Digital Backend Systems

2019 GBO/AO Single Dish Workshop

Ryan Lynch (GBO)

GREEN BANK OBSERVATORY

Outline

- Analog to digital conversion
- ADC calibration
- Sampling theorem
- Total power detectors
- Spectrometers
- Fourier Transform
- Additional data processing
- VEGAS at the GBT
- Next generation systems

Why Digital Backends?

- Signal processing over wide bandwidths can be achieved with modern computers at relatively low cost
  - Some components may be custom, many can be commercial off the shelf (COS)
- Provide flexibility, portability of data products, long term archiving, etc.
- Can take advantage of improvements in computing power, storage, etc.
- But need to discretely sample continuous signals without loss of information

EM Radiation

- Can be defined by
  - Direction of propogation
  - Amplitude
  - Frequency
  - Phase
  - Polarization
- We want to keep track of as much of this information as we can

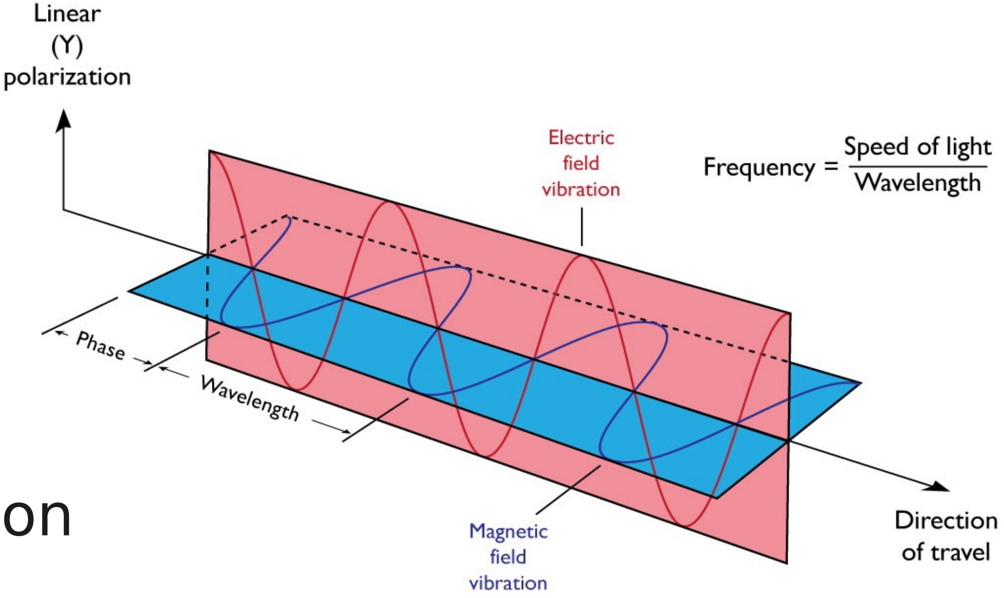

Continuous vs Discrete Signals

- Incoming radiation is a continuous change in electric field over a continuous range of frequencies
  - We do not typically think in terms of discrete photons in radio astronomy (with some exceptions)
- Digital systems work on discrete values that can be represented with some finite number of bits
- 1 bit = 2^1 values (0,1)
- 2 bit = 2^2 values (0 – 3)
- 8 bit = 2^8 = 256 values (0 – 255)

- The number of bits used to sample the signal defines the **dynamic range**
  - Smaller bit depth / resolution provides less granularity (1 bit = high or low)
  - Higher bit depth captures both weak and strong inputs
- This introduces some error, as perfect reconstruction is not possible with a finite number of bits
  - Bit depth chosen to keep quantization errors at or below an acceptable level

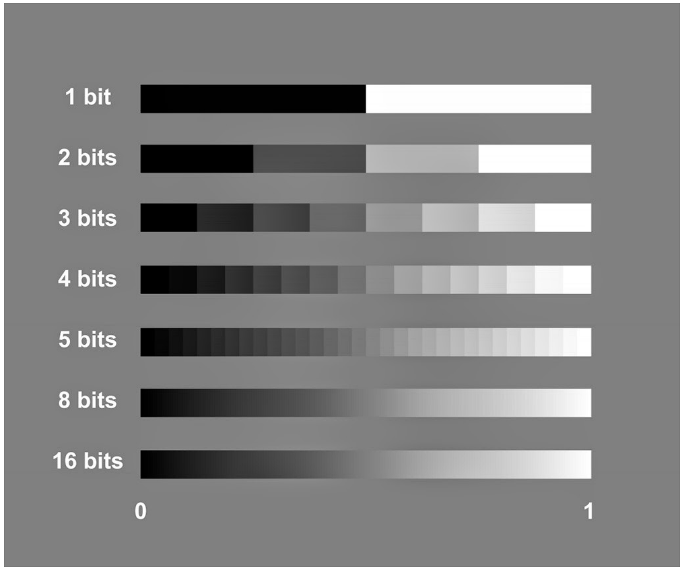

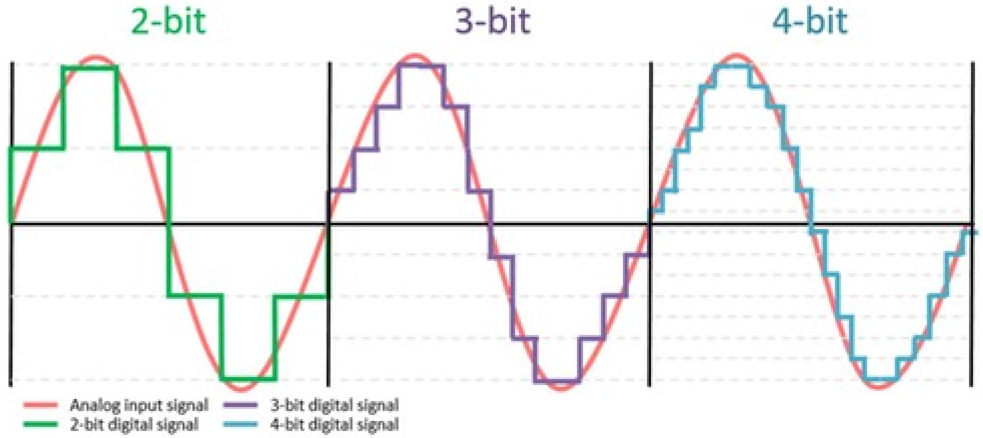

Analog to Digital Converters

- An **analog-to-digital converter** (ADC) is a device for converting continuous signal to discrete, digital signal
- ADCs use specialized circuits (**comparators**) to compare an input and reference voltage
  - Different types of ADCs exist with different schemes for using comparators
  - In simplest approach, 2^N comparators (N = number of bits) are used to successfully determine input level

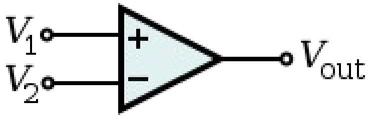

Analog to Digital Converters

- An **analog-to-digital converter** (ADC) is a device for converting continuous signal to discrete, digital signal
- Key features of ADCs are
  - Spurious free dynamic range (SFDR, related to resolution)
  - Sampling rate (determines bandwidth)
    - Why does sampling rate determine bandwidth??

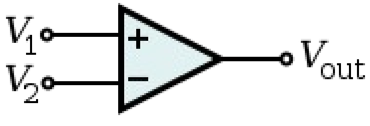

Nyquist-Shannon Sampling Theorem

- To perfectly reconstruct a time varying signal, we must sample at a critical rate, f_N, that is twice the highest frequency contained in the signal
  - A signal at a frequency $f > f_N$ will be **aliased** into our sampling band at a lower apparent frequency
- f_N is known as the **Nyquist frequency**

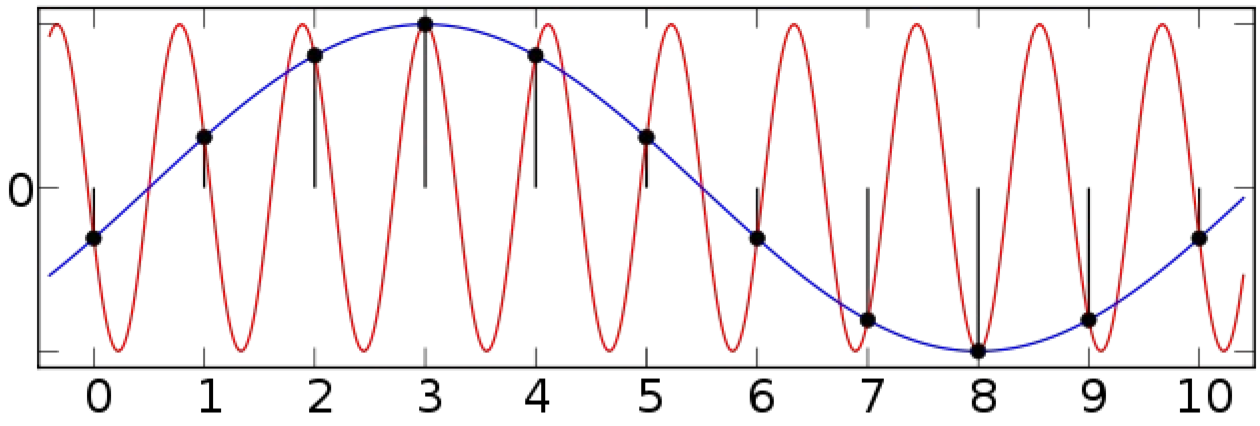

Nyquist-Shannon Sampling Theorem

- This is not just a time/frequency phenomenon
- Spatial variations can be decomposed into spatial frequencies
  - Sharp features contain higher frequency components (see Fourier transform later in talk)
- Nyquist sampling in spatial domain is important in mapping (see talks tomorrow by Larry Morgan and Dave Frayer)

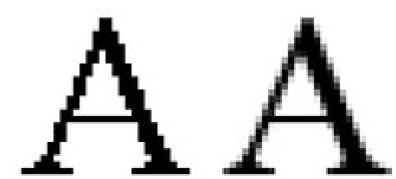

Image credit: Wikipedia

Sampling Rate and Bandwidth

- To avoid aliasing, we must apply an analog filter to suppress power outside some desired bandwidth B
- ADCs sample at a frequency $f_s = 2 \times B$
  - Example: We want to sample 800 MHz bandwidth
  - Downcovert to **baseband** and apply low-pass filter
  - Sample at 1.6 Gsps
- Remember: filters are not perfectly sharp
  - Filter roll-off needs to start below $f_s/2$ to ensure aliasing is kept below an acceptable level

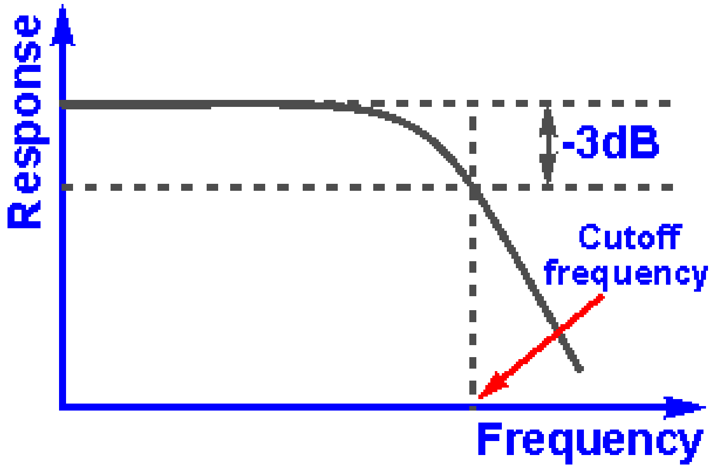

ADC Calibration

- To achieve high sampling rates, many ADCs are **interleaved**
  - Actually consist of several individual ADC **cores**, each clocked at a lower sampling rate than desired
  - Sampling is offset and a switch selects output from ADCs at appropriate times

