# GPU-enabled Imaging of an Irregular nonFlat Sky

## (with Warped Snapshot Imaging)

Lincoln Greenhill (Harvard / Smithsonian)
D. Mitchell, G. Bernardi, M. Clark, S. Ord

Abstract: science, warped snapshot imaging, ~~gridding~~, ~~deconvolution~~, GPUs, LEDA, ~~HERA, SKA~~ ? !

# Goal: Fill the Last Gap in the Cosmological Record

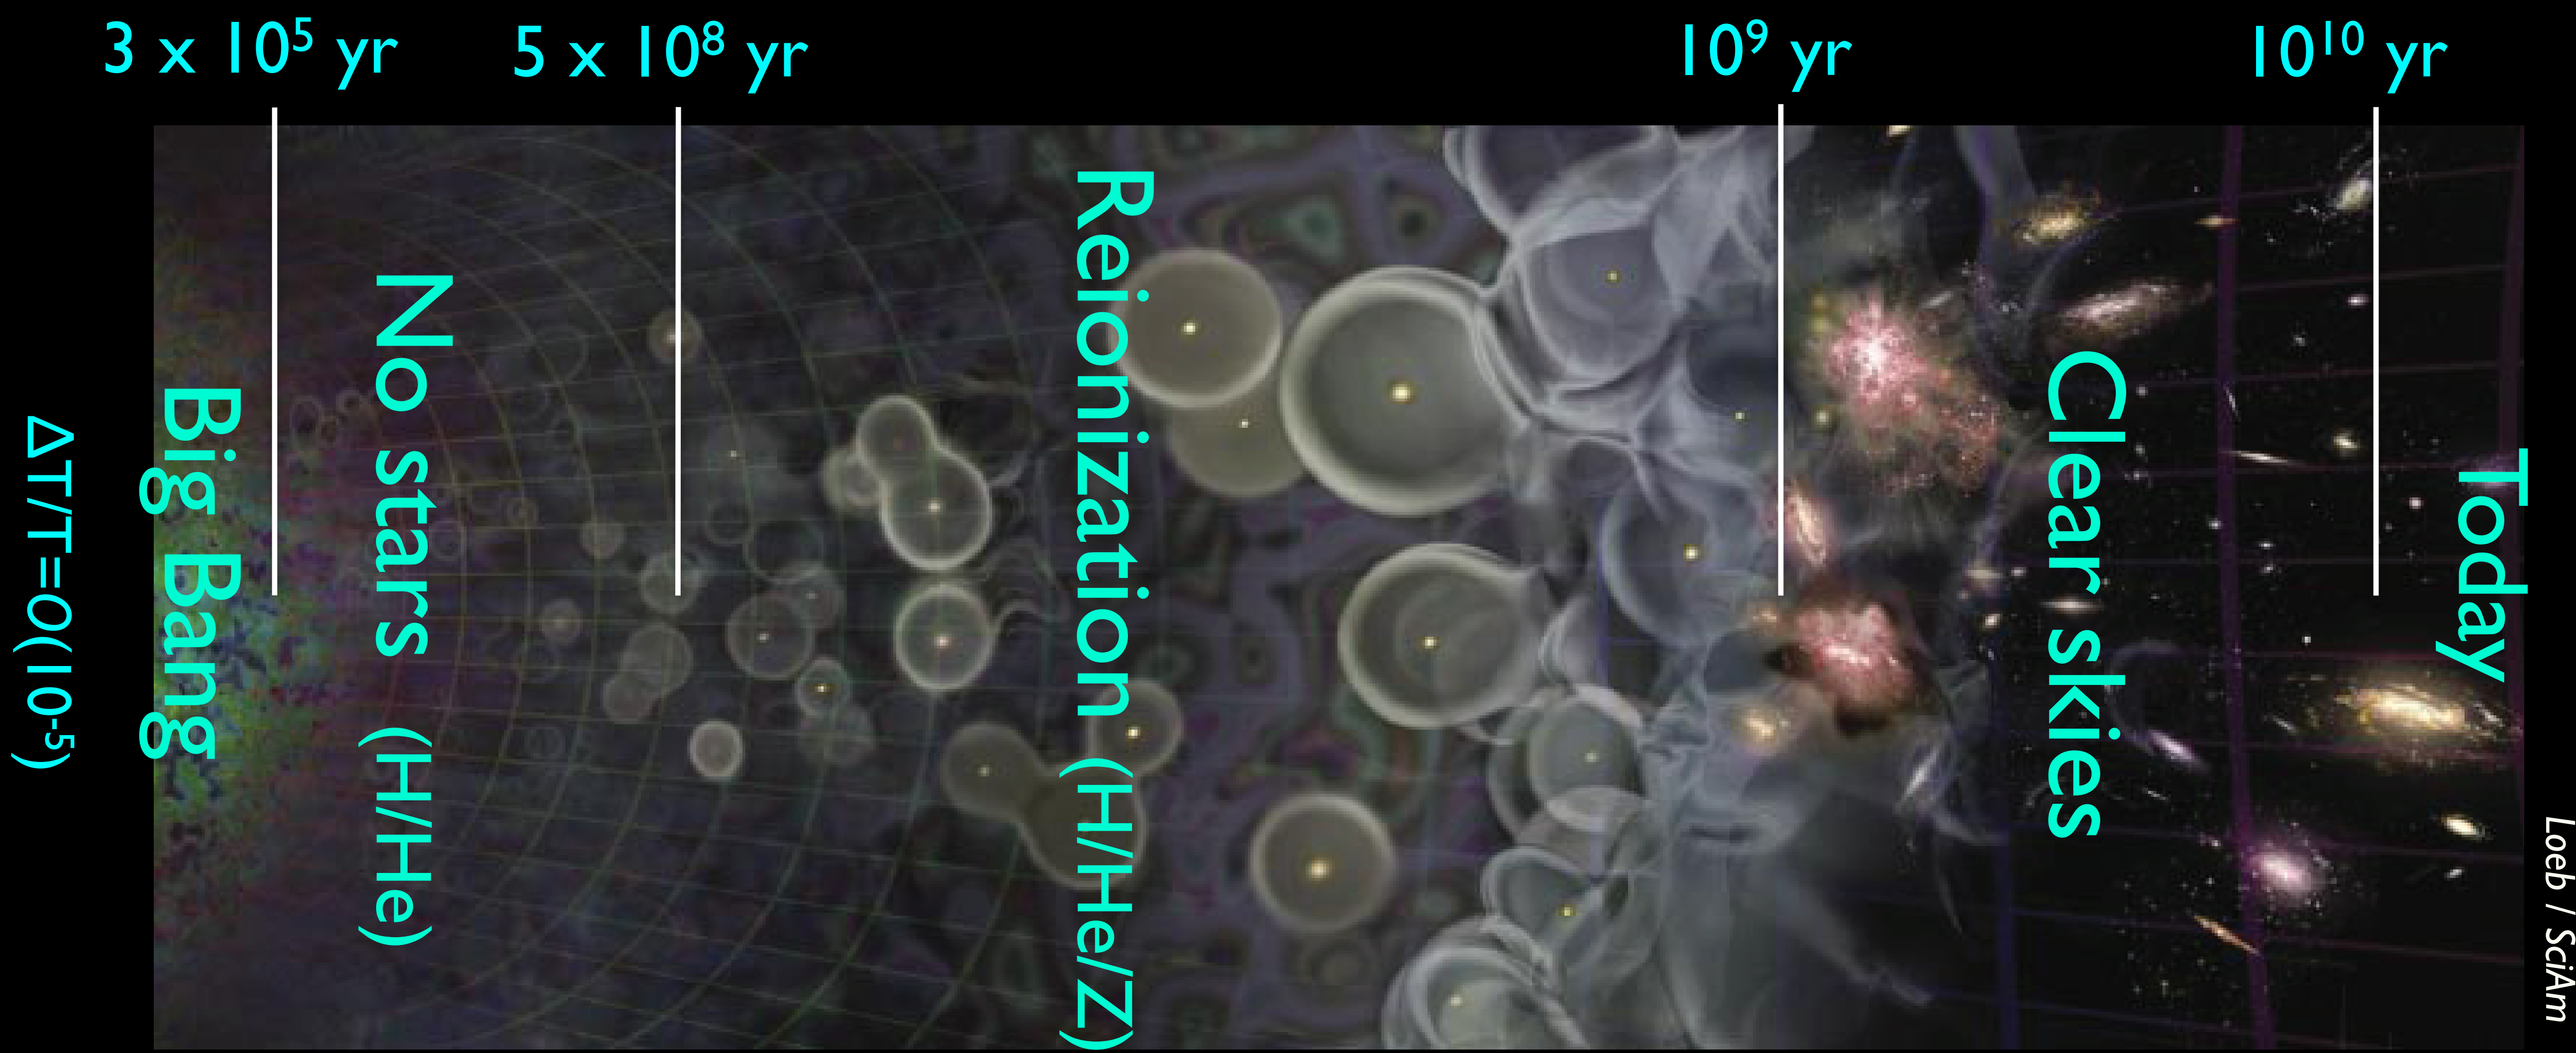


- When / how did the first stars & galaxies form?
  - the birth of the Universe we know today - ionized and metal rich
- The $\lambda$21cm hyperfine transition ($1^2s_{1/2}$) is a unique tracer
  - broadly distributed on the sky and in redshift (age)

