Challenges and opportunities for innovation in imaging and cosmology with the SKA

Nithyanandan Thyagarajan (CSIRO)

I acknowledge the Traditional Owners of the land, sea and waters, of the area that we live and work on across Australia. I acknowledge their continuing connection to their culture and pay my respects to their Elders past and present.

Challenges or opportunities?

- EoR studies using redshifted 21cm
- Time-domain and Data Intensive astronomy
- Imaging
- Calibration
- RFI

Epoch of Reionisation (EoR) studies

Global Average Spectrum

Fluctuations Power Spectrum

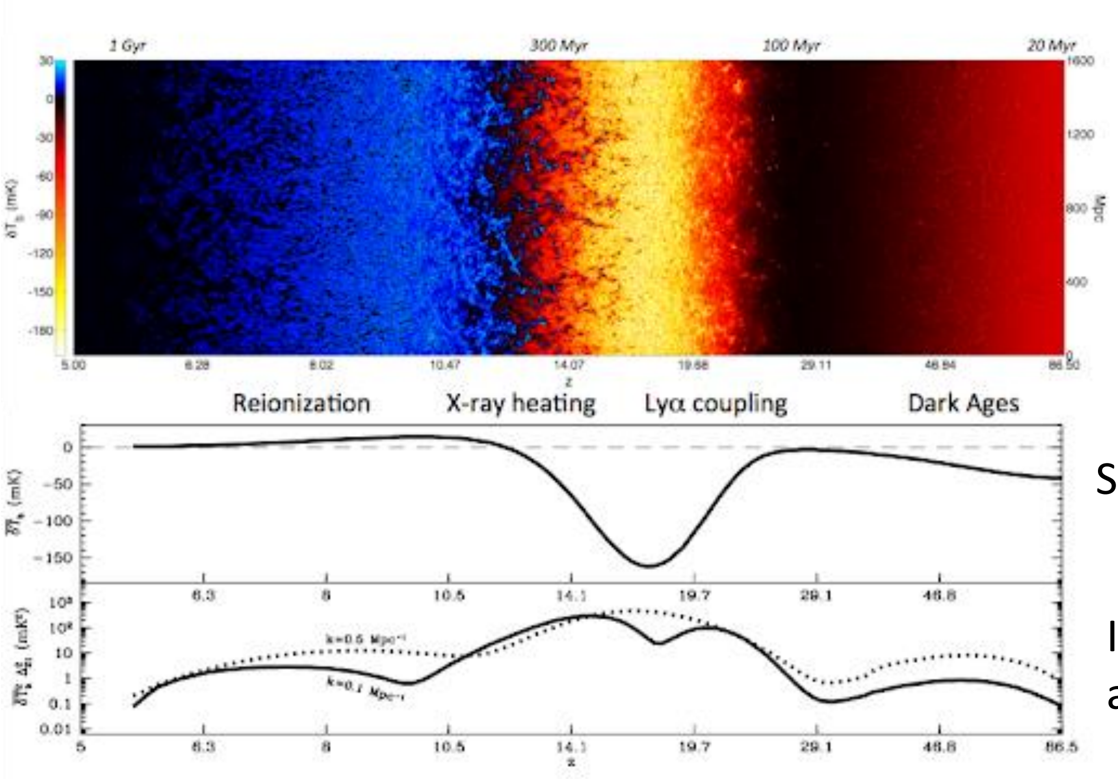

Single element

Interferometer arrays

The Potential of 21cm EoR experiments

Exquisite Astrophysical constraints

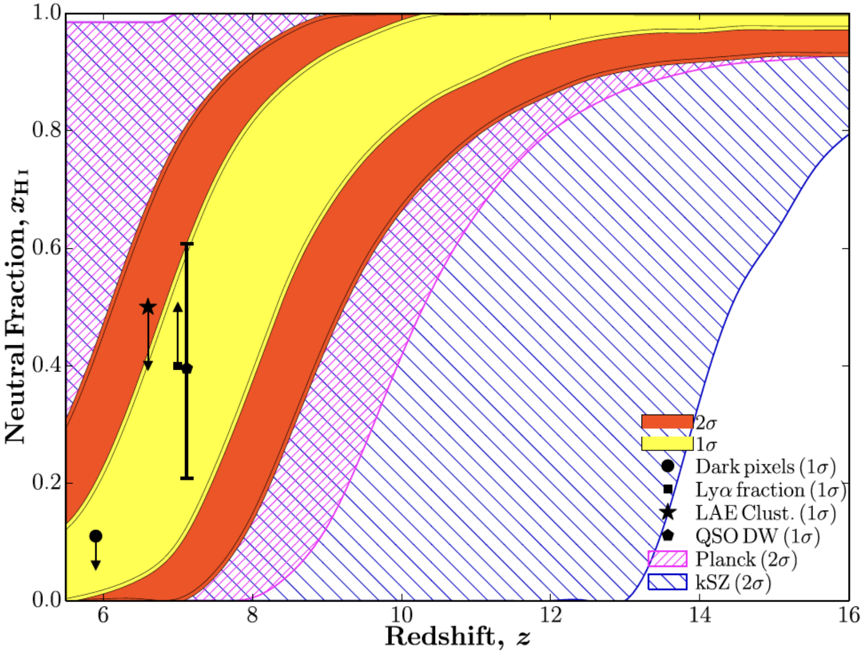

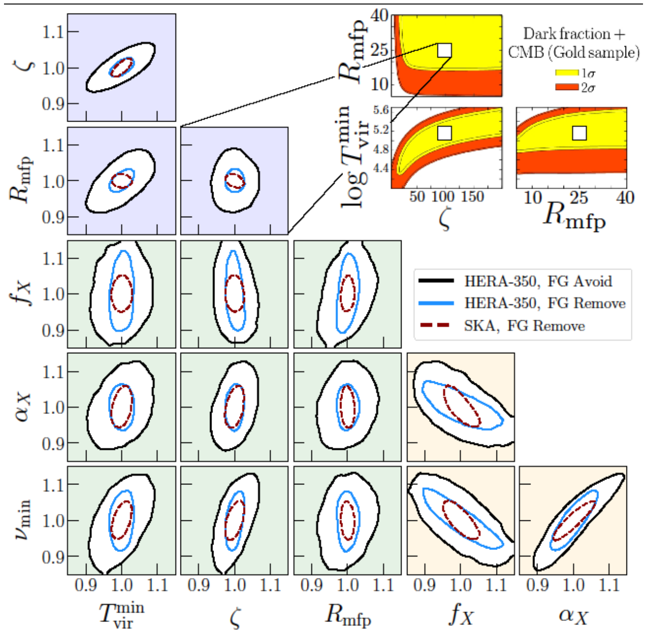

Courtesy: Bradley Greig & Andrei Mesinger

The Potential of 21cm EoR experiments

Disentangling cosmology constraints

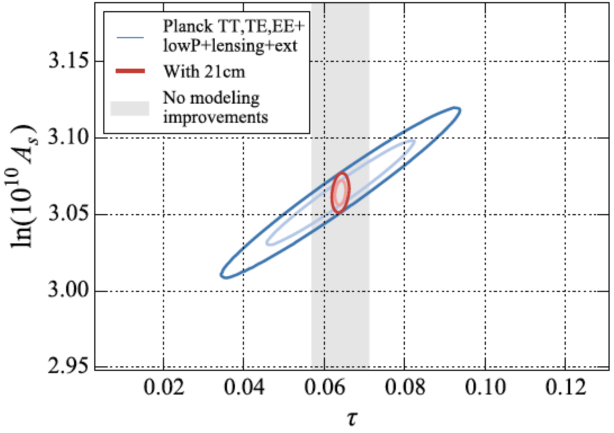

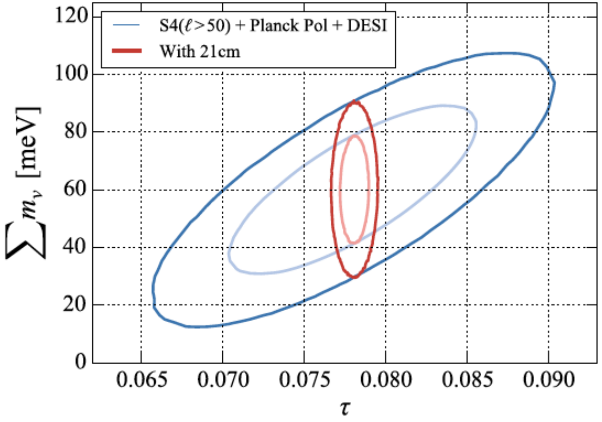

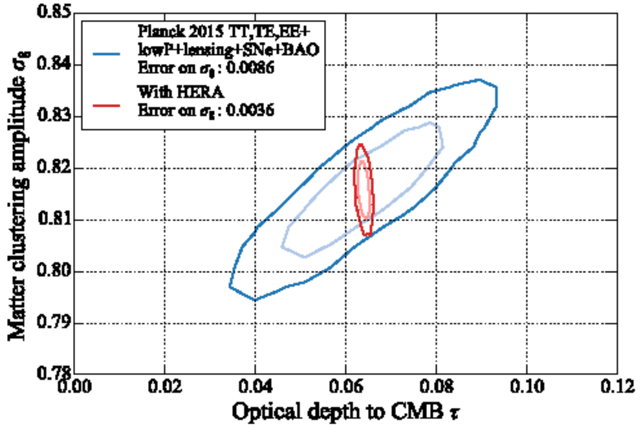

Liu+ (2016)