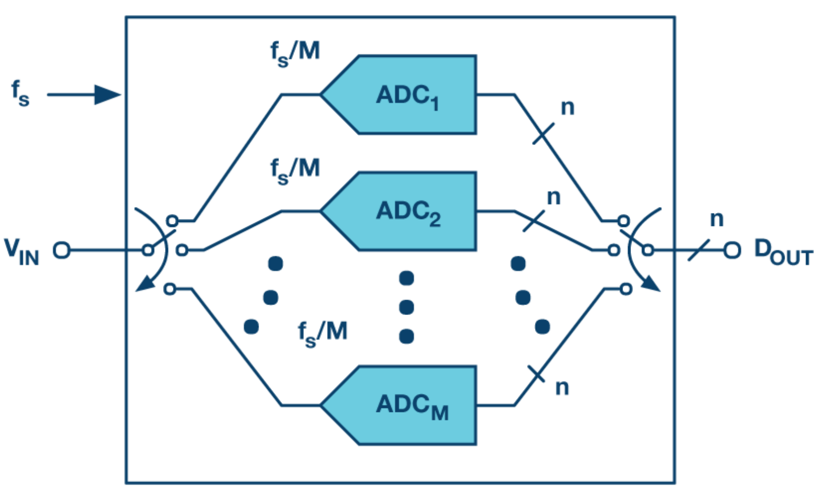

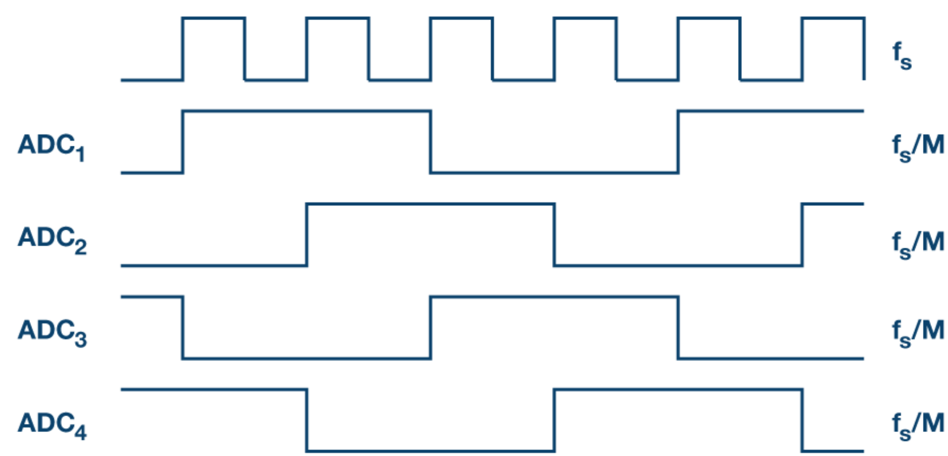

Image credit: analog.com

ADC Calibration

- Minor differences in the ADCs (e.g. gain, starting phase) leads to artifacts in the output data
- Proper calibration is important to keep spurious signals below an acceptable level

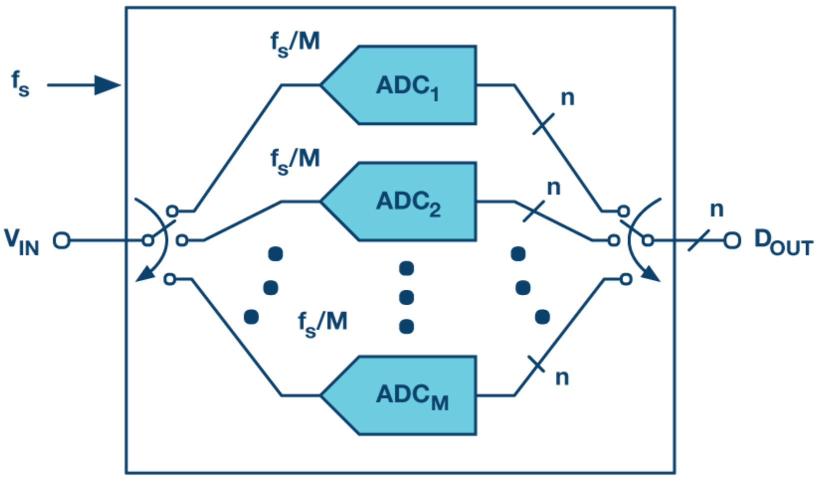

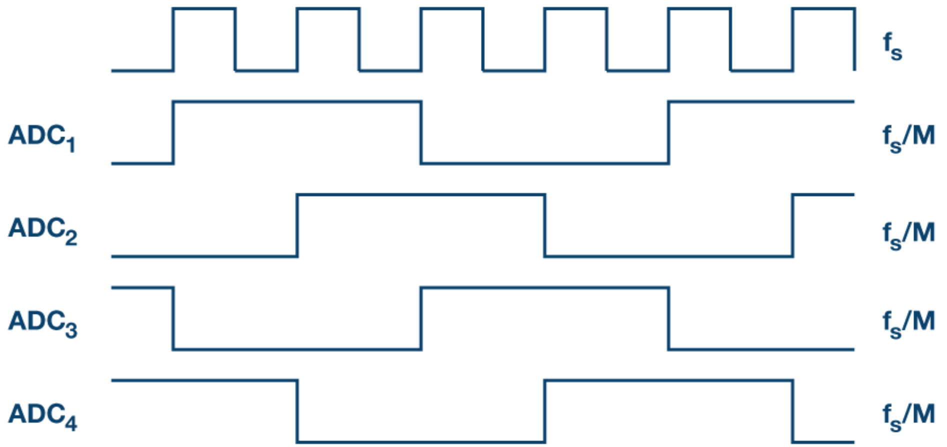

Image credit: analog.com

Dynamic Range for Wideband Systems

- Note that ADCs are **total power** devices
  - We have not yet sampled the power contributed at individual frequencies
- As the bandwidth goes up, so to do does the total power contributed by noise, RFI, and signal of interest
- Resolution / bit depth becomes increasingly important for wideband systems
  - Strong signals can push ADCs into non-linear regime
- Often use multiple digitizers to cover smaller portions of the total desired bandwidth
  - Analog filters may also be needed to remove portions of a band with strong RFI

Complex Voltage and Power

- Because incoming radiation is described by both an electric field amplitude and phase, it is convenient to represent it as a **phasor**
- This lends itself to using complex numbers to describe the voltage (recall Euler's formula)

$$A e^{i\theta x} = A[\cos(\theta x) + i\sin(\theta x)]$$

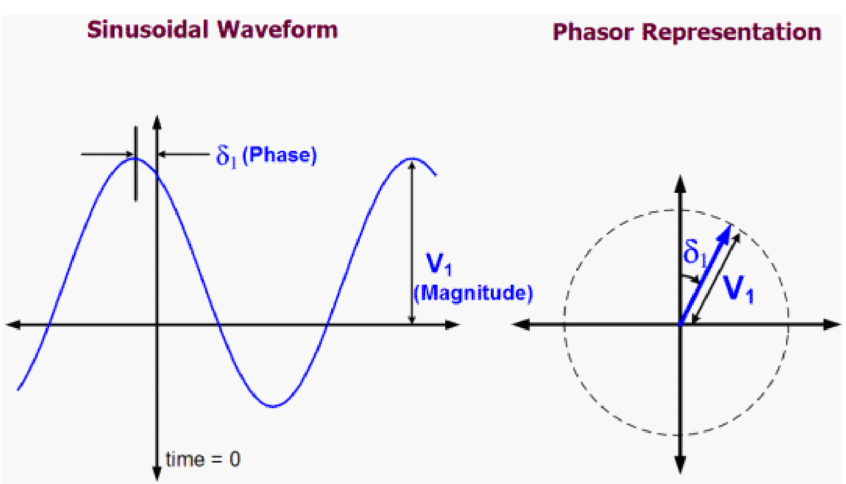

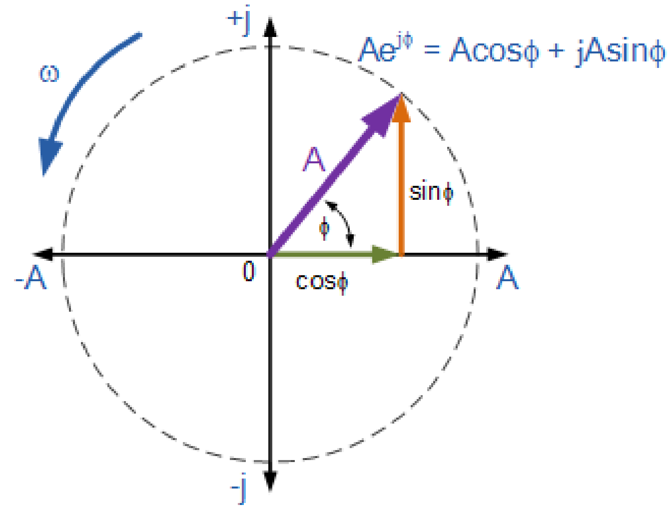

Complex Voltage and Power

- Note that when sampled at baseband the cosine and sine terms are often referred to as $I(t)$ and $Q(t)$ (i.e. I/Q values)
  - I corresponds to the real part of the complex voltage, and Q to the imaginary part
  - Don’t confuse these with the I and Q of Stokes parameters!

Complex Voltage and Power

- This allows us to represent the digitized signal with a real and imaginary part
  - Retains full amplitude and phase information so can be used for **coherent** processing
- In the final analysis we are usually interested in the **power** (which has non-zero mean), rather than the amplitude

$$P = \left| A e^{i \theta x} \right|^2$$

- This step is usually referred to as **detection**
  - If we sample two polarization states, we can form Stokes parameters or other polarization products prior to detection
  - Note that we lose phase information at this stage!

Spectrometers

- So far we have only been talking about time series data
- We are often interested in decomposing the time series into a **spectrum** that measures powers as a function of time *and radio frequency*
- Radio spectrometers do not rely on optical or electronic devices to spatially disperse different frequencies
  - Instead, take advantage of digital, phase coherent data and use Fourier transform

The Fourier Transform

- Recall that the ADC output are time samples of a band-limited signal containing power at many individual frequencies
- The (discrete) Fourier transform (DFT) is used to form a **power spectrum**, i.e. power measured at some discrete number of frequencies, or **channels**
  - Each discrete frequency is itself a measure of the total power within some finite channel bandwidth

The Fourier Transform

- Recall that the FT describes a function as a (finite, in our case) sum of sines and cosines

$$X_k = \sum_{n=0}^{N-1} x_n e^{2\pi i k n/N}$$

$$= \sum_{n=0}^{N-1} x_n [\cos(2\pi k n/N) + i \sin(2\pi k n/N)]$$

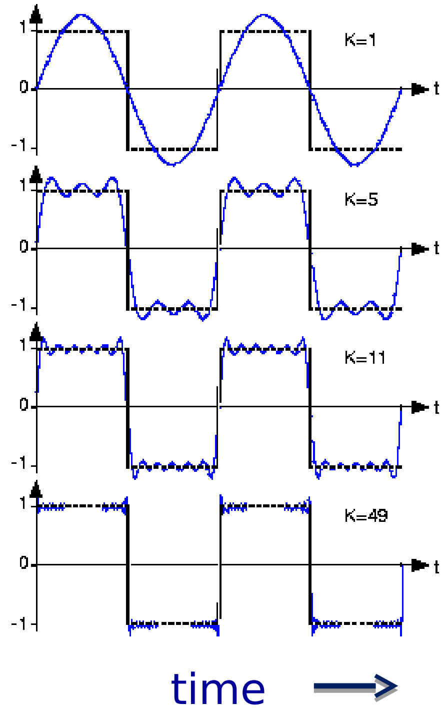

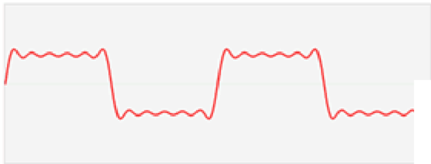

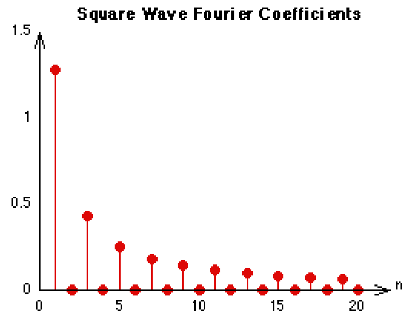

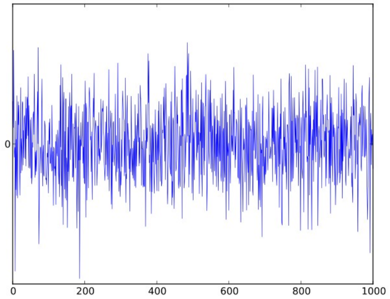

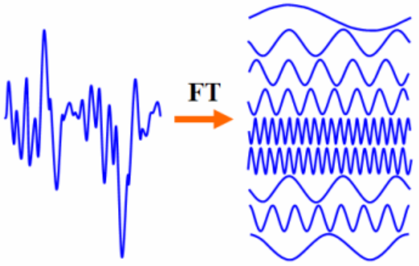

Weiner-Kinchin Theorem

- Relates the power spectrum to the **autocorrelation** of the incoming time series

$$S(f) = \int_{-\infty}^{\infty} r_{xx}(\tau) e^{-2\pi i f \tau} \, d\tau.$$

- r_{xx} is the autocorrelation, defined as

$$r_{xx} = \int_{-\infty}^{\infty} f(u) f^*(u-\tau) \, du$$

  - τ is known as the **lag**, and * denotes the complex conjugate
- In words, the power spectrum is the Fourier transform of the integral of the input signal multiplied point-wise by a time-delayed version of itself

Autocorrelation Spectrometer

- An **autocorrelation spectrometer** is highly flexible in terms of total bandwidth and channel bandwidth
  - The sampling interval $\Delta\tau$ and total number of lags *N* completely determine these parameters

$$B = \frac{1}{2\Delta\tau}$$

$$\Delta f = 1.2\frac{B}{N}$$

    The factor of 1.2 comes from the **windowing function**, which is simply a hard cutoff at t > Dt*N* (i.e. w(t) = 1 for t <= DtN, else 0)

  - The observed power spectrum is a convolution of the true spectrum with the Fourier transform of w

$$\tilde{S}(f) = S(f) \circ W(f)$$

Autocorrelation Spectrometer

- Because the Fourier transform of a top-hat is a sinc function, the channel shape of an ACS is itself a sinc, defined by it's FWHM
  - This is where the factor of 1.2 comes from
- While an ACS is flexible and easy to implement, this frequency response is undesirable
  - Power can leak into adjacent channels
  - For very strong signals, leakage can impact significant part of band
- Can we do better?
  - Yes!

