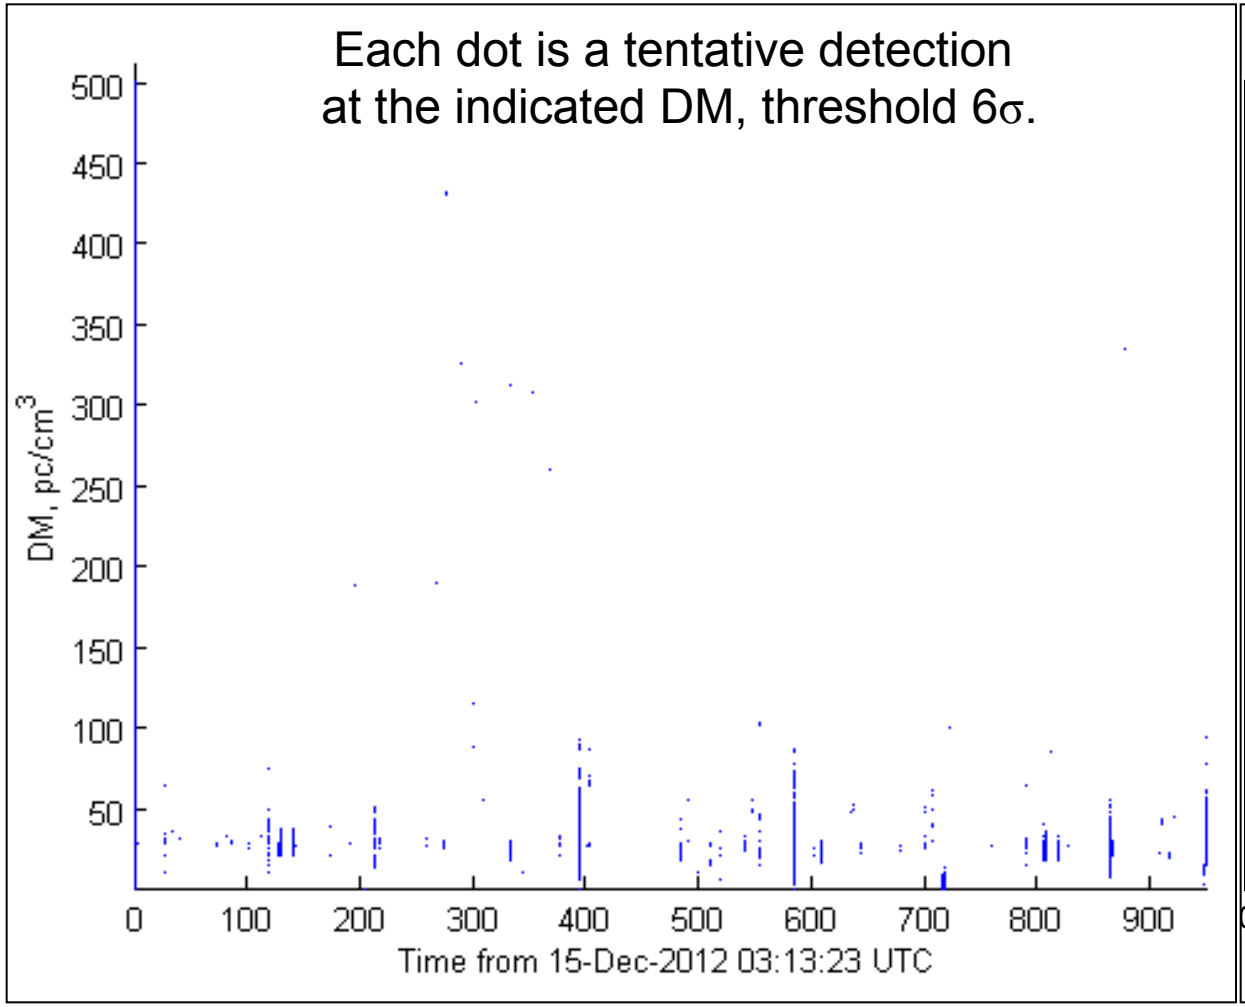

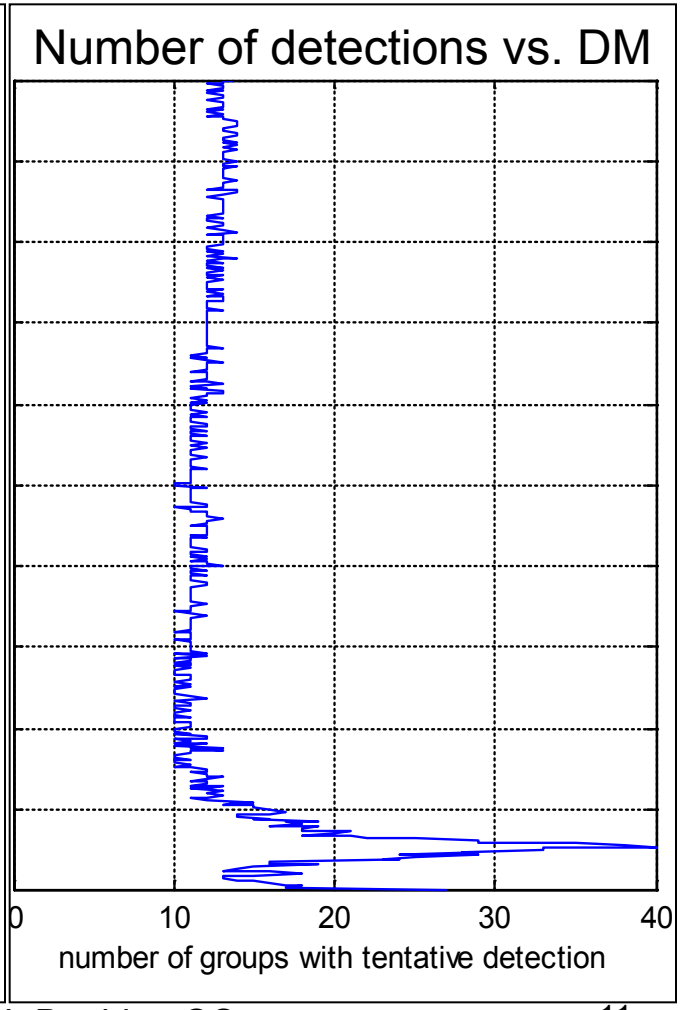

Test Results (2 of 2)