Challenges to 21cm experiments

21cm HI signature

Bright Synchrotron Foregrounds

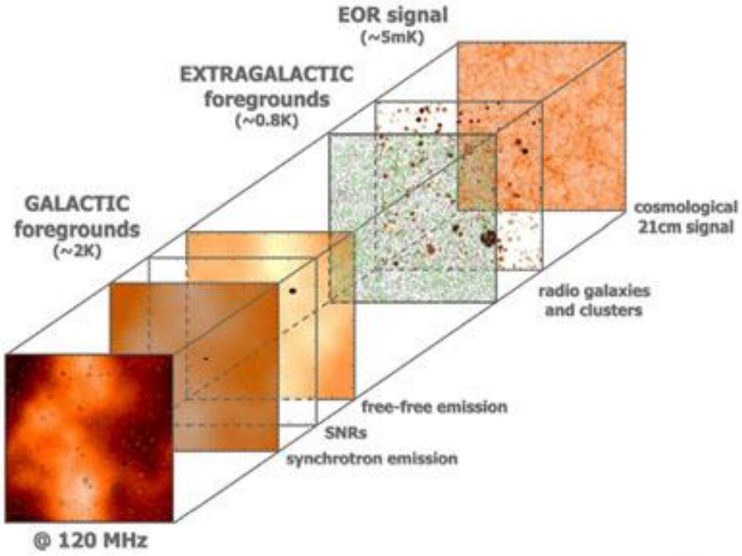

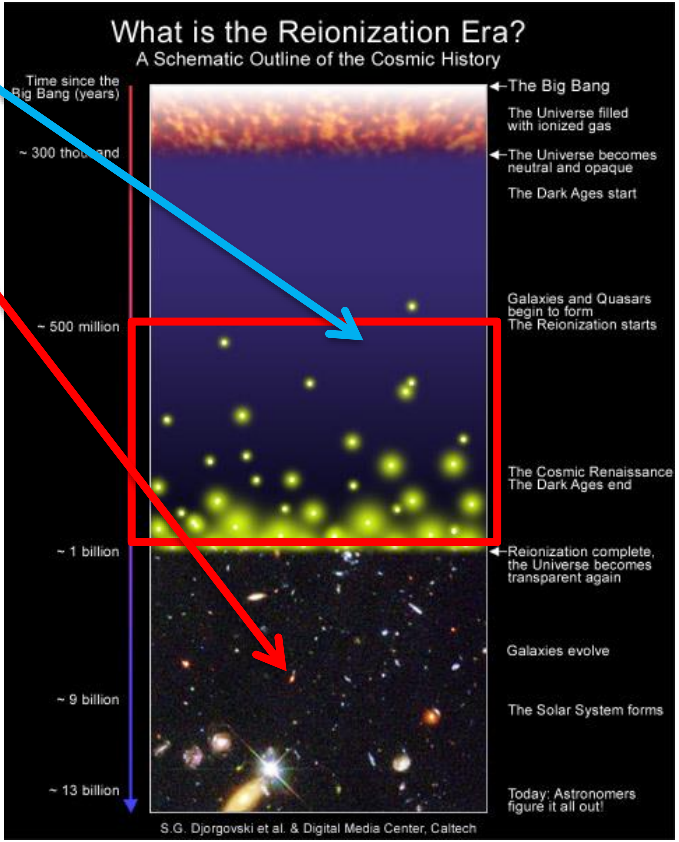

Calibration Challenges

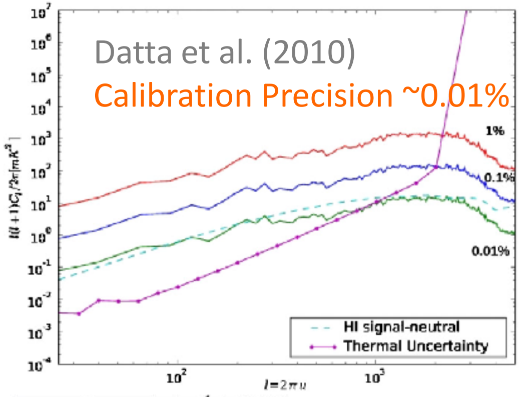

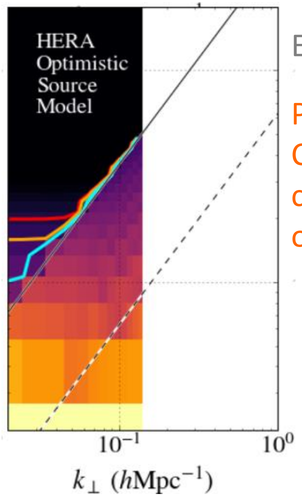

Ewall-Wice et al. (2017)

Power Spectrum Contamination from calibration errors in an optimistic case

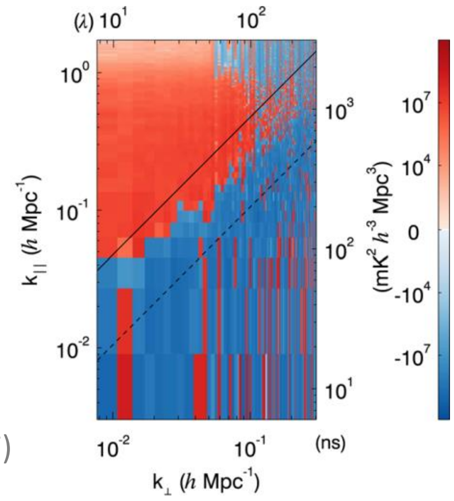

Barry et al. (2016)

Calibration Precision ~10^{-5}

Thorough knowledge of foregrounds required to achieve this precision

Similar conclusions from …

- Trott & Wayth (2016) for MWA and SKA
- Patil et al. (2017) for LOFAR
- …
- Sophisticated strategies being developed
- HERA redundant calibration (Dillon et al. 2017; Orosz et al. 2018; Byrne+ 2020)
- Sagecal, DDEcal for LOFAR
- Hybrid calibration?