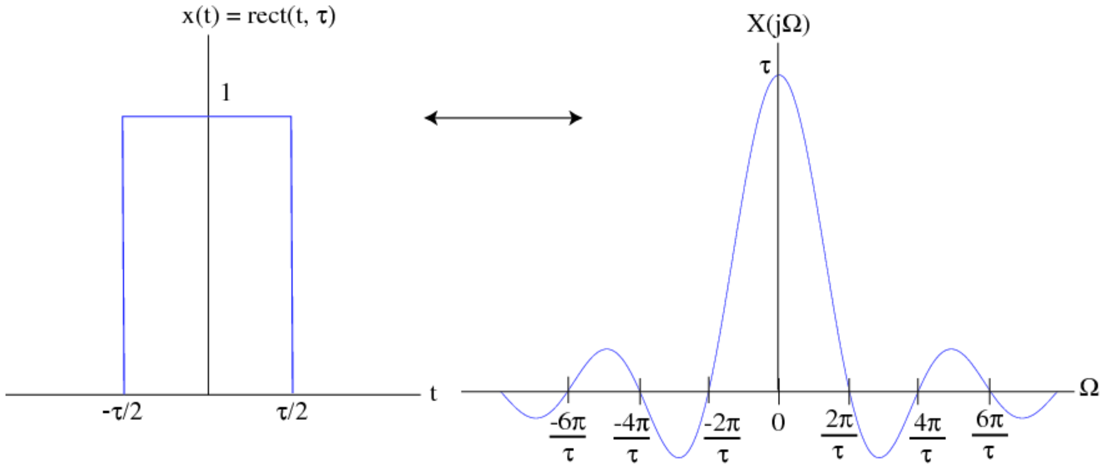

Polyphase Filterbank

- In a direct DFT we start with a rectangular windowing function (in time) and end with a sinc response (in frequency)
- We prefer to have a rectangular (i.e. flat) response in frequency across a channel
  - Use the Fourier inverse as the time-domain window, i.e. a sinc filter
- In practice, to obtain an N-point spectrum, use $M = N \times P$ points
  - P is the number of phases in the **polyphase** filterbank, also referred to as the number of taps

Polyphase Filterbank

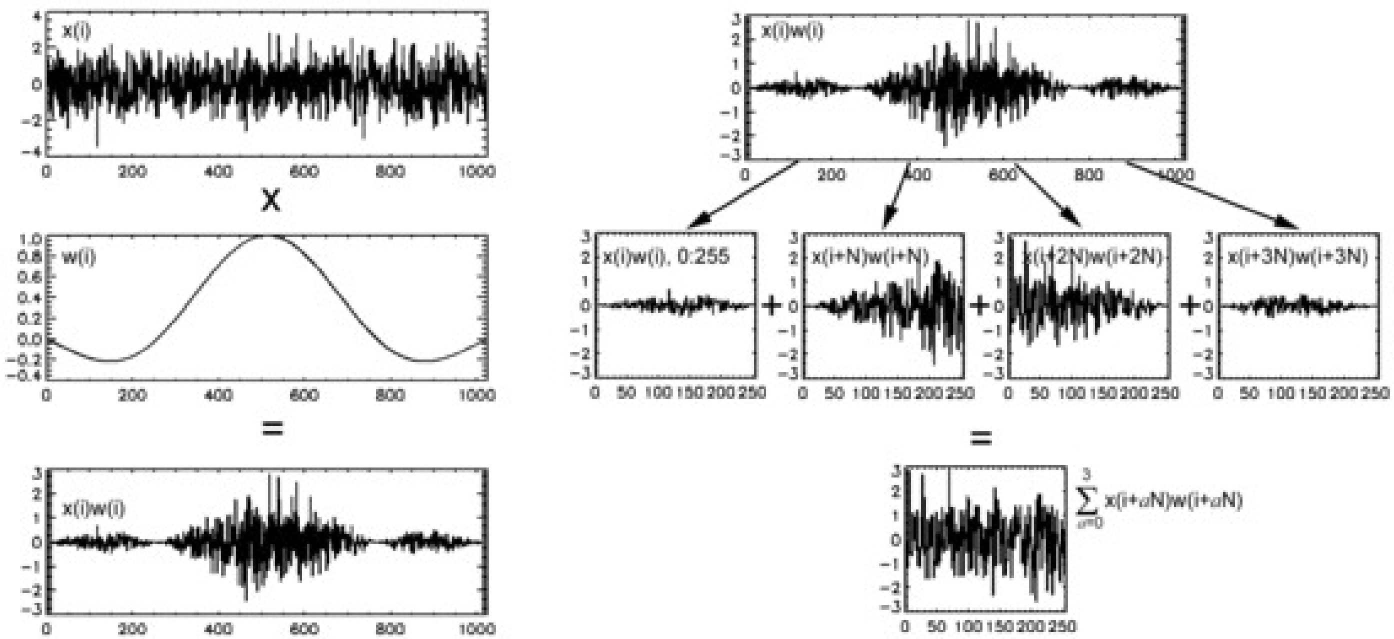

- After multiplication by an *M*-point filter, each phase is added to produce an *N*-point input to the DFT
- The DFT can now be taken, the result squared, and then accumulated to produce a power spectrum

Image credit: Dale Gary

Polyphase Filterbank

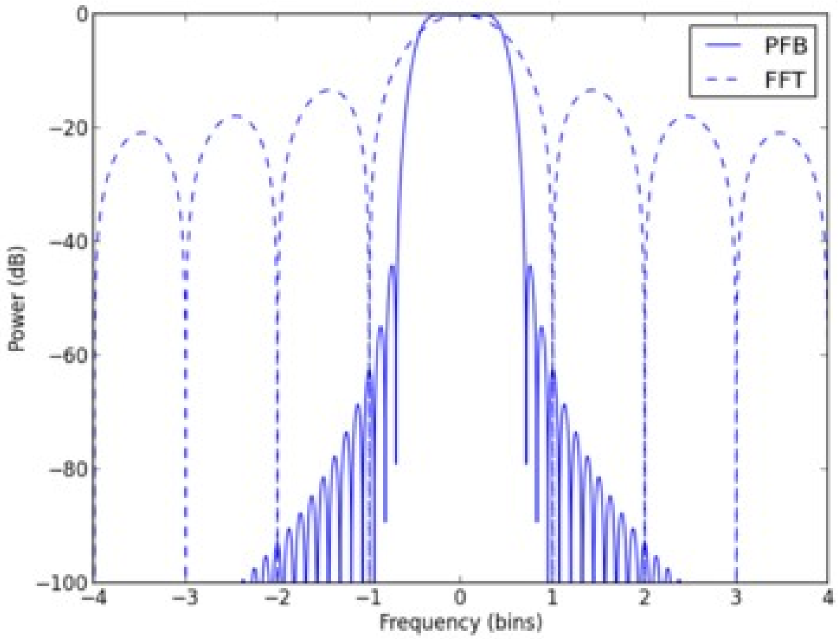

Image credit: Jayanth Chennamangalam

- Caveats
  - In pratice, the sinc window must be truncated so the frequency response is not perfectly flat
  - We typically multiply the sinc window by an finite impulse response (FIR) filter to improve frequency response
  - Using more taps also improves response

- PFB is more computationally intensive (~1.5x) than direct DFT but improved spectral response is usually worth the trade-off

Astronomical Spectrometers

- Note that the frequency resolution we obtain is determined by the number of points in the FFT
  - The sampling theorem is also relevant here: we need $2N$ time samples for N frequency channels
- This creates an inverse relationship between time and frequency resolution
- In typical spectral line observing, we are more concerned with frequency resolution than time resolution
- In pulsar observing we are usually more concerned with time resolution that frequency resolution

Astronomical Spectrometers

- The last* step is typically to detect and accumulate power spectra for some integration time
  - The choice of integration time depends on the stability of the instrument and scientific goals
  - Typically use ~0.1 – 10 s for spectral line observing to allow efficient excising of RFI
  - Typically use 10s μs in pulsar observing to retain sensitivity to fast pulsars

*Additiona signal processing often performed in pulsar observing (e.g. dedispersion, folding)

Polarization Products

- Most receivers sample two polarization states (typically linear [X/Y] or circular [L/R])
- Everything described above must be duplicated for each polarization channel
  - 2x ADCs, 2x spectrometer engines
- The polarization products that one records depends on science goals
  - Typically sufficient to record each channel's self-products independently (e.g. $|X|^2$ and $|Y|^2$)
- For strongly polarized sources, typically record Stokes parameters or self and cross terms

Polarization Products

- Stokes parameters allow complete recovery of polarized signal
  - For a linear basis:

    $I = |X|^2 + |Y|^2$ (total intensity)
    $Q = |X|^2 - |Y|^2$
    $U = 2\,\mathrm{Re}(X^*\,Y)$
    $V = 2\,\mathrm{Im}(X^*\,Y)$
  - $|V|$ = circular polarization
    $|L| = \sqrt{(Q^2 + U^2)}$ = linear polarization

- We may also record the self and cross terms directly, [i.e. $|A|^2$, $|B|^2$, $\mathrm{Re}(A^*\,B)$, $\mathrm{Im}(A^*\,B)$]

RFI Mitigation

- RFI is to radio astronomers as light pollution is to optical astronomers
- RFI almost only get’s worse with time, even in radio quiet zone

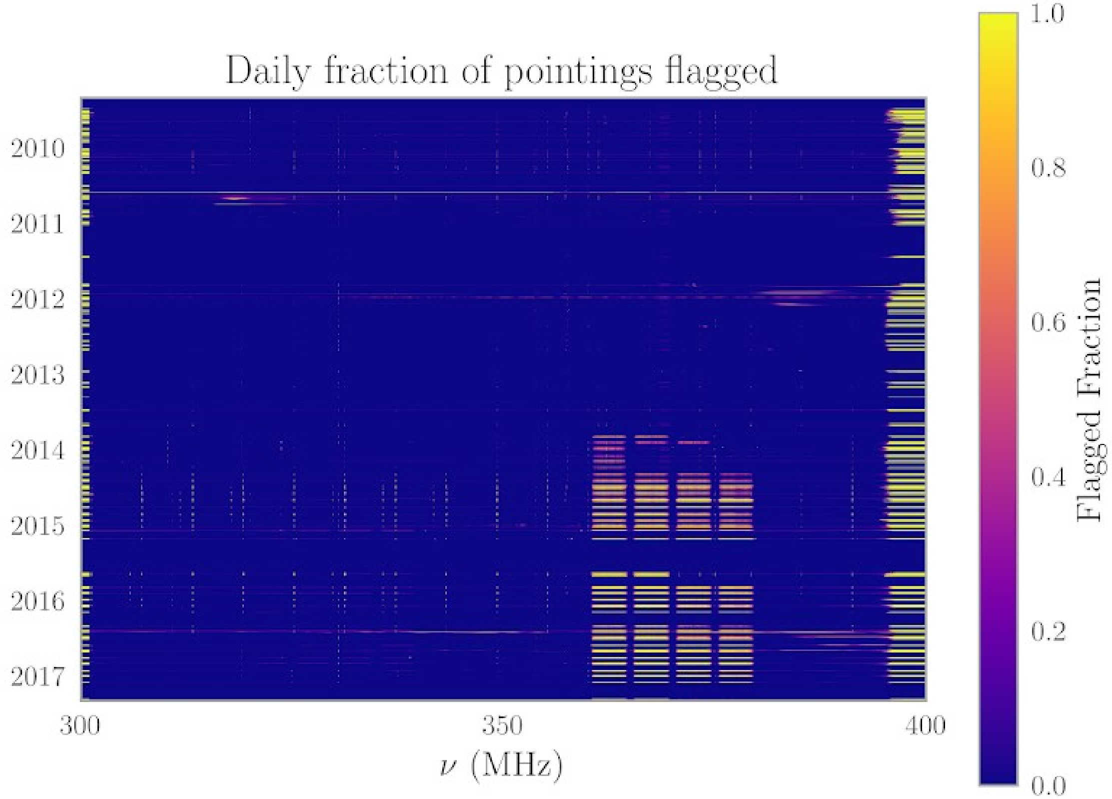

Image credit: Will Fiore (WVU)

