(GBT) Backends

Signal Processing Fundamentals - II

Richard Prestage

NAIC / NRAO Single Dish Summer School

July 2015

Atacama Large Millimeter/submillimeter Array
Karl G. Jansky Very Large Array
Robert C. Byrd Green Bank Telescope
Very Long Baseline Array

Overview of Talk

- Reminder of what we are trying to measure
  - ***Heterodyne*** instrumentation, not bolometers
- Review of existing GBT backends
- Backends – "old" versus "new"
- Detailed overview of GUPPI / VEGAS
  - Pulsar backend
  - Spectral line backend
- Brief summary of some other GBT CASPER-based backends

- Will NOT talk about interferometer backends (correlators)

- Instruments are GBT specific, but the concepts are generally applicable

Electromagnetic Waves

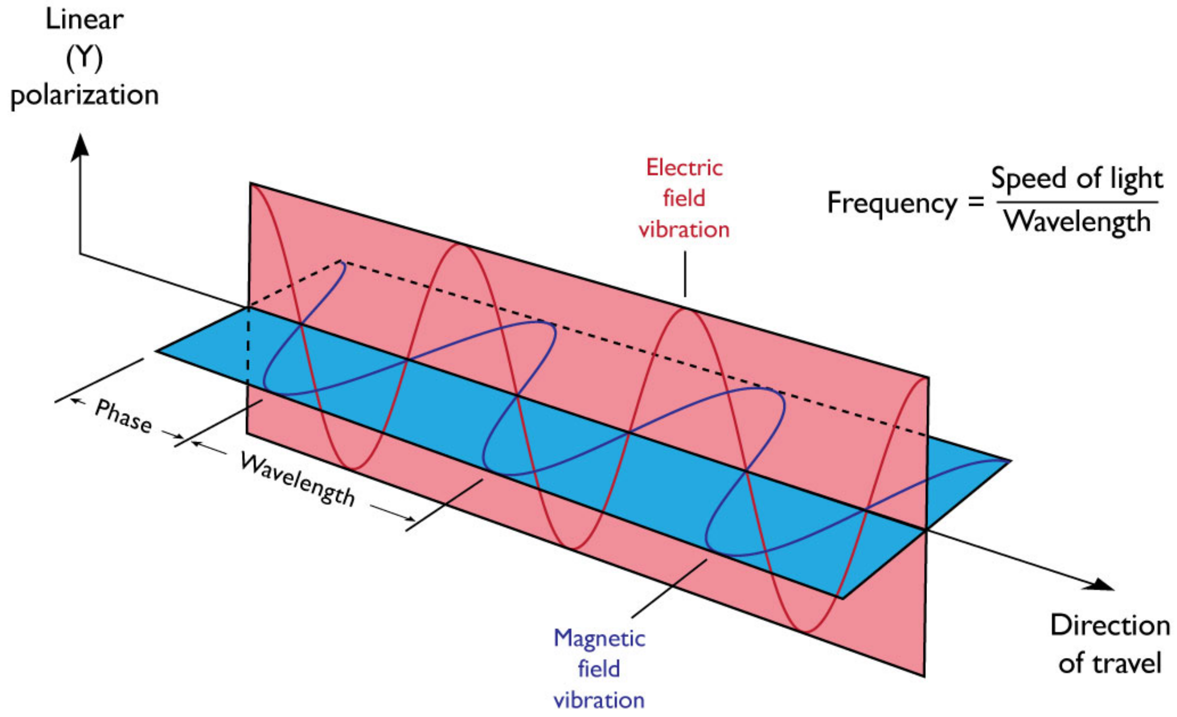

Only showing one polarization

Electromagnetic Wave Detection

- We want to measure ***everything*** we can know about the wave:
  - Direction of arrival
  - Frequency
  - Amplitude (Power)
  - Phase (not usually directly useful)
  - Polarization

  Full Stokes parameters

- Two approaches:
  - Coherent (heterodyne) detection
    - Treat the radiation like an electromagnetic wave
  - Incoherent (bolometer) detection
    - Treat the radiation like photons ("bundles" of energy)
- Complimentary Techniques:
  - Microwave ⇒ ~ few 100 GHz – heterodyne receivers / backends
  - ~ few 100 GHz ⇒ infrared – bolometers

Typical Astronomical Signals

- Typical astronomical signals are random Gaussian-distributed white noise.
- Sample in the ***time*** domain, Fourier transform to the ***frequency*** domain

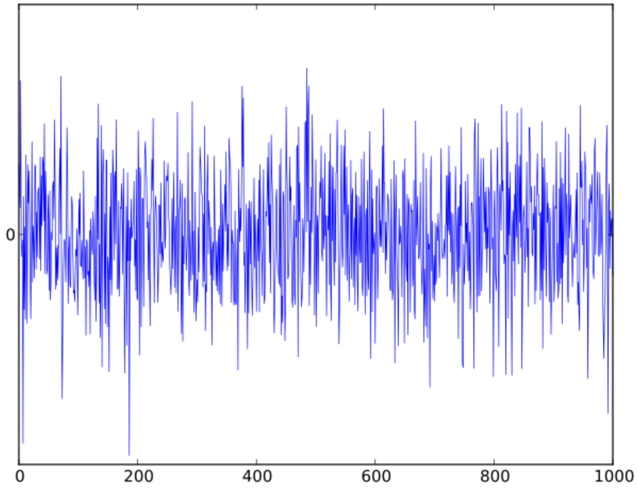

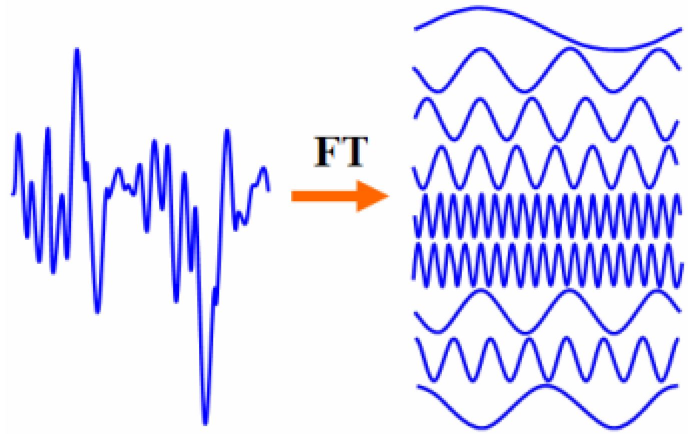

NRAO

Instrument Performance Specifications

- Bandwidth:
  - Extragalactic spectral lines, v ~ 3000 km s^{-1}
  - 900 MHz at 90 GHz ⇒ ~ 1 GHz
  - For blind redshift searches, spectral line surveys ⇒ ~ 10 GHz

- Spectral Resolution:
  - Maser lines v ~ 0.01 km s^{-1}
  - 55 Hz at OH 1665.4 MHz maser transition frequency

- Channels:
  - 5 km s^{-1} over 3000 km s^{-1} ⇒ 600 channels; round up to 1024
  - 22 GHz H2O masers, 0.3 km s^{-1} over 5000 km s^{-1} ⇒ 16384 channels

Performance Specifications

- Time Resolution:
  - ms pulsars, 10% duty cycle = 100 μs ⇒ ~ 1 μs (timing) - 100 μs (searching)

- RFI considerations / spurious signals:
  - For GBT, most difficult case is Prime Focus (< 1 GHz) operation
  - Typical (measured) RFI ~ $10^3\ Tsys \times \Delta\nu_m$
  - $SFDR = 10\ Log\left(10^{-4} \times \sqrt{\frac{\Delta v}{\tau}} \times \frac{1}{\Delta v_m}\right) = -89\ dBc$

    $\Delta v = 1\ kHz, \tau = 12\ hours$

  - Challenging spec, actual ADCs (e.g. VEGAS) = 56 dBc (decibels relative to the carrier)

NRAO

Specifications summary

Quantity	Specification
Bandwidth	~ 10 MHz to 10s of GHz
Spectral resolution	~ 10s of Hz to ~ MHz
Number of channels	1024 – 32768
Time resolution	~ 1 us to 10s of seconds
Number of IF inputs	8 x dual polarization
Number of spectral "windows"	~ 10 - 100
ADCs	8-bit
Spurious Free Dynamic Range	High!

- For most general purpose spectral line / continuum / pulsar observations (e.g. not Radar)

NRAO

Time / Frequency domain

- The Nyquist theorem states that a signal must be sampled at least twice as fast as the bandwidth of the signal to accurately reconstruct the waveform; otherwise, the high-frequency content will alias at a frequency inside the spectrum of interest (passband).
- Lowpass filter limits the frequency content of the input signal above the Nyquist rate.
- 1500 MHz bandwidth $\Rightarrow$ 3 Gsps sampling rate (0.3 ns durations)

- Frequency resolution $\Delta f = F_s / N$, F_s = Sampling Frequency, N = length of FFT
- At a fixed sampling rate, increasing frequency resolution decreases temporal resolution. e.g.:
  - 1.5MHz frequency resolution $\Rightarrow$ 0.3 μs time resolution (N = 1024)
  - 66 kHz frequency resolution $\Rightarrow$ 5.0 μs time resolution (N = 16384)

NRAO

Sampling referenced to Hydrogen Maser

- Hydrogen Maser generates an extremely stable atomic frequency reference. It uses the quantum transition between the hyperfine levels of the ground electron state of atomic hydrogen, corresponding to a frequency of 1.420 405 751 786 GHz
- Accuracy ~ 1×10^{-15} (femtoseconds)
- Generate 5 MHz, 100 MHz and 1 PPS (pulse per second)
- Used as reference for Local Oscillators, and to drive ADC and FPGA clocks
- 1 PPS used to trigger start of data acquisition
- Time, frequency and phase known to extremely high precision

A Menagerie of GBT Backends

Digital Continuum Receiver (DCR)