Recent improved HERA results

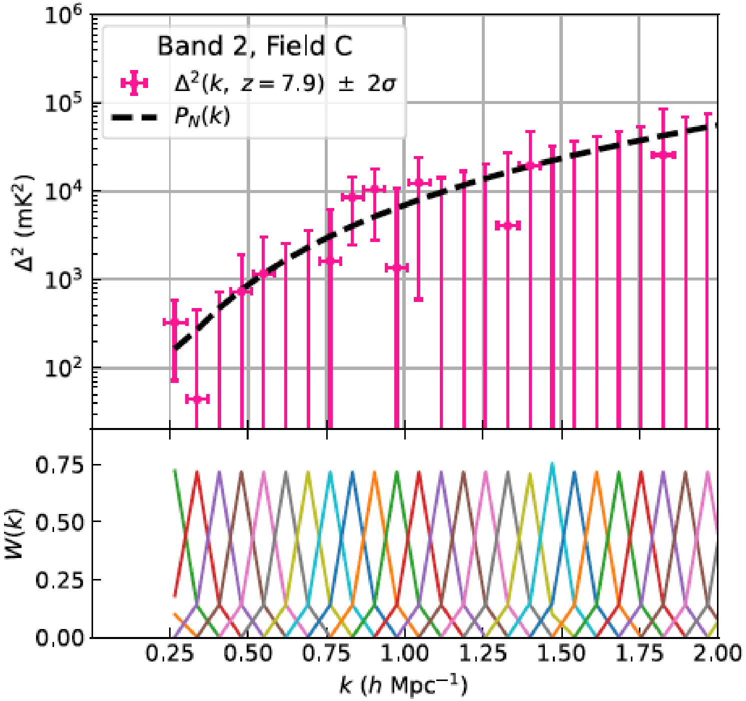

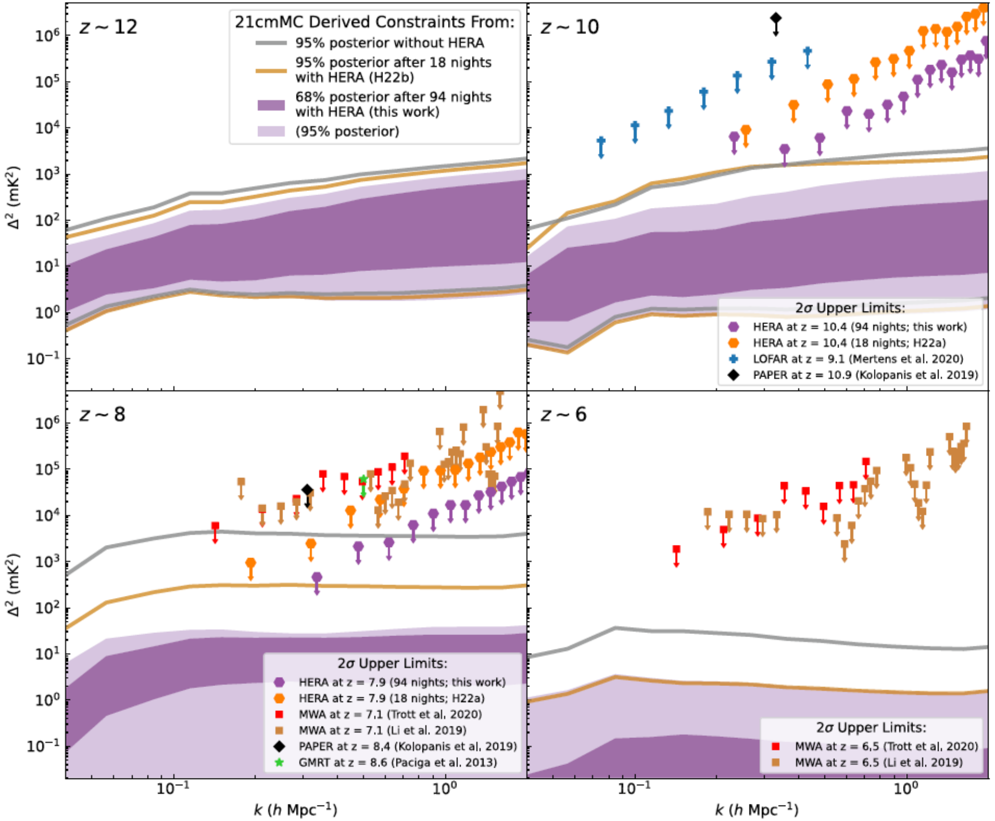

HERA et al. (2022)

Recent improved LOFAR results

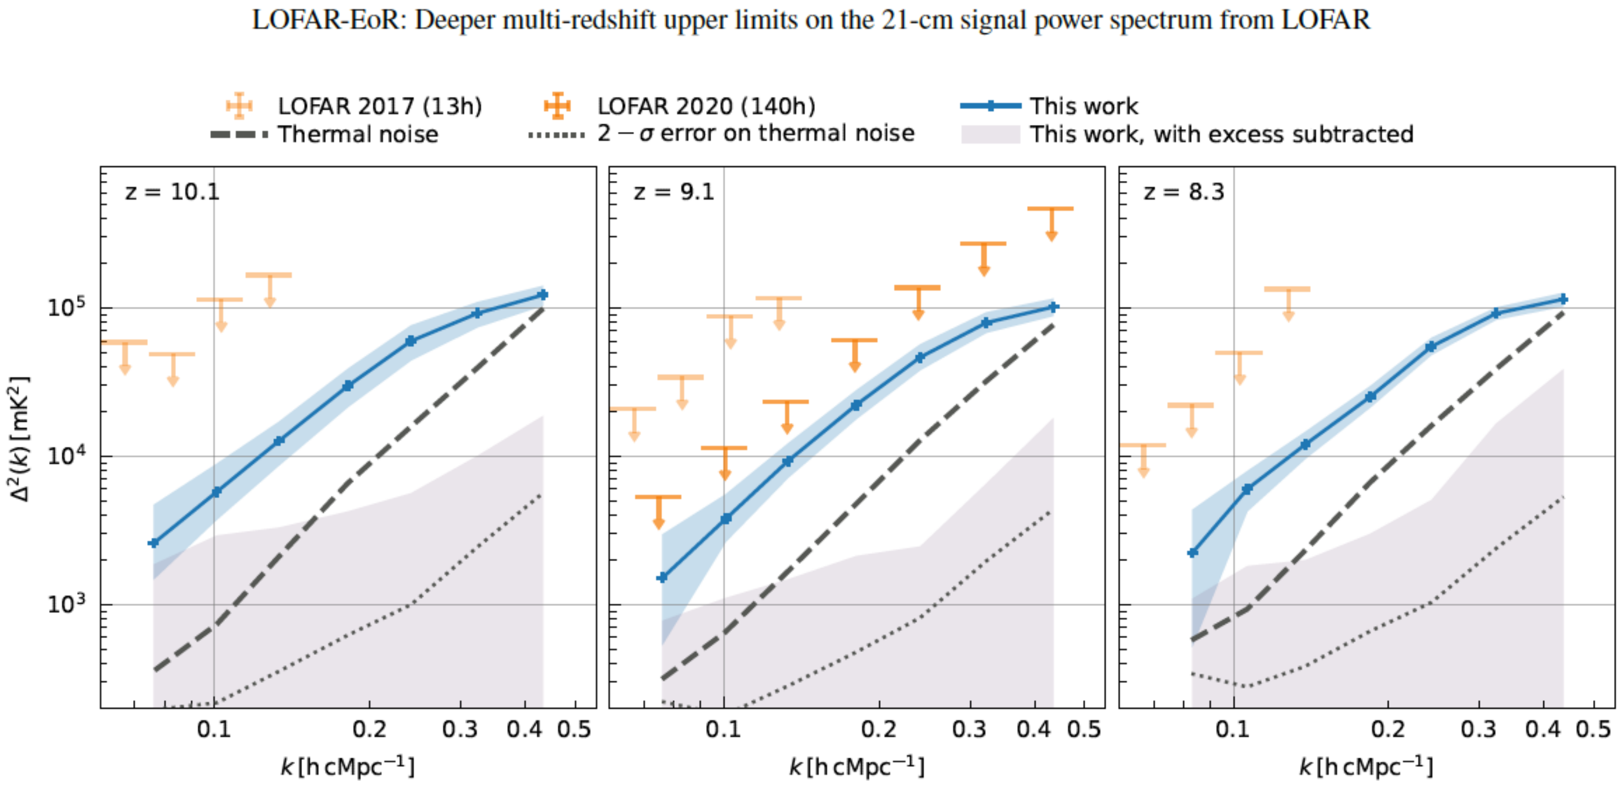

Mertens et al. (2025)

Improved MWA results (coming soon!)