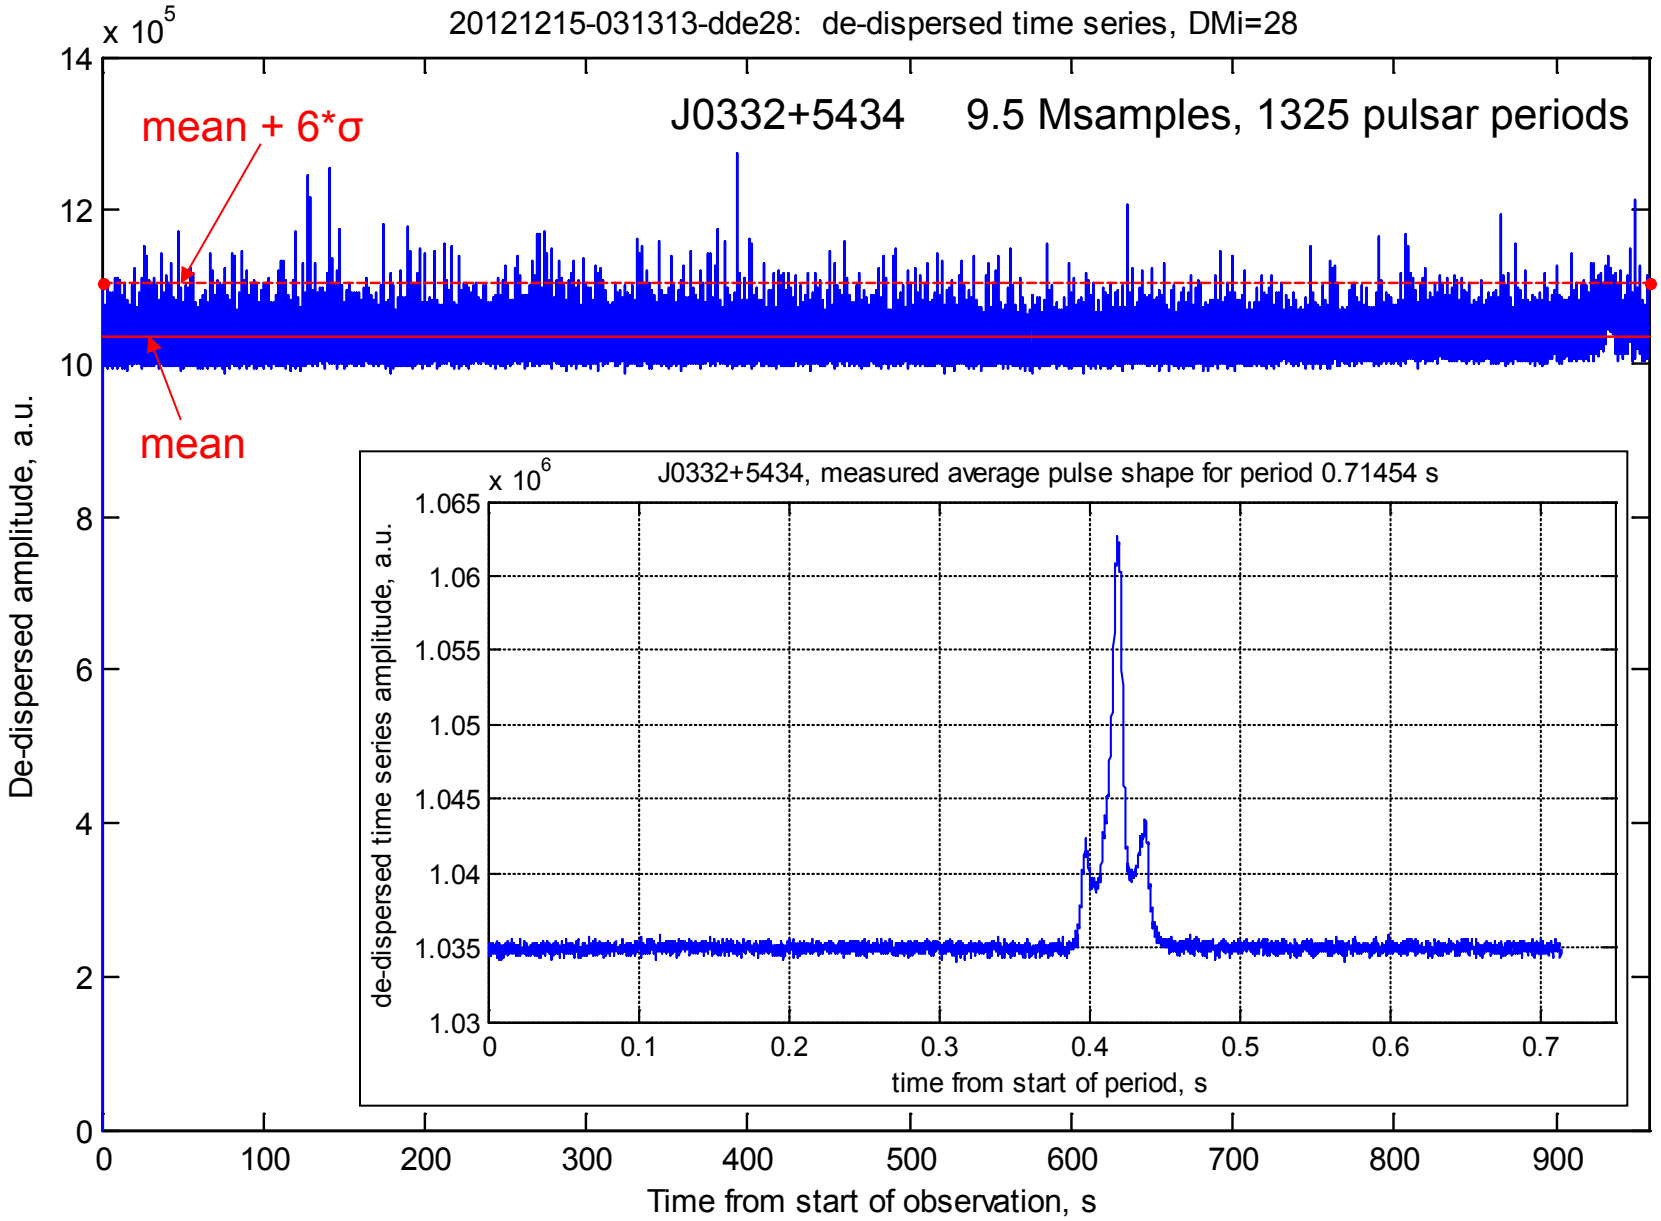

Conclusions and Further Work

- We have designed and implemented a radio telescope back end that detects transient pulses in real time after incoherent de-dispersion at ~500 DMs and for multiple beams simultaneously.
- Two versions so far: one tailored to the ASKAP telescope (36 beams) and one for single-dish telescopes (up to 6 beams).
- Detection latencies ~30 ms or less are achieved, allowing capture of raw voltage samples around the detected pulse.
- A new incoherent de-dispersion algorithm is used in order to maximize SNR over a wide range of DMs. Distribution of trial DMs across the search range is also optimized.
- Desirable further work:
  - Implement sample capture buffer in the ROACH spectrometer at DSS13
  - Implement software to discriminate against non-astrophysical tentative detections (RFI, noise)
  - Automate operation of the DSS13 installation
  - Possible deployment of Tardis-SD to other telescopes
  - Eventual deployment of Tardis-ASKAP to the ASKAP telescope