RFI Mitigation

- RF techniques for RFI mitigation are “notch” filters that remove affected band
  - Degrades Tsys but may be necessary for strong, persistent RFI
- Digital techniques can be passive or active
- Passive techniques
  - Flag/mask small numbers of channels/integrations from downstream processing
  - Preserves original data at expense of losing all information in a flagged channe/integration
  - Adds a (potentially expensive) step to post processing

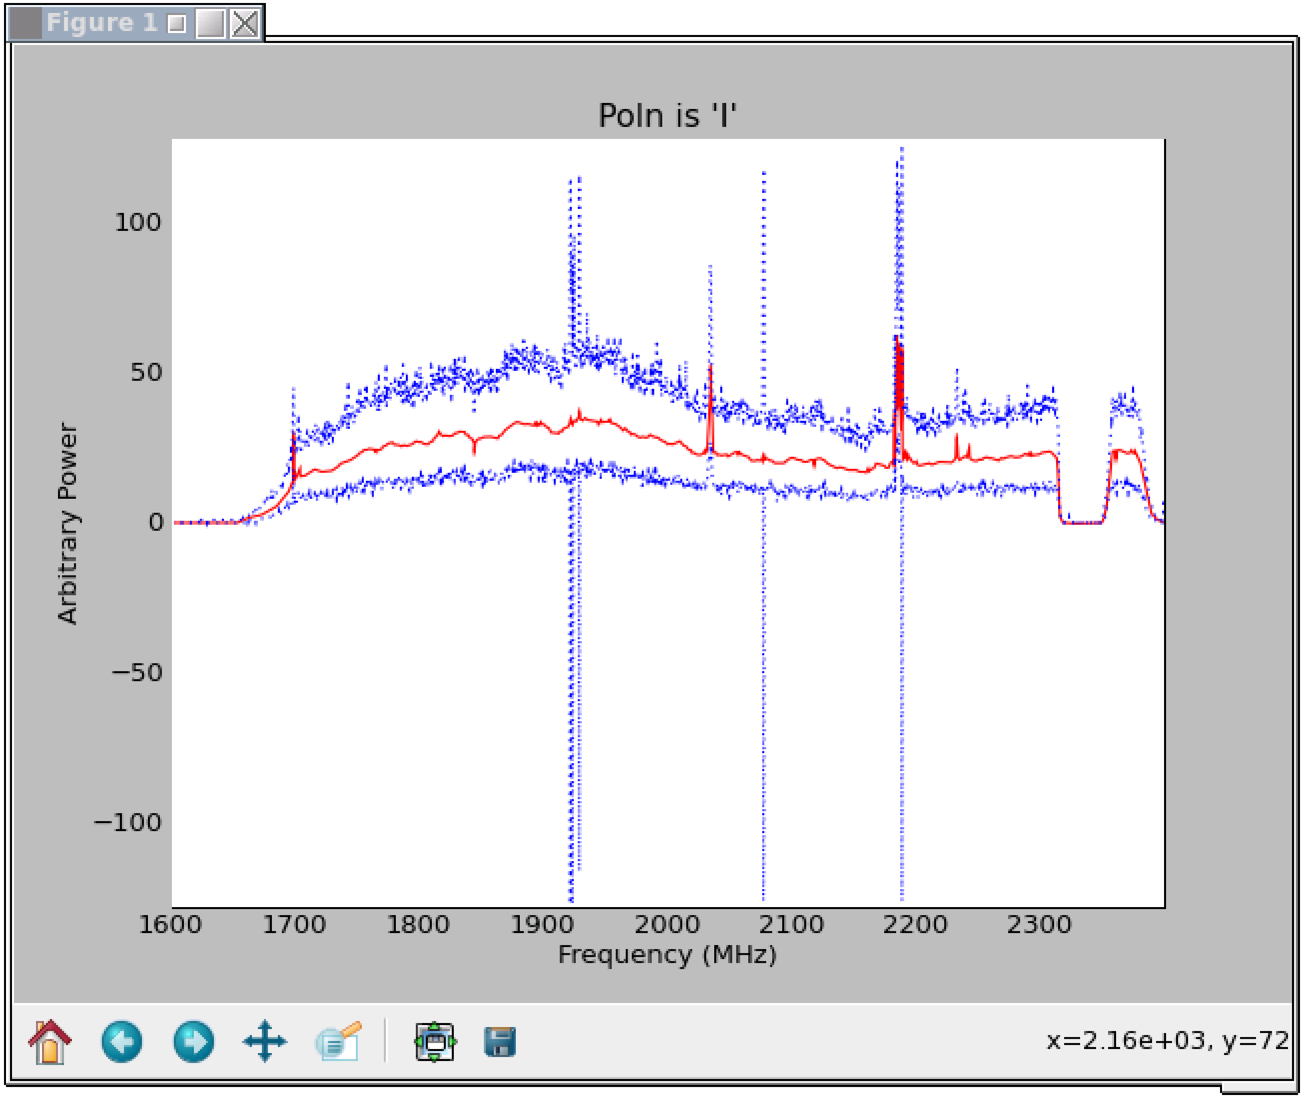

Figure 1
Poln is 'I'
100
50
0
−50
−100
Arbitrary Power
1600
1700
1800
1900
2000
2100
2200
2300
Frequency (MHz)
x=2.16e+03, y=72

RFI Mitigation

- Active mitigation may include
  - Subtraction of reference antenna signal
    - Complicated by differences in gain, beam shape, etc.
    - Must be done on complex data
  - Statistical flagging/replacement pre-detection
    - Look for statistical outliers in voltage data
    - Replace with zeros, Gaussian noise with same statistics as unaffected data, etc.
- Removes RFI closer to the source
- Alters original data in unrecoverable way (unless a second copy is made)
- Statistical flagging/replacement being investigated at GBO

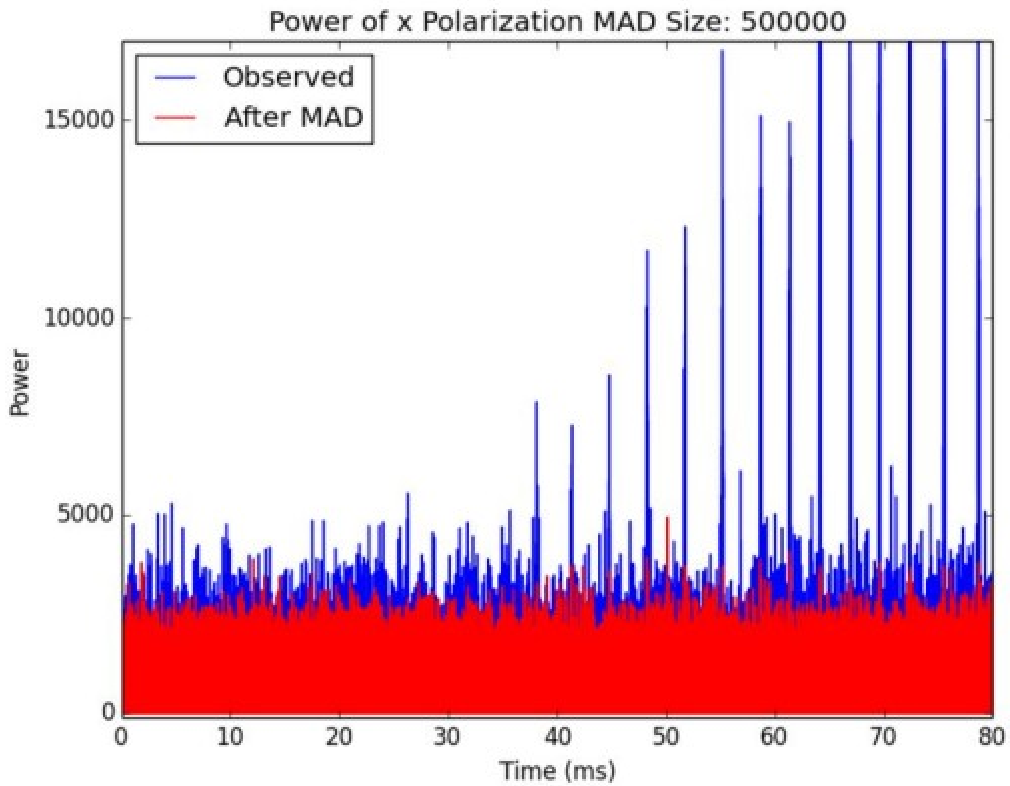
Power of x Polarization MAD Size: 500000
Observed
After MAD
Power
Time (ms)
0
5000
10000
15000
0
10
20
30
40
50
60
70
80

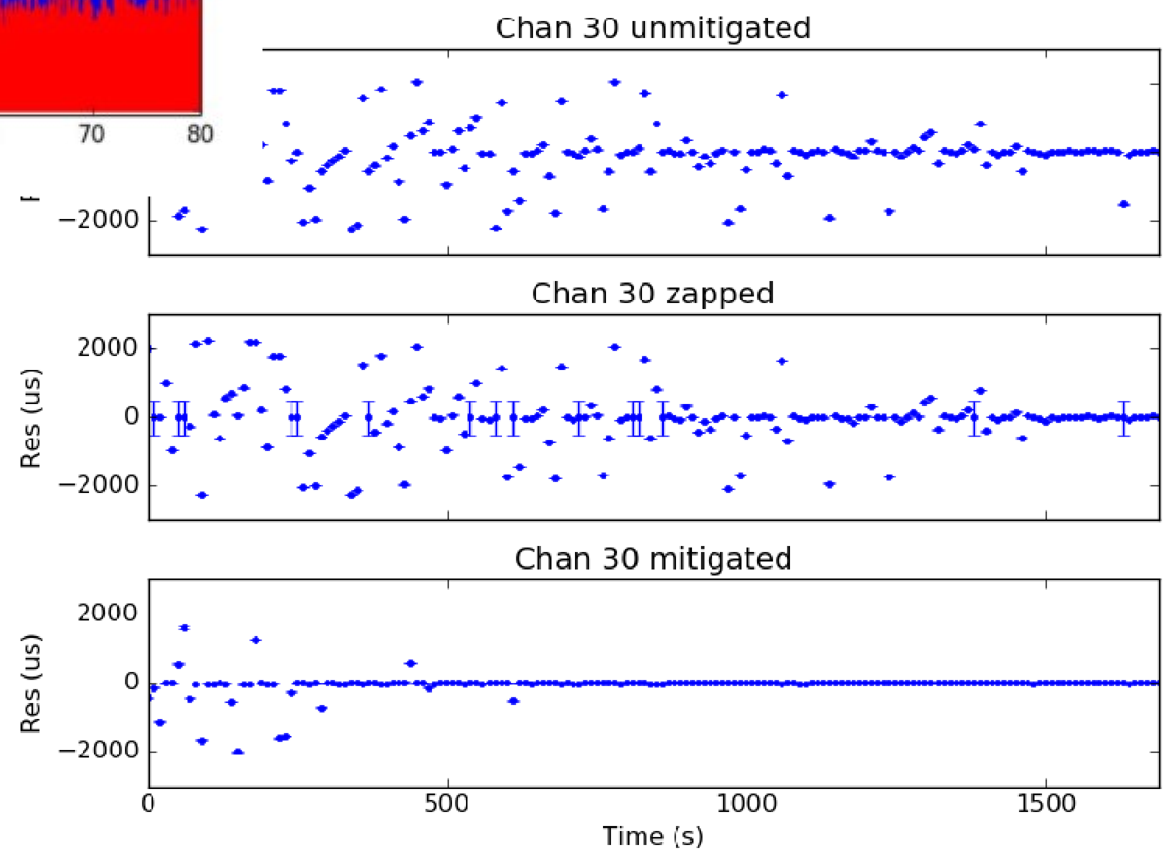
Chan 30 unmitigated
−2000
Chan 30 zapped
Res (us)
2000
0
−2000
Chan 30 mitigated
Res (us)
2000
0
−2000
0
500
1000
1500
Time (s)