- MWA results using 268 hours
- Redshifts z=6.5 to 7

Gradual but steady improvements

LOFAR is improving its results

- ❑Main limitation has been rotation of sources around North Celestial Pole and associated sidelobe removal due to beam errors

HERA is also improving its results

- ❑Has set very competitive results to date (depends on redshift, k-bin, etc.)
- ❑Main limitation is inter-antenna coupling effects that introduce spectral structure

MWA is also improving its results

- ❑Main limitation is sharp features in bandpass structure

21cm forest: Independent, alternate EoR measurement

- Have to be lucky for bright background source
- Statistical detection has more prospects
- Sample of high-z RLQSOs growing

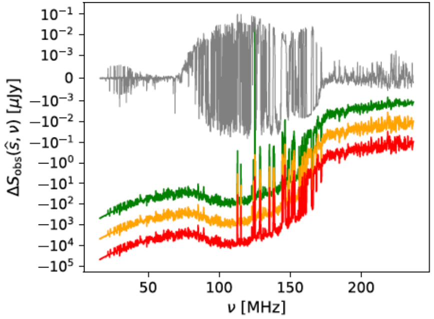

(a) EoS model 1 (BRIGHT GALAXIES)

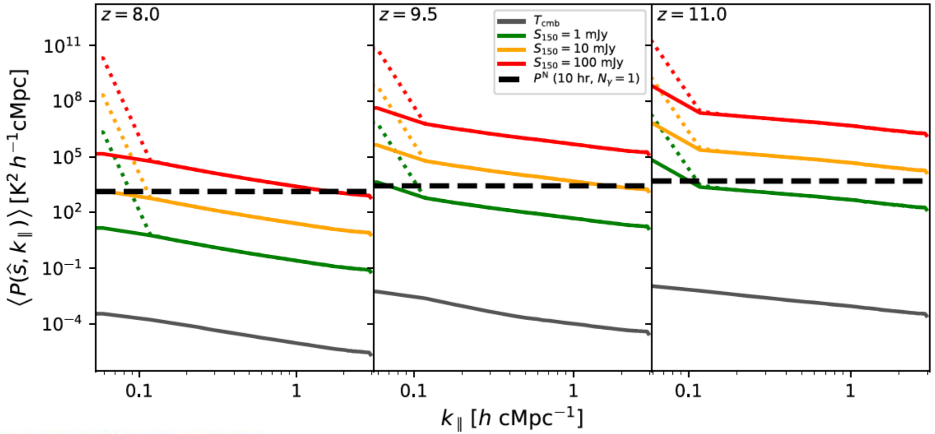

Thyagarajan (2020)

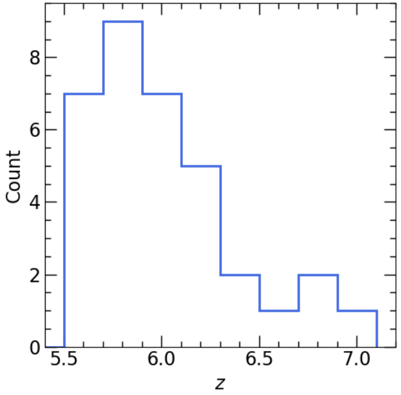

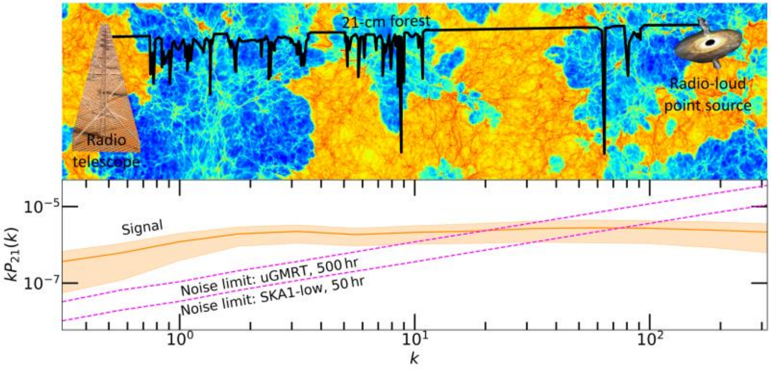

Soltinsky et al. (2025)

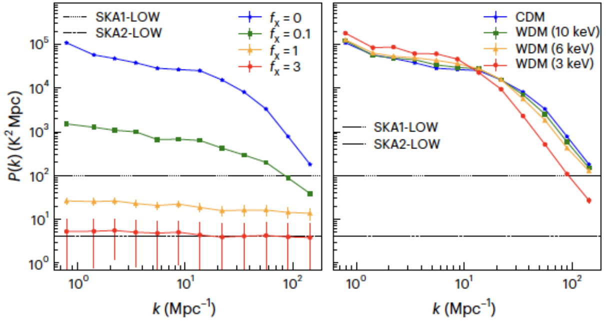

Shao et al. (2023)

Closure phase spectrum (avoids calibration systematics)

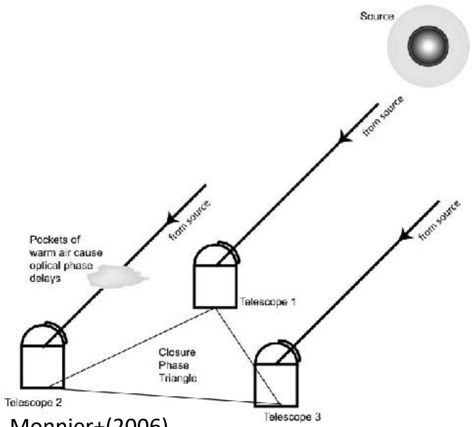

Monnier+(2006)

Visibility Fluctuations

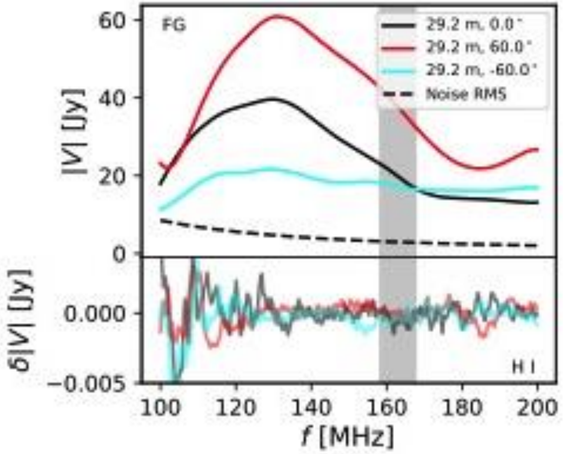

Closure Phase Fluctuations

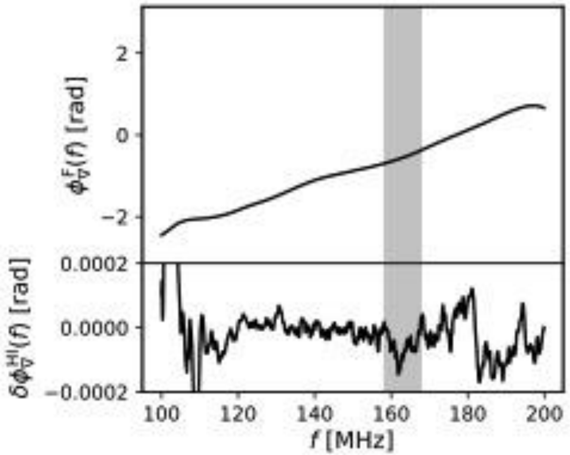

Thyagarajan et al. (2020): PRD 102, 022002

Good correspondence between fluctuations.

Shape, Dynamic range, Sensitivity, etc.