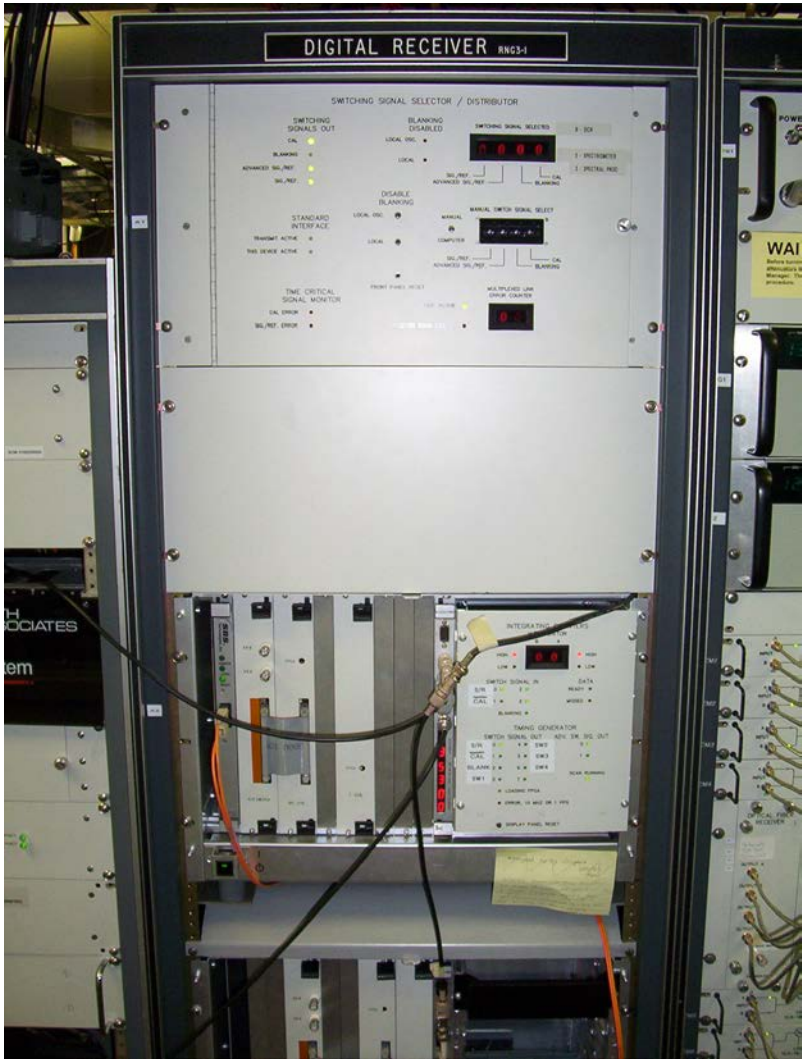

Caltech Continuum Backend

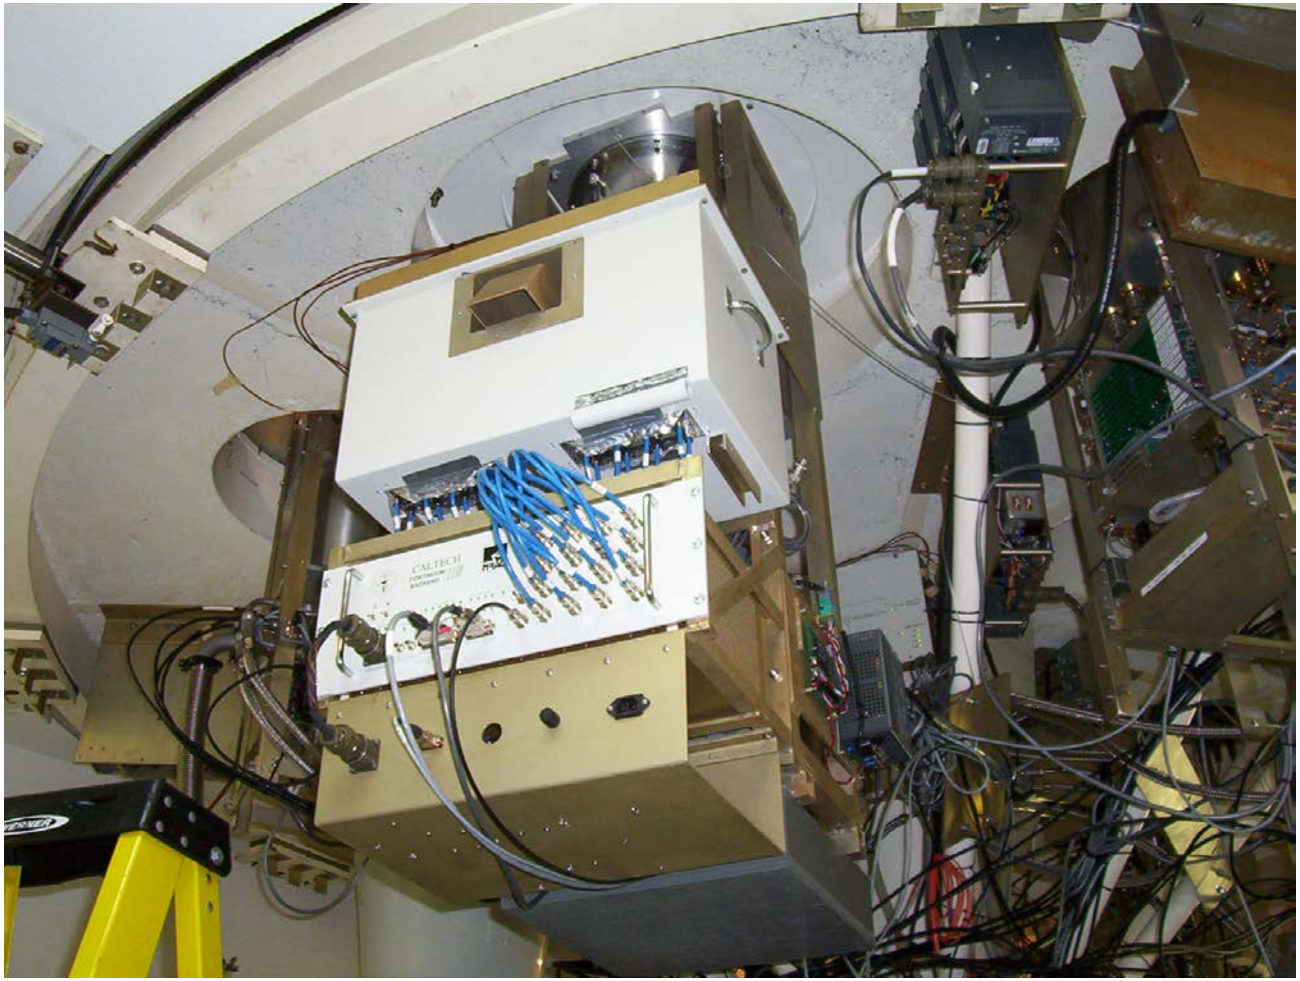

Spectral Processor

Spectrometer

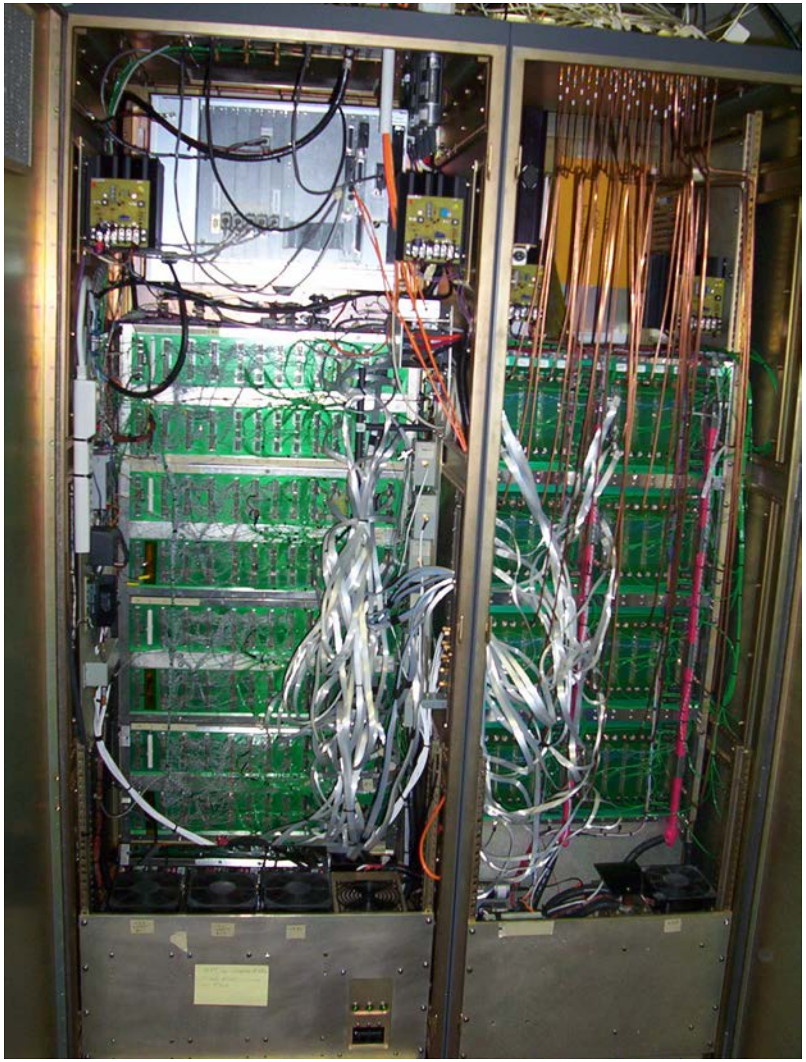

Zpectrometer

VEGAS

GUPPI (Green Bank Ultimate Pulsar Processing Instrument)

VLBA Mark 5b

GBT Backends - I

- DCR – workhorse utility (pointing, focus) backend
  - Voltage-to-frequency converters into 28-bit counters
- CCB (Caltech Continuum Backend)
  - Sensitive wideband backend for use with Ka-band receiver
  - Optimized RF detector circuits and fast beam-switching
- Spectral Processor (disassembled)
  - FFT Spectrometer, pulsars, polarization measurements
  - High dynamic range $\Rightarrow$ good immunity to RFI
- Spectrometer (retired)
  - GBT workhorse spectrometer until ~ 16 months ago
  - Autocorrelation Spectrometer
  - Application Specific Integrated Circuit (ASIC) technology
  - 256k lags, flexible observing modes

NRAO

GBT Backends - II

- Zpectrometer (retired) – Harris et al.
  - Analog log correlation spectrometers (WASP2)
  - 14 GHz bandwidth
  - Used with Ka-band for CO at z ~ 3.4 - 8
- GUPPI
  - Current workhorse pulsar backend
  - 8-bit sampling, full stokes, up to 800 MHz BW / 40.96μs min. int. time
  - Incoherent and coherent modes.
- VLBA Backends (Roach Digital Backend)
  - Standard VLBA backend

GBT Backends - III

- MUSTANG 2
  - Bolometer array with Roach-based readout electronics
- FLAG Beamformer
  - Roach / GPU based beamformer for 19 element phased-array feed (PAF)
- Portable Fast Sampler Radar Backend (Margot et al.)
  - Programmable Logic Device (PLD) based system
  - 60 Msps (16ns time resolution)
  - 2,4 or 8-bit sampling
- JPL Radar Backend

GBT Backends - IV

- VEGAS
  - Versatile GBT Astronomical Spectrometer
  - Single Spectrometer which combines the features of all previous GBT Spectrometers
  - Also provides Continuum and Pulsar capabilities

Backends – the "old" days

- > 10 different backends
- Each backend custom-designed
  - Usually for a specific type of science
  - Designed to optimize a specific property (time resolution, dynamic range)
- Each backend unique
  - Custom hardware / software
- Often, lengthy development time
  - GBT Spectrometer ~ 10 years

- Many still in use at many Observatories

Backends – the "new" days

- Use CASPER-based technology
- Standard processors - FPGAs, GPUs, multi-core CPUs
- Standard (or no) backplanes
- Small number of commercial hardware components (ADC boards, etc)
- 10 (40) GbE network switches used for interconnect / corner turning
- Open source gateware / firmware / software
- In use at GBT since ~ 2007 (early adopter)
- Other countries have similar approaches
  - ASTRON (NL) UniBoard I / II

- Specifics will change, concepts are the most important

- Radio Telescopes becoming reconfigurable high-performance computers with antennas as peripherals

Backends – the "new" days

- In principle, a single high time/frequency resolution, wide bandwidth spectrometer could do everything
- Not quite there, but VEGAS is close
- 8 banks provide up to 10 GHz bandwidth
- Capable hardware, reprogram to achieve different functionality:
  - Continuum
  - Spectroscopy
  - Pulsar searching
  - Pulsar timing
- Multi-pixel instruments $\Rightarrow$ multiple spectrometers
- Beamformers add another complexity

GUPPI (Green Bank Ultimate Pulsar Processing Instrument)

First Generation CASPER hardware (IBOB, BEE-2)

Searching: 800MHz of BW, 4096 chan, and RFI-resistance (from polyphase filterbank) makes GUPPI a "Super-Spigot" that will be unbeatable for searches from 1-5 GHz (previously the BCPM and Spigot)

Ultra-High Precision Timing: 8-bit sampling, 800MHz of BW, coherent de-dispersion and RFI resistance allows unprecedented timing precision from 1-3 GHz for millisecond pulsars (i.e. the NANOGrav project and the search for nHz gravitational waves; previously GASP/CGSR2)

Other High Dynamic-Range Studies: 8-bit sampling, high spectral resolution and Full Stokes makes GUPPI perfect for scintillation studies, HI absorption, Faraday rotation measurements, polarization studies, and singlepulse studies (previously the Spectral Processor)

NRAO

Slide after one by Scott Ransom

GUPPI Specifications

- 8-bit sampling
- 2 polarization and full stokes parameters
- Bandwidths: 100, 200, 800 MHz
- Spectral Channels: 64, 128, 256, 512, 1024, 2048, 4096
- Minimum integration time 40.96 μs (20.48 μs with on-line folding)
- Search mode and coherent de-dispersion mode