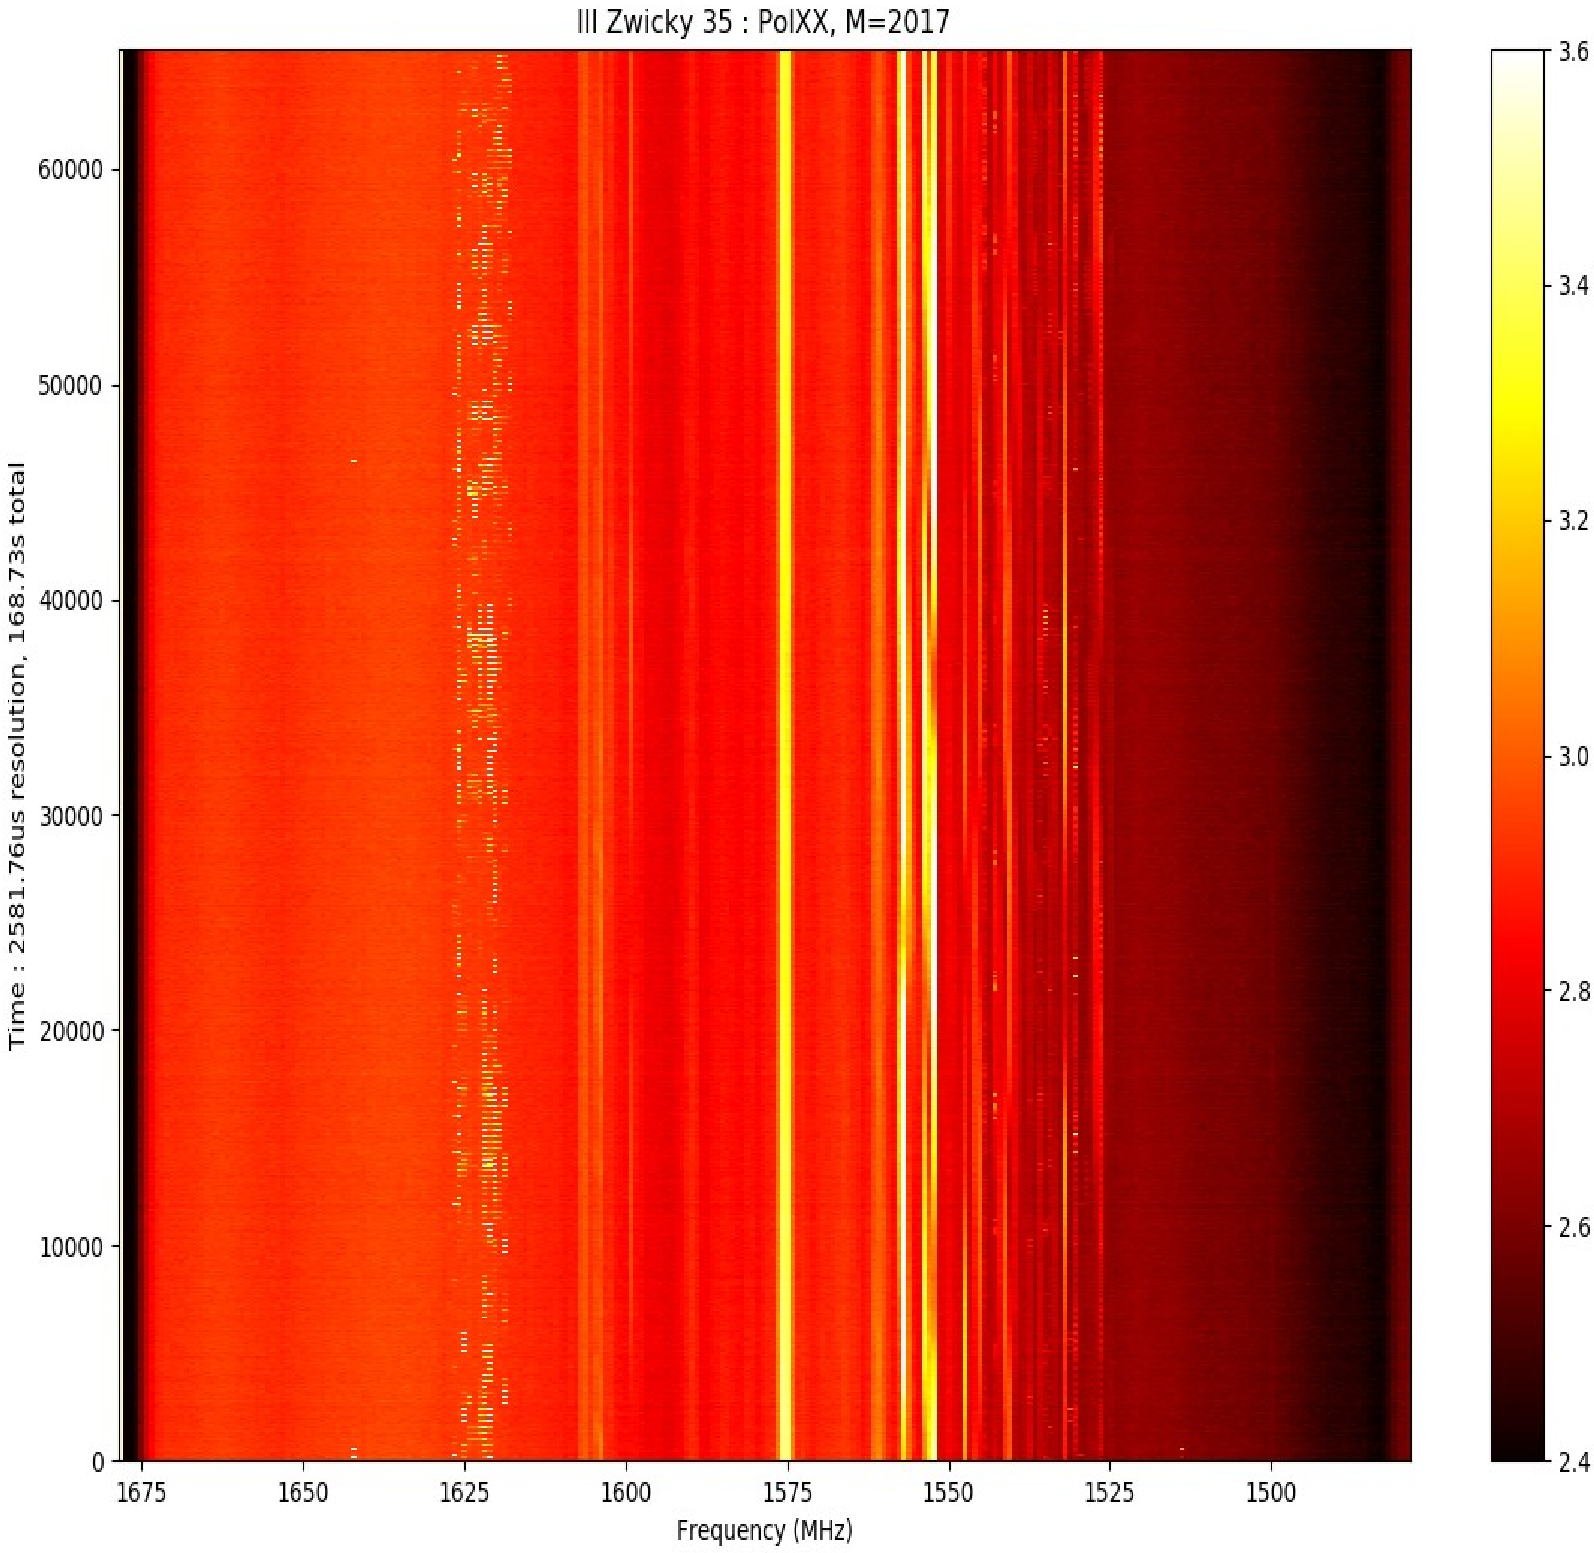

Observation of OH megamaser in III Zwicky 35

Signal at 1622MHz, underneath Iridium Satellite RFI

---

200 MHz bandwidth centered at 1581.5 MHz, 256 channels

One GUPPI raw file pictured, with 2017 spectra averaged at a time

Leading to 2581μs time resolution on the plot and 168.73s / 2.81 min total

Slide by Evan Smith (WVU)