Closure Phase Power Spectrum Results provides independent constraints

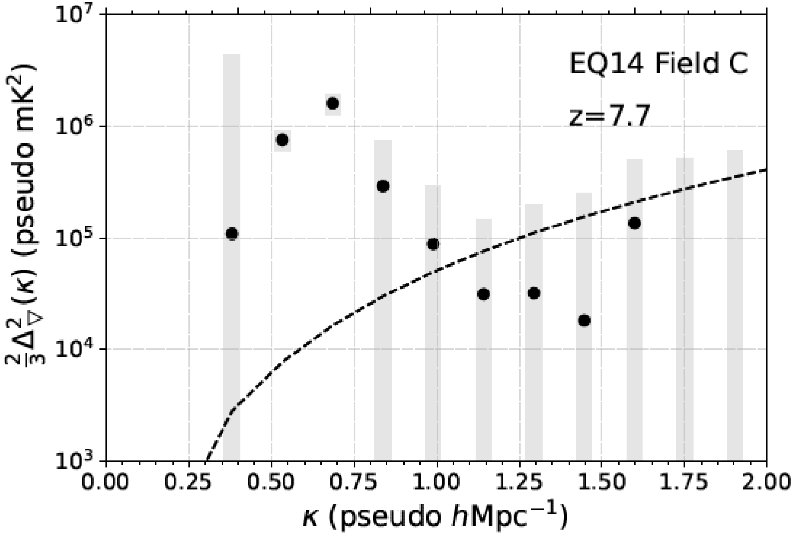

Keller+ 2023 using HERA

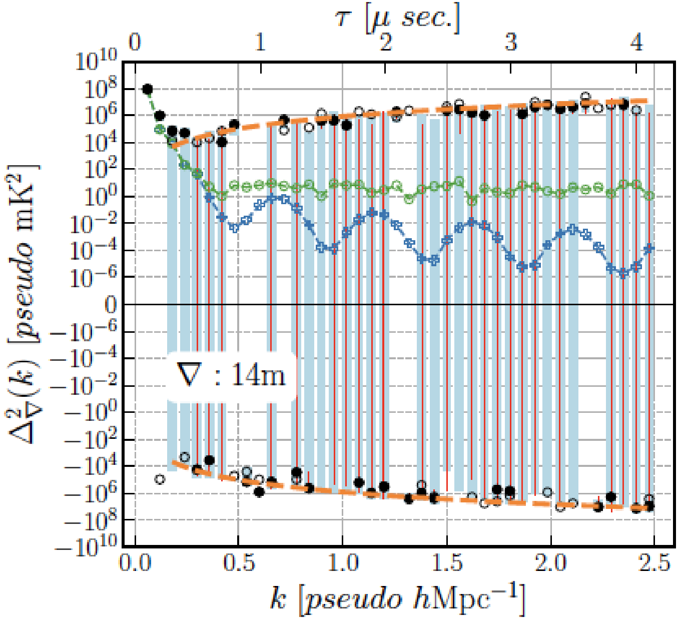

Tiwari+ 2023 using MWA

Time domain and Data Intensive Astronomy

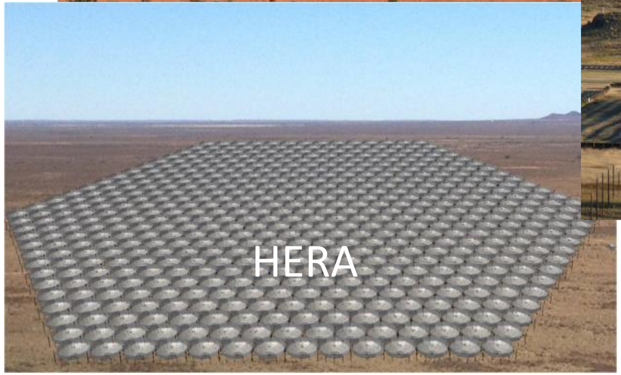

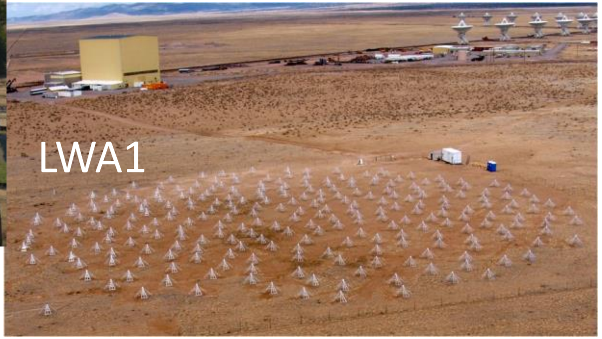

- Aperture elements
- Aperture array stations
- Multi-scale array of stations

Need for new paradigm: Extreme time-domain phenomena

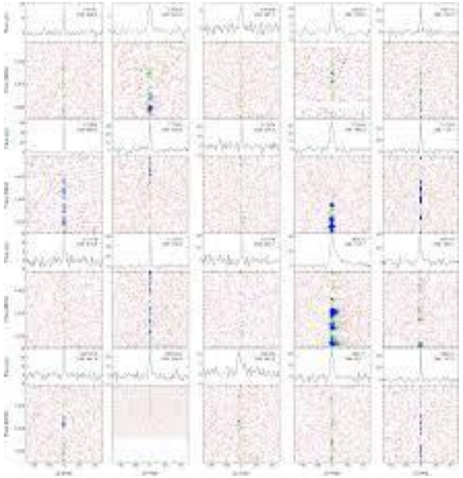

ASKAP FRBs: Shannon et al. (2018)

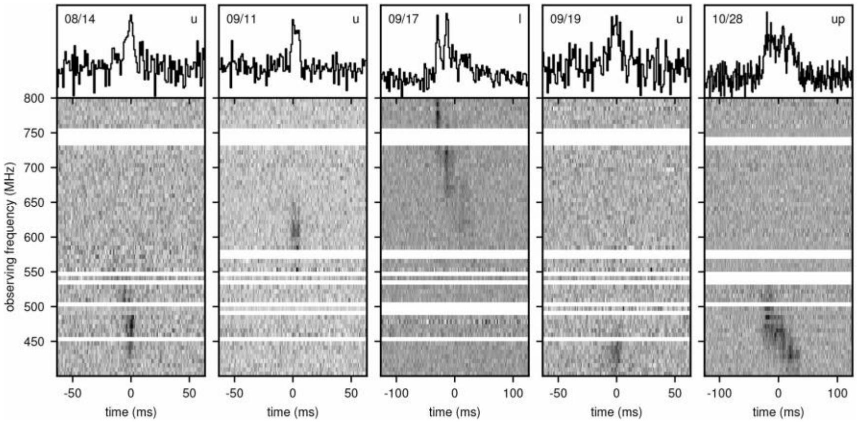

The CHIME / FRB collaboration (2019)

Fast transients

- Extreme explosive astrophysical phenomena
- Very short duration (< 1 ms)
- Unknown but very distant origin in the Universe
- Probes for Dark matter, matter under extreme densities

Time domain requires…

- Wide fields of view
- Large number of antennas for collecting area
- High time cadence, fast writeouts

Observational Cosmology like 21cm EoR also favours dense, large-N arrays

E-field Parallel Imaging "Correlator" (EPIC)

Playing the role of Nature more directly: Direct Imaging

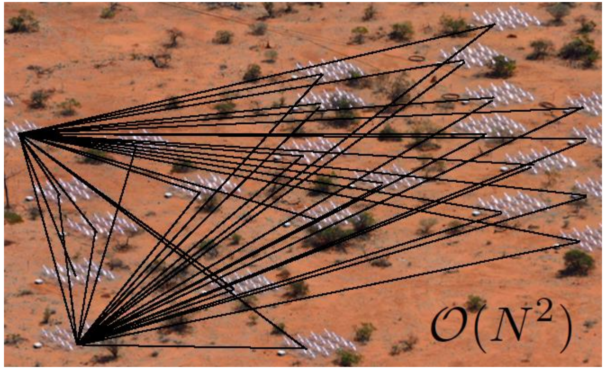

Image credit: Adam Beardsley

$$I(l,m) = \iint \mathcal{V}(u,v) e^{2\pi i(ul+vm)} du dv$$

Spatial correlation: Computationally expensive for Large-N telescopes

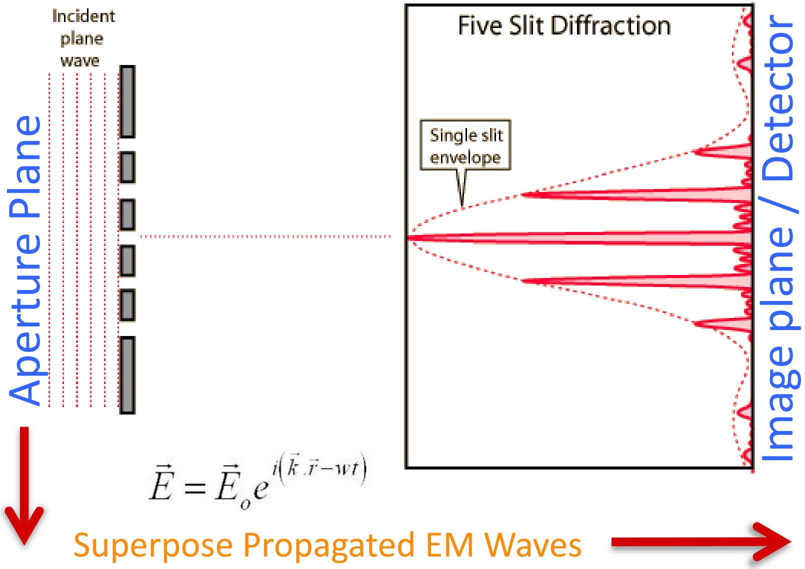

Superpose Propagated EM Waves
= Fourier Transform

Convolution theorem of Fourier Transform:

FT(Correlation) ←→ FT(.) x FT(.) = |FT(.)|^2

Traditional / Correlation-based

Direct Imaging (FT and square)

EPIC implementation of Direct Imaging

- FX

Traditional / Correlation-based

FT(Correlation)

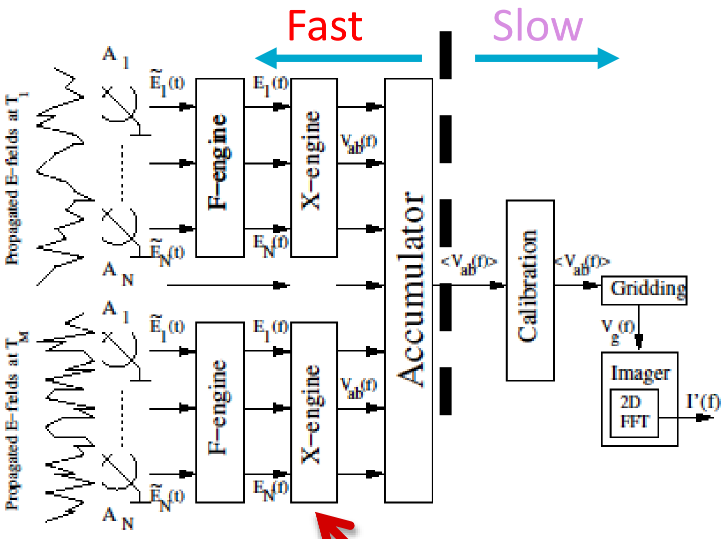

O(N_a^2)

Thyagarajan et al. (2017)

- EPIC

Direct Imaging

FT(.) x FT(.)