Fragmentation of dark matter into halos
Coalesence of gas, driven by gravity...
Collapse into stars...
Supernova explosions of stars...
Distribution of heavy elements through the cosmos...
Recycling of material into new stars, etc

2009: Ralf Kaehler, Marcelo Alvarez & Tom Abel

Fragmentation of dark matter into halos
Coalesence of gas, driven by gravity...
Collapse into stars...
Supernova explosions of stars...
Distribution of heavy elements through the cosmos...
Recycling of material into new stars, etc

2009: Ralf Kaehler, Marcelo Alvarez & Tom Abel

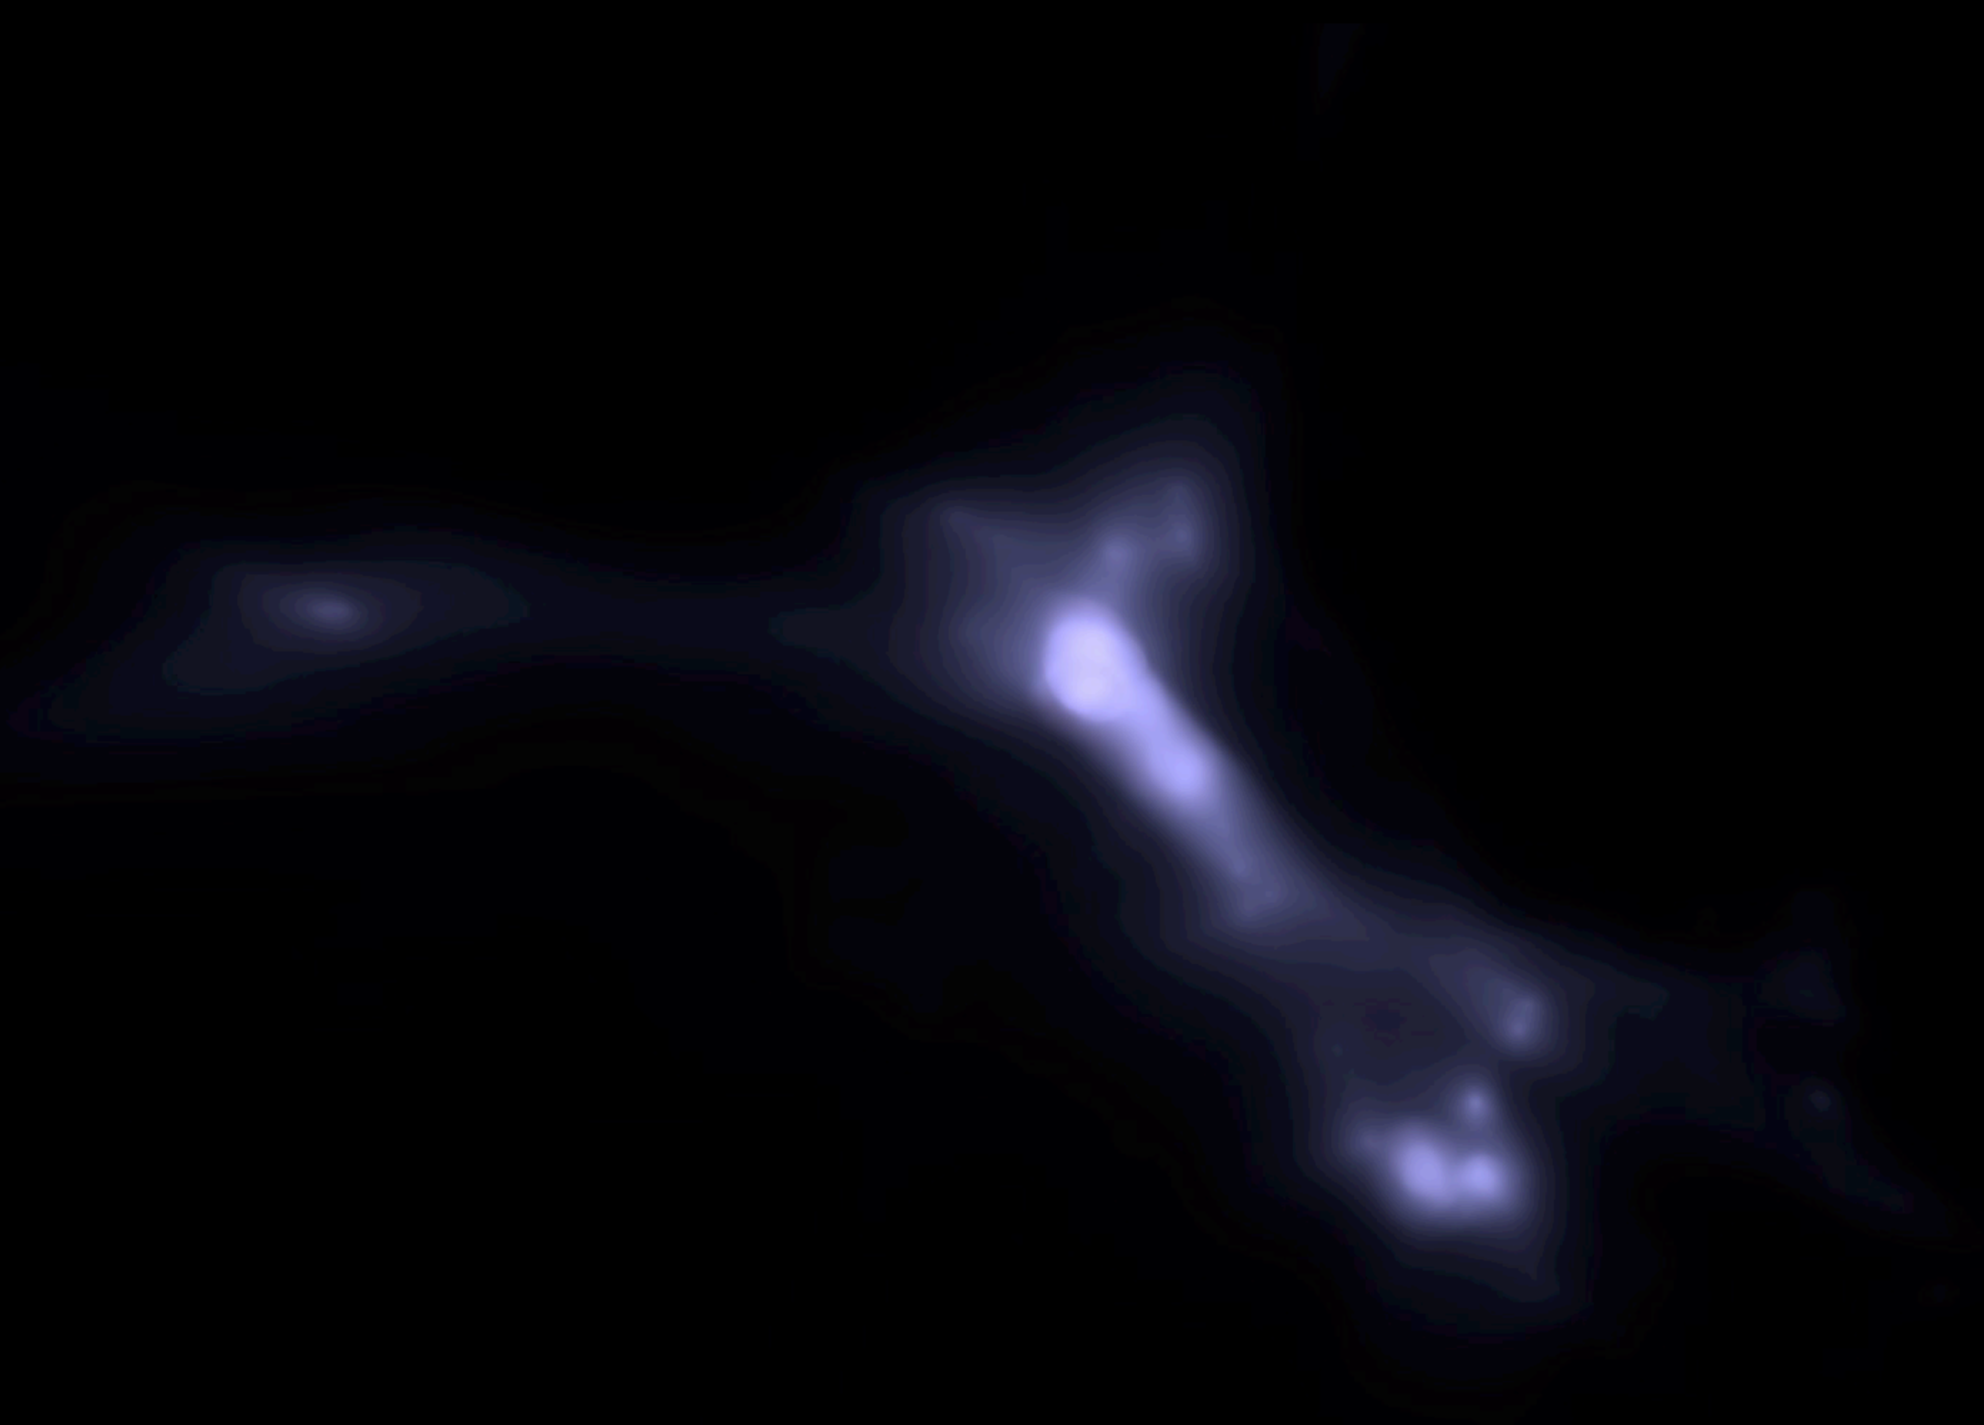

2009: Ralf Kaehler, Marcelo Alvarez & Tom Abel

# Simulation of Reionization

Neutral H →

Evolution in time
and frequency

2009: Ralf Kaehler, Marcelo Alvarez & Tom Abel

# Simulation of Reionization

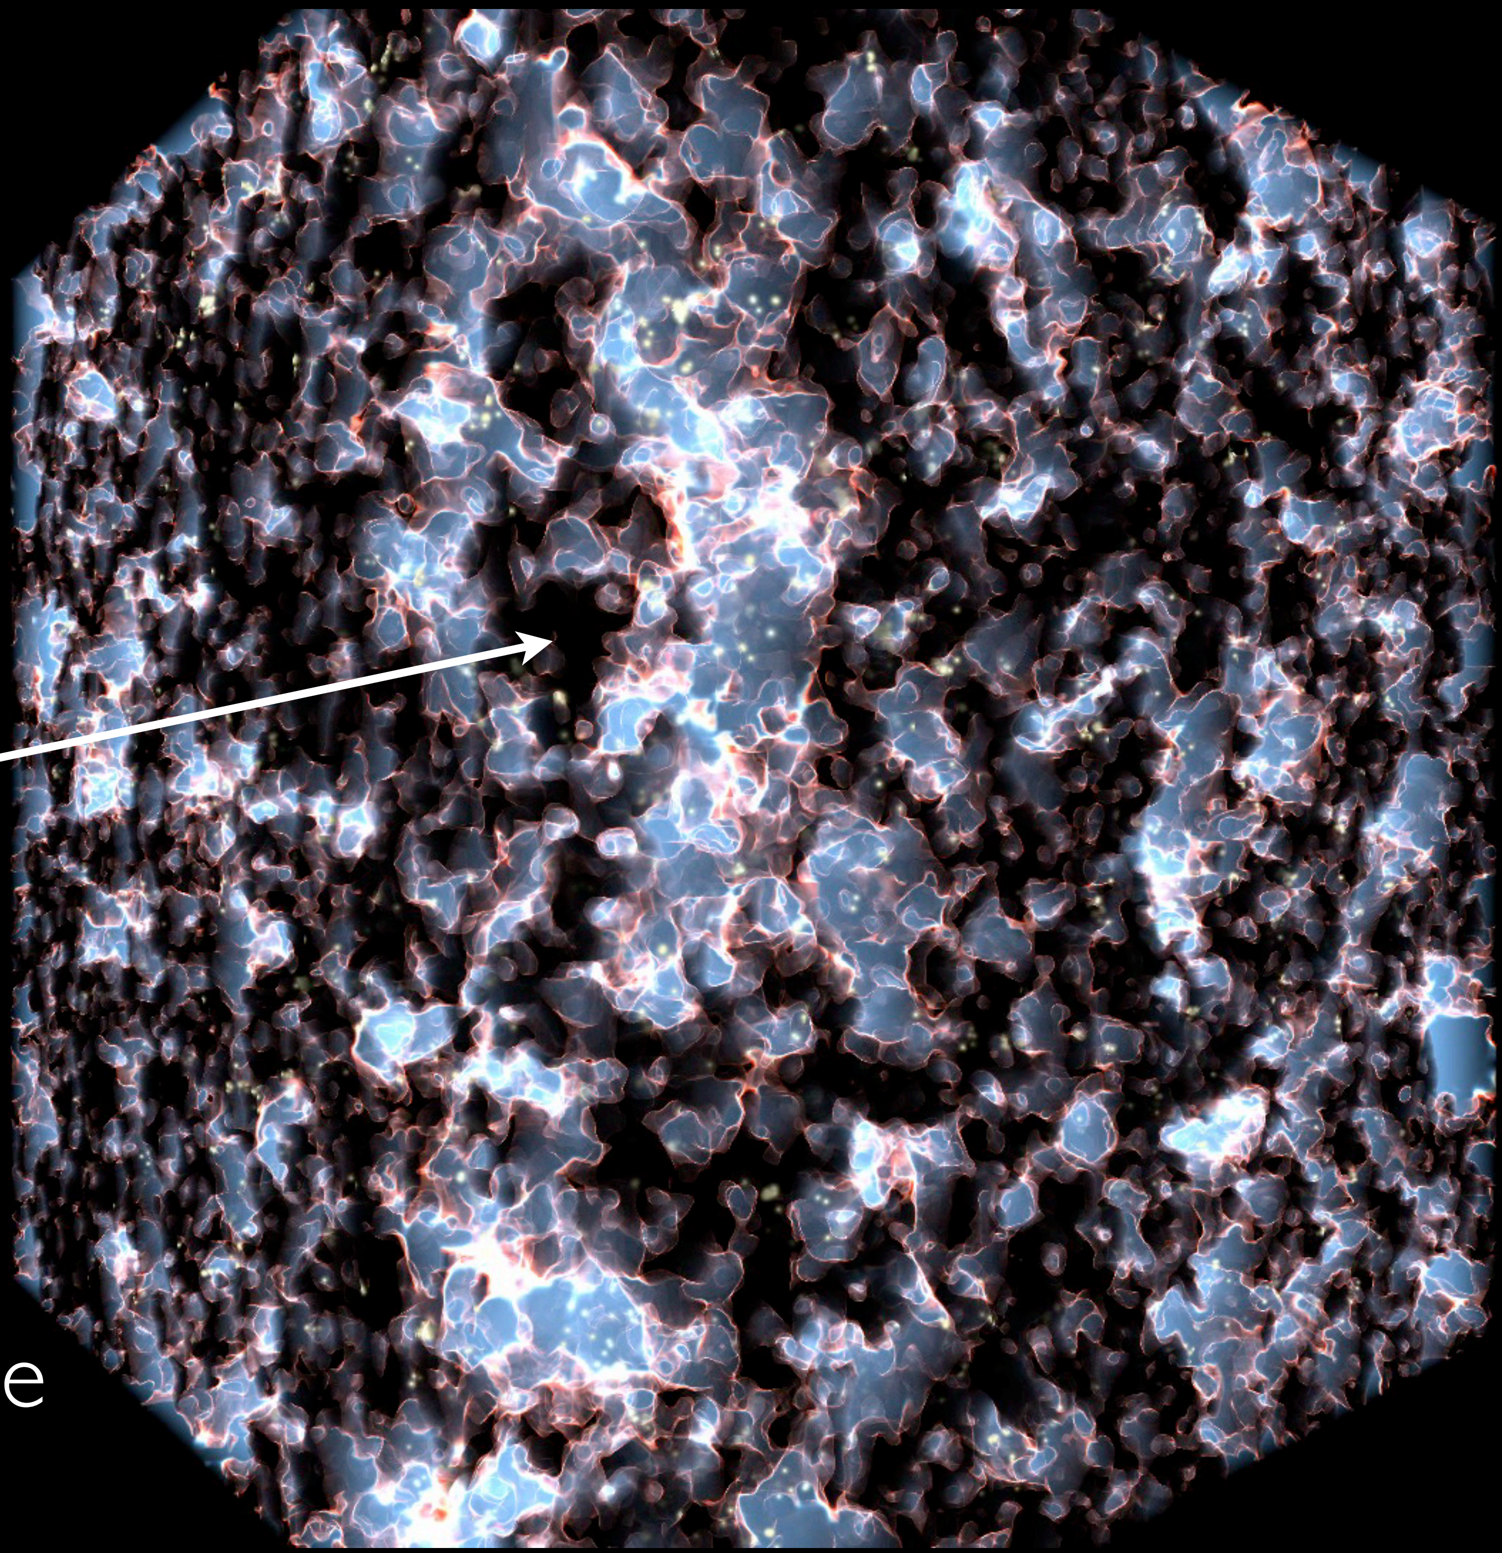

Neutral H

Evolution in time and frequency

2009: Ralf Kaehler, Marcelo Alvarez & Tom Abel

# Challenges

- "extreme" requirements
  - intricate structure in HI emission distribution ... large-N interferometer*
  - emission is distributed over the sky ... v. wide field of view
  - high redshifted: 1420 MHz $(1+z)^{-1}$... VHF (30- 200 MHz)
  - weak signal (1 : $10^{4\text{-}5}$) of foreground ... big $A_e$, good calibration
- dipole receptors appear well suited
  - but introduce fundamental challenges
    - non-flat sky
    - continuous brightness distribution on many angular scales
    - ionospheric distortion
    - direction dependent gains
    - variation element to element
    - low gain ➟ large numbers of receptor elements – $O(10^2$ - $10^4)$
      - » $O(N^2)$ compute & storage challenges

*** Single ideal antenna adequate to provide angle-averaged spectrum ➟overall temperature.**

# Challenges

- "extreme" requirements
  - intricate structure in HI emission distribution ... large-N interferometer*
  - emission is distributed over the sky ... v. wide field of view
  - high redshifted: 1420 MHz $(1+z)^{-1}$... VHF (30- 200 MHz)
  - weak signal ($1 : 10^{4\text{-}5}$) of foreground ... big $A_e$, good calibration
- dipole receptors appear well suited
  - but introduce fundamental challenges
    - non-flat sky
    - continuous brightness distribution on many angular scales
    - ionospheric distortion
    - direction dependent gains
    - variation element to element
    - low gain ➟ large numbers of receptor elements – $O(10^2 - 10^4)$
      - » $O(N^2)$ compute & storage challenges

