

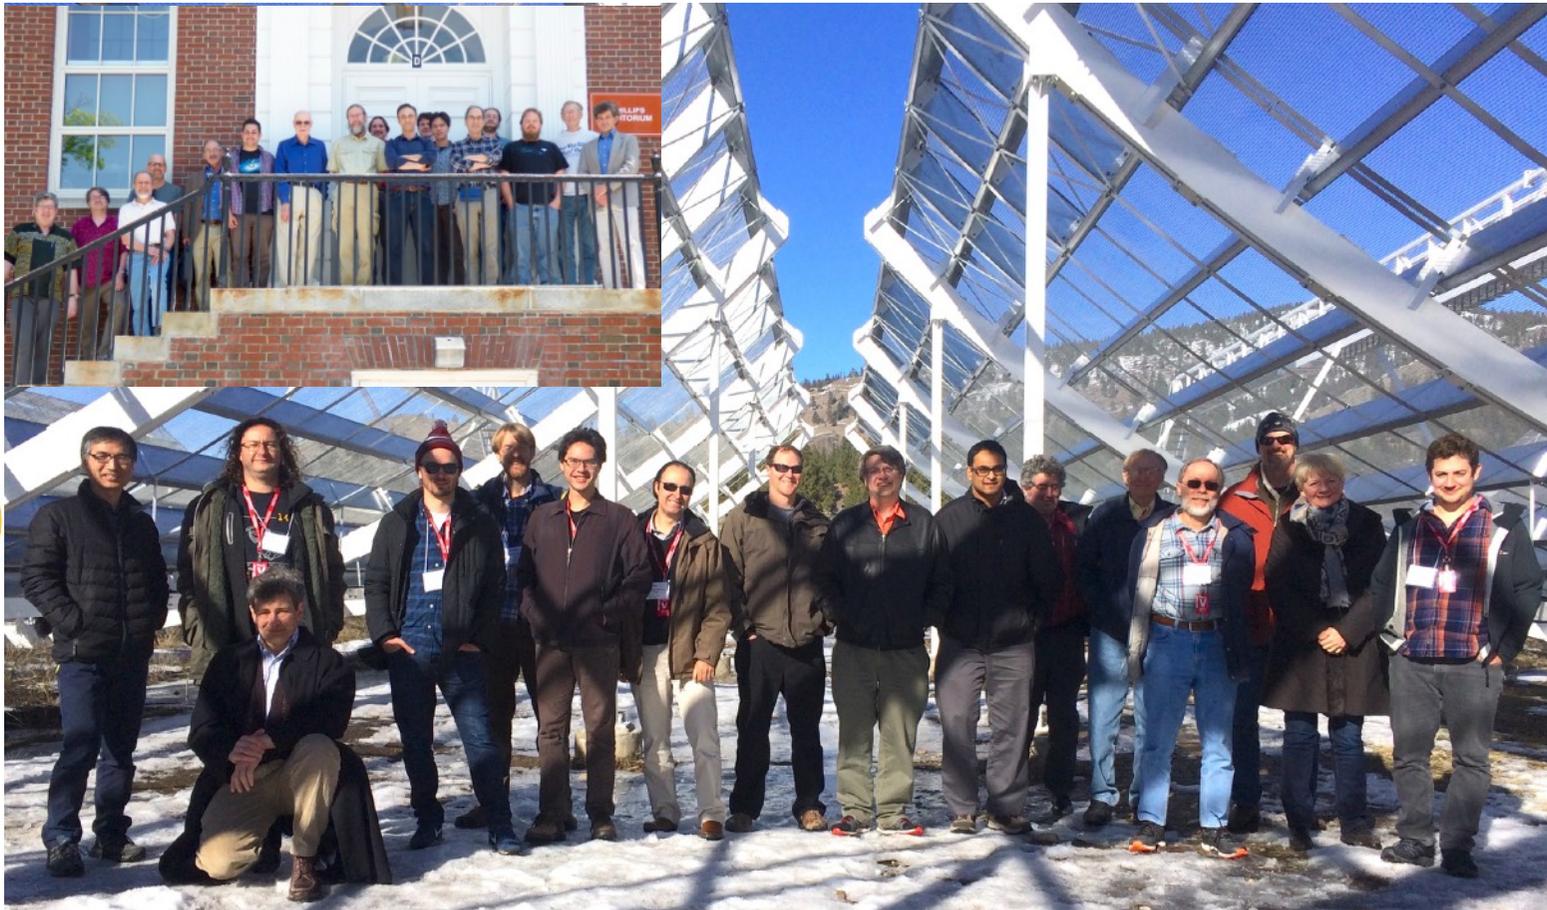
The Next Generation ALMA Correlator and Phased Array