Gridding makes it generic

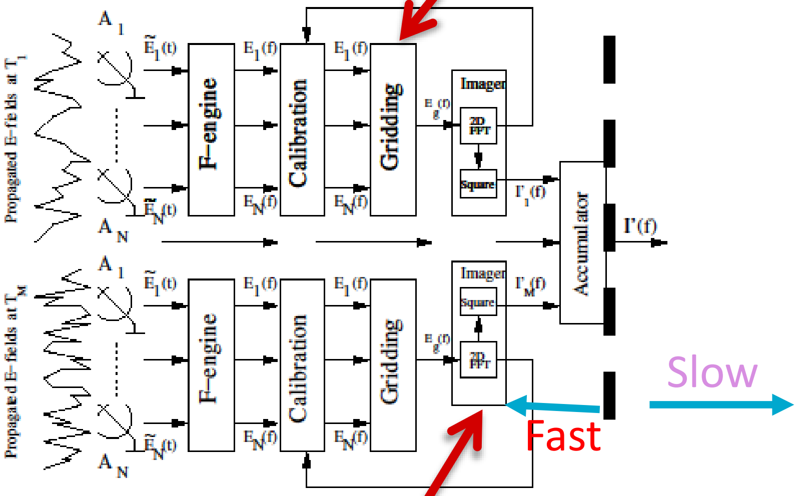

O(N_g log N_g)

Fast Imaging
(With calibration)

Deployment on LWA-Sevilleta (NM, USA)

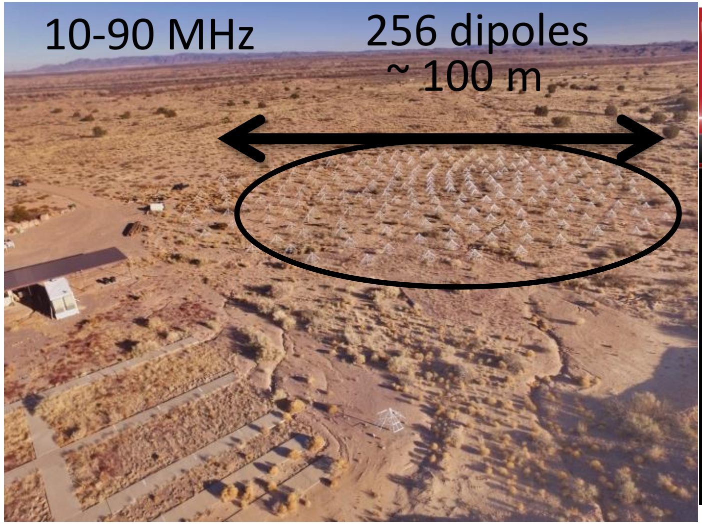

Image Credit: Greg Taylor (PI: LWA-SV)

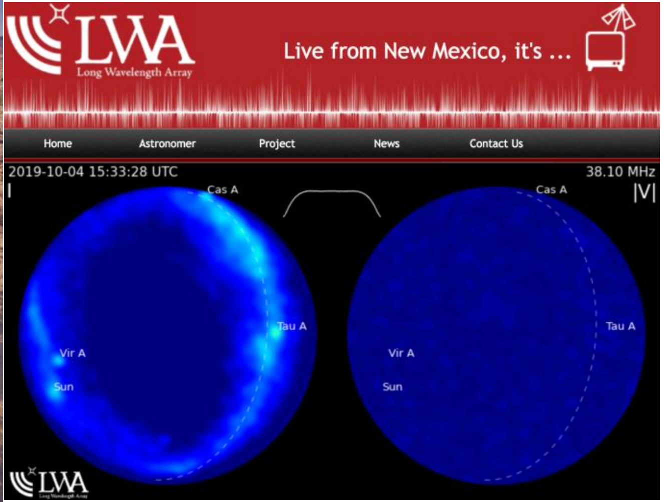

Significant Time-domain Capability added to LWA-SV
Before EPIC: > 1 s, After EPIC: ~ 80 ms

Parameters	GTX 980	RTX 2080Ti
Number of cores	2048	4352
GPU clock (MHz)	1127	1350
Streaming multiprocessors	16	68
Memory bandwidth (GB/s)	224.4	616
FP32 performance (TFLOPS)	4.891	13.45
FP64 performance (GFLOPS)	155.6	420.2

64x64 images, 3.6 MHz, 80 ms cadence on two GPU cards
EPIC on commensal mode for users

Comparison with other architectures

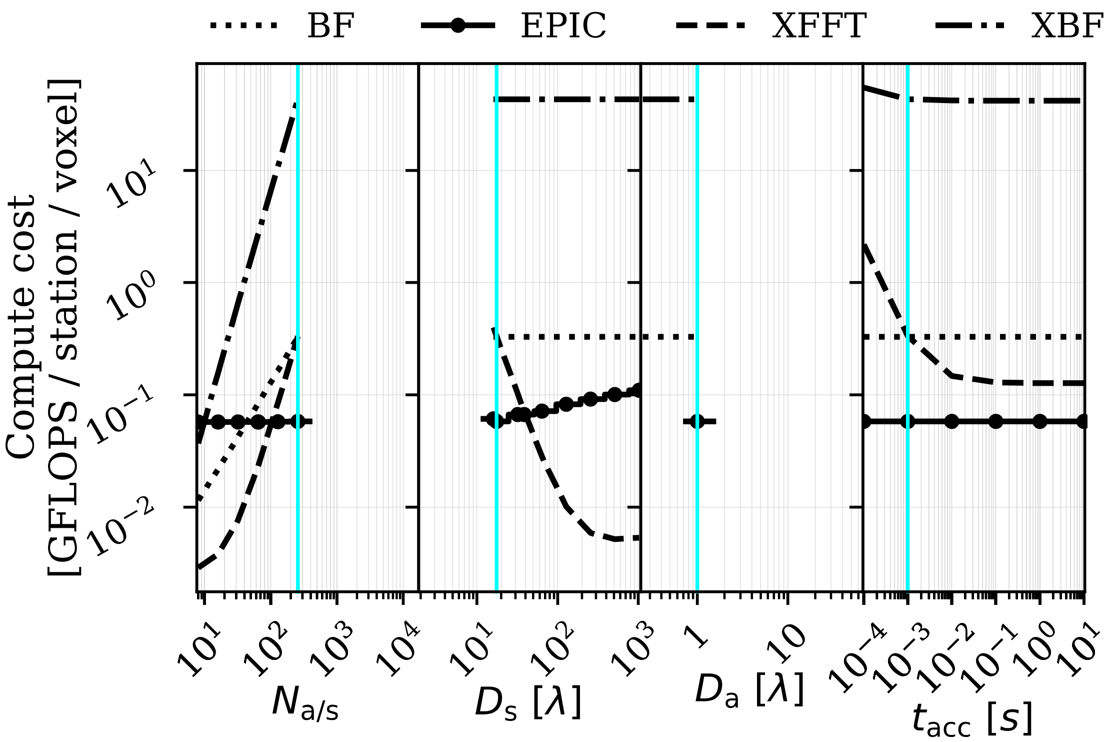

SKA-low or similar

CSIRO

EPIC on SKA-low

- 224 stations in the SKA-Low core
  - That's 57344 antennas !
  - 1 km diameter
- Many options for EPIC
  - Single station : 64x64 pixels
  - Sum multiple stations for greater sensitivity
  - Or image with the entire core e.g. 1024x1024 pixel all-sky images at kHz update rates

See everything that flashes in the radio sky!