GUPPI

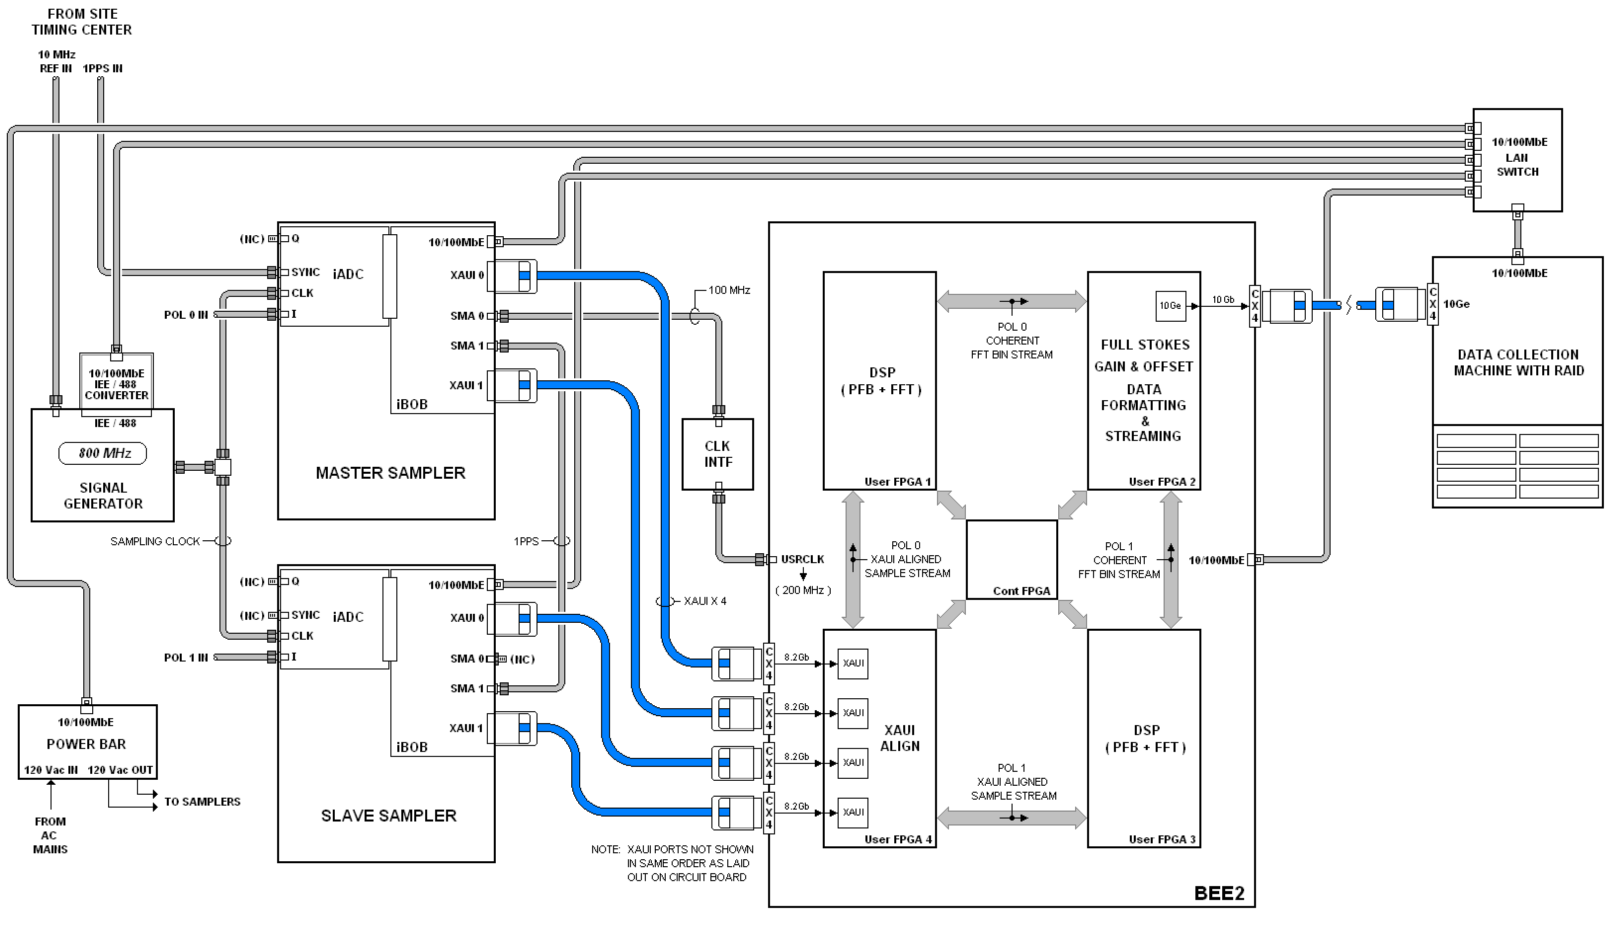

FIGURE 1

GUPPI Hardware

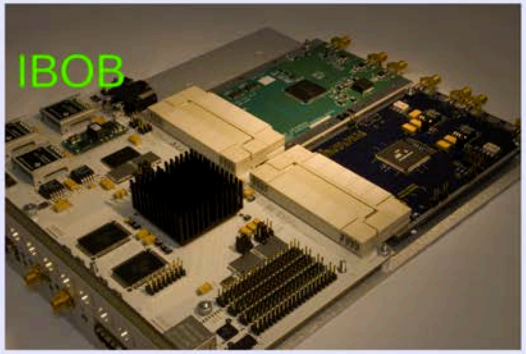

GUPPI architecture:
~1 MHz PFB in FPGAs
Coherent dedisp in GPUs

XAUI

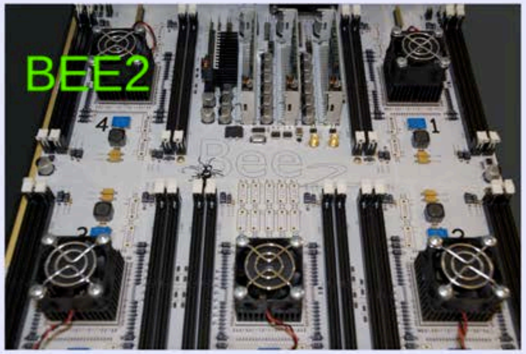

10 Ge
switch;
24 Gb/s

Slide courtesy of Paul Demorest

Search Mode

- Fairly conventional low frequency resolution high time-resolution spectrometer

Coherent De-dispersion Mode

- Pulse is dispersed during propagation through interstellar medium
- Can be described as the work of a "phase-only" filter, or ***transfer function*** H
- In frequency domain:

$$V(\nu_0 + \nu) = V_{int}(\nu_0 + \nu)\, H(\nu_0 + \nu)$$

- Where $V(\nu), V_{int}(\nu)$ = Fourier Transforms of measured and intrinsic complex voltages.

- Delay can be represented as phase rotations that depend on frequency and path length travelled, $\Delta\Phi = -k(v_0 + v)d$. Transfer function becomes:

$$H(\nu_0 + \nu) = e^{-ik(v_0+v)d}$$

Coherent De-dispersion Mode

- Recall ISM can be modelled as a low-density electron plasma, with:

$$v_p = \sqrt{\frac{n_e e^2}{\pi m_e}}$$

- The wavenumber $k(v)$ is given by:

$$k(v) = \frac{2\pi}{c}\sqrt{v^2 - \ v_p{}^2}$$

- Using first terms in Taylor expansion:

$$k(v_0 + v) = \frac{2\pi}{c}\ (v_0 + v)\left[1\ -\ \frac{v_{p^2}}{2(v_0 + v)^2}\right]$$

Coherent De-dispersion Mode

- Which leads to the transfer function:

$$H(v_0 + v) = exp\left[2\pi i \frac{DM}{2.4 \times 10^{-16}} \frac{v^2}{(v + v_0)\, {v_0}^2}\right]$$

- (v in MHz)

- Finally:

$$V_{int}(\nu_0 + \nu) = V(\nu_0 + \nu)\, H^{-1}(\nu_0 + \nu)$$

- Usually combine inverse transfer function with a taper, T, so:

$$C(v_0, v) \equiv T(v) H^{-1}(v_0 + v)$$

C is known as the ***chirp function***

Implementation

- Obtain in-phase (I) and quadrature (Q) baseband voltage samples (I,Q are ***not*** Stokes parameters)
- Take a set of data of length n samples
- (n must be at least $2t_{DM}\Delta f$ samples long, where t_{DM}= dispersion delay across bandpass Δf).
- Compute a discretely sampled chirp function for n samples
- Fourier transform the data set of n points, and multiply Fourier components with chirp function
- Inverse Fourier transform back into time domain
- Repeat with next set of samples

- In GUPPI, 100 MHz chunks of bandwidth are de-dispersed in each of 8 GPUs.

NRAO

VEGAS (Versatile GBT Astronomical Spectrometer)

- Designed to make maximum use of GBT K-band focal plane array (7 pixels) – 8 x dual-polarization inputs
- General purpose replacement for old spectrometers (GBT Spectrometer, Spectral Processor)
- Replacement / Upgrade for GUPPI pulsar capabilities (in progress)

New Spectrometer	Old GBT Spectrometer
Polyphase Filter Bank for RFI rejection	Autocorrelation Spectrometer
256 level samplers	3 level samplers
16 high-speed samplers	8 high-speed samplers
1350 MHz bandwidth/sampler	800 MHz bandwidth/sampler
800 MB/s max disk I/O rate	< 25 MB/s max disk I/O rate
New technology (FPGA based, Virtex-{5,6}, GPUs)	Designed and built in early 1990's

VEGAS Spectral Line Modes

Table 1: VEGAS Modes

The configurations supported by each of the 8 spectrometers of VEGAS are listed in the table.

Mode Name	Number of spectral windows/pol	Bandwidth of spectral window (MHz)	Number of channels per spectral window per pol	Spectral resolution (kHz)	Minimum int time[b] (msec)	Max data rate[a] (MB/sec)	Data rate supported by software[c] (MB/sec)	Max data rate from 8 spec (GB/sec)
Single Sub-band modes[d]								
Mode 1 (H1K/HBW)	1	1500	1024	1465	0.5	32	32	0.3
Mode 2 (H16K/HBW)	1	1500	16384	92	1.4	187	32	1.5
Mode 3 (H16K/HBW)	1	1080	16384	66	2.0	130	32	1.0
Mode 4 (L1/LBW1)	1	187.5	32768	5.7	9	52	32	0.4
Mode 5 (L1/LBW1)	1	187.5	65536	2.9	19	52	32	0.4
Mode 6 (L1/LBW1)	1	187.5	131072	1.4	30	69	32	0.6
Mode 7 (L1/LBW1)	1	100	32768	3.1	10	51	32	0.4
Mode 8 (L1/LBW1)	1	100	65536	1.5	20	51	32	0.4
Mode 9 (L1/LBW1)	1	100	131072	0.8	30	69	32	0.6
Mode 10 (L8/LBW1)	1	23.44	32768	0.7	5	93	32	0.8
Mode 11 (L8/LBW1)	1	23.44	65536	0.4	11	93	32	0.8
Mode 12 (L8/LBW1)	1	23.44	131072	0.2	27	75	32	0.6
Mode 13 (L8/LBW1)	1	23.44	262144	0.1	44	93	32	0.8
Mode 14 (L8/LBW1)	1	23.43	524288	0.04	67	125	32	1.0
Mode 15 (L8/LBW1)	1	11.72	32768	0.4	5	93	32	0.8
Mode 16 (L8/LBW1)	1	11.72	65536	0.2	11	93	32	0.8
Mode 17 (L8/LBW1)	1	11.72	131072	0.1	33	62	32	0.5
Mode 18 (L8/LBW1)	1	11.72	262144	0.04	44	93	32	0.8
Mode 19 (L8/LBW1)	1	11.72	524288	0.02	89	93	32	0.8
8 Sub-band modes[d]								
Mode 20 (L8/LBW1)	8	23.44	4096	5.7	5	96	32	0.8
Mode 21 (L8/LBW1)	8	23.44	8192	2.9	10	96	32	0.8
Mode 22 (L8/LBW1)	8	23.44	16384	1.4	30	64	32	0.6
Mode 23 (L8/LBW1)	8	23.44	32768	0.7	40	96	32	0.8
Mode 24 (L8/LBW1)	8	23.44	65536	0.4	75	104	32	0.9
Mode 25 (L8/LBW1)	8	16.9	4096	4.1	7	72	32	0.6
Mode 26 (L8/LBW1)	8	16.9	8192	2.1	14	72	32	0.6
Mode 27 (L8/LBW1)	8	16.9	16384	1.0	39	48	32	0.4
Mode 28 (L8/LBW1)	8	16.9	32768	0.51	54	72	32	0.6
Mode 29 (L8/LBW1)	8	16.9	65536	0.26	100	80	32	0.7