Smithsonian



CASPER



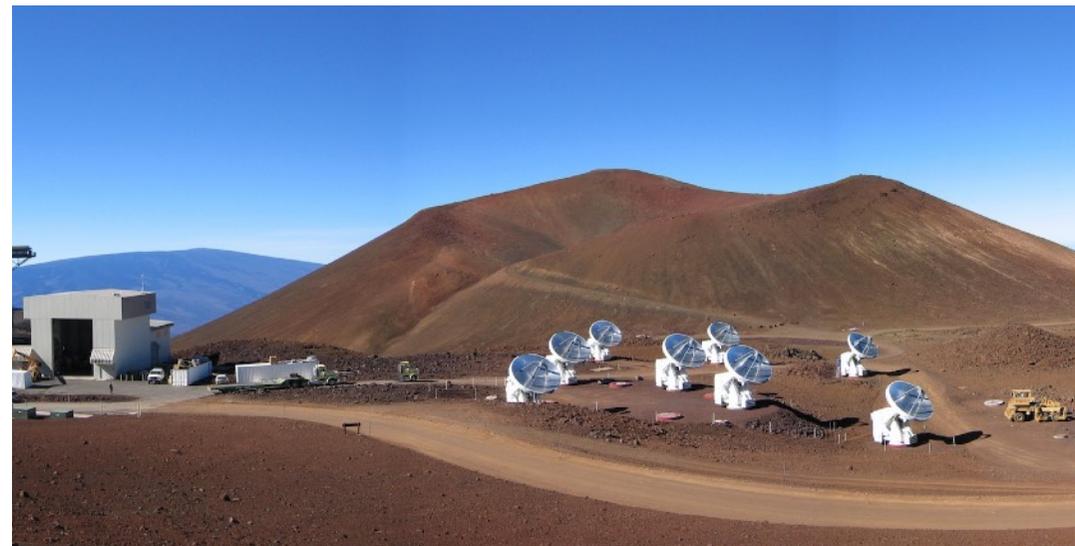
Jonathan Weintroub

Smithsonian Astrophysical Observatory

on behalf of the the Cycle 3 ALMA Development Study Team

USNC-URSI National Radio Science Meeting

Boulder, Colorado, 5 January 2018



ALMA Memo 607
SAO Cycle 3 Development Study Closeout Report
Digital Correlator and Phased Array Architectures for Upgrading ALMA

Alain Baudry, Lindy Blackburn, Brent Carlson, Geoff Crew, Shep Doleman,
Ray Escoffier, Lincoln Greenhill, Daniel Herrera, Jack Hickish, Rich Lacasse, Rurik Primiani,
Michael Rupen, Alejandro Saez, Jonathan Weintraub (PI) & André Young

December 22, 2017

Abstract

This Closeout Report documents the outcome of a SAO-led ALMA Development Study of a next generation combined correlator and VLBI phased array to take greater advantage of fundamental scientific capabilities, such as sensitivity, resolution and flexibility. ALMA already represents a huge advance in collecting area and frequency coverage making it the dominant instrument for high frequency radio astronomy. We have studied processing architectures that maximize bandwidth, and thus sensitivity, allow flexible ultra high resolution spectral processing, and supports other operational modes, such as VLBI. The ALMA Science Advisory Committee (ASAC) studies *Pathways to Developing ALMA* and *A Road Map for Developing ALMA* (both referenced as Bolatto et al., 2015) comprehensively describe the community view of ALMA upgrades and their key science impact.

The methodology of the Study was to examine a variety of technologies, algorithms, balancing costs and timelines against potential benefits. The scientific impact for the proposed study derives from several key new areas of enhanced capability. The Study is divided into eight technical work packages. This Outcomes Report gives a concise summary of each, and eight detailed appendices are provided. A top-level conceptual framing of the full installation, including specifications and rough equipment costing and schedule, is presented as Phase III of three suggested design phases. Phase I is this Study, now complete.