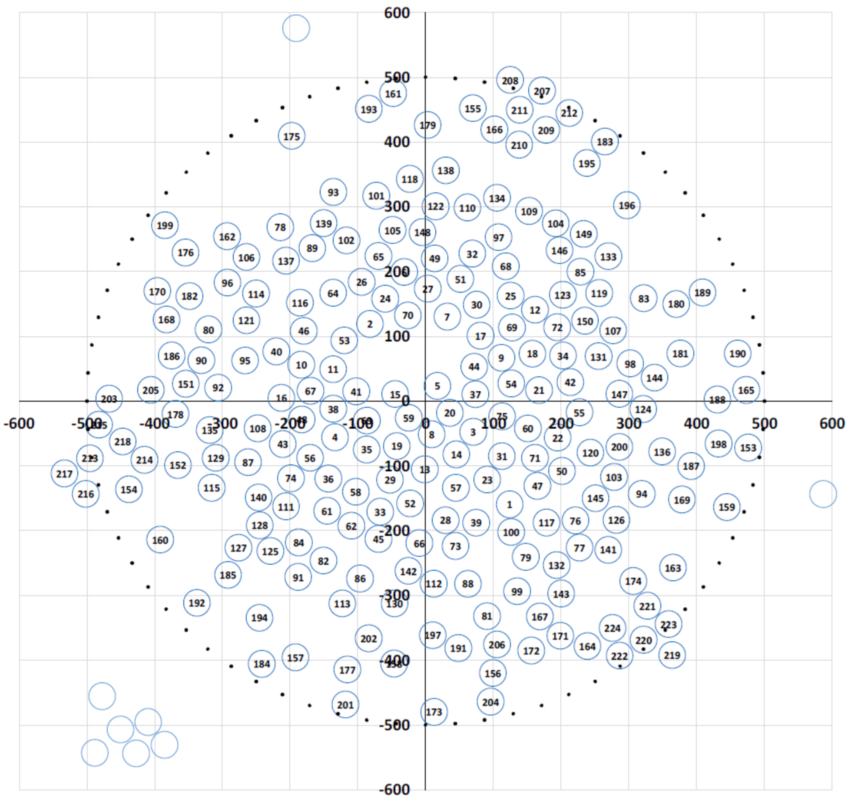

LWA EPIC GPU Implementation

LWA performed a quite detailed analysis of EPIC GPU performance

Memory bandwidth limited

GPU	SM count	Shared Memory per SM (KiB)	Image size	Achievable bandwidth[a] (Channels/MHz)	Polarization	Base [Boost] Clock (MHz)	Compute capability	Bus interface	Estimated price
RTX 2080 Ti	68	64	64x64 128x128	132/3.3 MHz 68/1.65 MHz	Dual Single	1350 [1545]	7.5	PCIe 3.0 x16	$1k
RTX A4500	56	128	64x64 128x128	112/2.8 MHz 56/1.4 MHz	Dual Single	1050 [1650]	8.6	PCIe 4.0 x16	$1.5k
RTX 3090 Ti	84	128	64x64 128x128	132/3.3 MHz 84/2.1 MHz	Dual Single	1560 [1860]	8.6	PCIe 4.0 x16	$1.6k
RTX 4090	128	128	64x64 128x128	132/3.3 MHz 128/3.2 MHz	Dual Dual	2235 [2520]	8.9	PCIe 4.0 x16	$1.8k
RTX 4090 Ti	142	128	64x64 128x128	132/3.3 MHz 132/3.3 MHz	Dual Dual	2235 [2520]	8.9	PCIe 4.0 x16	$2k[#]
				Data Center GPUs					
L40	142	128	64x64 128x128	132/3.3 MHz 132/3.3 MHz	Dual Single	735 [2490]	8.9	PCIe 4.0 x16	$7.6k
A100 PCIe 40 GB	108	192	64x64 128x128	132/3.3 MHz 108/2.7 MHz	Dual Dual	765 [1410]	8.0	PCIe 4.0 x16	$11k
H100 PCIe	114	256	64x64 128x128	132/3.3 MHz 114*/2.85 MHz	Dual Dual	1095 [1795]	9.0	PCIe 5.0 x16	$35k

https://github.com/epic-astronomy/Memos/blob/master/PDFs/009_EPIC_Code_Optimizations.md

Fully parallel FFT

- 64 point FFT computed every clock cycle
- Every path in the FFT signal flow graph is implemented in the hardware
- Can be constructed entirely with adders
- The U55C Alveo can implement over 1,000,000 parallel 8-bit adders.

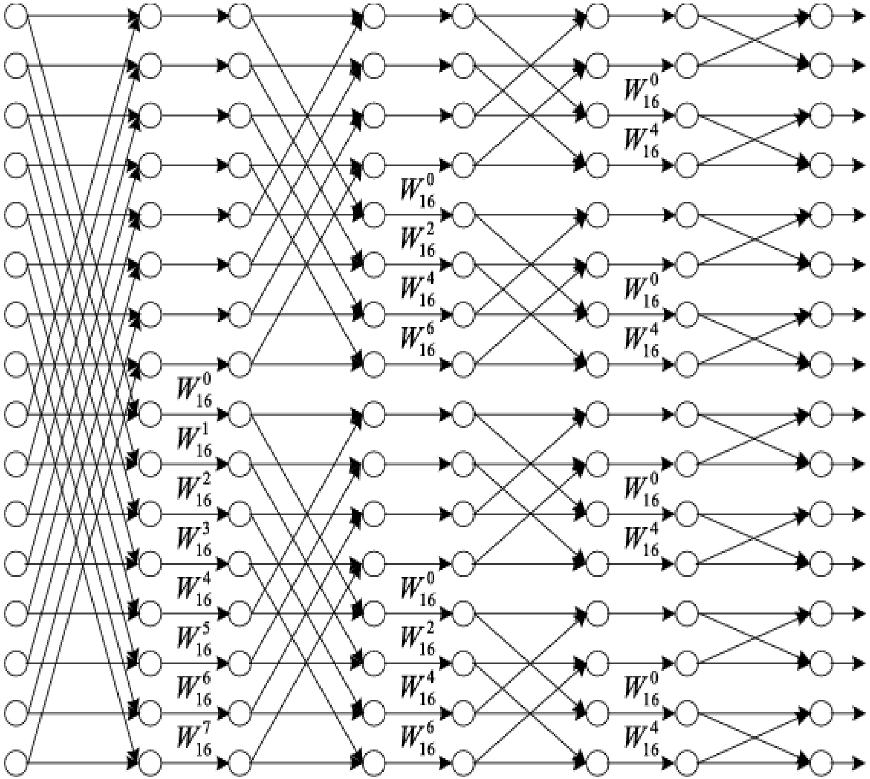

Signal flow graph for a 16-point FFT

FPGA/GPU comparison

- The LWA GPU implementation is severely memory bandwidth limited
  - gridded voltages and the 2D FFT doesn't fit in GPU register space.
  - Short FFTs are also not particularly efficient in the GPU

- LWA GPU implementation ~ 3 MHz of bandwidth at 300W power (Nvidia RTX 3080/4080)

- The FPGA allows construction of a processing pipeline that uses on chip block RAM with only the final time/frequency integrated image written back to external memory.
  - FFTs achieved with adders and managed with BRAMs
  - expandable to 1024x1024 grids, with thousands of antennas (with significantly reduced bandwidth per FPGA).

- U55C consumes ~120W power for ~60 MHz processing bandwidth
- ~50x bandwidth/power efficiency

Images produced by EPIC on FPGA would be exported to a GPU for further processing, e.g. transient search, de-dispersion etc.