VEGAS Bands / Sub-bands

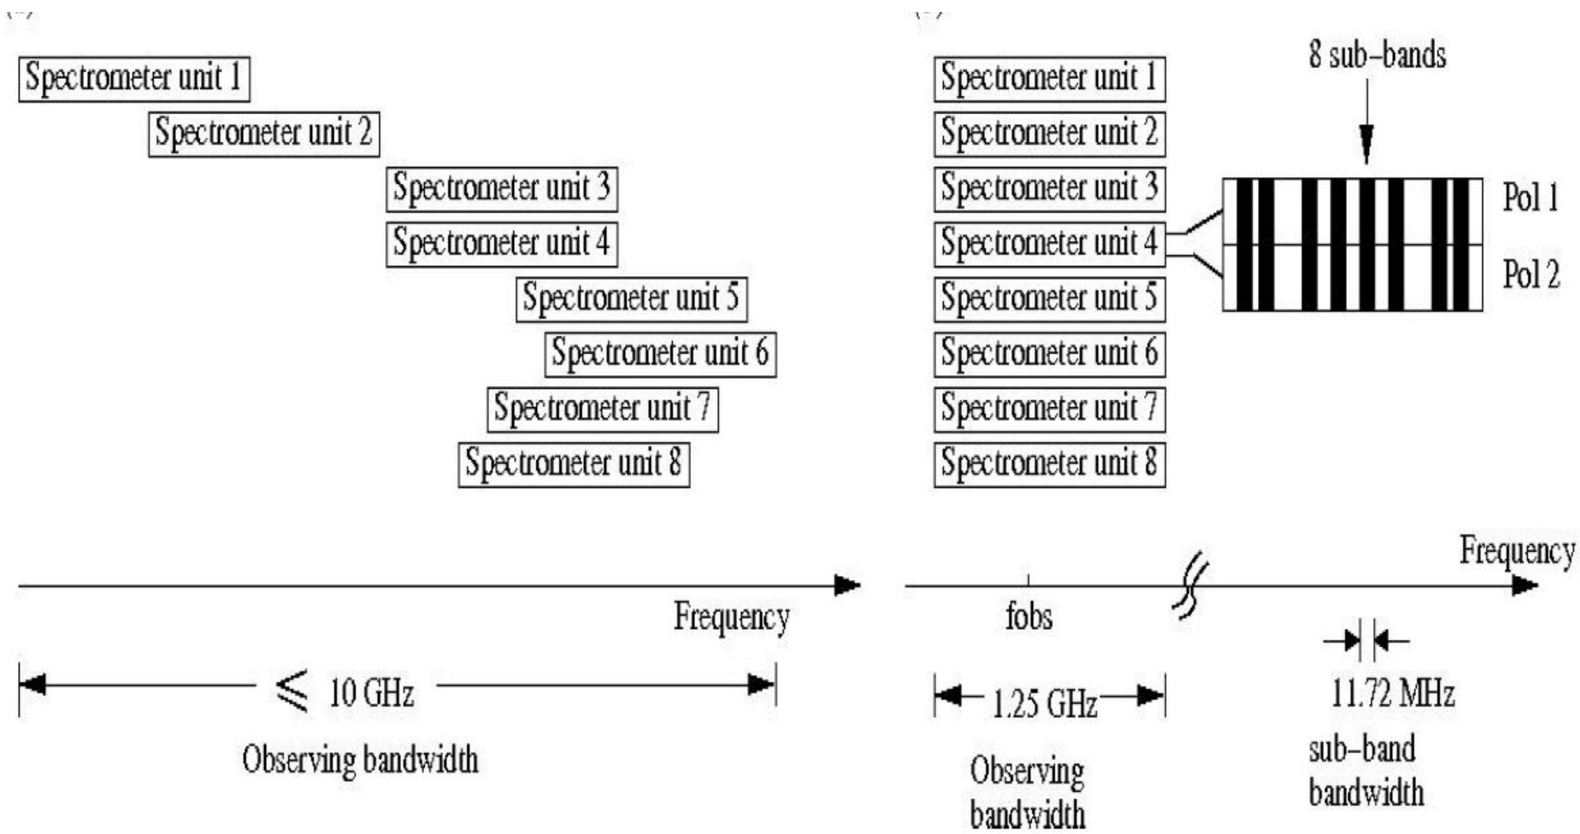

Design

- 5 Gsps max sampling and PFB/FFT calculations
- Heterogeneous Computing Approach
  - Divide processing into two components
  - Use FPGAs to fully process bandwidths greater than 400 MHz
  - Use FPGA front-ends to pre-process, split and packetize data, then GPUs to provide fine channelization on narrower chunks
- Software Design
  - Adapt concepts and code from GUPPI
  - Fully Integrate into the GBT Monitor and Control system