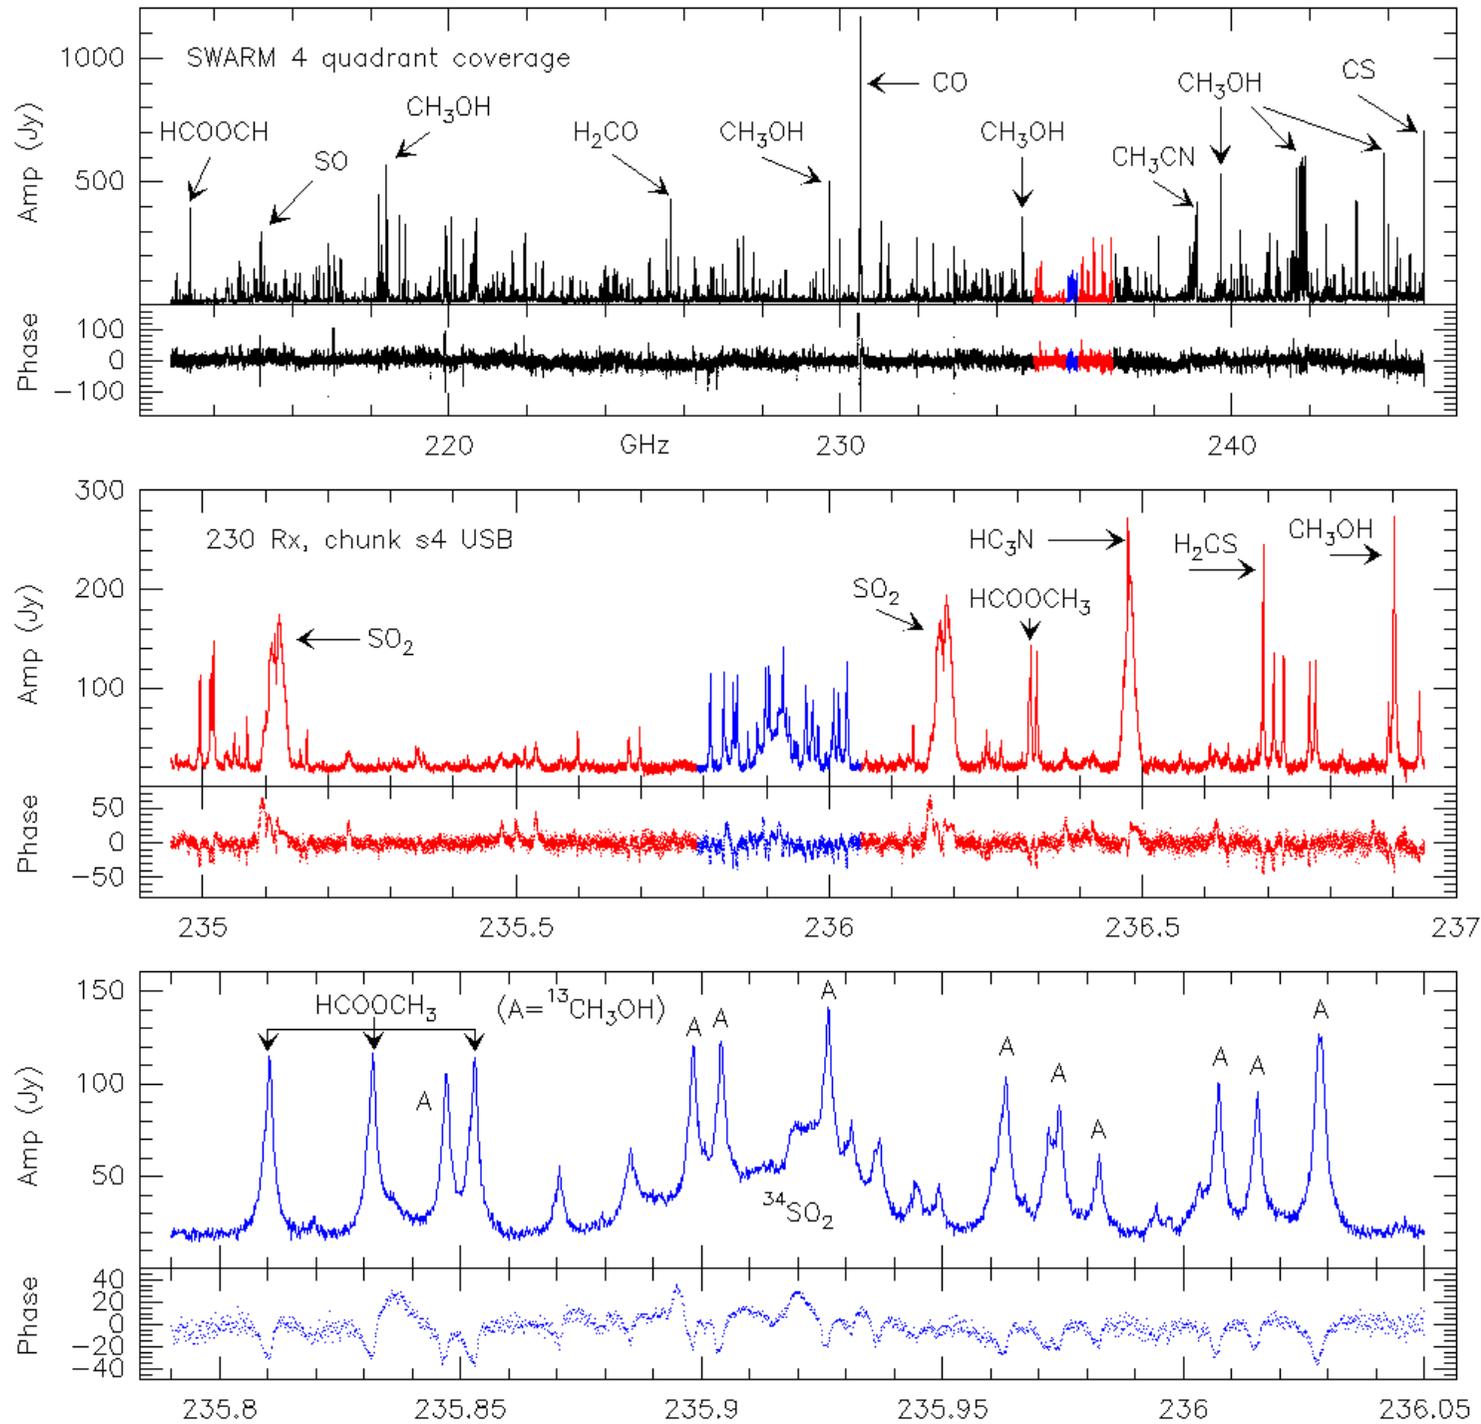


Summary

- Our international team have developed a ultra-wideband correlator/phased array design concept, using *proven* high performance packetized FX technology
- It leverages remarkable recent industry advances in high performance computing and high speed data communications, for a full COTS solution, at modest cost and power
- It is a powerful, flexible and transformative upgrade to BW (4X), resolution, phased array, 4-bit math for 99% efficiency: & supports the *full* ALMA2030 science vision
- The concept has a number of benefits including:
 - Scalable, extensible (e.g. to array feeds), flexible and supports user instruments
 - Small, so it can be assembled while present operational system is running
 - Siting at OSF? increase uptime, reduce maintenance, easier cooling—and so leveraging *full ALMA investment*
 - Native phased array, with very low latency in phase-feedback look for efficiency
 - 4-bit arithmetic in all modes translates to an effective 22% time savings
- If started now, we could produce a commissioned system by ~2027
- 8 GHz “BBC” blocks, so initially could support 8 GHz/SB/Pol (16 GHz goal)
- *System Design for ngALMA to support this? Antennas, Receivers, ADCs, DTS, Archive demands, online calibration, are examples of open issues*