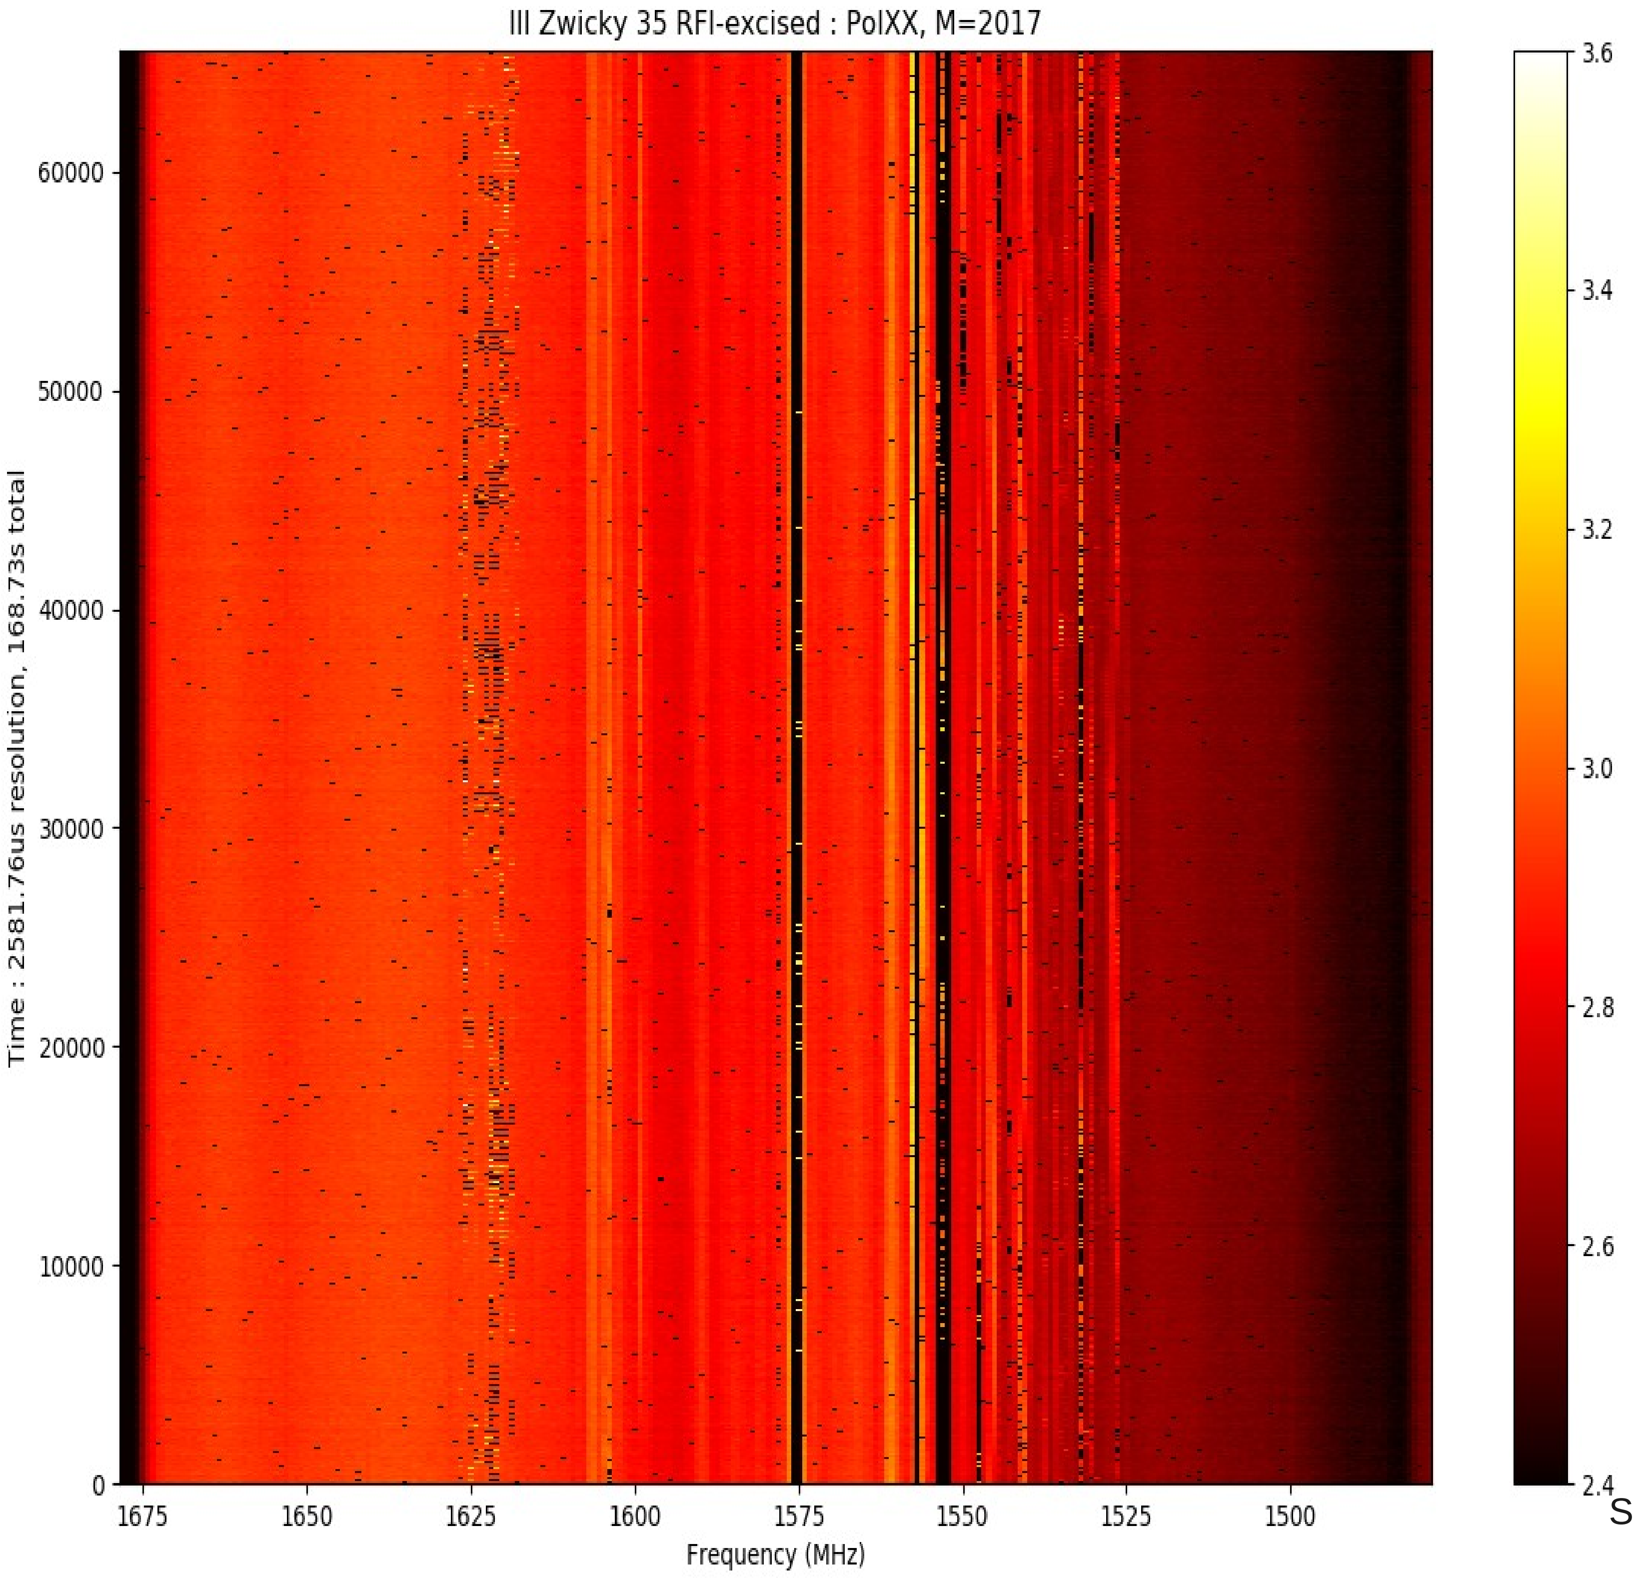

Observation of OH megamaser in III Zwicky 35

Spectral Kurtosis excision applied

---

200 MHz bandwidth centered at 1581.5 MHz, 256 channels

One GUPPI raw file pictured, with 2017 spectra averaged at a time

Leading to 2581μs time resolution on the plot and 168.73s / 2.81 min total

Slide by Evan Smith (WVU)

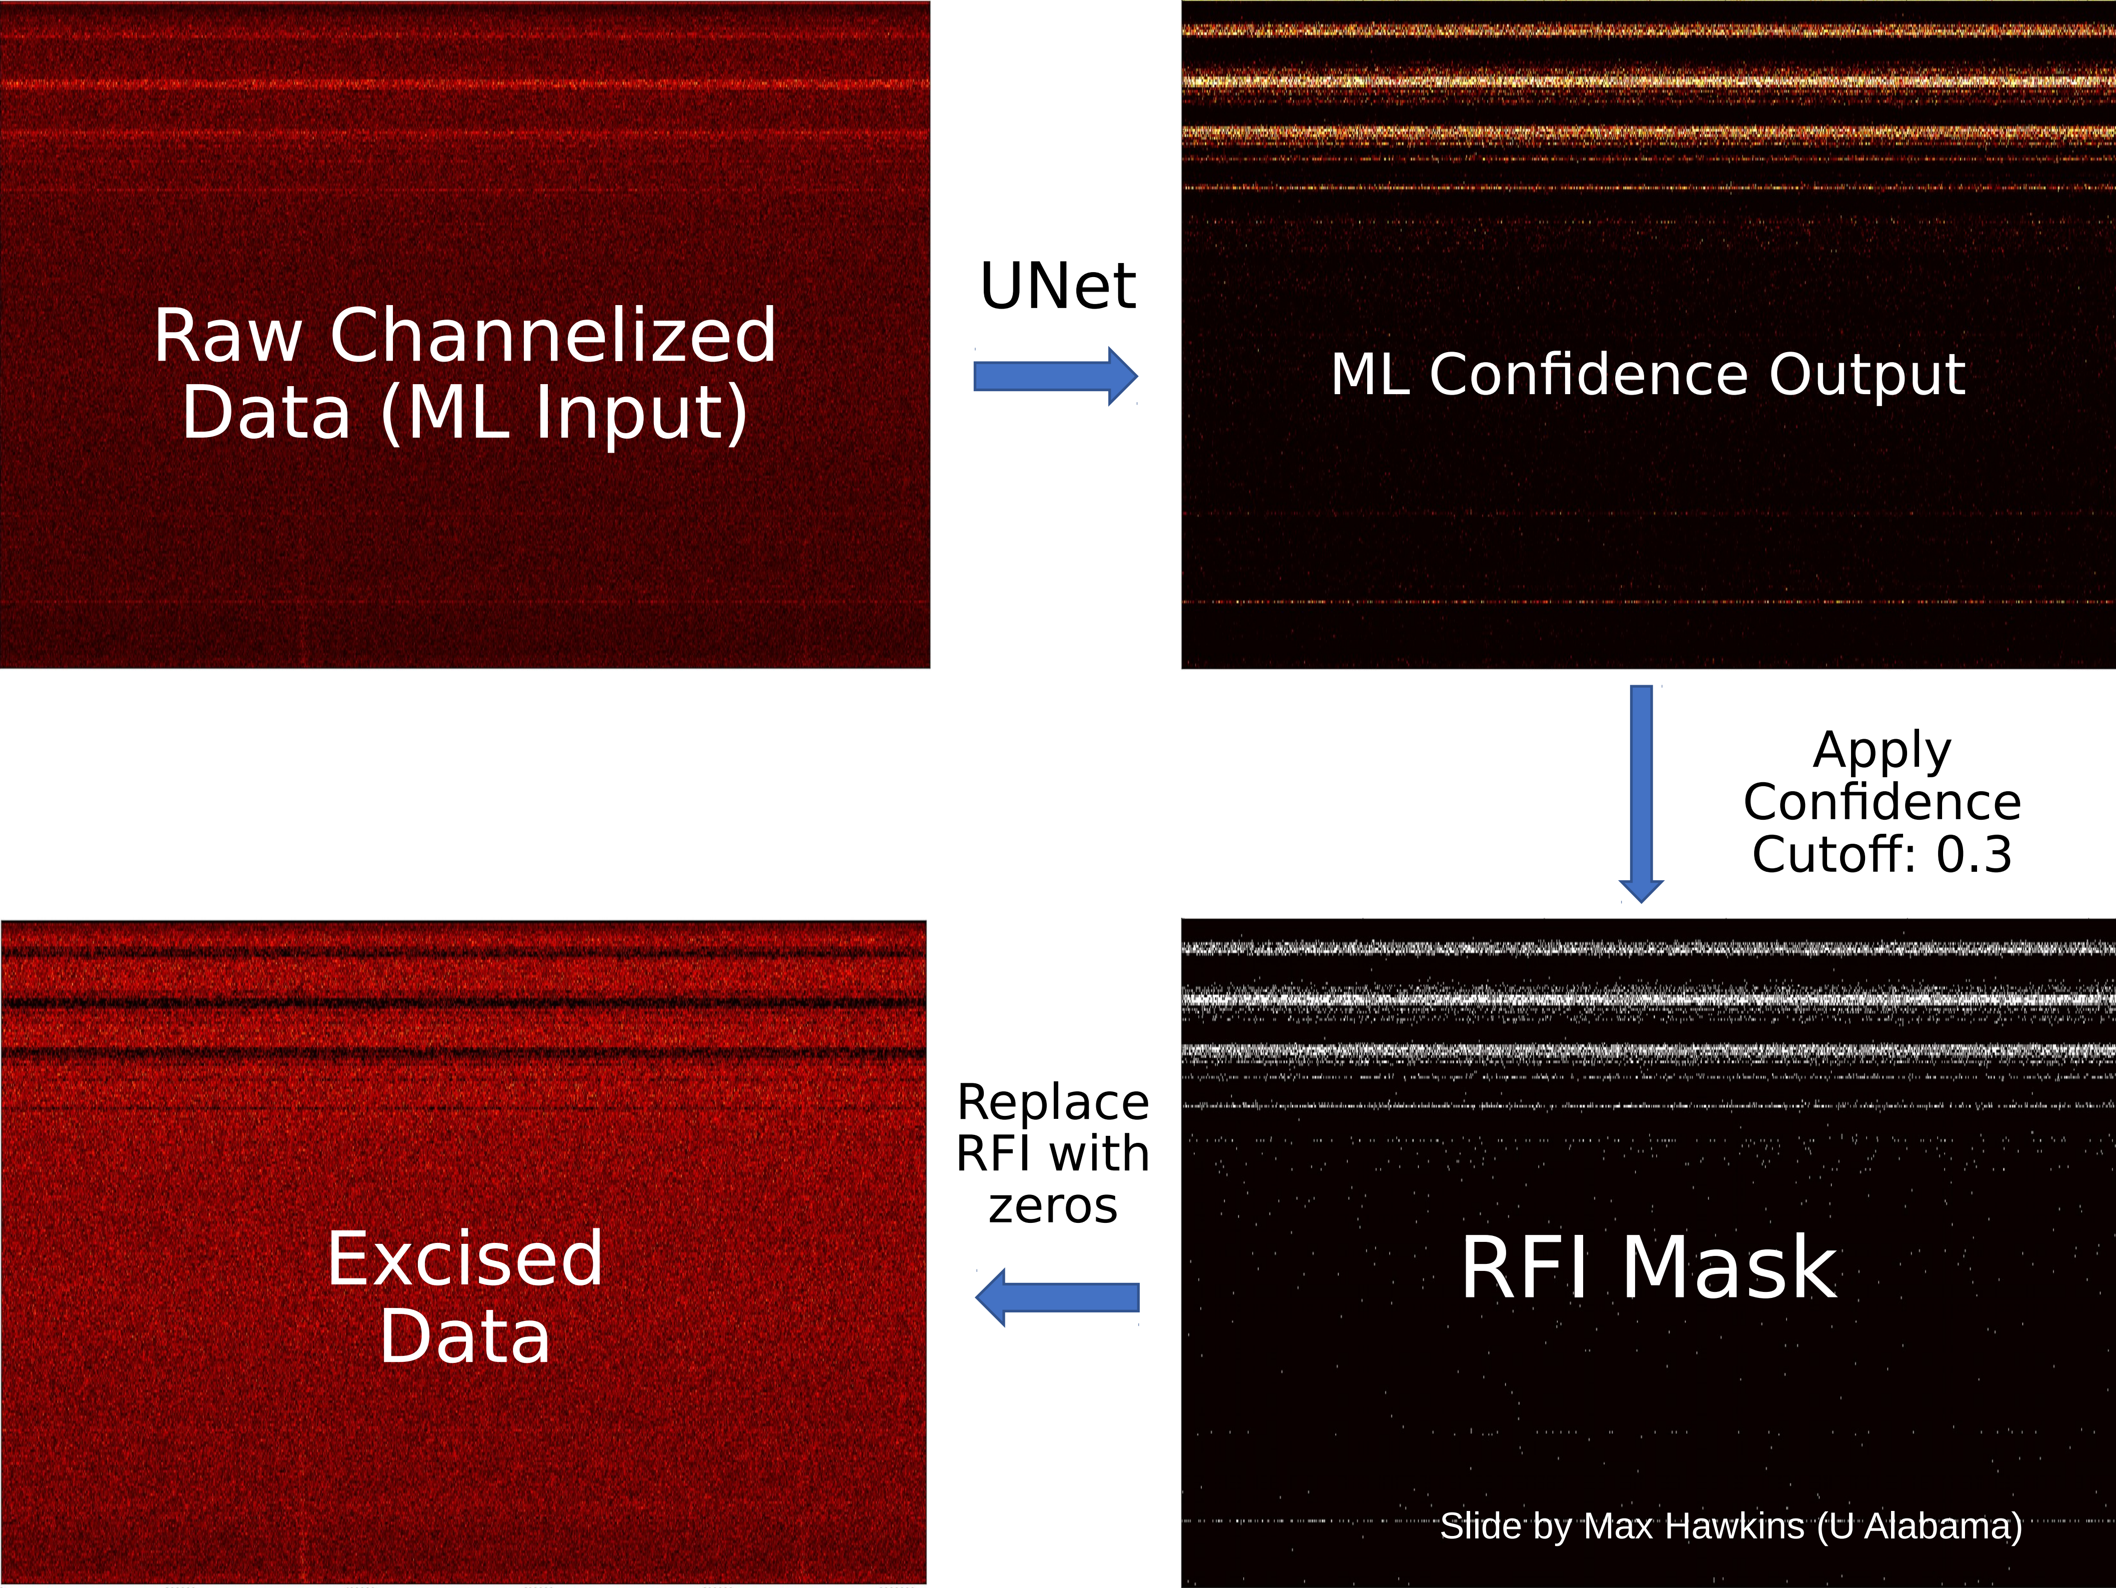
Raw Channelized Data (ML Input)
UNet
ML Confidence Output
Apply Confidence Cutoff: 0.3
RFI Mask
Replace RFI with zeros
Excised Data
Slide by Max Hawkins (U Alabama)