Scaling to larger N, larger images

- FPGA resources required roughly scales as the image size
  - E.g. 256x256 image = 16x the number of pixels compared with 64x64, so around 4 MHz of bandwidth could be processed per FPGA
  - Allows arrays up to about 8192 antennas.
  - 8192 antennas = about 800 Gbps of voltage data input
- Scales to images of ~1000x1000 = 1MPixels
- Beyond 1000x1000, performance will drop dramatically due to on-chip memory constraints

SKA-Low - EPIC: 64 MHz bandwidth

- EPIC is well suited to SKA-low core

Configuration	Image Size (pixels)	Antennas Used	Estimated Alveos (Versal) required, 64 MHz bandwidth
Single SKA-Low Station	64x64 = 4096	256	1
Entire core, incoherent sum	64x64 = 4096	224x256 = 57344	224
Entire SKA-Low Core	1024x1024 = 1M pixel	224x256 = 57344	448

The science cases for EoR, pulsar search, and general transient search needs large field of view, not high angular resolution.

cost(448 Alveos) ~$5,000,000 = small fraction of total SKA-low cost

Imaging

Significant developments in last 2 decades

- Non-coplanarity, w-terms
- Heterogeneity, A-terms
- Large sky curvature within FoV
- W-projection, w-stacking, WSCLEAN, etc.
- Increasing use of RML methods for sparse apertures

New developments

- Image Domain gridding
- Spherical Harmonic Imaging
- M-mode imaging
- ML methods

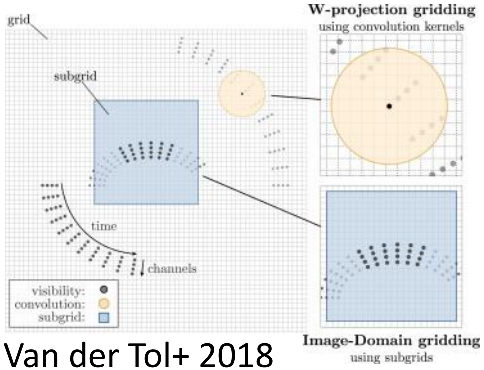

Van der Tol+ 2018

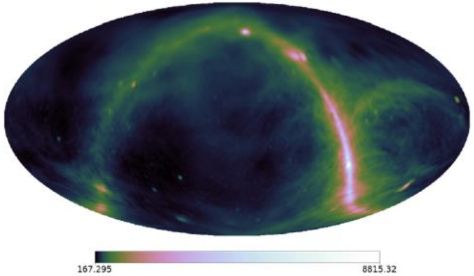

M-mode imaging
Kriele+ 2022

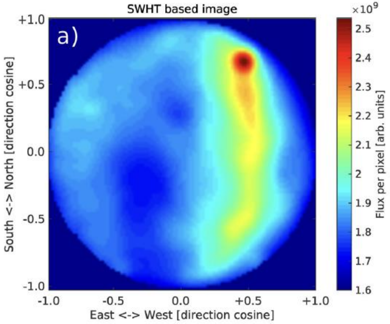

Spherical imaging
Carozzi 2015

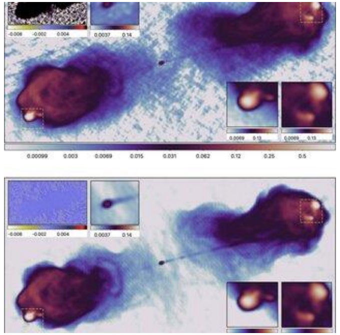

Deep Learning imaging
Image credit: Yves Wiaux

Calibration

Challenging when:

- Signal is weak compared to overwhelming nuisance signal e.g. EoR
  - Applies to low-frequency, dense aperture arrays
- Measurements on sparse aperture coverage, low S/N
  - Applies to VLBI, for example

Closure Invariants: Calibration-independent Imaging

- Closure invariants are calibration-independent
- E.g. Closure phases and amplitudes, but polarimetric invariants too
- Information ratio = $(N^2-3N+1)/(N^2-N) \sim 1-2/N$ when N is large
- Imaging possible without calibration without much loss of information

EHT layout

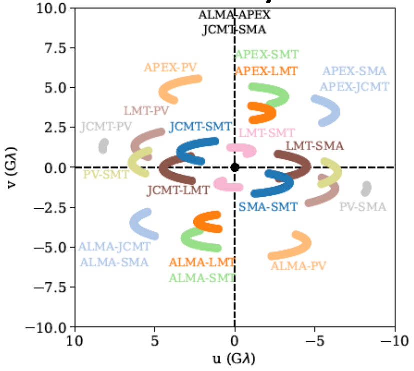

Closure Invariants + Deep Learning

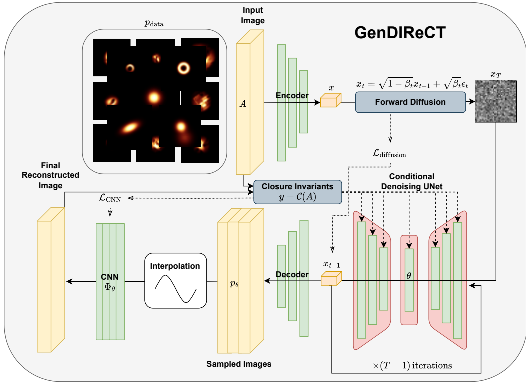

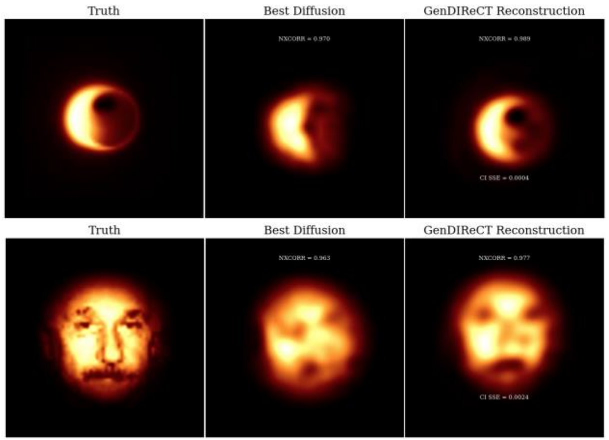

Lai et al. (2025)

RFI

- Site selection critical for RFI
- Site options shrinking with satellite RFI
- Lunar far side??

- Explore Spatial nulling in beamforming
- Online RFI detection and excision
- Explore pre-correlation RFI detection and flagging at high time- and spectral resolution becoming important

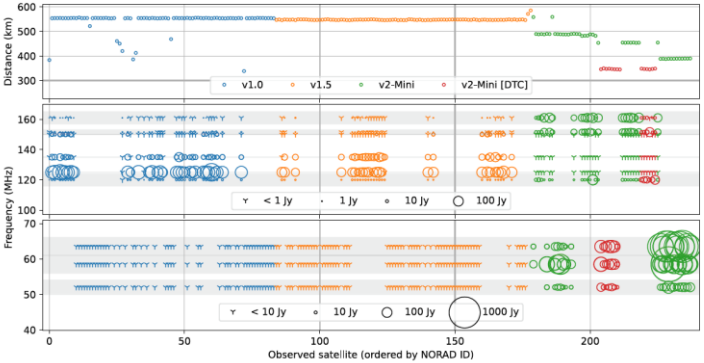

Bassa+ 2024

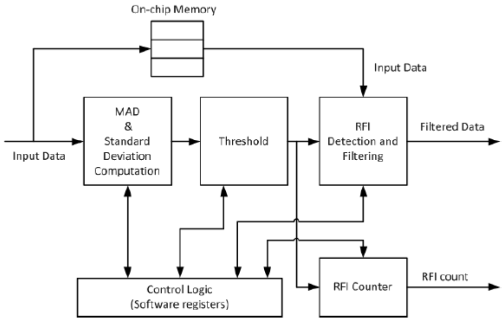

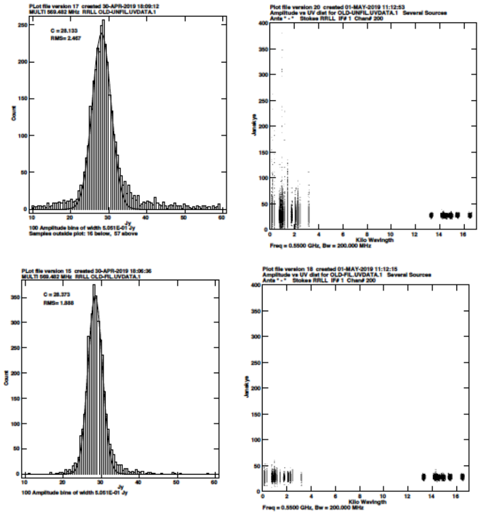

Real-time RFI excision in uGMRT
Buch+ 2019

Summary

- SKA is made possible after decades of innovation
- But requirements have grown due to
  - Date intensive science
  - Dynamic range
  - Sensitivity
  - Sparse measurements of information
  - Time and angular resolution
  - Field of view
- Difficult for current techniques, algorithms, and instrumentation to scale sustainably
- Must start thinking beyond SKA as it is full of opportunities