VEGAS

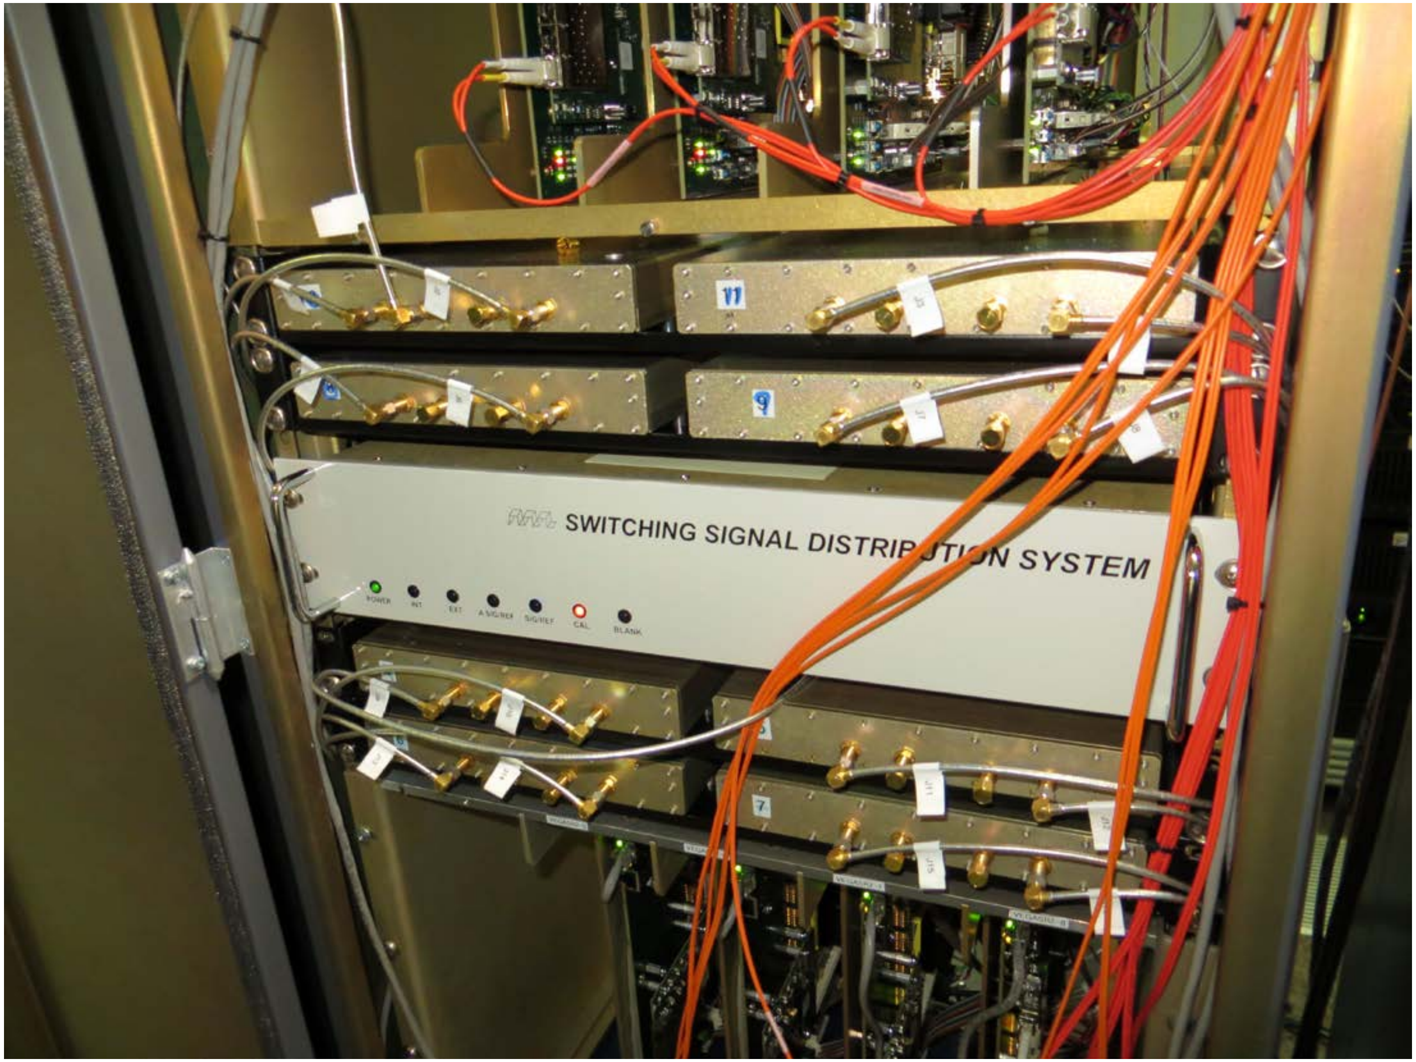
SWITCHING SIGNAL DISTRIBUTION SYSTEM
POWER
INT
EXT
CAL
BLANK

VEGAS

VEGAS Simplified Block Diagram

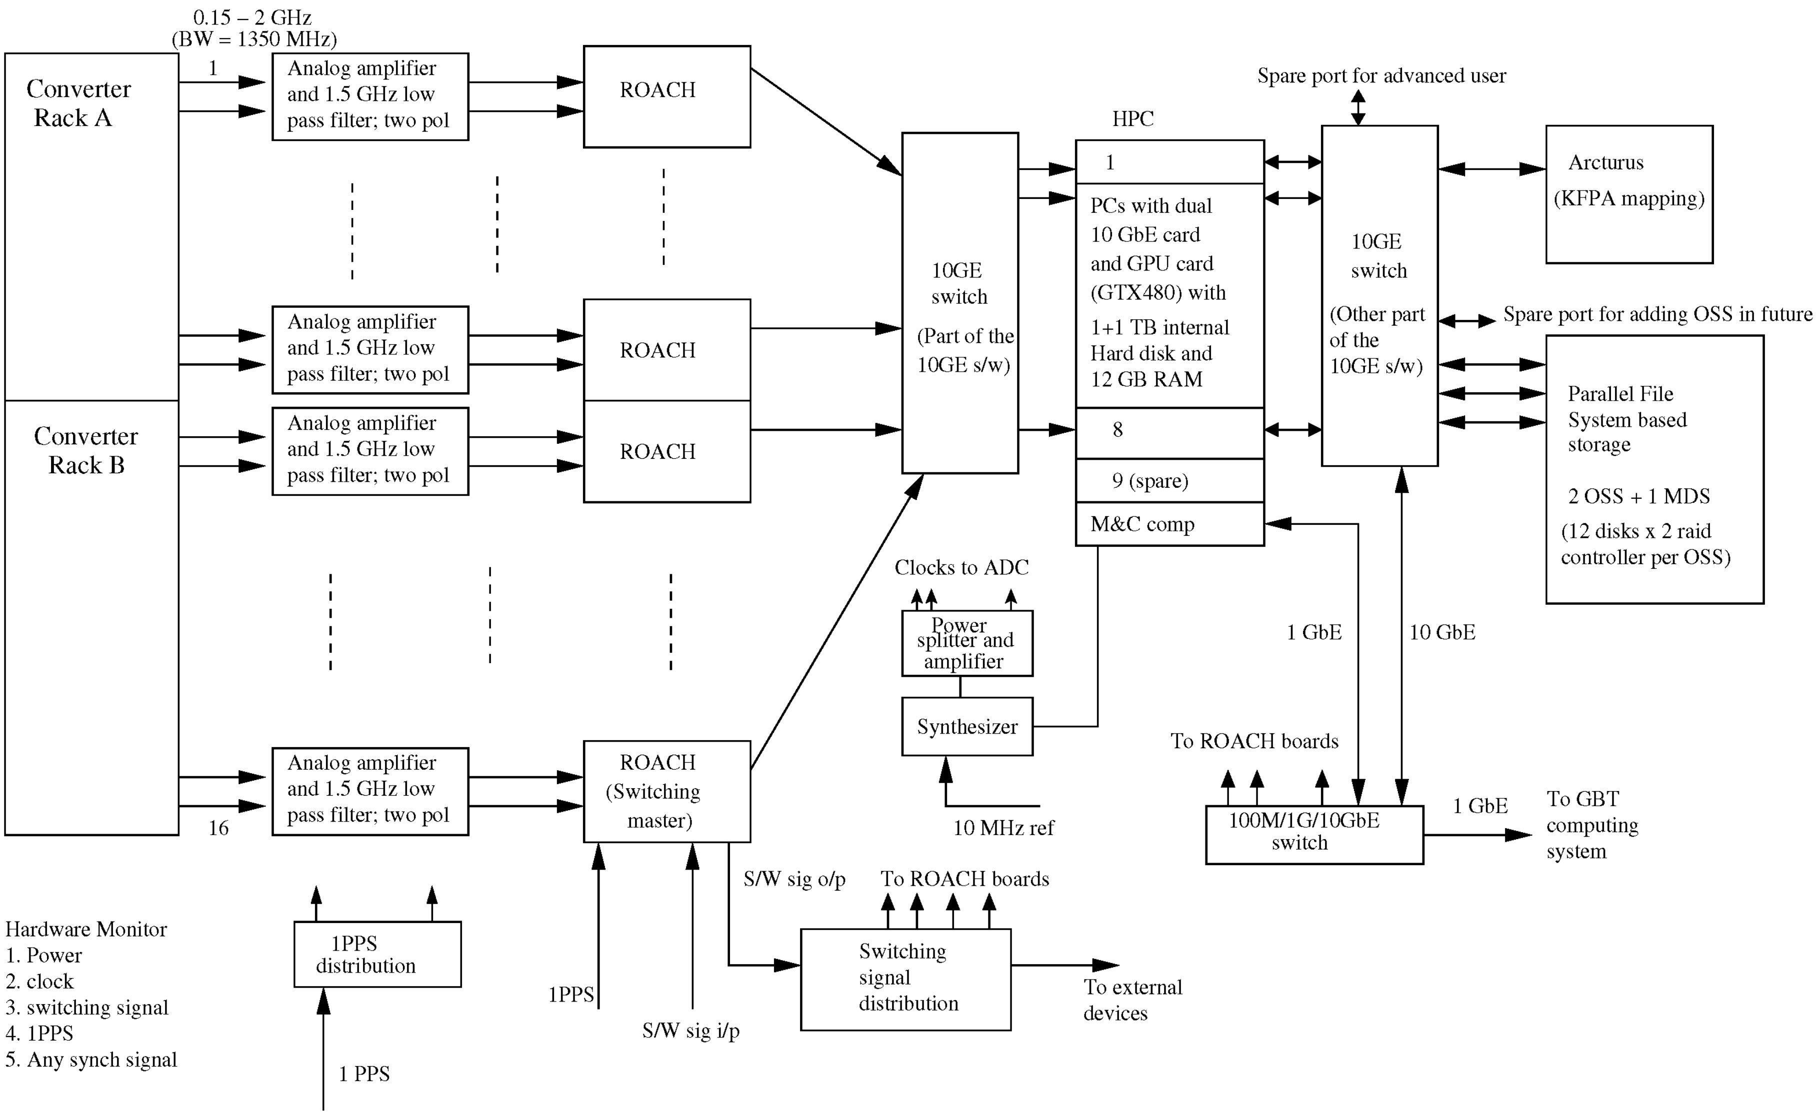

“High bandwidth” modes (FPGA only)

- Max Bandwidth 1.5 GHz at 1024 Channels.
- Spectral Resolution 32768 channels
- DDC, PFB, Stokes, Accumulate, buffer, packetize

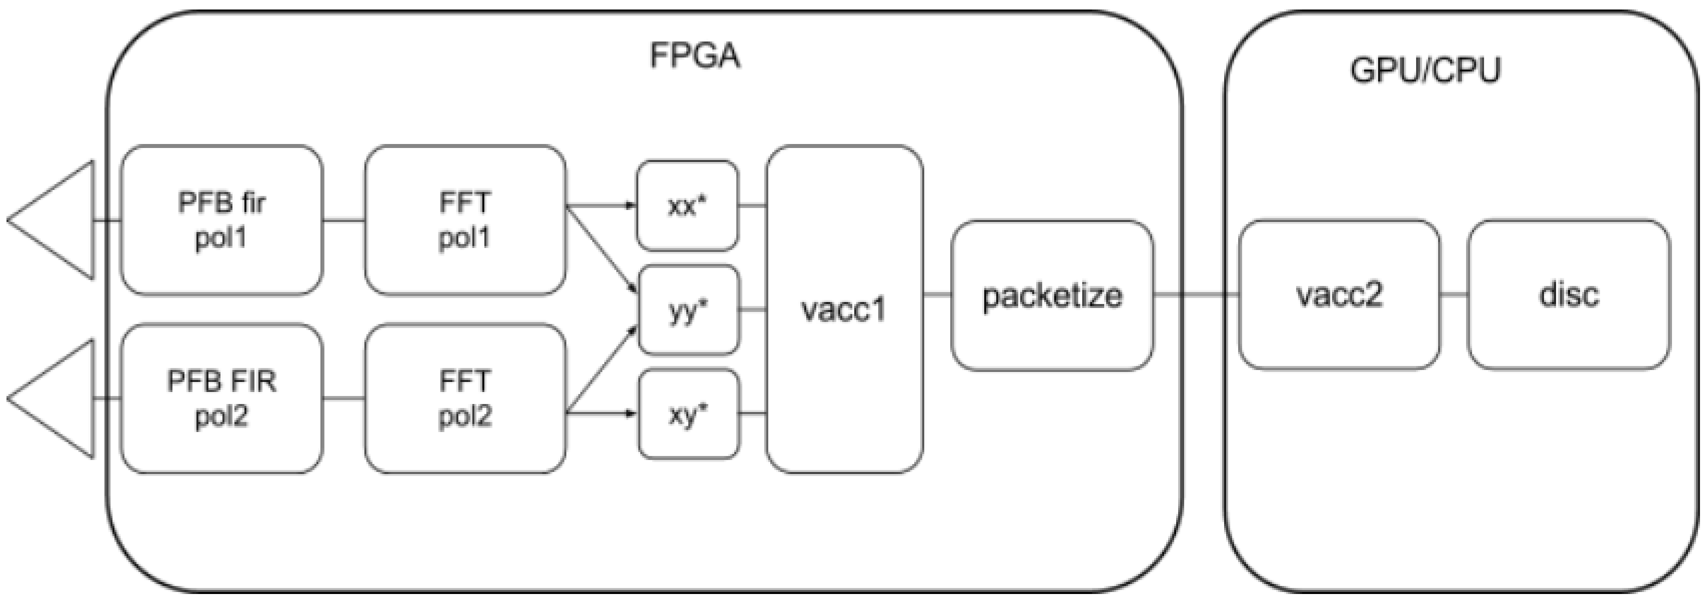

“Low Bandwidth” Modes (FPGA+GPU)

- 2x 8 Sub IFs with 1 - 30 MHz bandwidth
- Spectral resolution ranging from 0.2 - 7.3 KHz
- DDC mixes at arbitrary frequencies with 10KHz resolution
- Uses computationally efficient CIC filter (H.Chen)