*** Single ideal antenna adequate to provide angle-averaged spectrum ➟overall temperature.**

# Low-Frequency Foreground Sky

model

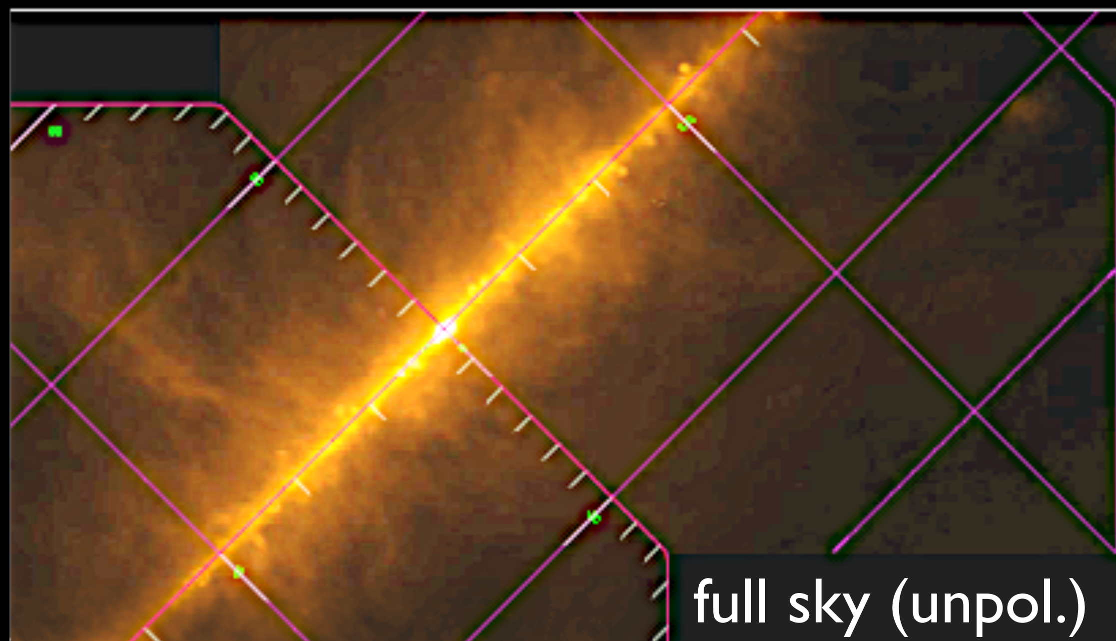


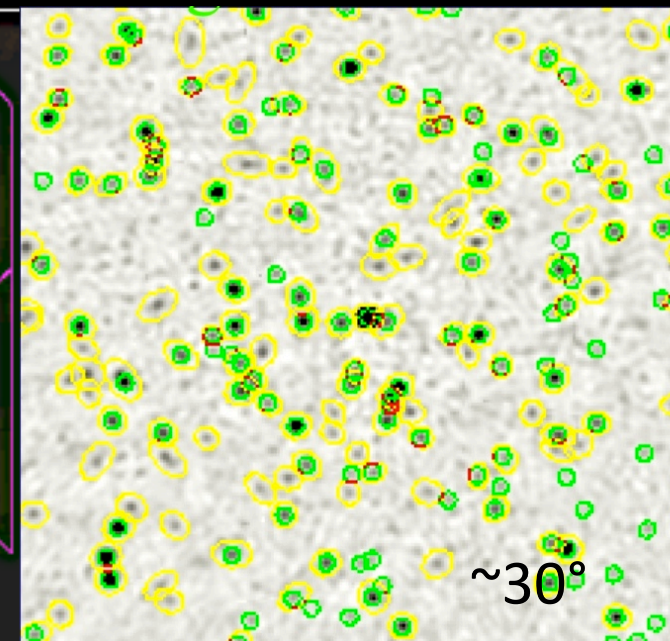


~150 MHz data

Briggs (ANU)

model

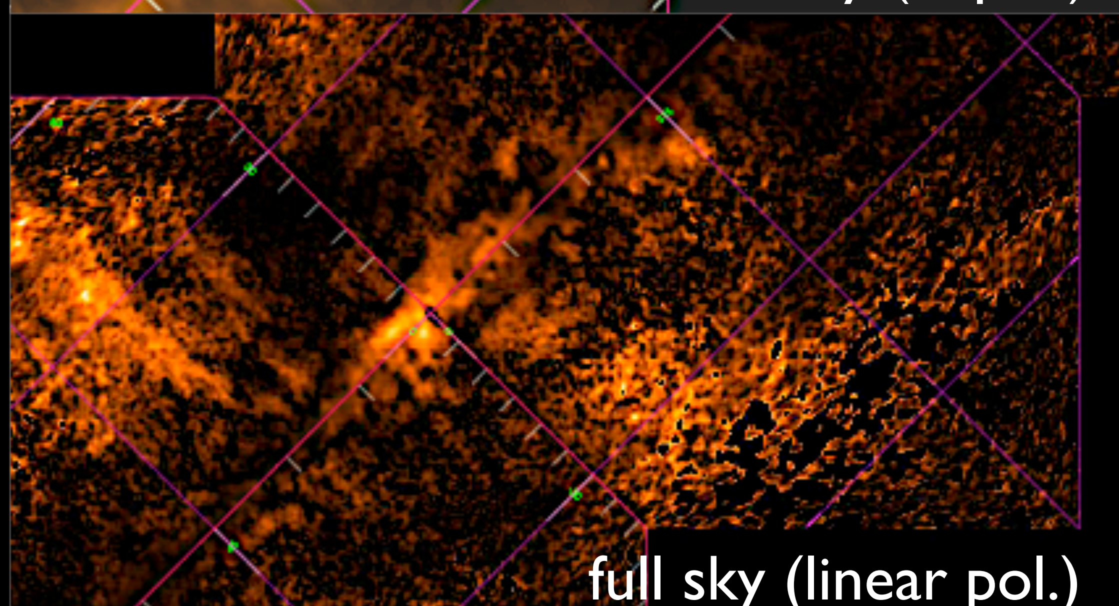


Ord (Curtin)

Emission is everywhere. Cal. errors ➟ pol. mode mixing

Cold sky: $10^{4\text{-}5}$x science signal

LOFAR


PAPER


*Parsons, Bradley, et al.*

LWA
LEDA
new start
15/8/11


Warped snapshot application

*Greenhill*
*Werthimer*
*Taylor*
*Ellingson*

LOFAR


PAPER


*Parsons, Bradley, et al.*

LWA
LEDA
new start
15/8/11


Warped snapshot application

*Greenhill*
*Werthimer*
*Taylor*
*Ellingson*

# Wide-field Approaches

- c.f., Cornwell et al., arXiv:0807.4161
- 3D transforms
  - FFT (sparse volume)
  - DFTs (expensive)
- 2D transforms
  - image-plane facets
  - uvw-space facets
  - warped snapshots
  - w-projection
- Combinations (e.g., peeling and segmenting)

courtesy D. Mitchell

# Wide-field Approaches

- c.f., Cornwell et al., arXiv:0807.4161
- 3D transforms
  - FFT (sparse volume)
  - DFTs (expensive)
- 2D transforms
  - image-plane facets
  - uvw-space facets
  - warped snapshots
  - w-projection
- Combinations (e.g., peeling and segmenting)

*W.S. appeal for compact filled low-freq arrays*

- coplanar-array snapshots obey sky-vis FT relation
- exact correction for wide-field distortion
- fixes ionospheric distortion (refractive approx.)
  - snapshot imaging + image resampling
  - correction in image plane
  - imaging cadence set by ionospheric τ-scales
- enabled by
  - good snapshot beam (large-N, filled aperture)

courtesy D. Mitchell

# Fourier sampling

- snapshot imaging requires
  - filled array
  - minimum redundancy array
  - subtraction of bright objects in Fourier plane
- LEDA example (right)
  - optical telescope-like response
  - subtraction ↔ occultation