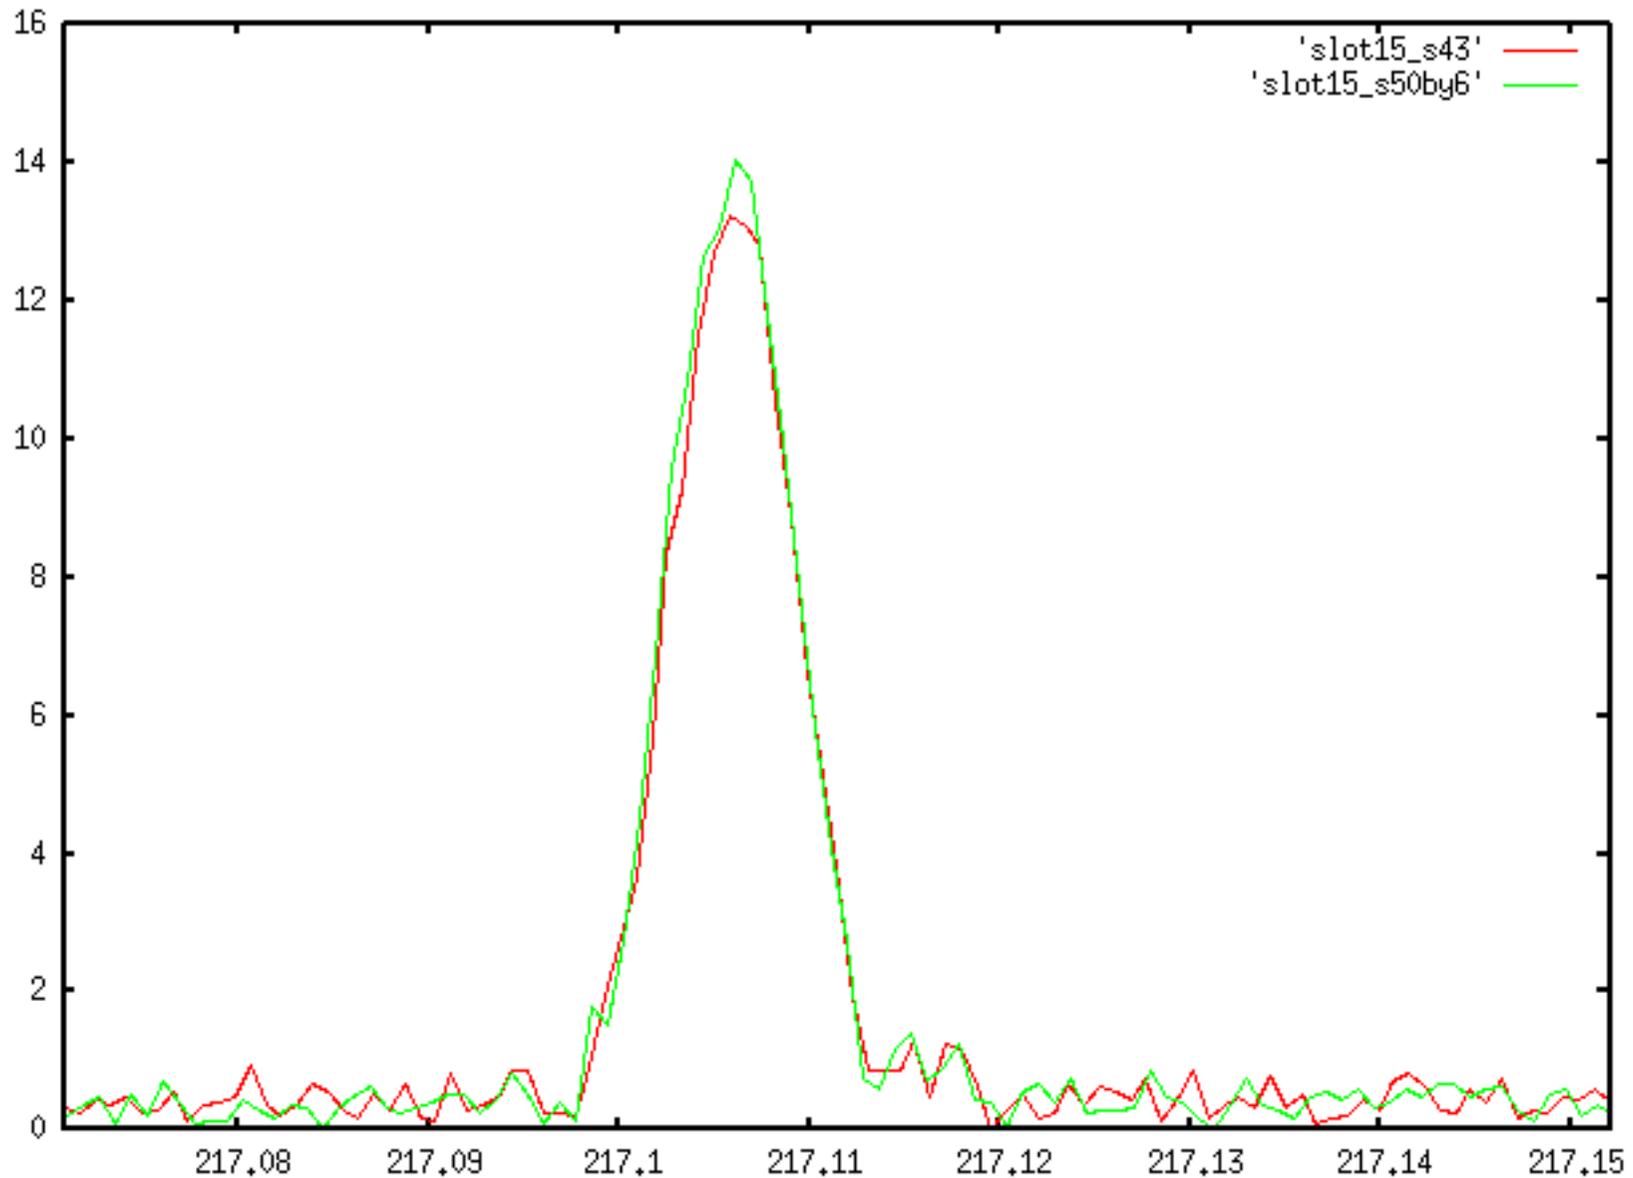
SWARM: SMA Wideband Astronomical ROACH2 Machine

Orion BN/KL



A 4-bit correlator is more efficient than 2-bit

SiO maser in R-Cas was used to measure the ratio of SWARM/ASIC SNR.



217.139, 6.48203 ASIC SiO line is red. SWARM line is green.