Backup Slides Follow

End of Presentation

ASKAP Processing – Simplified View

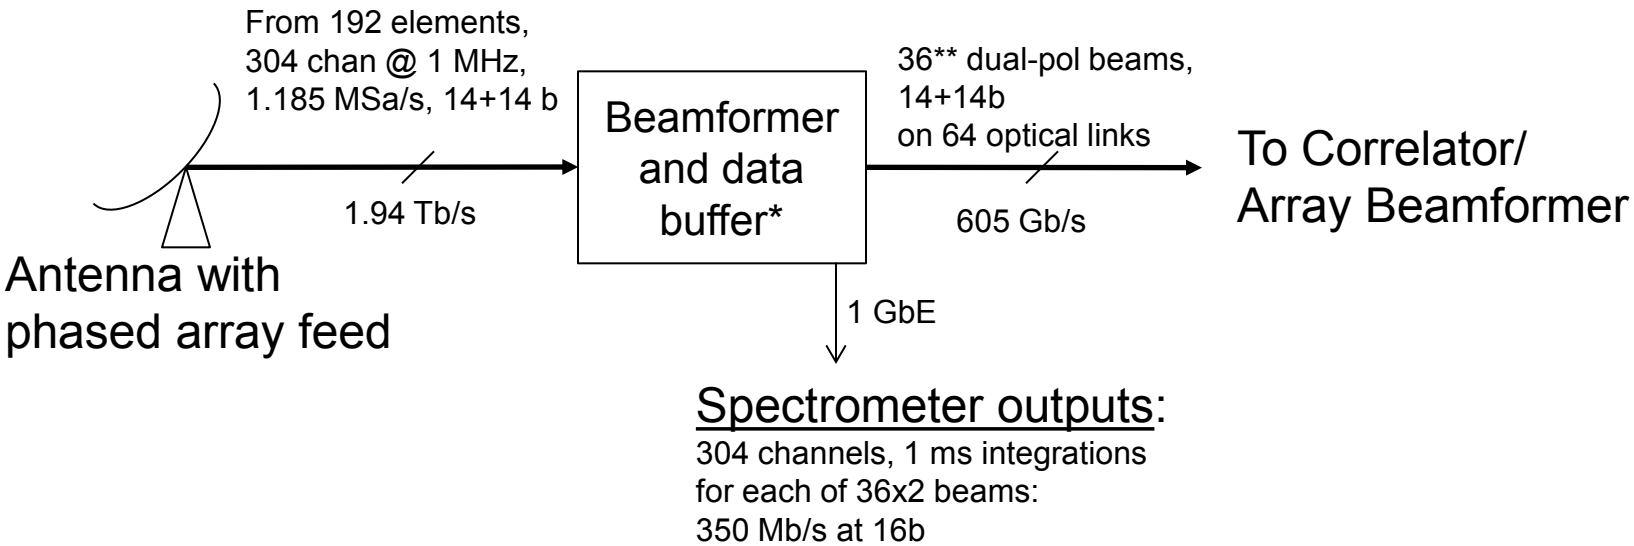

Notes:
* **Beamformer contains circular buffer for capturing input samples.**
** Maximum of 36 beams at high frequency end of tuning range, fewer at low frequency end.