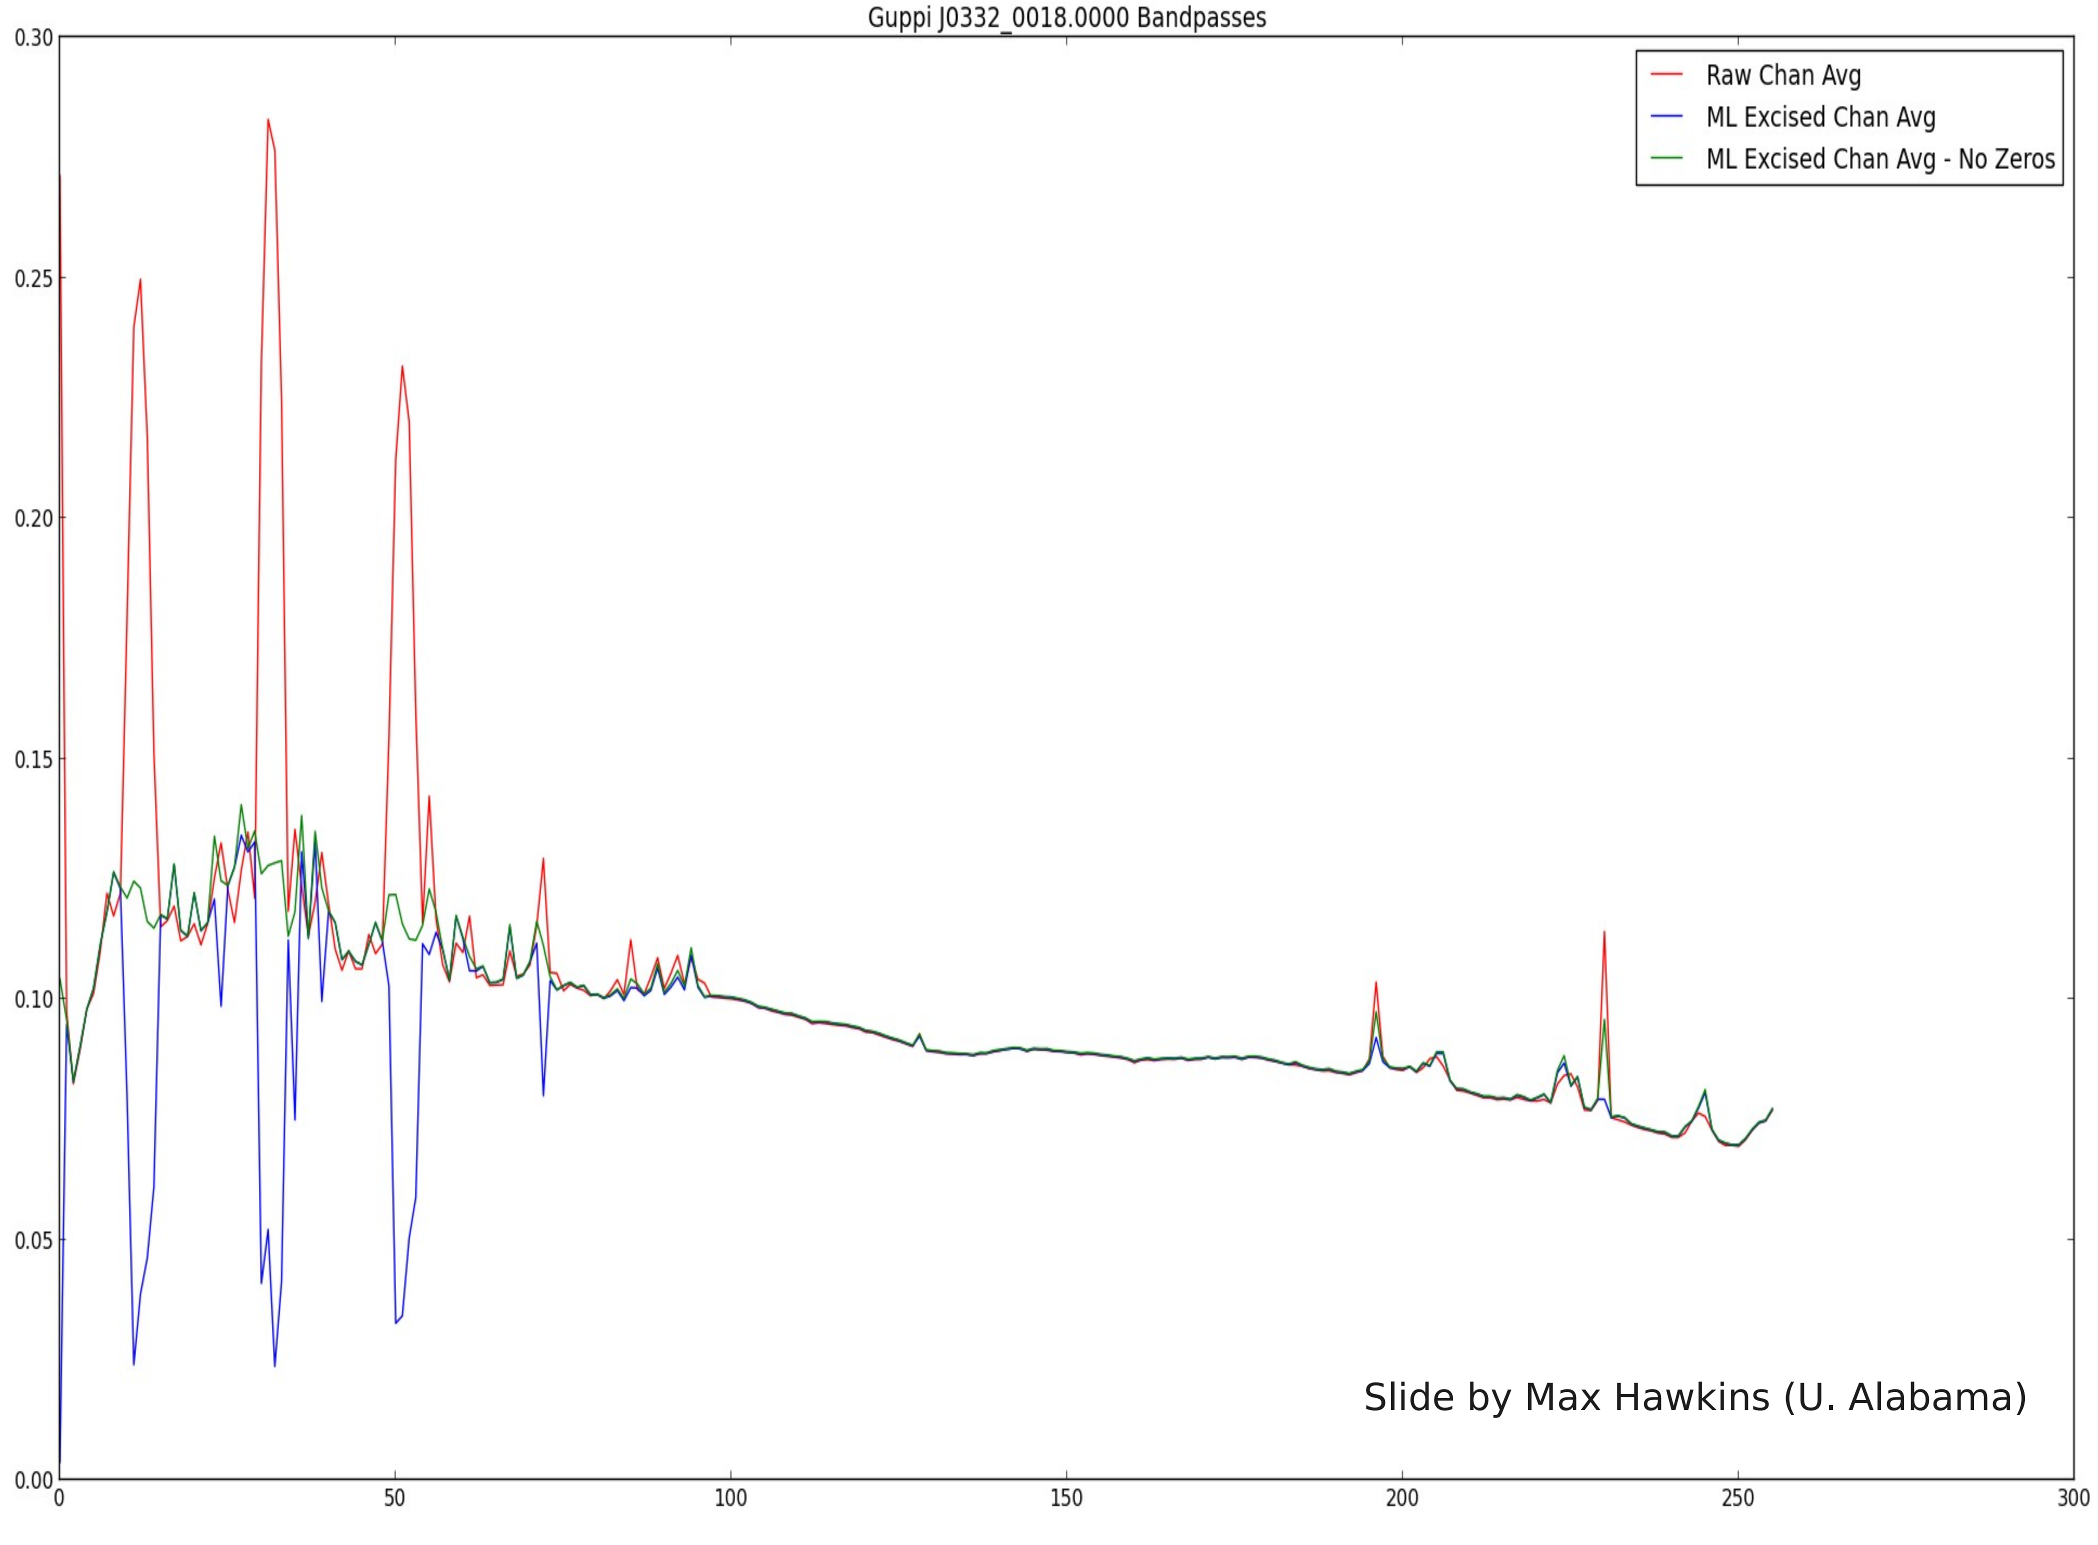

Slide by Max Hawkins (U. Alabama)

A Note on Complex Voltages

- There are some applications in which it is desirable/necessary to record pre-detection complex voltages
  - Very long baseline interferometry requires phase information for correlation
  - Offline analysis may be needed to form spectra with different resolutions for different applications
- This comes at the expense of very high data rates, requiring lots of storage

Hardware for Modern Digital Backends

- Modern systems are typically implemented with a combination of field programmable gate arrays (FPGAs) and GPU-equipped high performance computers running specialized digital signal processing software
- GBT currently uses five primary backends
  - Digital continuum receiver
  - Mark V (now Mark VI) VLBI baseband recorder
  - GUPPI (pulsar observing – being retired)
  - VEGAS (spectral line/pulsar observing)
  - Breakthrough Listen (baseband recording for SETI, etc.)