By measurement, SWARM SNR is 11 +/- 3% better

Yields about 70 extra days in a year of operations

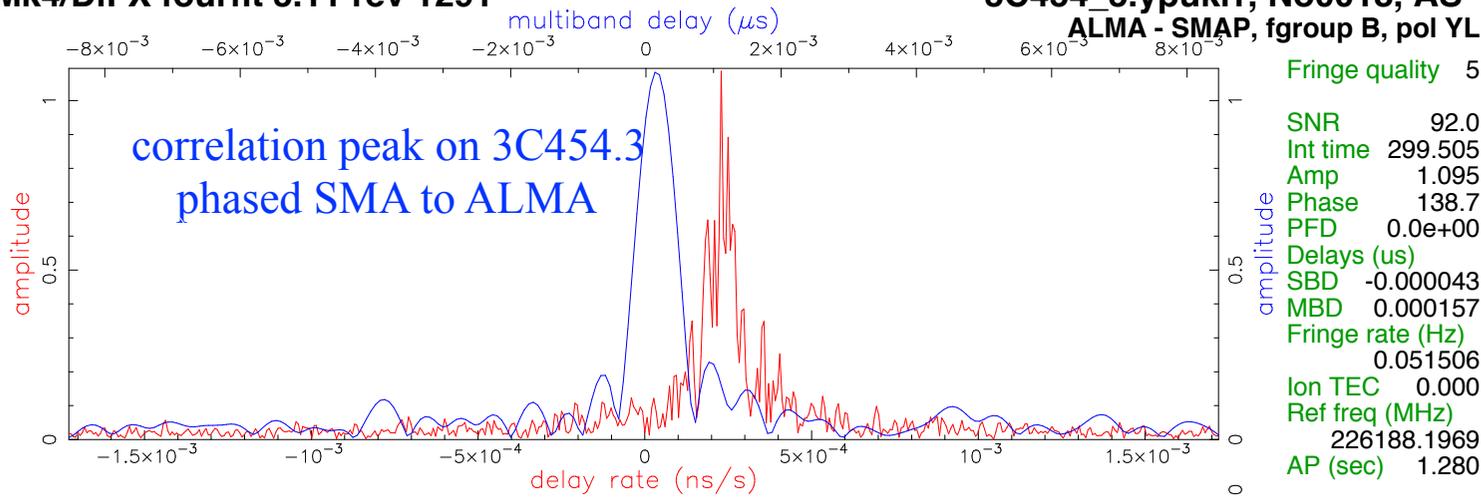
SWARM has a native phased array for VLBI/EHT

Phased ALMA to SMA Fringe, 22 Jan 2016

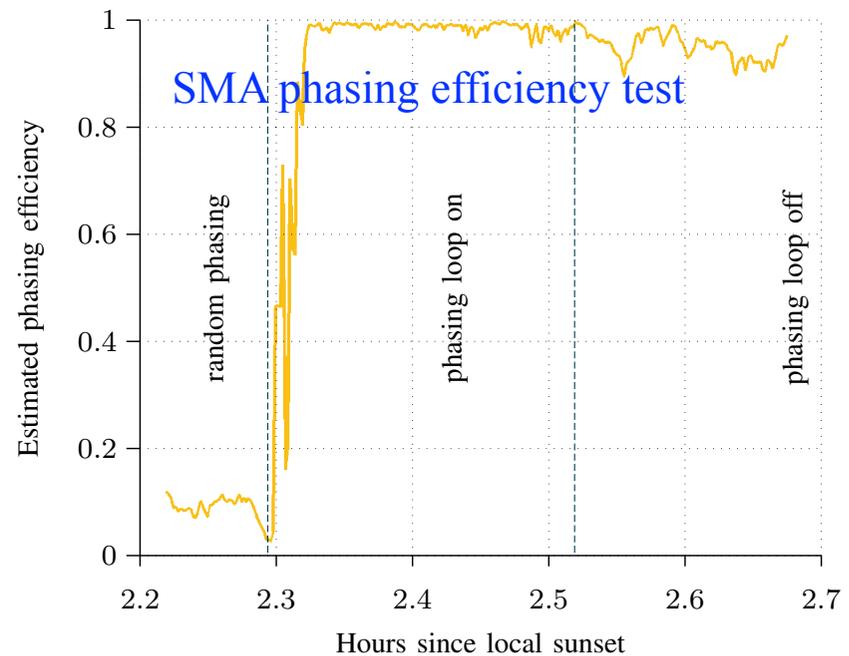
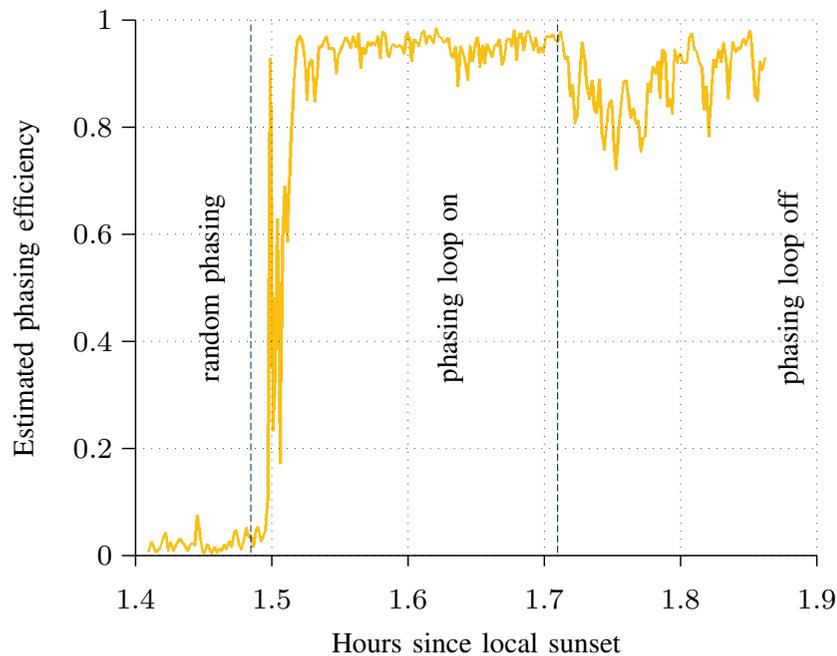
Mk4/DiFX fourfit 3.11 rev 1291

3C454_3.ypukrr, No0018, AS

ALMA - SMAP, fgroup B, pol YL



Geoff Crew, Mike Tttus,
Roger Cappallo, Adam Deller,
ALMA phasing Geoff Crew and
Lynn Matthews, and many more!



