Title: EHT Correlation

Date: Nov 13, 2014 02:55 PM

URL: http://pirsa.org/14110103

Abstract: Over the next few years the Event Horizon Telescope will greatly expand its capabilities from the current 8 Gbps at a few sites to 64 Gbps at a much larger number of sites. The good news is that this processing can proceed through 4 independent (16 Gbps) processing stages of 2 GHz of bandwidth. The bad news is that EHT stations come in multiple "flavors" each posing its own issues for correlation with DiFX (the current correlation option). This talk will discuss some of the issues and the road ahead.

EHT Correlation

Geoffrey B. Crew

MIT Haystack Observatory

EHT 2014 - Perimeter Institute

November 13, 2014

Acknowledgements

- A great many individuals have brought us to this point:
  - A. Deller, W. Brisken & the DiFX developer/user community
  - R. Cappallo, M. Titus & Haystack Correlation/Post-processing team
  - APP Collaborators (especially W. Alef, H. Rottmann & N. Pradel for some of the materials presented in this talk).

Overview

- EHT Correlation Requirements
- Haystack and Bonn Correlators
- Technical issues for DiFX/Post-processing
  - Mark6 playback
  - Multiple data streams
  - Polarization
  - Channels/FFT/DFT/zoom-mode/resampling
  - ALMA sub-band tuning
  - Planned work
- Logistics

EHT Correlation Requirements

- Use DiFX software correlator
- Expansion to ~8 stations over 5-year period
- Bandwidth:
  - Previously 8 Gbps (2 x 512-MHz band, dual pol.)
  - March 2015: 16 Gbps (2-GHz band, dual pol.)
  - Target (ALMA): 64 Gbps (4 x 2-GHz band, dual pol.)
- More than one campaign per year (2? 3?)
- Timely reduction of data (for media reuse)
- ~400 DiFX processing cores required

Current Bonn VLBI Korrelator

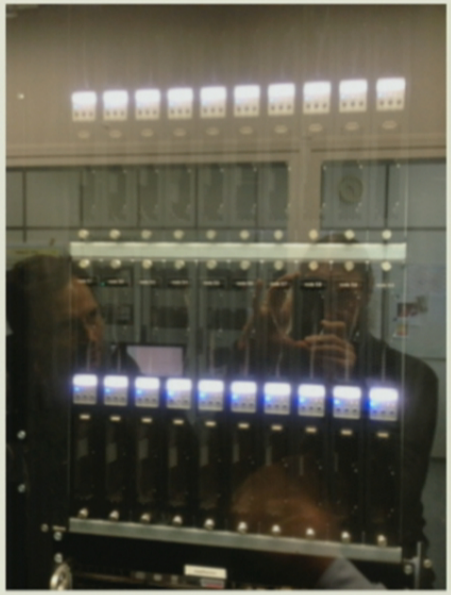

Current Haystack VLBI Correlator

Cluster Upgrade at Haystack

- Phased upgrade strategy:
  - Minor upgrade to support APP commissioning/ March 2015 campaign is in progress (have funds; equipment procurement has begun)
  - Significant upgrades for March 2016/2017 based on lessons learned from March 2015 production using MSIP funds
  - Target ~8 site playback units with 64 Gbps processed with 4 x 16 Gbps passes

Mark6 Disk Testing (8x6TB He)