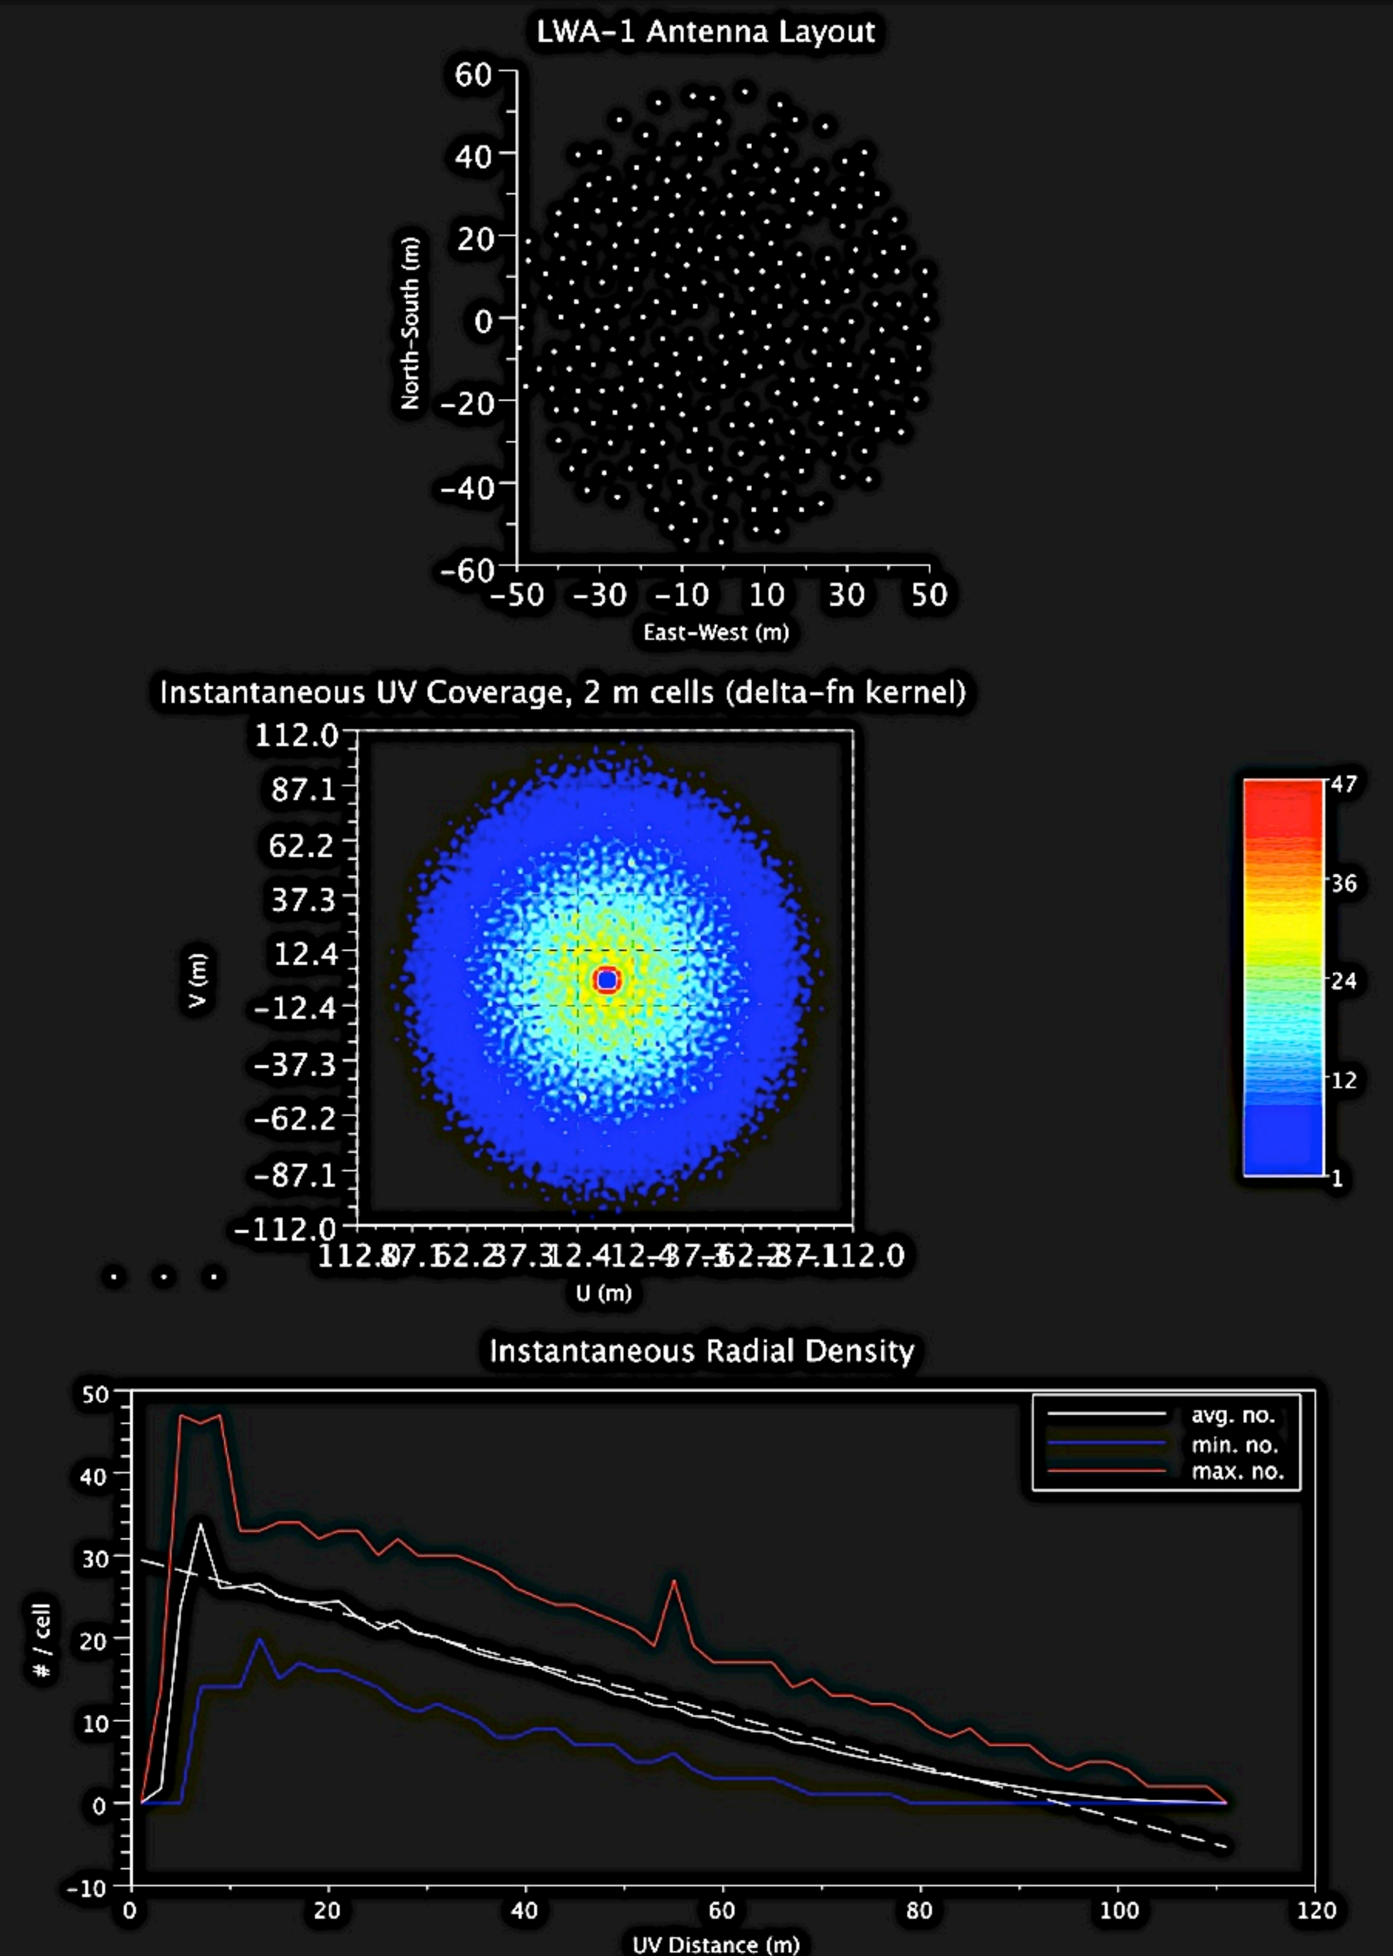

# Warped Snapshots

## Co-adding snapshots mimicks interferometric synthesis

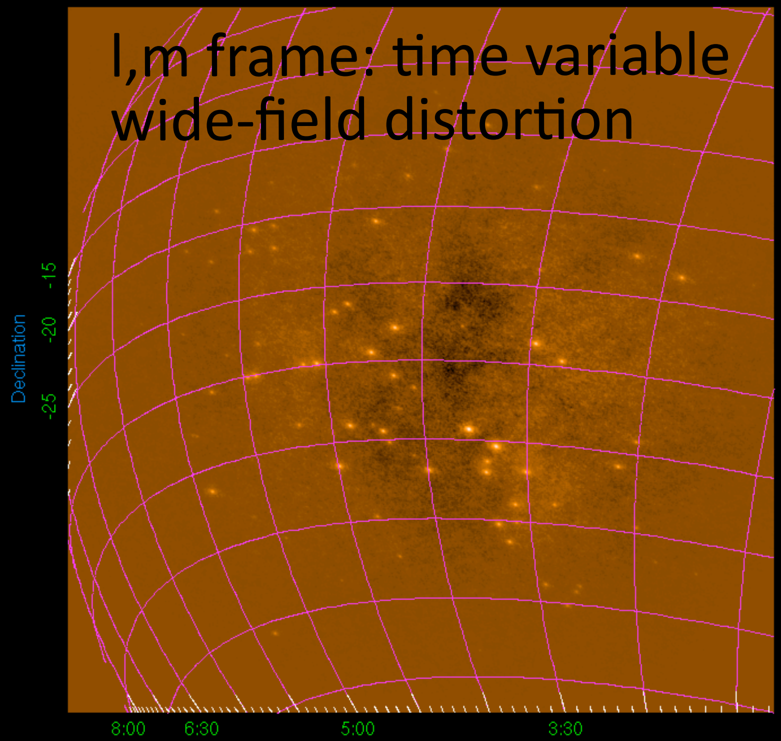


### Simulated data

- $7^h$ sky motion (HA=$\pm 3.5^h$)
- Interpolate snapshots to constant frame
- co-add
- weighted by (gain pattern)$^2$, c.f. A-projection, mosaicing
- pipeline (mostly) running on GPUs