Science case: *please see:*

- Opening presentation by Al Wootten, this session
- ALMA Memo 607: <http://library.nrao.edu/public/memos/alma/main/memo607.pdf> (section 2.1 and first appendix on requirements)
- *A Roadmap for Developing ALMA*: https://science.nrao.edu/facilities/alma/alma-dev/PathwaystoDevelopingALMA.pdf/at_download/file
- *Pathways to Developing ALMA*: https://science.nrao.edu/facilities/alma/science_sustainability/Pathways_finalv.pdf

A ROAD MAP FOR DEVELOPING ALMA

ASAC recommendations for ALMA 2030

Alberto D. Bolatto (chair), John Carpenter, Simon Casassus, Daisuke Iono, Rob Ivison, Kelsey Johnson, Huib van Langevelde, Jesús Martín-Pintado, Munetake Momose, Raphael Moreno, Kentaro Motohara, Roberto Neri, Nagayoshi Ohashi, Tomoharu Oka, Rachel Osten, Richard Plambeck, Eva Schinnerer, Douglas Scott, Leonardo Testi, & Alwyn Wootten

From the above document, the ASAC recommended focus areas are:

- Improvements to the ALMA Archive: enabling gains in usability and impact for the observatory
- *Larger bandwidths and better receiver sensitivity: enabling gains in speed*
- *Longer baselines: enabling qualitatively new science*
- *Increasing wide field mapping speed: enabling efficient mapping*

PATHWAYS TO DEVELOPING ALMA

A document to inform the scientific discussions leading to the development of a roadmap for improvements in ALMA

ALMA DEVELOPMENT WORKING GROUP REPORT

Alberto Bolatto -- Chair, ASAC
Chair, General Coordination

Stuartt Corder -- Deputy Director, JAO
Reliability & Efficiency lead

Daisuke Iono -- EA Project Scientist
Resolution, FOV, and Imaging Quality & Calibration lead

Leonardo Testi -- EU Project Scientist
Sensitivity, Spectral Coverage, and Flexibility lead

Alwyn Wootten -- NA Project Scientist
Simultaneous Frequency Coverage, and Usability lead

Some examples of well-established science use cases for increased BW, spectral density, phasing.:

- x10 increase in spectrally surveyed star forming regions and extragalactic sources
- Higher cosmic volumes for intensity mapping
- rapid high-z redshift surveys
- time domain observations of GRBs, comets.
- efficient VLBI/phasing capability for ultra-high resolution and pulsar studies

Study structure: 8 Work Packages...

- WP2.1 Scientific requirements & specifications
- WP2.2 Identify DSP F-engine platform
- WP2.3 Determine F-engine architecture given chosen DSP platform
- WP2.4 Identify corner-turn platform
- WP2.5 Identify DSP X-engine platform
- WP2.6 Determine optimal X-engine architecture
- WP2.7 Determine design of VLBI capability
- WP2.8 Staging of new correlator and phased array

...under certain baseline assumptions

1. Correlator architecture will be FX.
2. Future available bandwidth will be 16 GHz per sideband per polarization, or 64 GHz total usable instantaneous bandwidth, a quadrupling of the current ALMA processed bandwidth.
3. Even larger bandwidths still can be handled by modular replication
4. Samplers will remain at the antennas with digital data sent over fiber
5. Samplers will digitize 8 GHz bandwidth per baseband channel (BBC) at 4-bit resolution
6. The number of observation modes of the new digital system will be minimized.
7. A maximum number, 72, antennas will be supported over baselines extending to 300 km.

Assumed Requirements

(We do not presume to set system requirements for ALMA2030. Our study needed goals. Those we assume are truly transformational.)