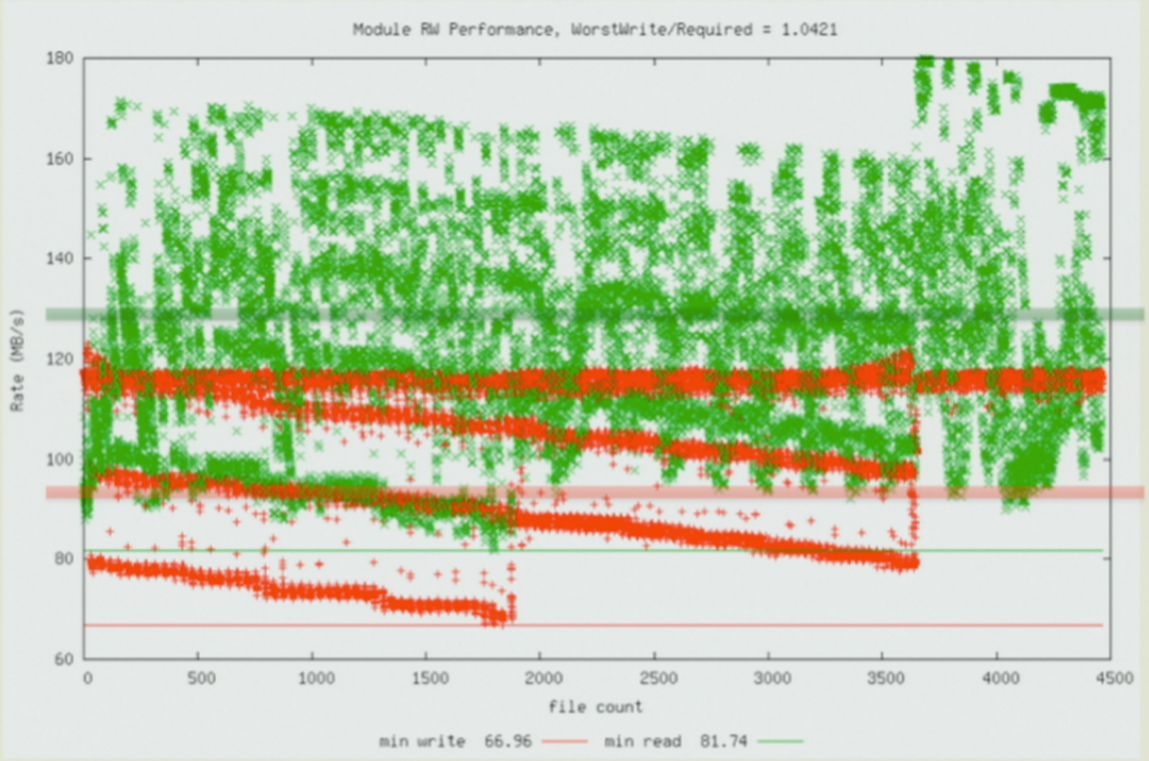

Read

Write

Technical Issues, 1

- Mark 6 Playback
  - 10 GbE limit is 40 Gbps (2 dual-ported NICs)
  - Individual disk limit is ~125 MB/s or ~8 Gbps/mod (required record rate is ~4 Gbps/module)
  - "robust" scatter-gather file system is recorded
  - ***vdifuse*** layer has 10-20% overhead for de-scatter "gather" process. Essential functionality exists now, but more capability is planned.
  - Mark6 playback into DiFX correlation was tested in ALMA recorder mini-cluster.

Technical Issues, 2

- Multiple data streams
  - (Eventually) 8 x 8 Gbps input streams presented to 4 Mark 6 recorders
  - Recording to groups of 4 modules, but each input stream is on one 2-module sub-group
  - DIFX ***mpifxcorr*** supports multiple data streams in principal (but not yet [conveniently] in practice)
  - Changes to ***difxio/vex2dif*** would need to be modified (MPI, perhaps in 2015)
  - changes for the new VEX 2.0, perhaps