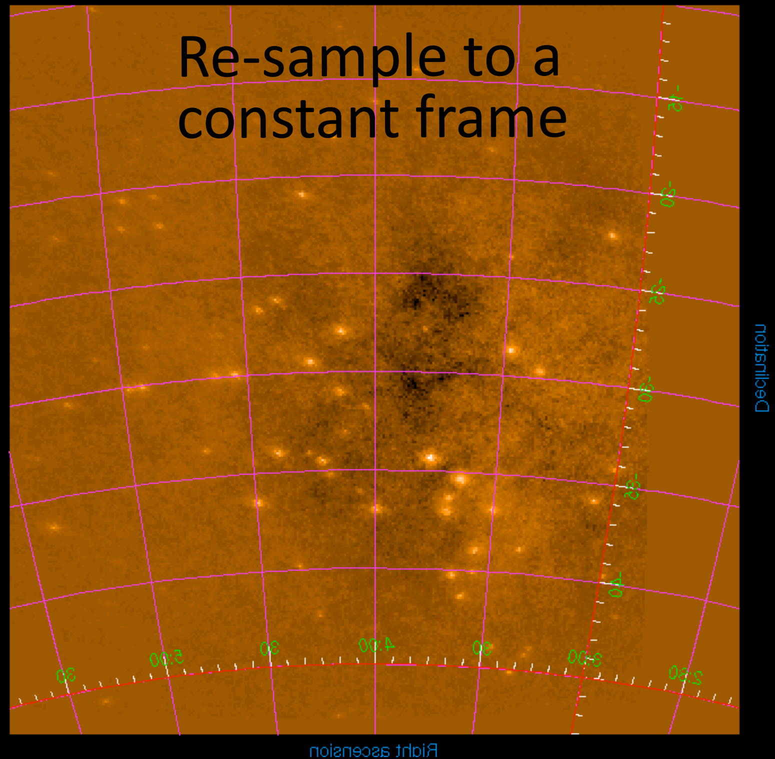


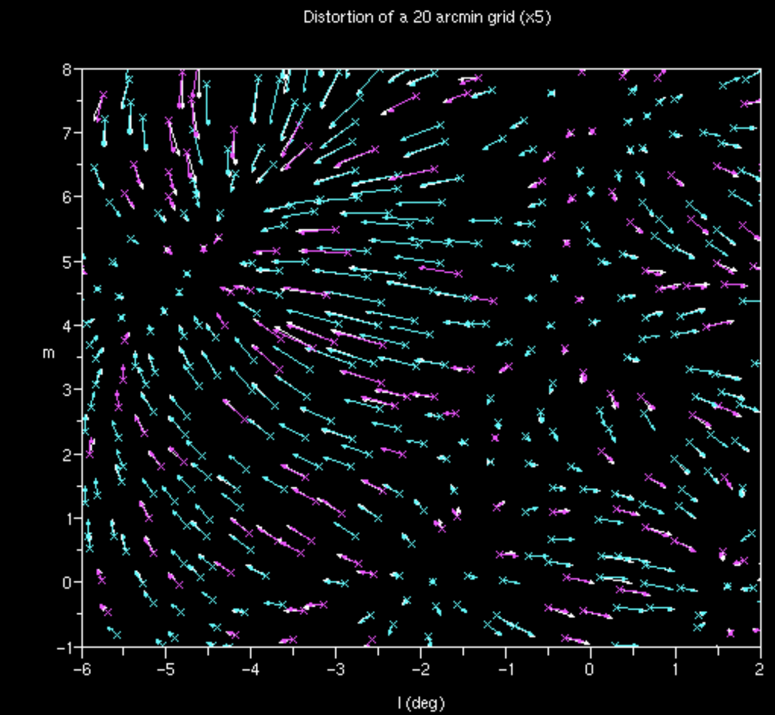


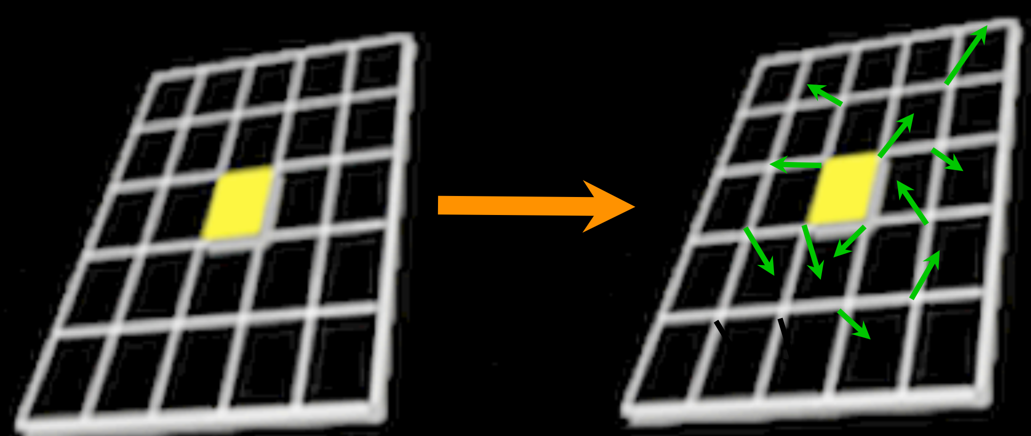

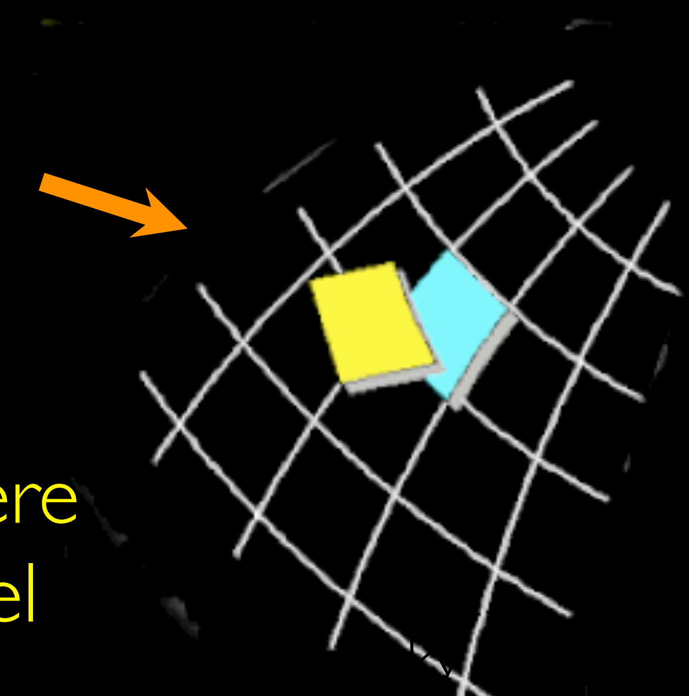

Rubbersheet approx. of ionosphere repaired by resampling, in parallel

# Computation - Motivating GPU Use

- Dipole arrays generate stringent requirements
  - owing to large-N & irregularity; downside of simple receptors
  - also stems from complex source distribution
- Present generation instruments @ O(100) TFlop $s^{-1}$
  - calibration/imaging (MWA)
    - N=512 x 2 polarizations x 1000 channels ($10^5$ Fourier samples / dump)
    - M=$10^6$ image pixels
    - interpolation onto all-sky frame
    - O(10) TF $s^{-1}$ per pass through data
      - peeling but not iteration of sky / instrument model
  - Forward modeling of sky & instrument is at least as great
    - involves iterative execution of same calibration/imaging system
- Scaling $O(N^2M^2F)$ to next generation (10 PF/s) within reach

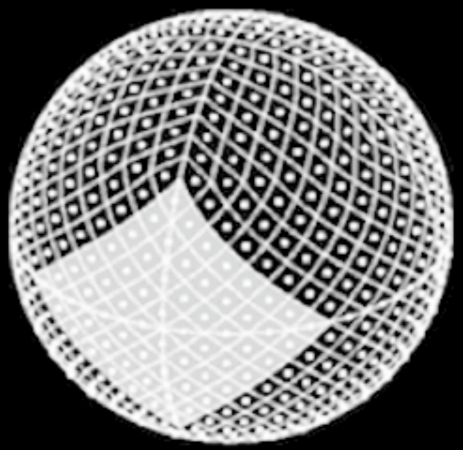

HealPix frame

# Application of GPUs

- Tailored to high arithmetic intensity (Flops/byte)
- Tiered memory
  - likely feature of exascale machines
- Some rules of Thumb
  - attempt end-to-end processing on GPU
  - engineer parallelism to hide transfer latency
    - on/off chip AND
    - within chip
    - use async. bilateral data transfer
  - thread carefully; not necessarily naturally
    - by Fourier sample, image pixel, time bin, ...
    - stay aware if access divergence & race conds.
  - identify units of data that fit hardware
    - register/cache size & access patterns
  - avoid synchronization & atomics
  - off-load tasks to CPU (heterogeneous comp.)