Conceptual Block Diagram

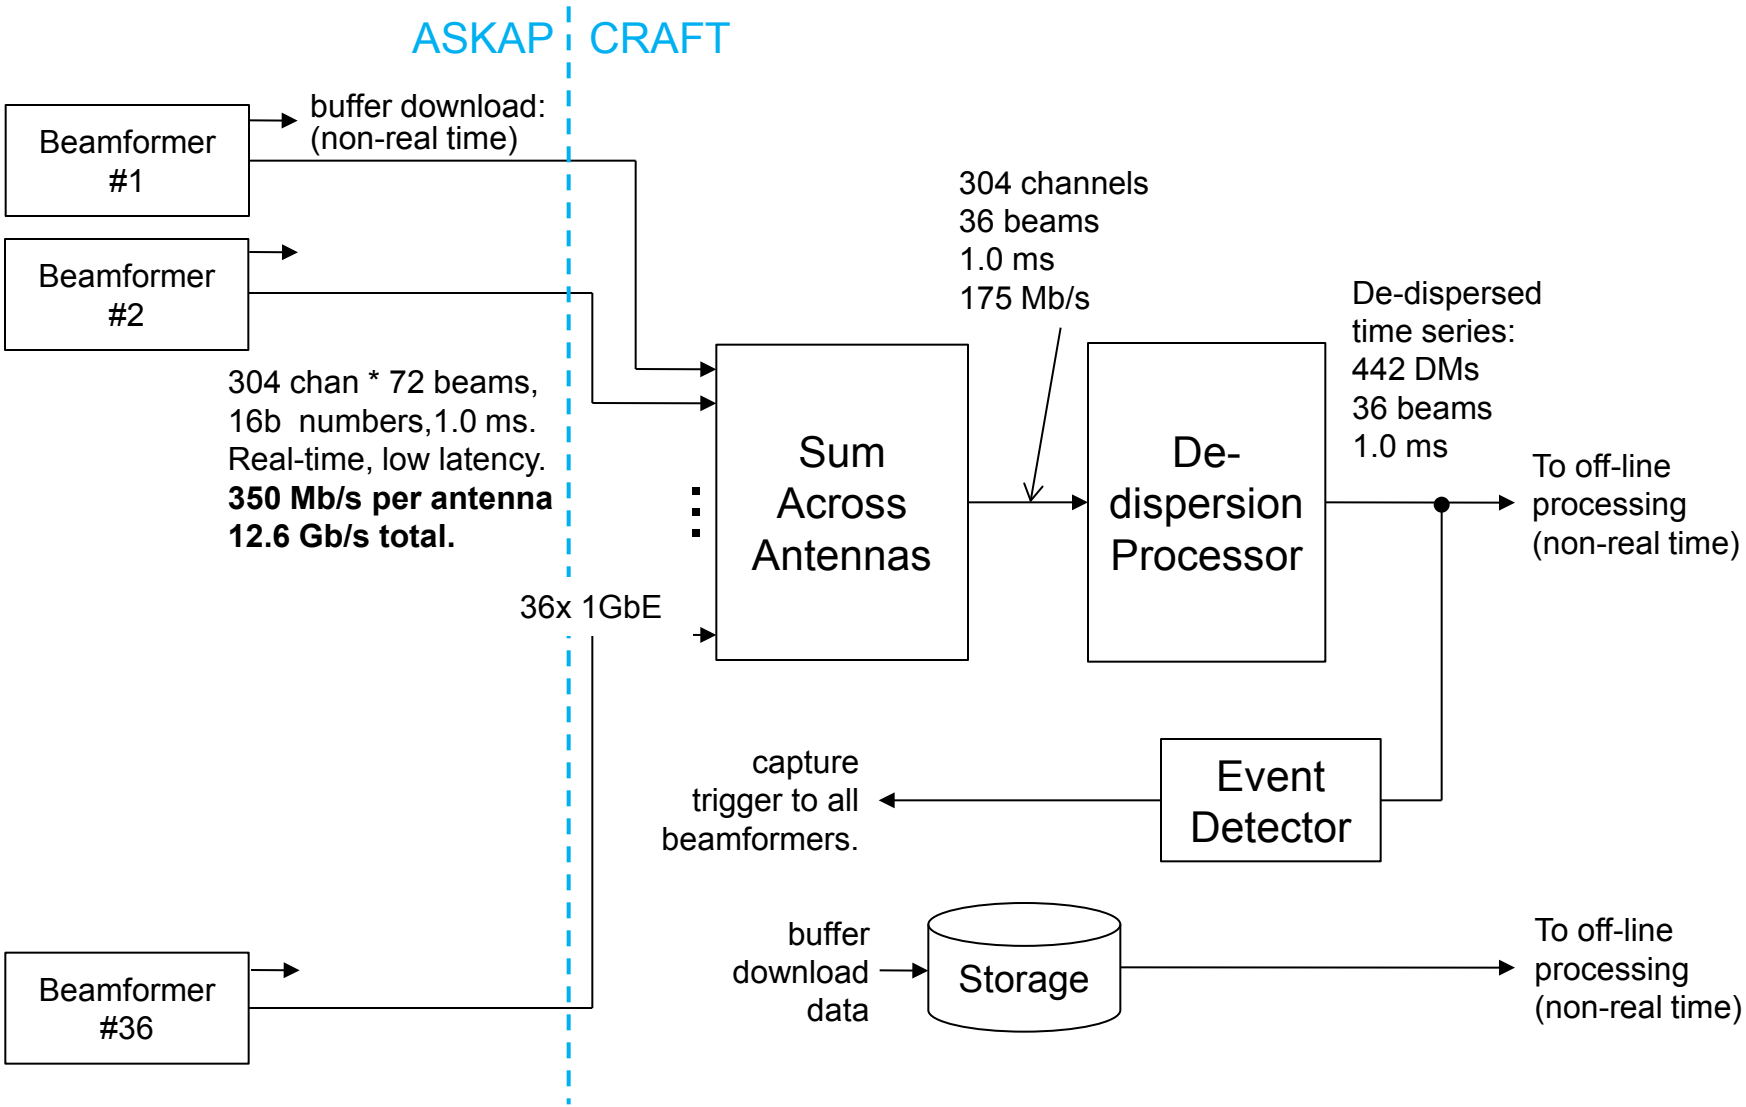

Performance Summary

Tardis Back-End Specifications

	Sym.	ASKAP	SD	
Maximum number of inputs (beams)	b	36	6	
DMs searched	D	442	512	
Integration time per spectrum, minimum	τ	0.9	0.1	ms
Frequency channels per spectrum	C	304	1024	
Group size, integrations	J	16	64	
Transient detection latency ($2J\tau$ + 1ms)		33	14	ms

Telescope Specifications

	ASKAP	DSS13		
RF band	700-1004	2200-2304	8200-8620	MHz
	1500-1804			MHz
IF band	424-724	200-304	100-520	MHz
Effective sampling rate	910.22	325	1300	MHz
Channels	304	1024	1024	
Useful channels	304	681	677	
Channel width	1.00	0.159	0.635	MHz
Integration time	999.8	100.825	100.825	μs
Beams (all polarizations)	72	1	1	