DiFX *mpifxcorr* Architecture

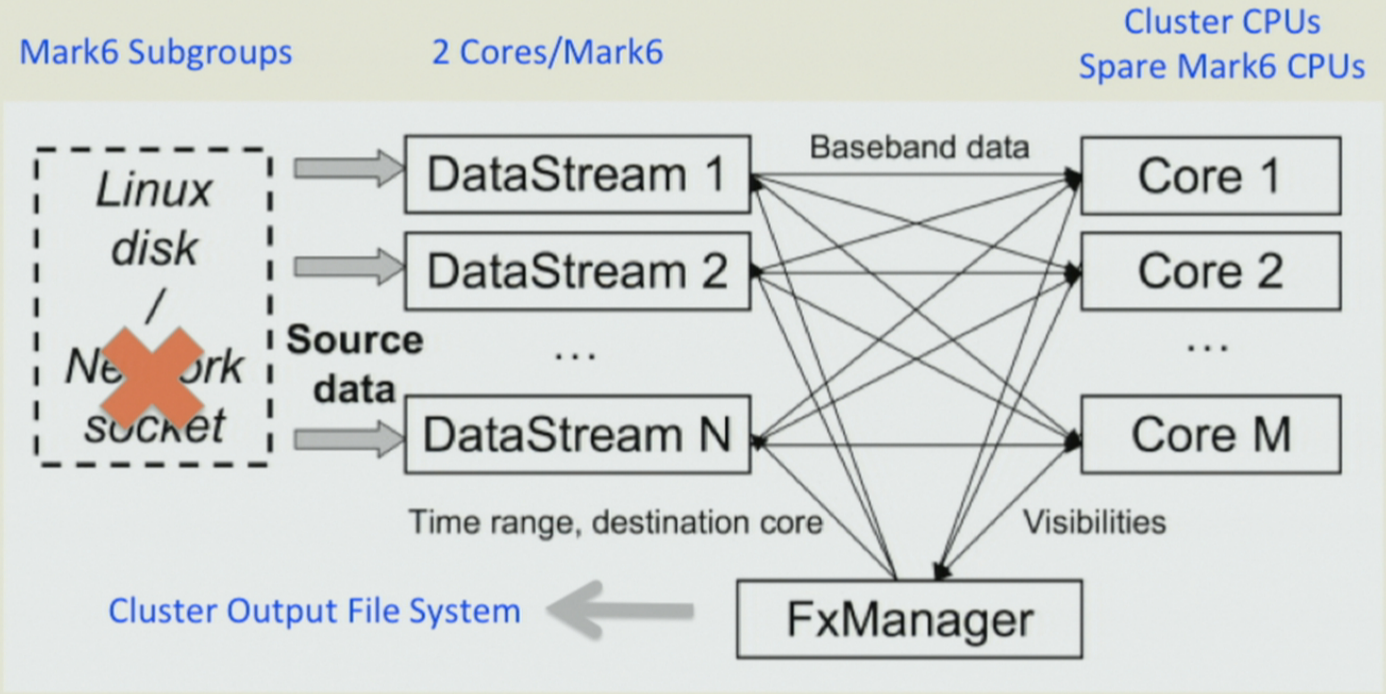

Technical Issues, 3

- Polarization
  - ALMA is X/Y, not R/L (quarter-wave plates for 66 antennas was not an allowed option)
  - Lie: DiFX doesn't need to know it's X/Y not R/L
  - ***polconvert*** tool(s) planned to "clean up" the DiFX mess (all that lying comes to no good end)
  - Calibration data (D terms) needs to be taken
  - For details, see Ivan Marti-Vidal's talk, later today

Technical Issues, 4

- Channels/FFT/DFT/zoom-mode/resampling
  - R1DBE: 15 x 32-MHz channels (64 Msps) per IF/pol
  - ALMA: 32 x 62.5-MHz channels (TFB -- 125 Msps)
  - R2DBE(now): 1 x 2048-MHz channel (4x1024 Msps)
  - Initial development was aimed for 32/64 v 62.5 MHz
  - R2DBE(future): 2^n-MHz channels? (if needed)
  - R2DBE(future): $2^n 5^m$-MHz channels? (if needed)
  - DBBC3 (in development) will be even more capable
  - ALMA flexibility/1-channel R2DBE allow other options
  - Accommodations to ***difx2mark4*** and **fourfit** remain

VDIF simulator: ***enoise***

- ***enoise*** can simulate non-0 baseline
  - Any frequency
  - Any baseline
  - Any bandwidth
- Delay model acc. ~100ps
- Remaining problems
  - Simulations are very long to compute (~0.1 day per channel per s in APP case)
  - No parallel computing
  - Long duration simulations untested