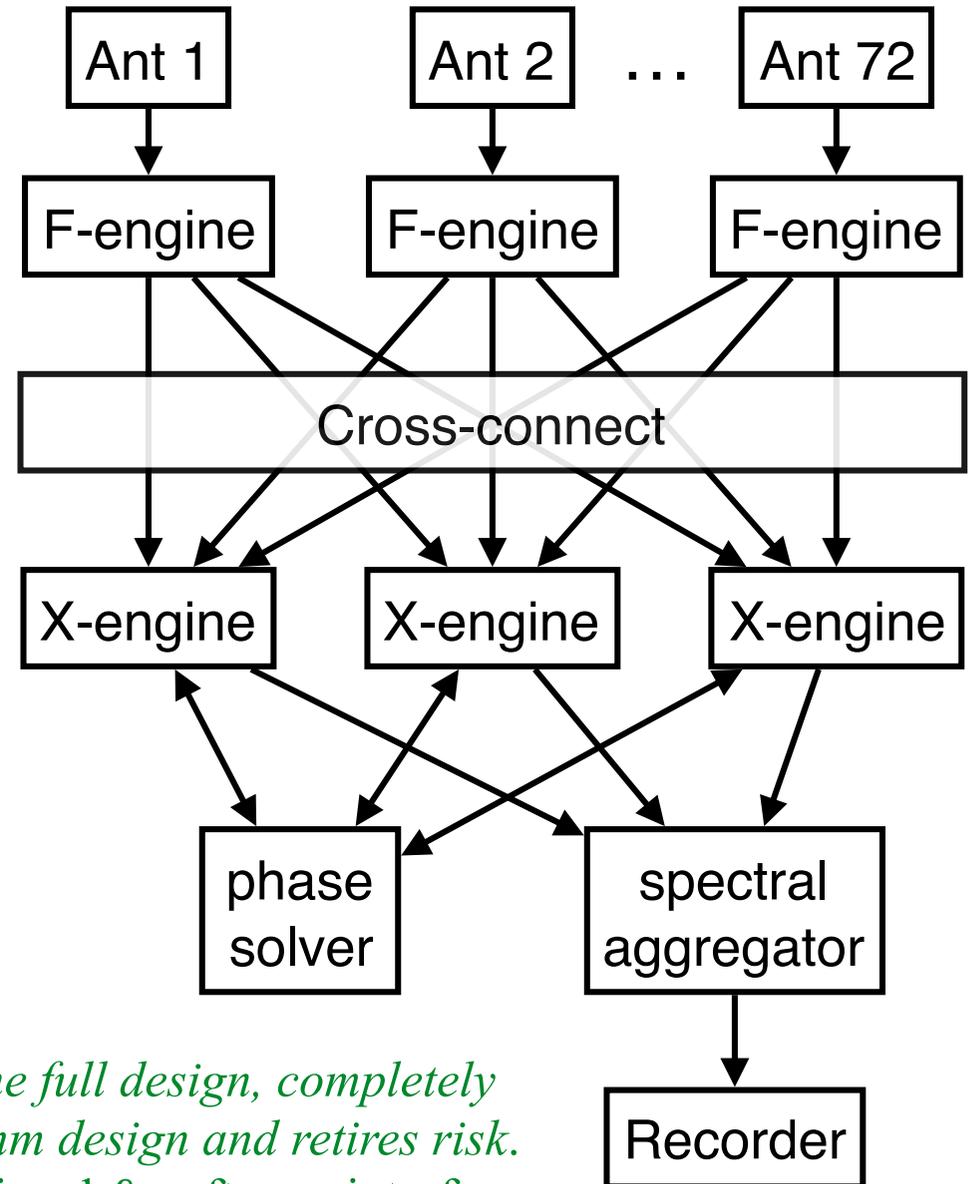
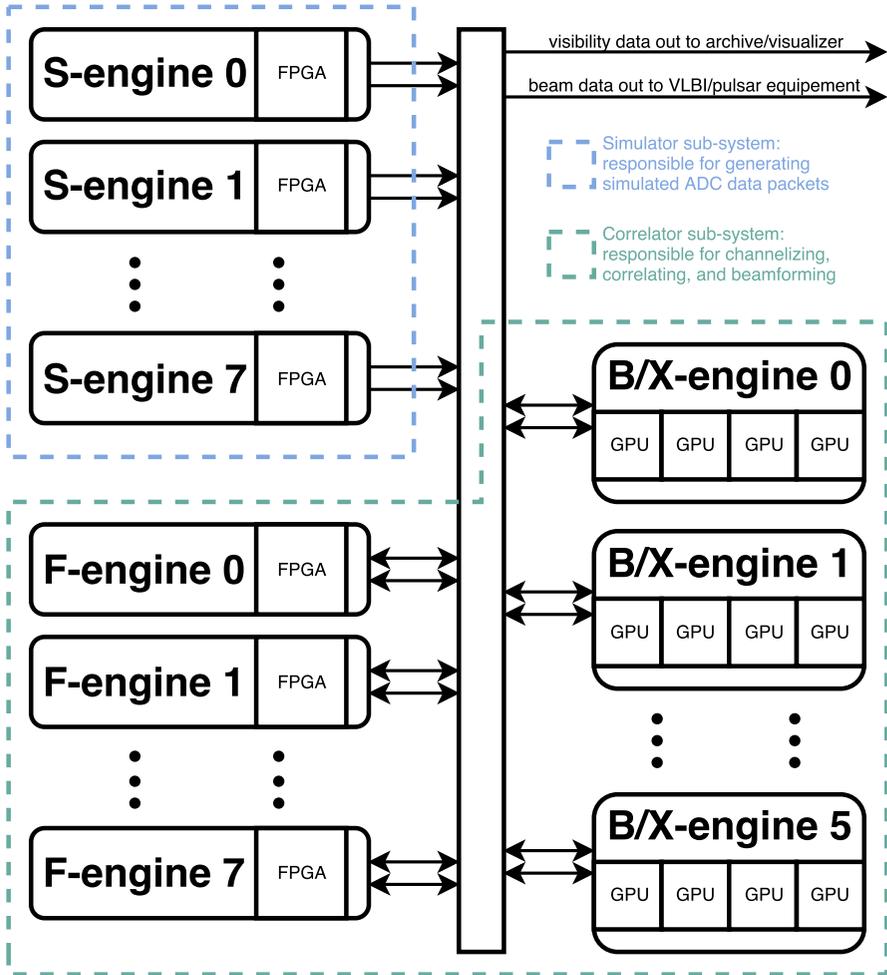
Parameter	Specification	Remarks
# antennas	72	configuration
Max. baseline length	300 km	Sets max delay
Instantaneous BW	64 GHz	16 GHz/SB/Pol
Baseband (BBC) BW	8 GHz	single ADC block
Finest spectral resolution	0.01 km/s:1 kHz	band 1, cold cores
Effective bits	4	99% digital efficiency
Spec. dynamic range	10,000:1	weak lines near strong
Spec. dynamic range	1,000:1	lines on continuum
readout interval	16 ms	for x-correlations
reconfiguration time	1.5 seconds	agile mode change
VLBI mode	phased sum out	two subarrays

The combination of these specifications drive output data rates to a impressive degree, with implications for processing and archive

Three phases of development, Phase I is completed Study

Phase II: Bench prototype (8 antennas, antenna simulator)