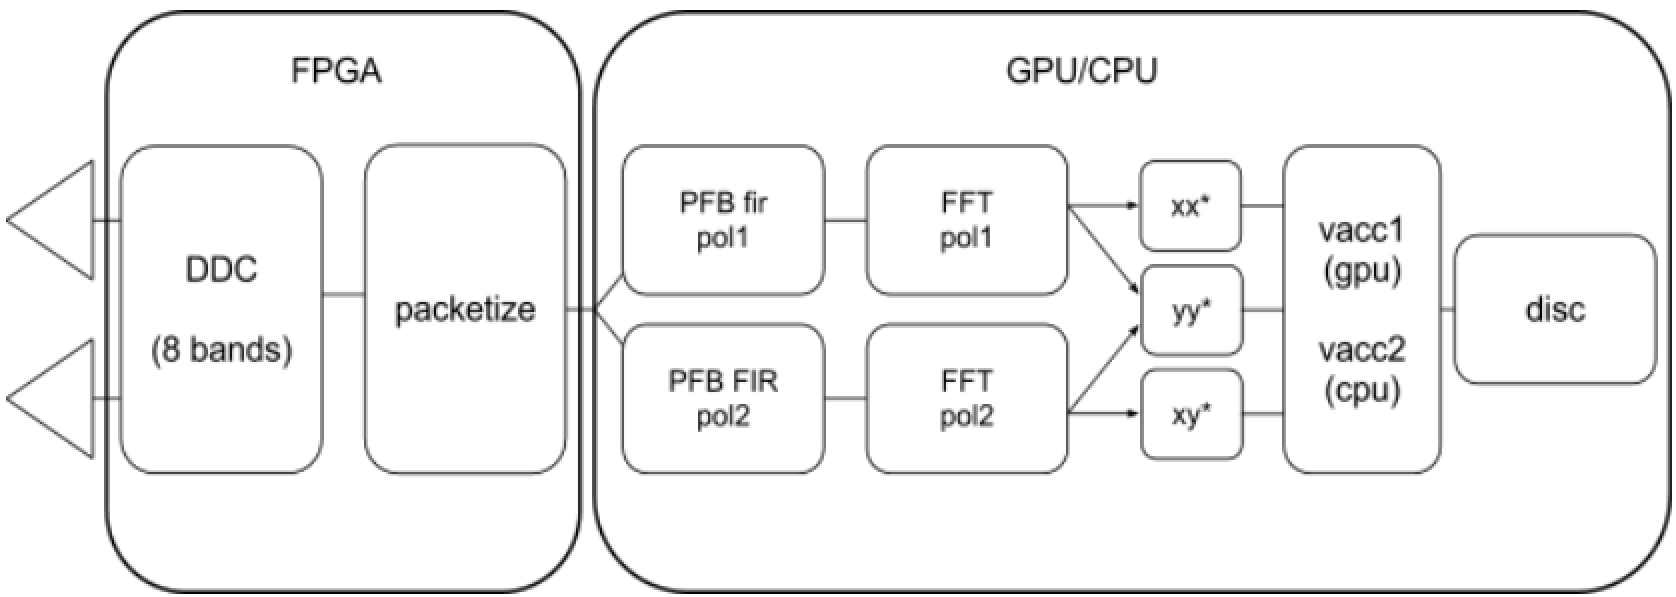

HPC Data Flow: High BW Modes

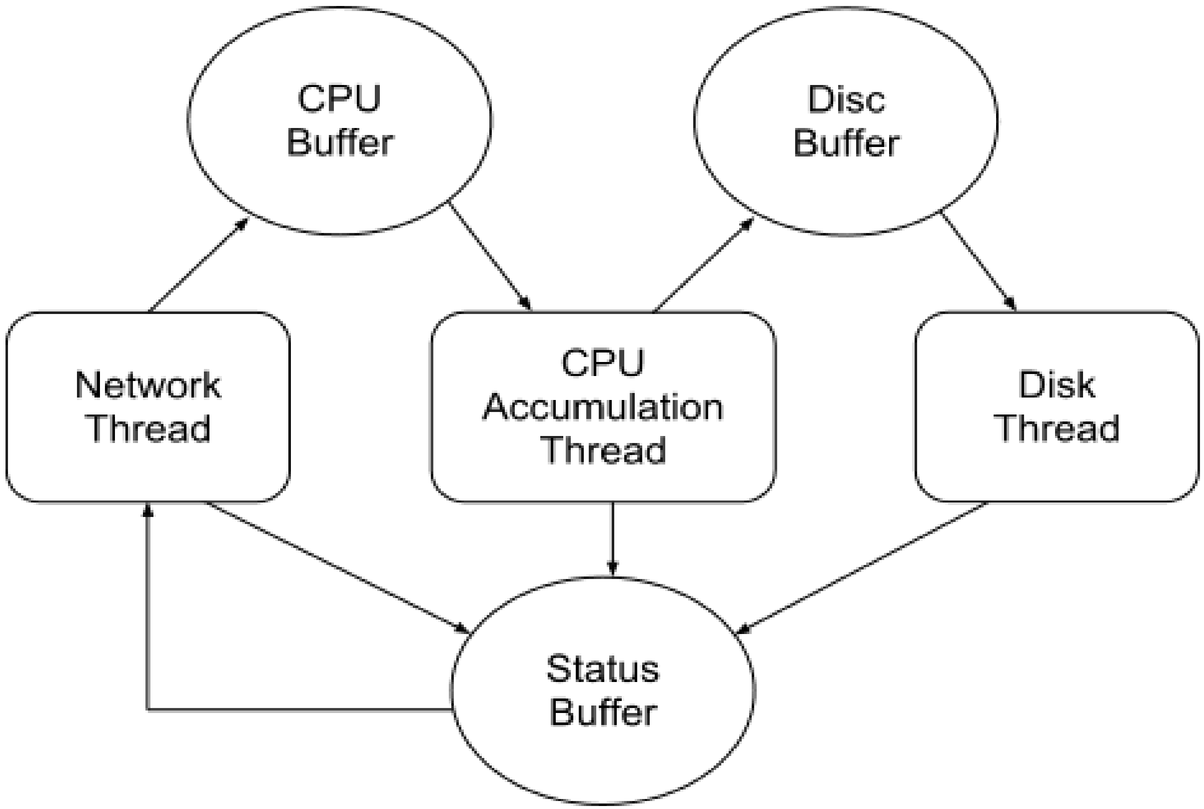

HPC Data Flow: Low BW Mode

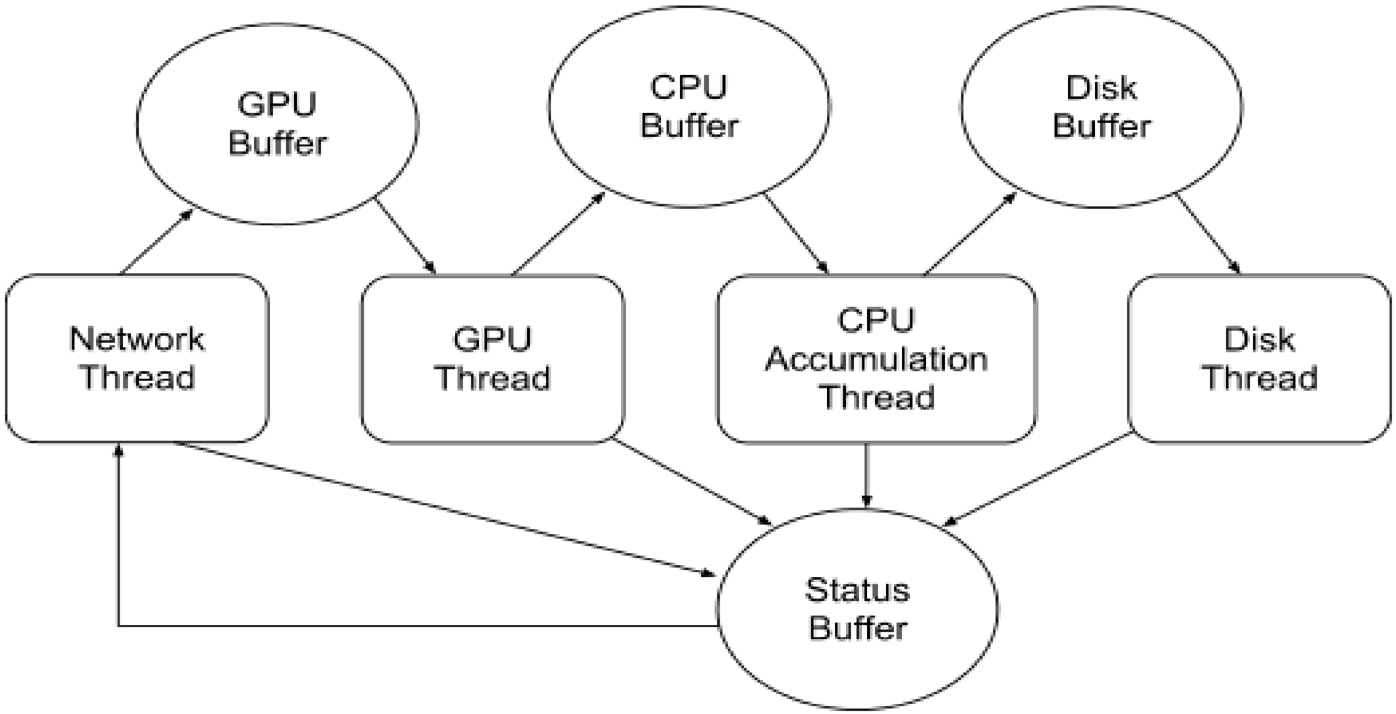

VEGAS Pulsar Modes

- Use identical Hardware (ADCs, FPGAs, GPU machines)
- Add port of GUPPI firmware / software
- Well underway
- Incoherent Modes:
  - 2 GHz bandwidth (4Gsps sampling)
  - 64 – 8192 channels, full Stokes
  - ~ 50 µs dump rate
- Coherent Modes
  - etc

Observing with VEGAS

- Choose line rest frequency(ies) and velocity
- Specify bandwidth and spectral resolution (low, medium, high)
- Specify number of spectral windows (1 – 64)
- Decide switching scheme
- Decide integration time
- "Configure()", "Balance()" and Observe.

- Piece of cake!

NRAO

What the control software is doing

- Deciding what FPGA "personality" (firmware) to load into the FPGA
- Setting the FPGA / ADC clock (sampling frequency)
- "Mixed Mode Clock Management" (calibrating the FPGA clock phase relative to the ADC inputs, to avoid glitches on data capture)
- Adjusting the offset, gain and phase of the four ADC cores
- Setting the (analog) filters and LOs upstream of VEGAS, and the VEGAS digital filters and LOs as appropriate for the selected bandwidth, etc
- Setting the FPGA "requantization gains" depending on input power level
- Starting to obtain ADC samples synchronized to the Observatory 1 pps (pulse per second) tick
- Packetizing the data and sending it off to the GPU for further processing