Hardware for Digital Backends

- GUPPI, VEGAS, BTL developed through CASPER (Collaboration for Astronomical Signal Processing and Electronics Research)
- VEGAS uses 8x ROACH2 boards and NVIDIA GPUs
  - Integrated ADCs, FPGAs, 10 gigabit ethernet, serial communication ports, onboard flash memory perform initial conditioning, supply channelized data or I/Q values
  - Additional spectral line / pulsar processing performed on GPUs/CPUs
  - Data stored on lustre distributed filesystem
  - 8 independent spectrometer banks for maximumize frequency coverage/flexibility

Hardware for Digital Backends

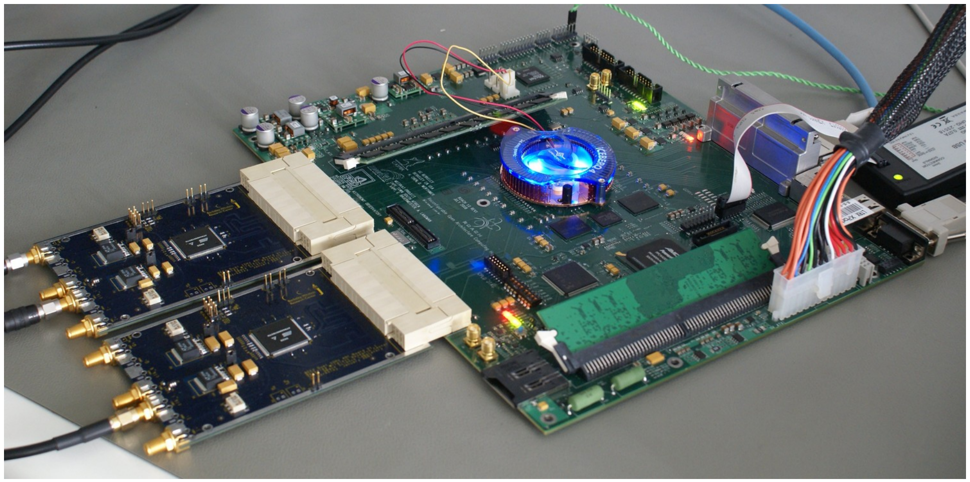

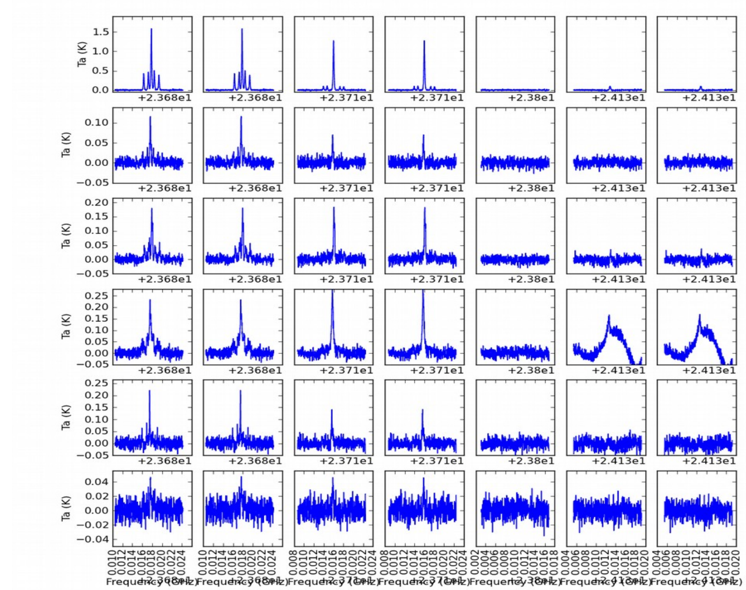

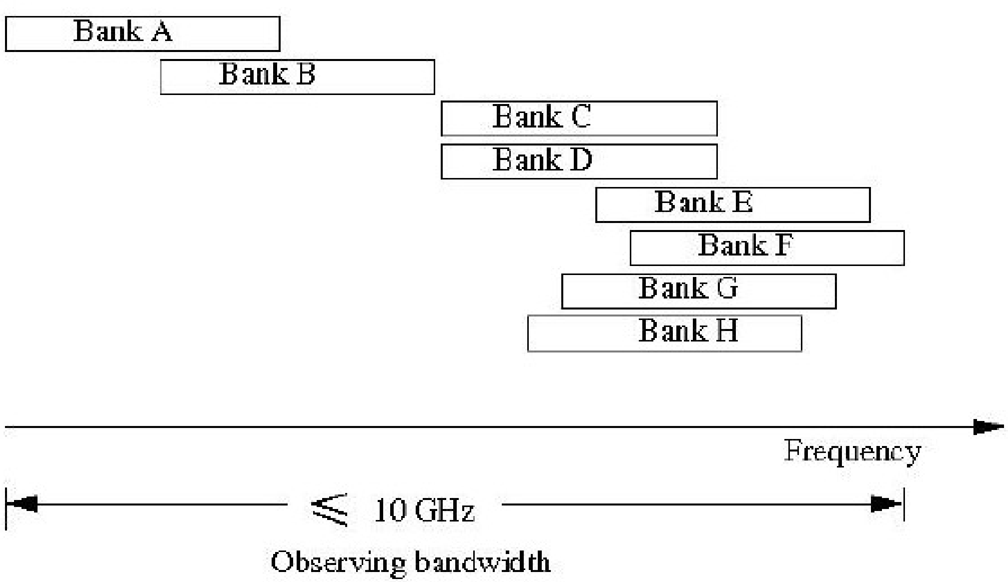

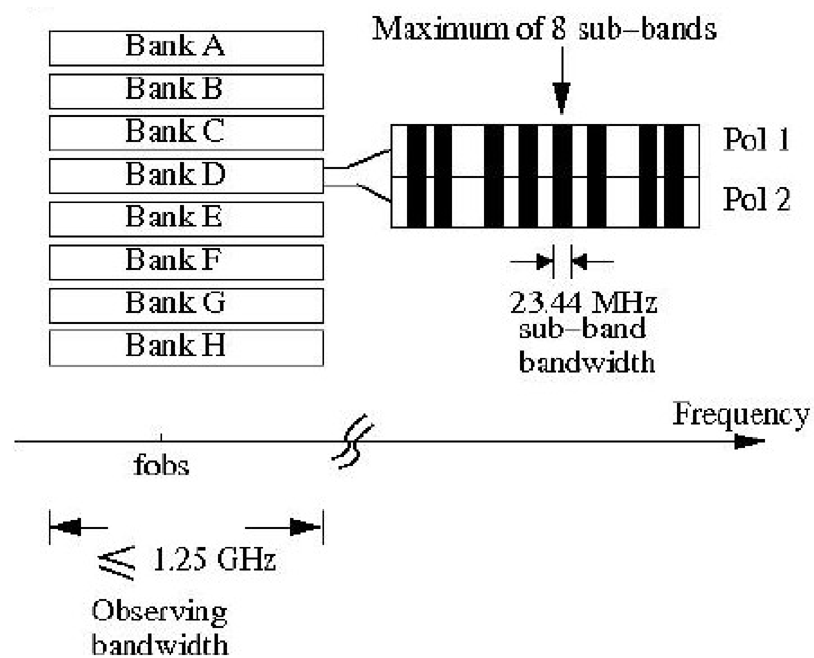

Next Generation Backends

- Fast (i.e. wideband) ADCs with lots of bits are relatively new
- ADCs can also be noisy (i.e. generate RFI) and produce a lot of heat
- Historically, most radio telescopes convert to IF and use lots of analog components
  - Lots of components can be impacted by RFI

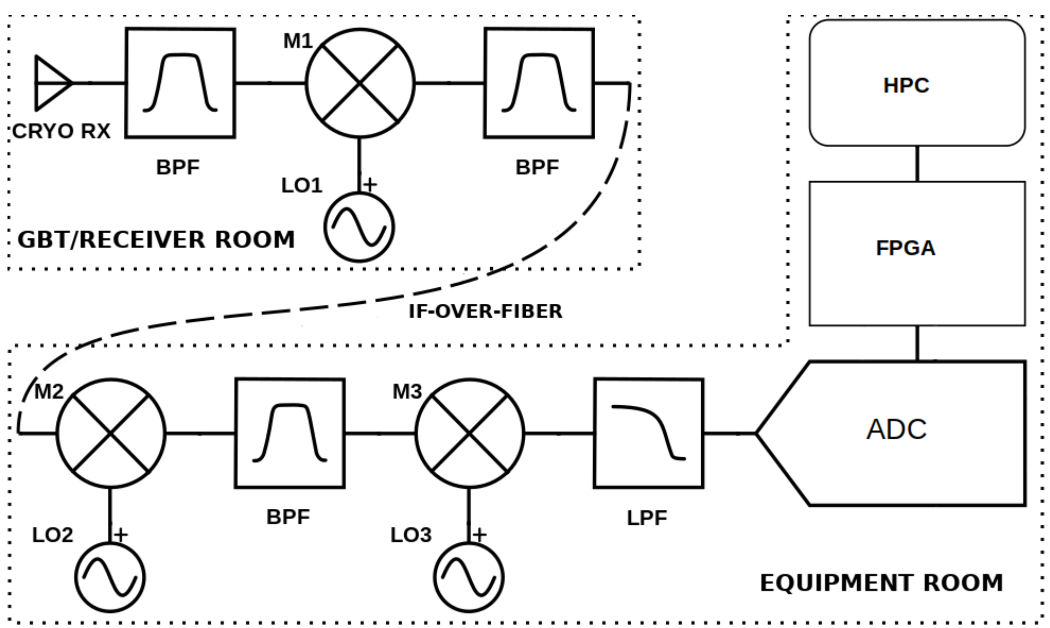

Next Generation Backends

- New ADCs and FPGAs make it feasible to directly sample wide bandwidths at RF
  - Already being done for some instruments (e.g. UWBL receiver at Parkes, FLAG at GBT, and others)
- GBO has just started a new NSF-funded R&D project to develop integrated RF sampling for GBT ultrawideband receiver

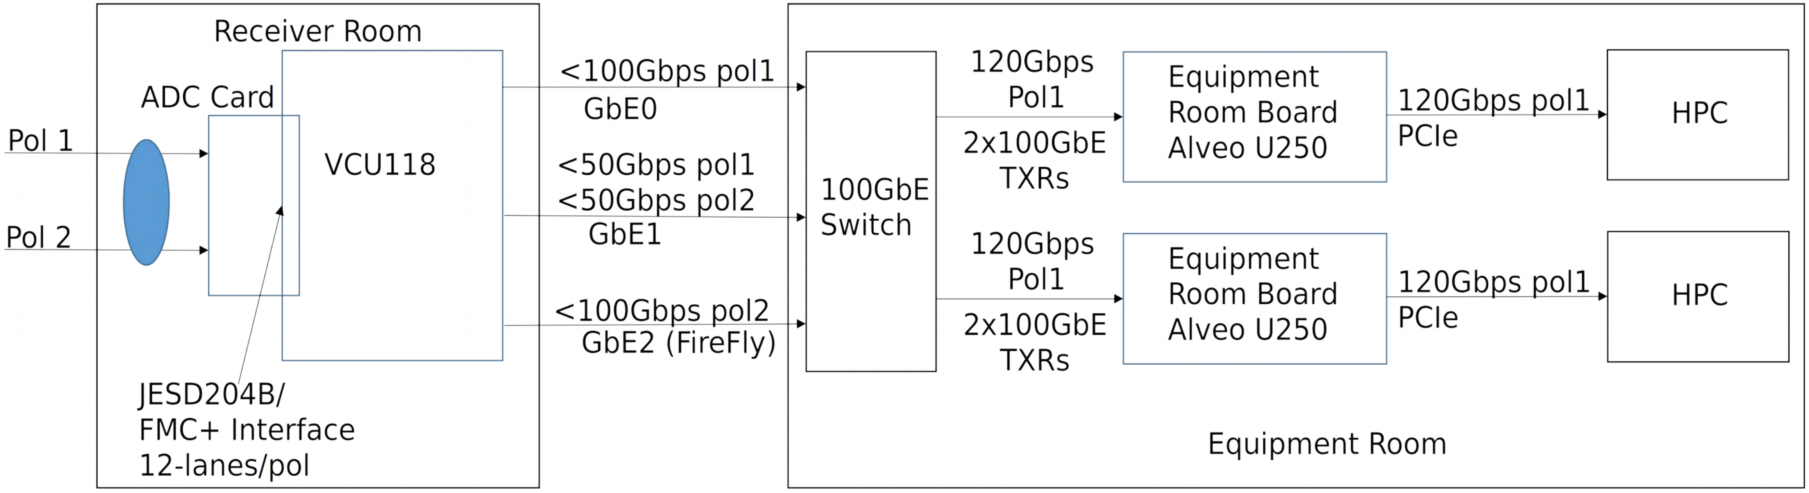

Next Generation Backends

- GBO has just started a new NSF-funded R&D project to develop integrated RF sampling for GBT ultrawideband receiver
- Goal is use 10 Gsamp/s, 12-bit ADCs to sample 4 GHz of bandwidth
  - Data rates > 240 Gbits/s
- Will also offer optional active RFI excision
- R&D phase over next two years
  - Would then build new spectrometer

Questions?

GREEN BANK OBSERVATORY

greenbankobservatory.org

The Green Bank Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.