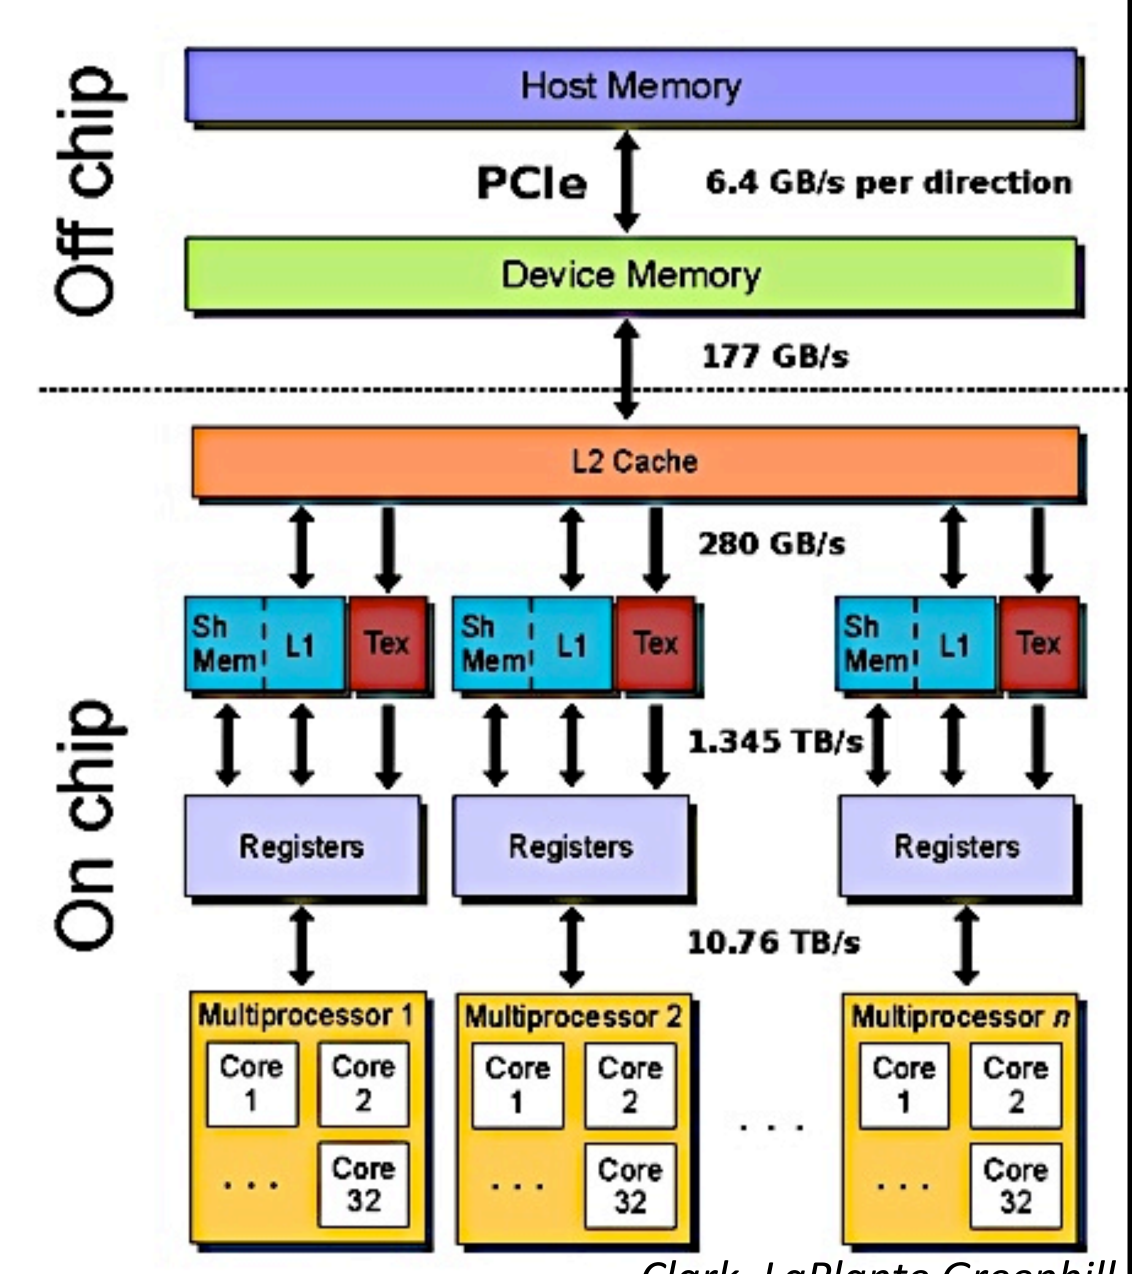


Clark, LaPlante Greenhill

# Application of GPUs

- good applications, e.g.
  - FT
  - gridding & convolution (structured as "gather" op.)
  - re-sampling
  - transposes
  - large matrix / vector problems
  - Texturing operations (e.g., type conversion)
  - cross correlation - X-engine @ 80% theoretical
- O(1) TFlop $s^{-1}$ per present day GPU is feasible
  - depends strongly on quality of adaptation to GPU env.
  - bird in the hand vs future many-core CPUs
    - GPU development environment is advanced
      - partial convergence of hardware coming
        - » CPU & GPU cores on one die; limited by Watts