VEGAS results - continuum

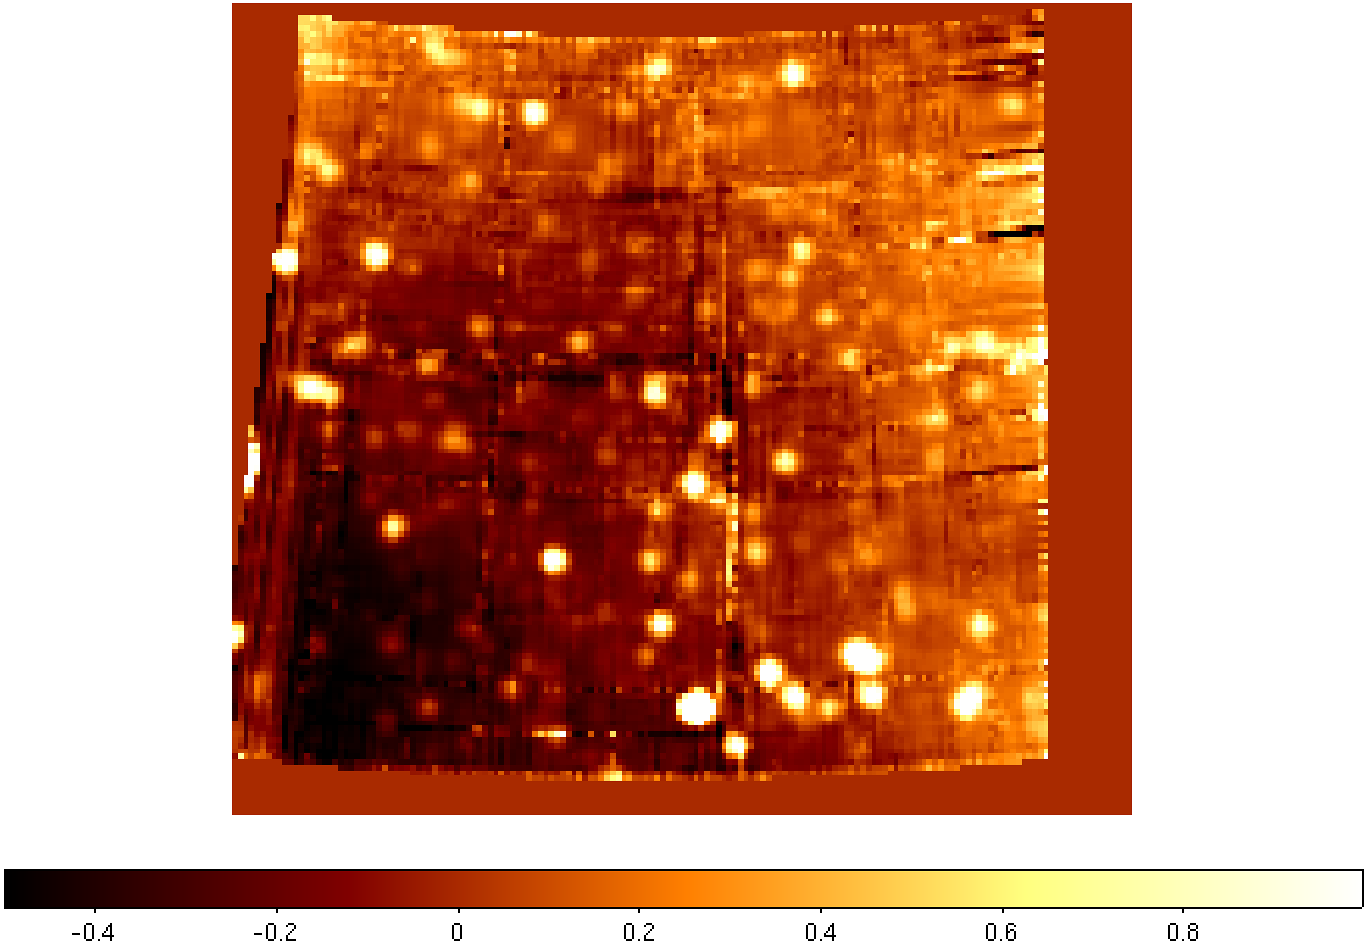

Image courtesy of Brian Mason

VEGAS results – spectral line

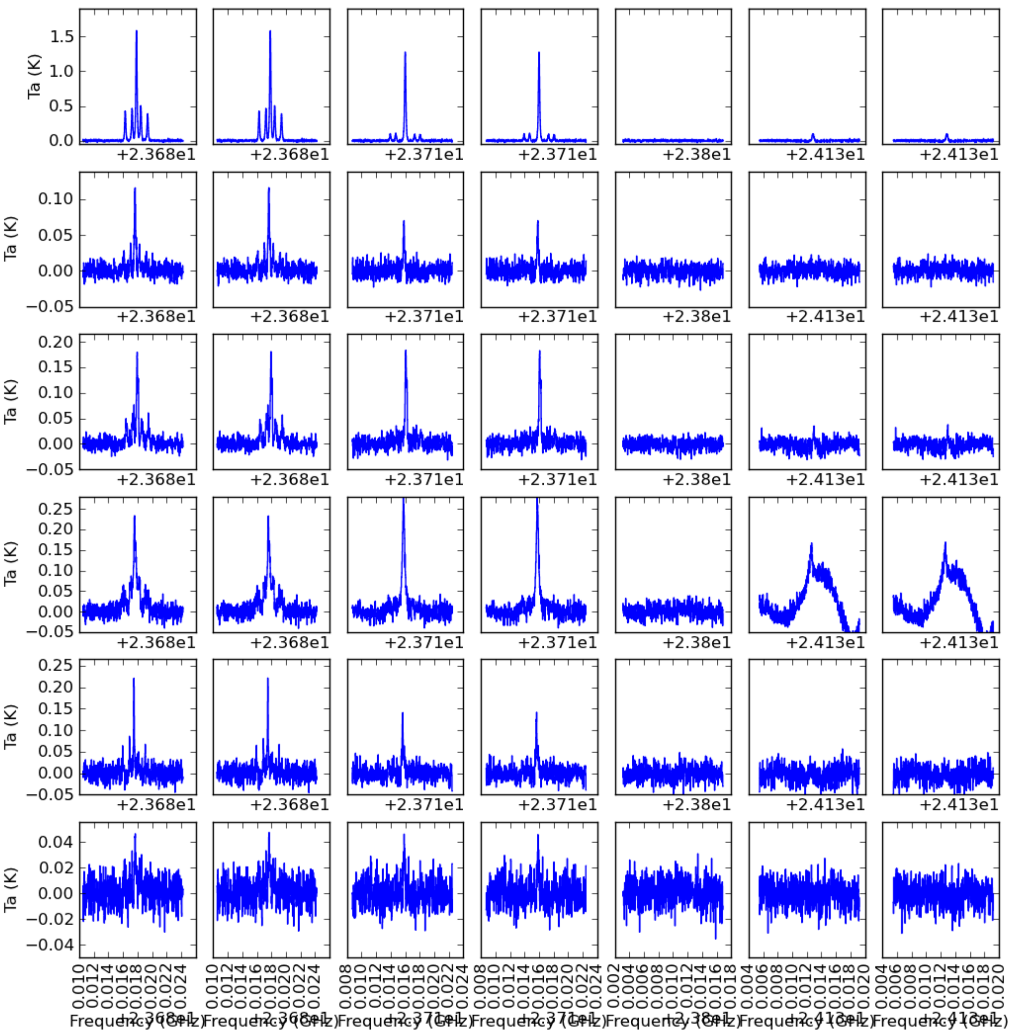

Image courtesy of Amanda Kepley

Other CASPER backends - RDBE

- Roach Digital Backend
- 4 Gsps Digital VLBI backend
- NRAO / MIT Haystack / CASPER / KAT collaboration
- Records to Mark 5C recorder

Other CASPER backends - RDBE

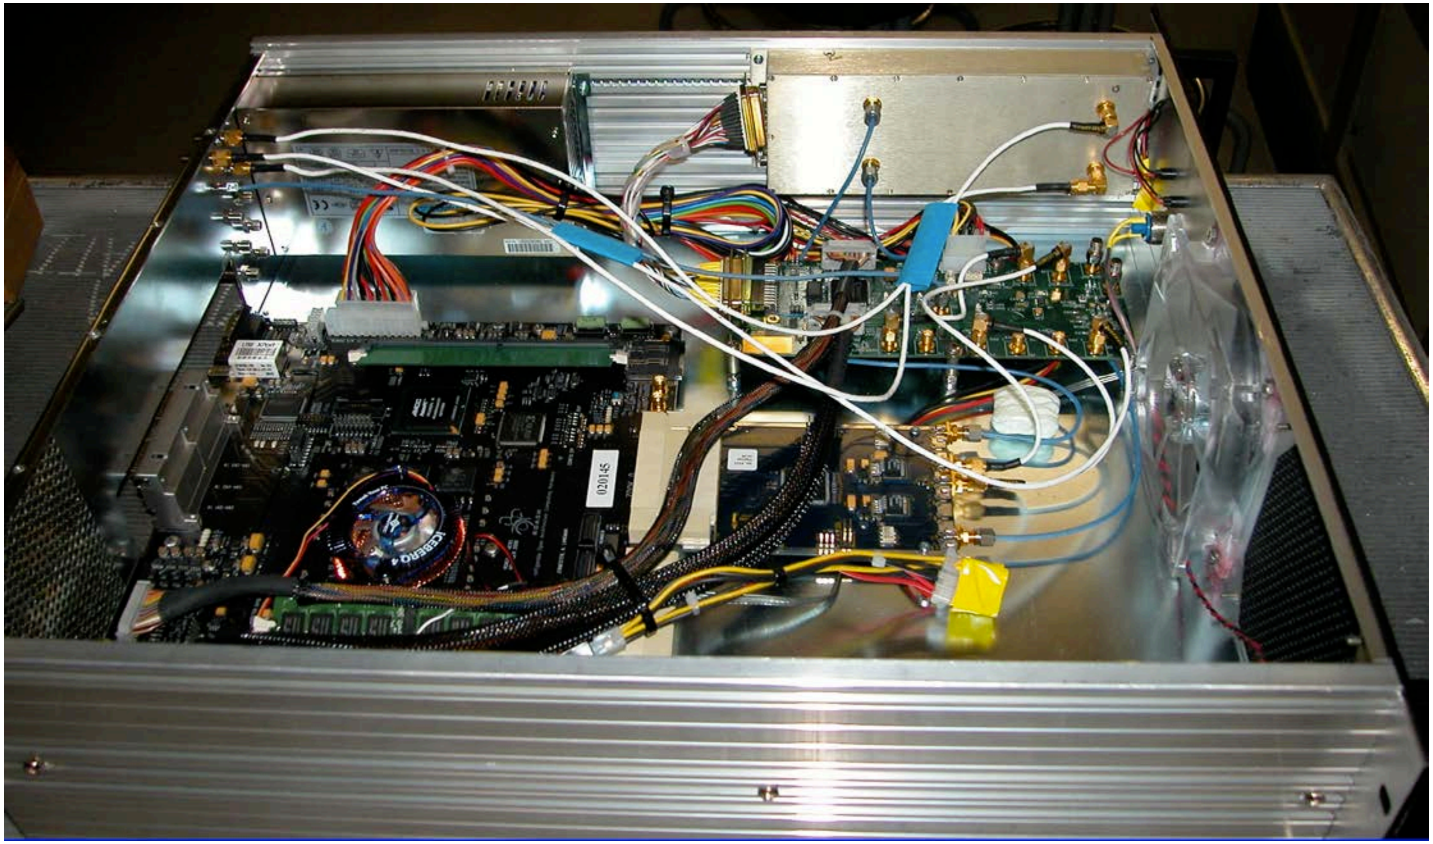

FLAG Beamformer

- CASPER – based beamformer for 38 element FLAG L-band (1.4 GHz) PAF
- Intended primarily for HI mapping, and pulsar searching
- 150 MHz bandwidth (set by OC48 SONET digital downlink)
- Real-time "Fast Radio Burst" output spigot
- 500 coarse channels, 150 MHz bandwidth
- 1600 fine channels, 15 MHz bandwidth
- Implemented using:
  - RF digitized at antennas, and transmitted digitally over fiber
  - 5 Roach IIs
  - 5 dual-GPU HPC Machines
  - Mellanox 40GbE network switch