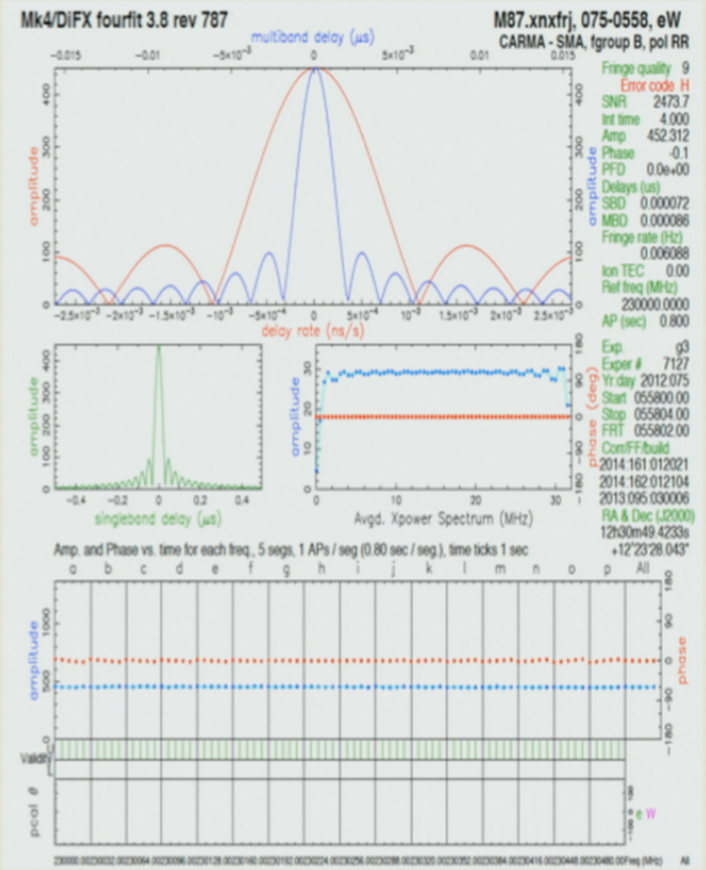

ALMA Quadrant-1 vs Quadrant-3

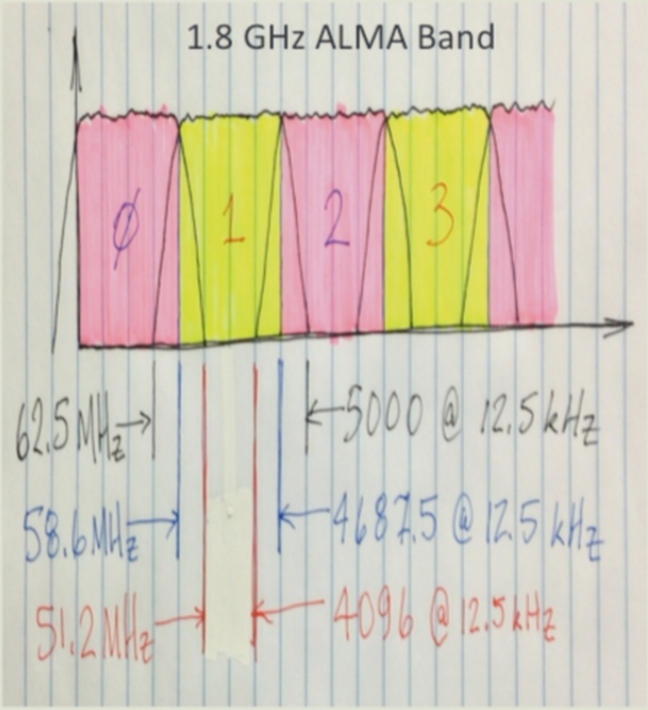

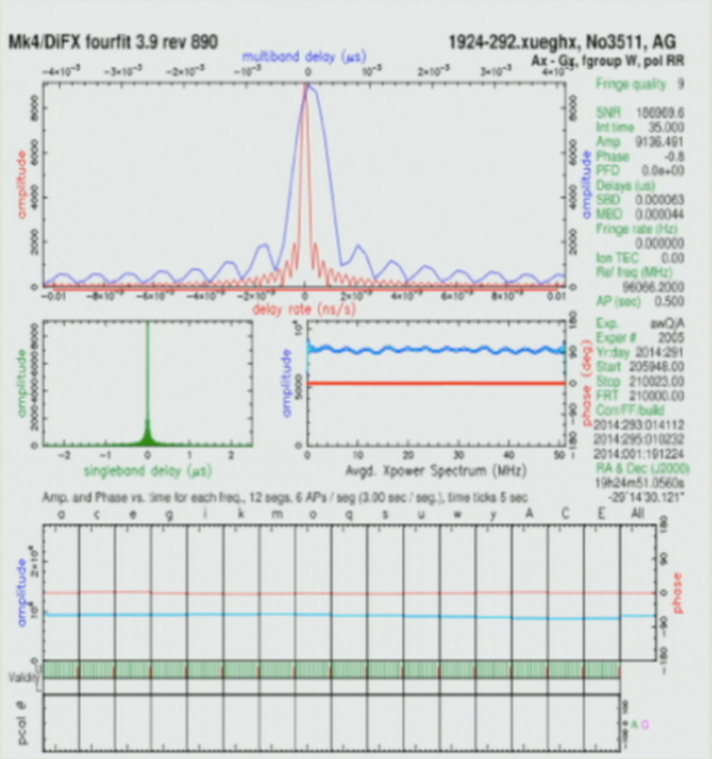

Technical Issues, 5

- ALMA sub-band tuning
  - Tunable filter banks tune in 2.0 GHz/2^{16} steps (30.517578125 kHz)
  - Per-sub-band LO offsetting is required in DiFX in order to compensate for the frequency mismatch
  - DiFX support with zoom mode by A. Deller based on work with ASIAA ***enoise*** simulator
  - Small LO offset tested with SPT equipment and the Wf/Gg geodetic R1DBE (v1.4)

Technical Issues, 6

- Planned work
  - ***enoise*** enhancements (until we have data)
  - Possible ***difx2mark4*** enhancements to better accommodate DFT vs FFT and zoom channels (*e.g.* round up # channels to nearest power of 2)
  - Modify ***fourfit*** (and other HOPS tools) to use DFT rather than FFT processing (i.e. relax assumptions of 2^n) and modify internal data structures for significantly larger data sets.