Gridding

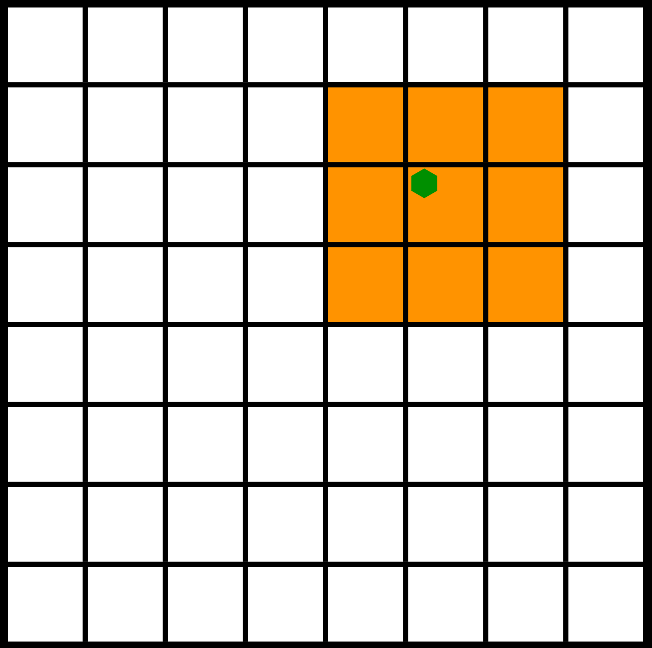

Scatter structure - BAD

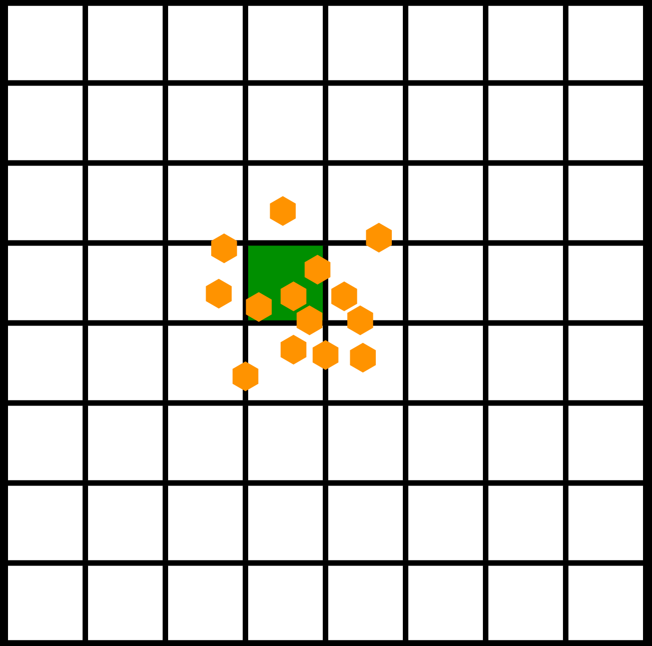

Gather structure - GOOD

*R. Edgar NVidia GTC09*

# GPU R/T Delivery of Fourier Samples
## (LEDA correlator)

Rules of thumb deliver
record performance for real-science application

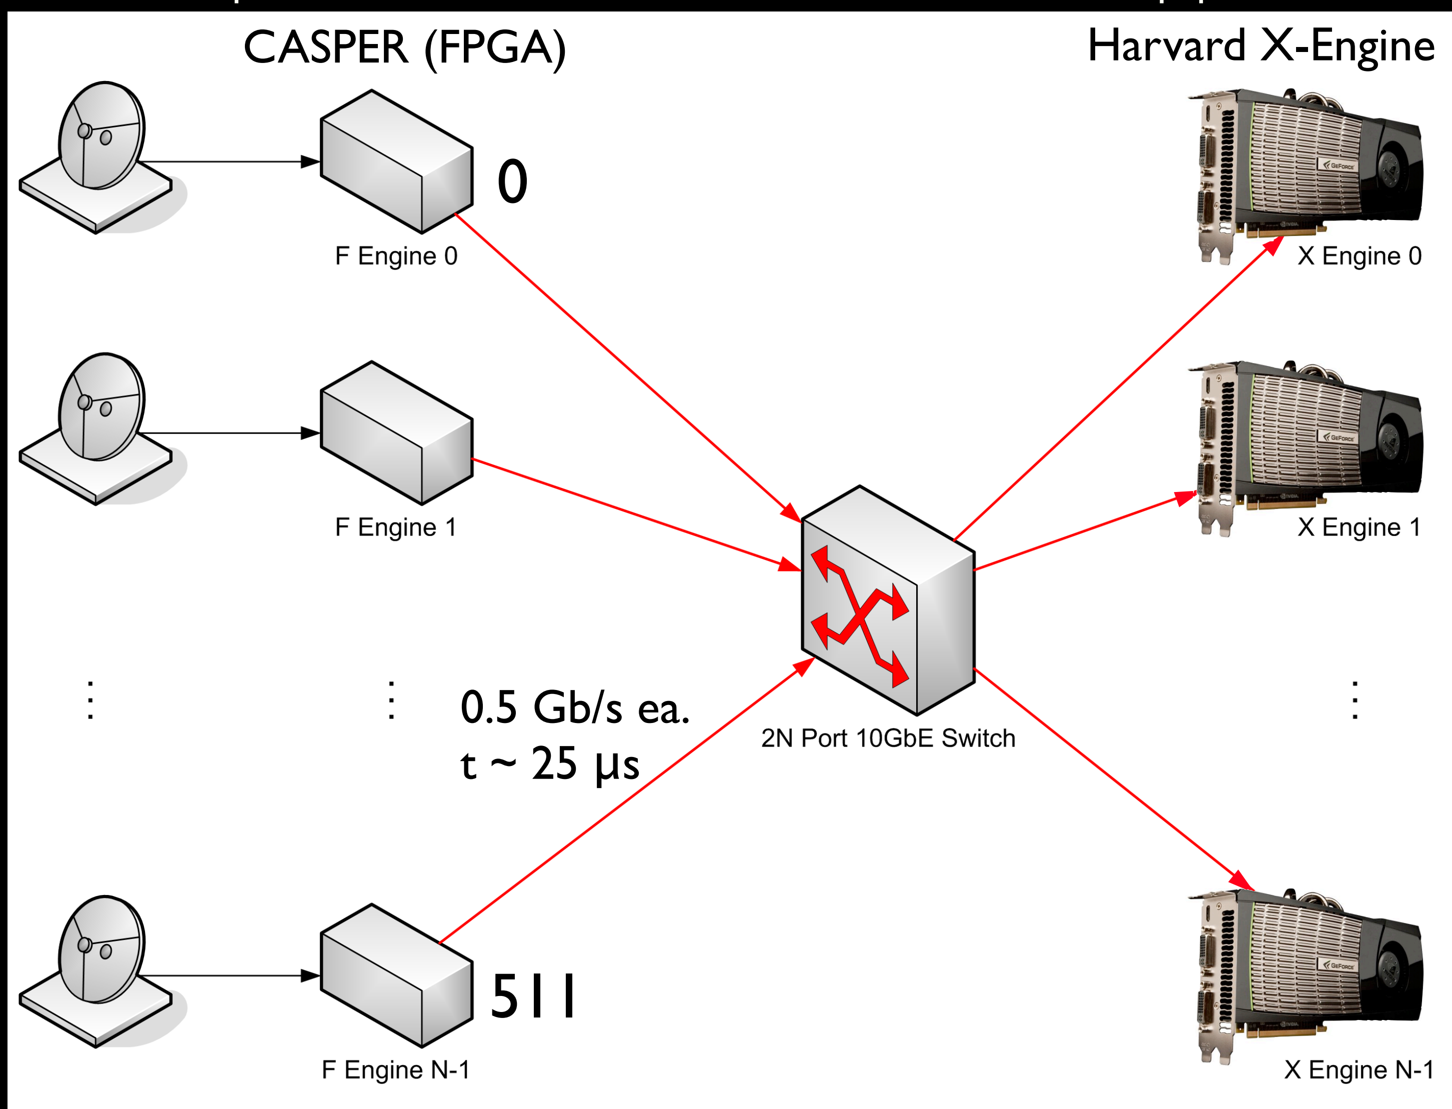


Clark, LaPlante, & Greenhill, arxiv:1107.4264 (source: /github.com/mikeaclark/xGPU)

Performance X-engine vs problem size

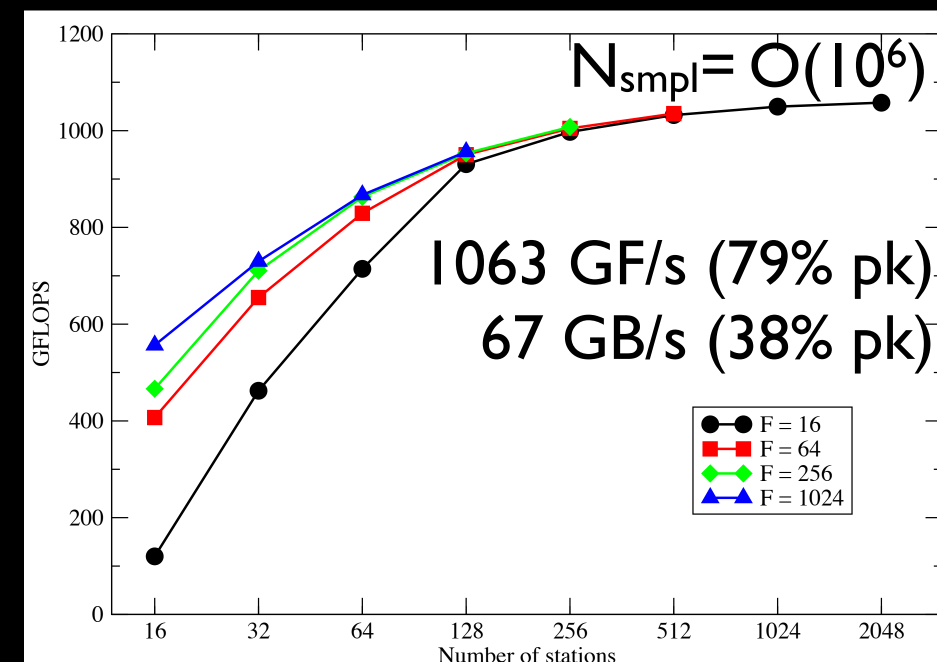


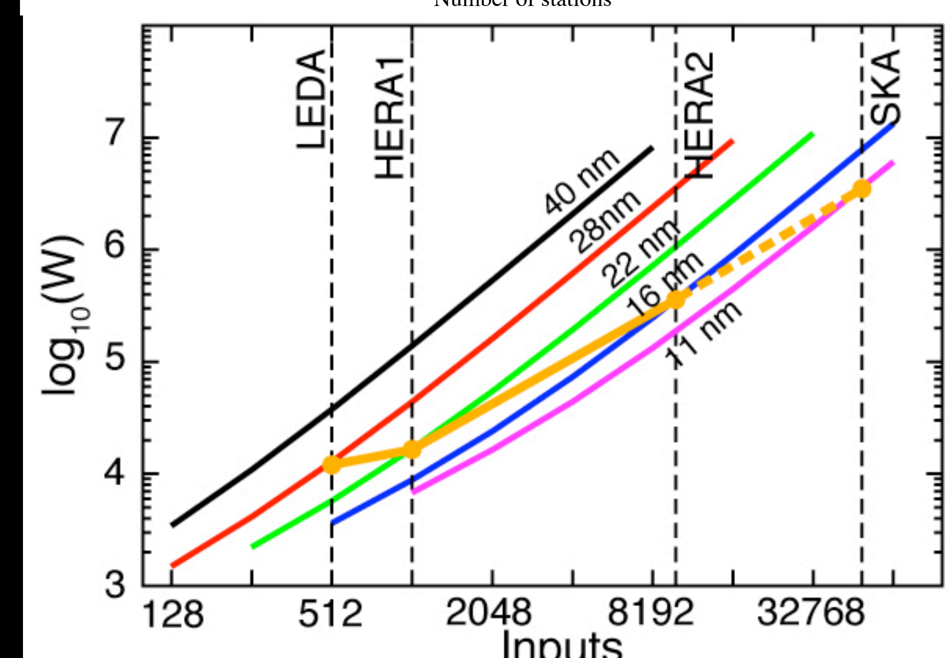


Time series ➡ FFT spectrum ➡ Cross-multiply

High-efficiency F/X GPU for noise limited signals in progress

# Summary

- Frontier cosmology drives large-N dipole array dev.
- Dipoles are simple, dipole arrays are not
  - cutting edge interferometry and computation make up balance
- Warped snapshot imaging promising for v. wide FOVs
  - counterpart to W/A-projection
  - best approaches are context dependent
- GPU computing is adaptable
    - maintain high arimethic intensity
    - threading problems & structuring data may not be natural
- for applications that depend on residual images
  - gridding kernels that maximize SNR may not minimize source confusion/contamination
- multiple gridding schemes w/in one application may be appropriate

*-end-*