NRAO

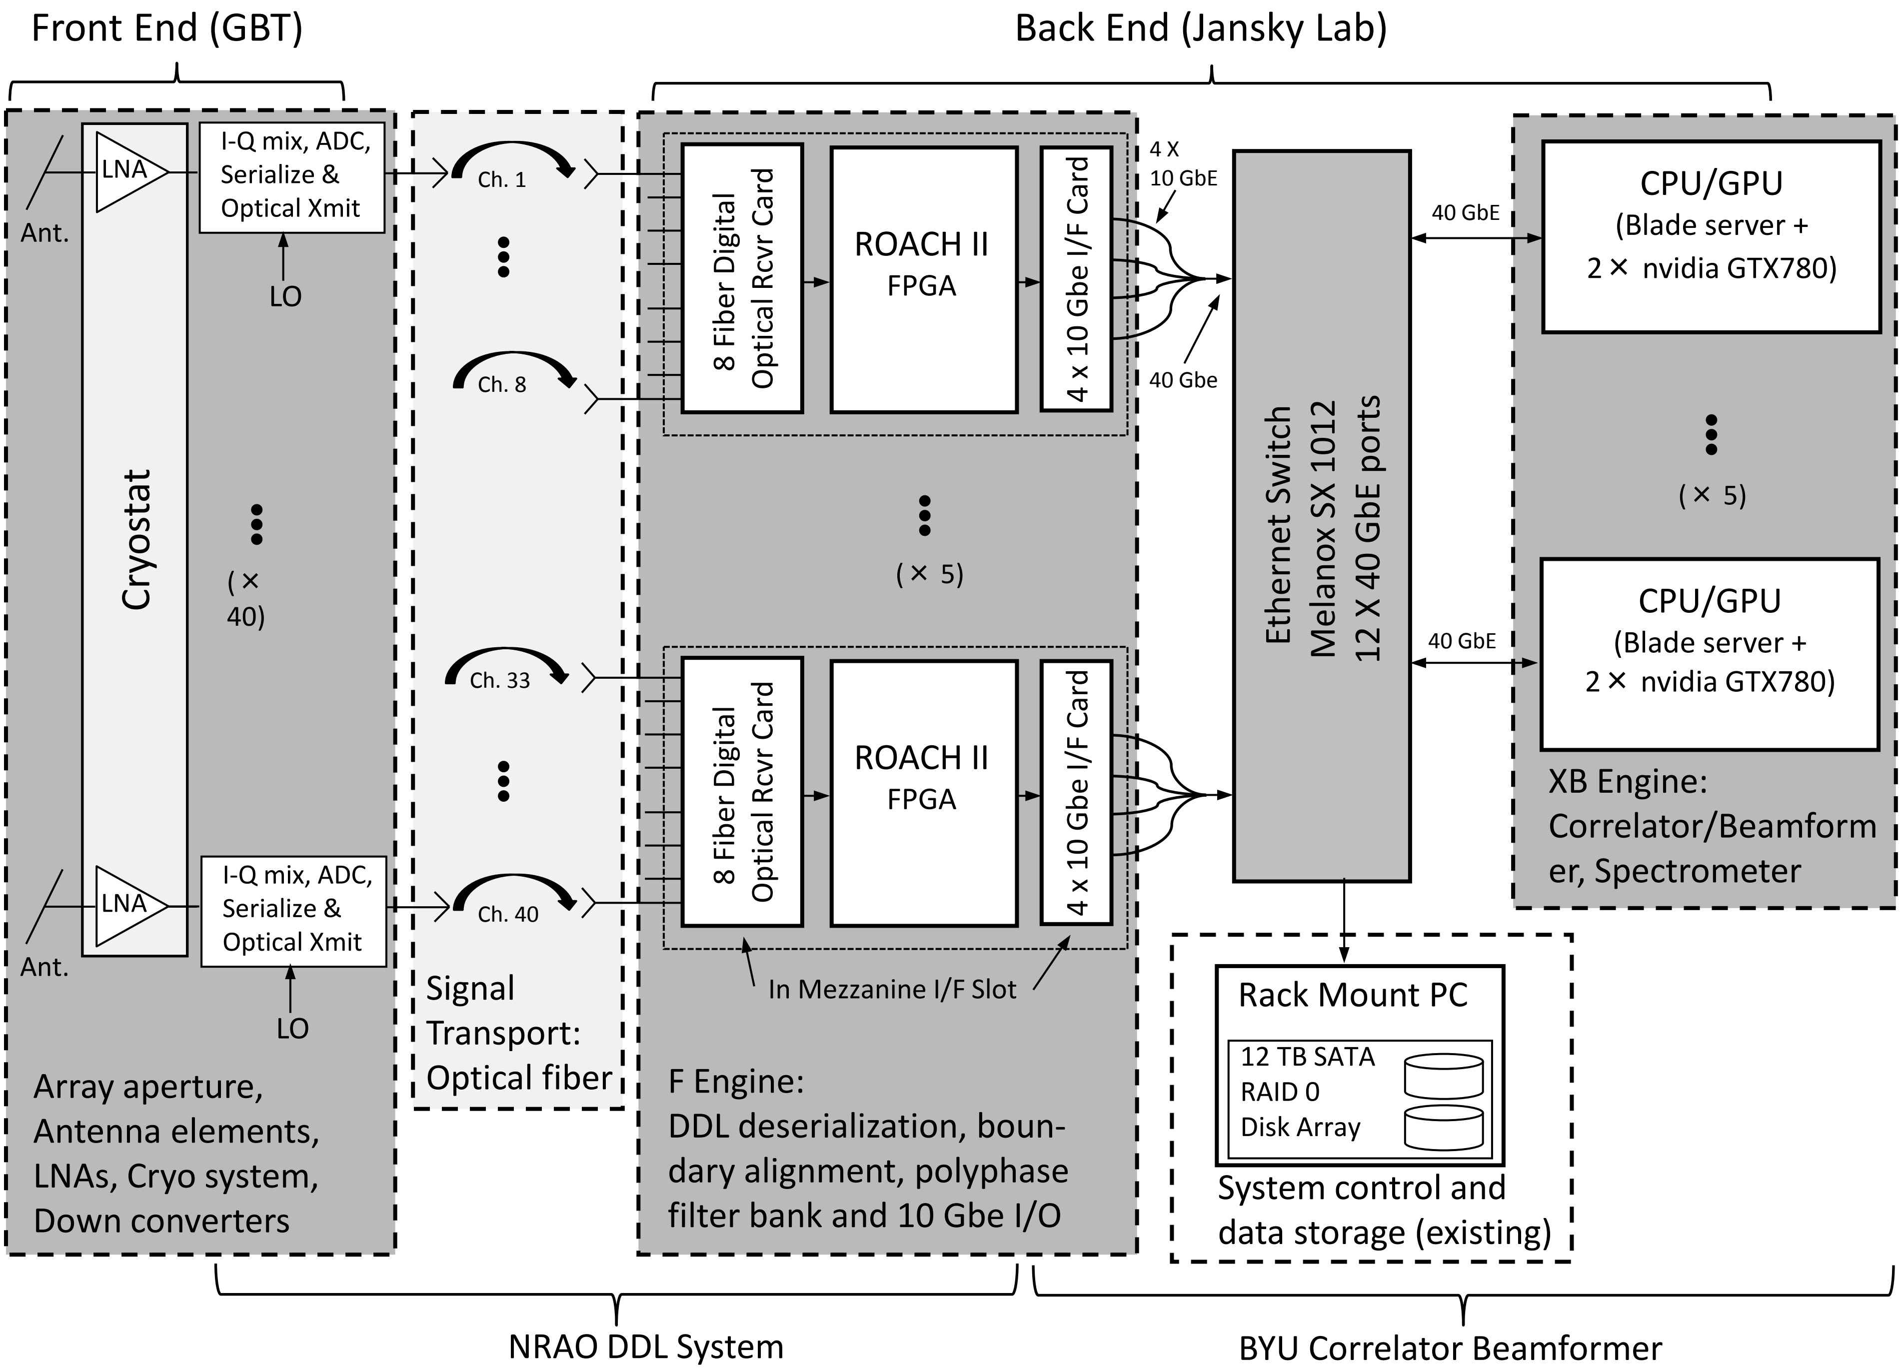
Front End (GBT)
Back End (Jansky Lab)
Ant.
LNA
Cryostat
I-Q mix, ADC, Serialize & Optical Xmit
LO
(× 40)
Ant.
LNA
I-Q mix, ADC, Serialize & Optical Xmit
LO
Array aperture, Antenna elements, LNAs, Cryo system, Down converters
Ch. 1
Ch. 8
Ch. 33
Ch. 40
Signal Transport: Optical fiber
8 Fiber Digital Optical Rcvr Card
ROACH II FPGA
4 x 10 Gbe I/F Card
(× 5)
8 Fiber Digital Optical Rcvr Card
ROACH II FPGA
4 x 10 Gbe I/F Card
In Mezzanine I/F Slot
F Engine: DDL deserialization, boundary alignment, polyphase filter bank and 10 Gbe I/O
4 X 10 GbE
40 Gbe
Ethernet Switch Melanox SX 1012 12 X 40 GbE ports
40 GbE
40 GbE
CPU/GPU (Blade server + 2 × nvidia GTX780)
(× 5)
CPU/GPU (Blade server + 2 × nvidia GTX780)
XB Engine: Correlator/Beamformer, Spectrometer
Rack Mount PC
12 TB SATA RAID 0 Disk Array
System control and data storage (existing)
NRAO DDL System
BYU Correlator Beamformer

Target Design: XB-engine Processing Functional Block Diagram (1 of 10 GPUs shown)

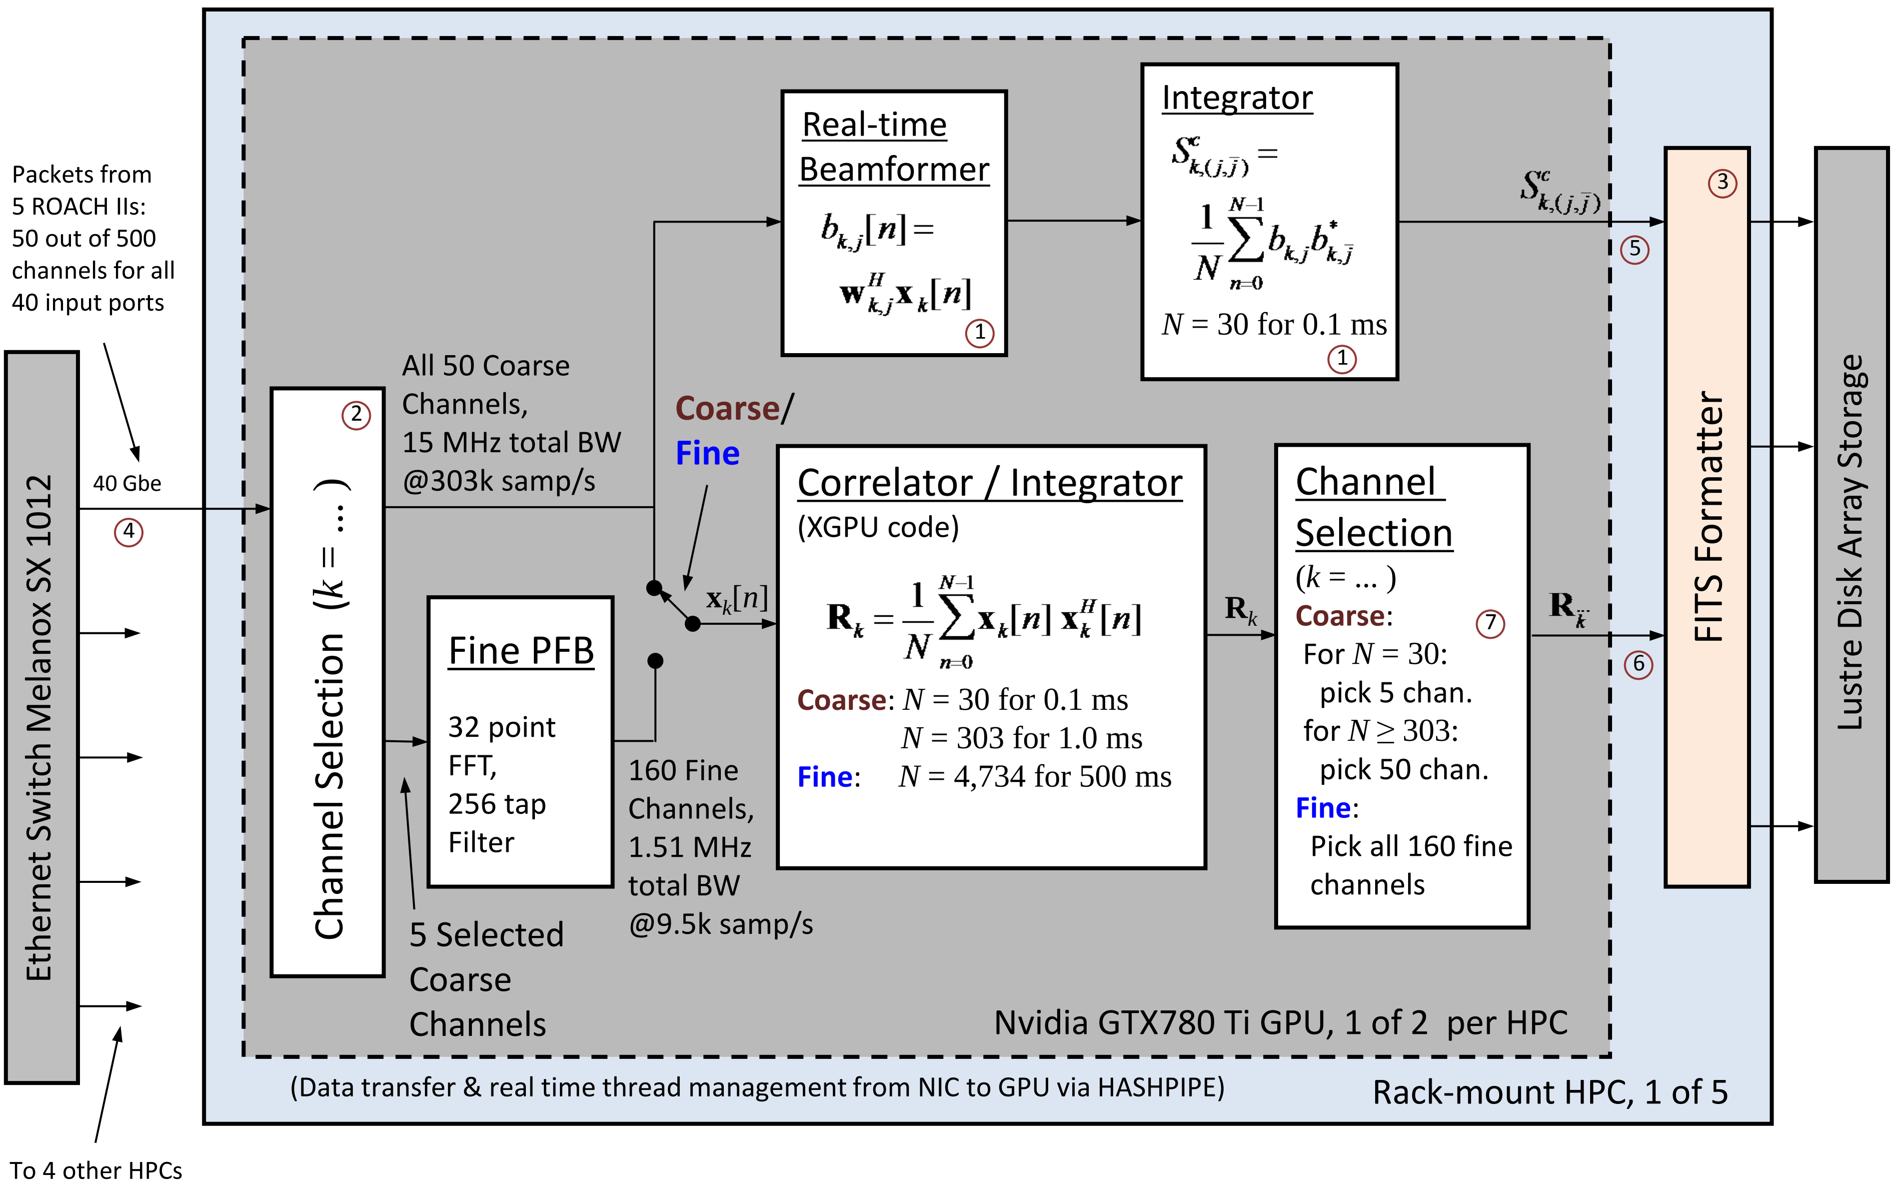

Legend: ① ... ⑦: See notes on following pages.

The End!