Phase III: Full deployment



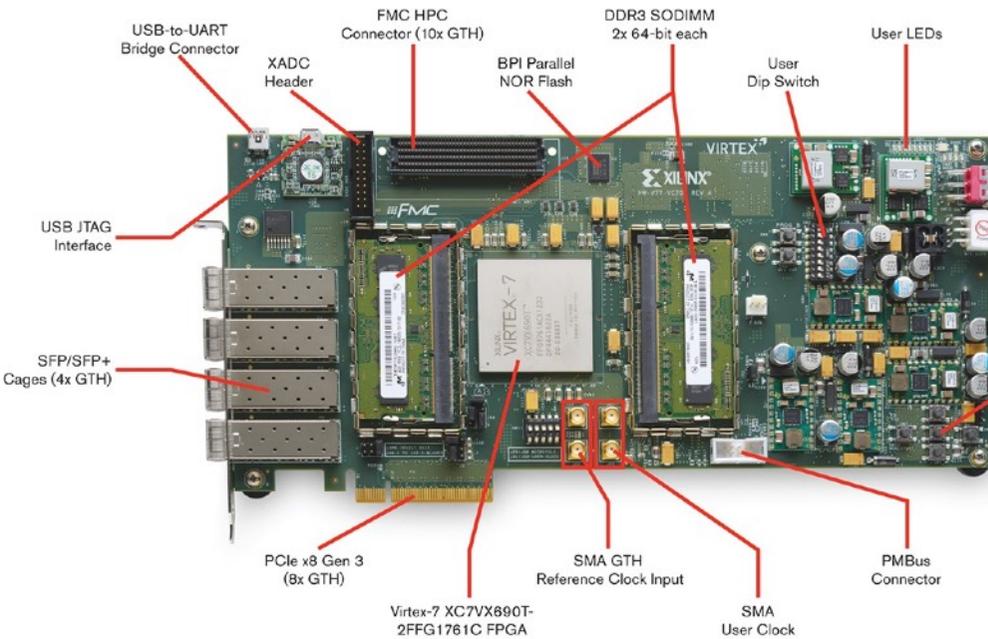
Phase II: *A full featured scaled down version of the full design, completely demonstrates feasibility of concept, details algorithm design and retires risk. Also used as opportunity to resolve ICDs, timing signal & software interfaces.*

Re BBC BW: SAO 20 GS/s ADC based on Hittite HMCAD5831LP9BE

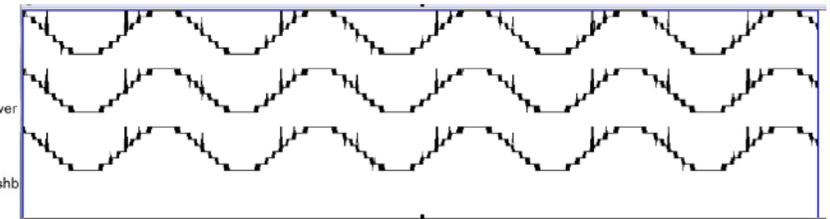
(Weintroub and Raffanti, ISSTT, 2015)

Funded by SI Competitive Grants Program for Science

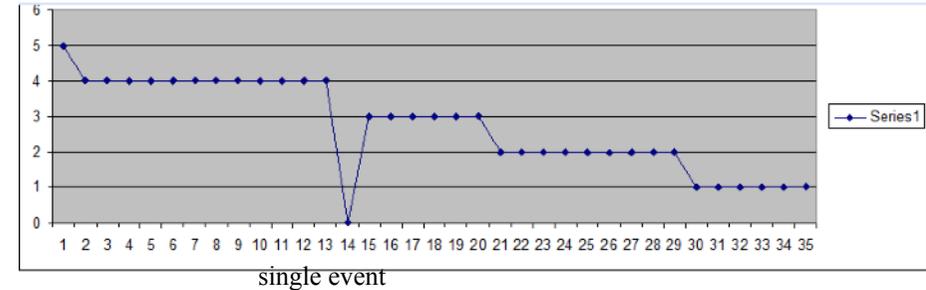
“sparkle codes” ~ 1 in 1000 frequency



VC709 Xilinx V7 Evaluation Board
bargain basement price: \$4995



persistence display



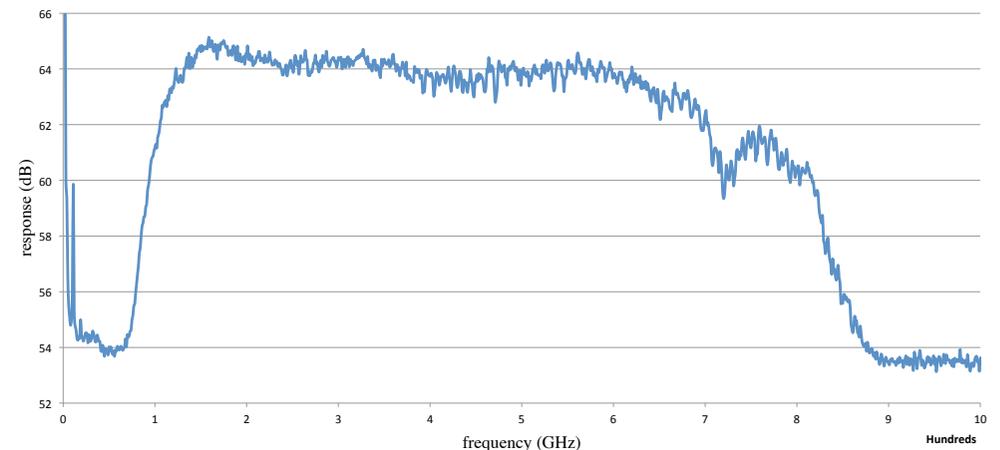
SERDES Transceivers:
GTP, GTX, **GTH**, GTZ

2 to 8.2 GHz Noise passband

GTH features:

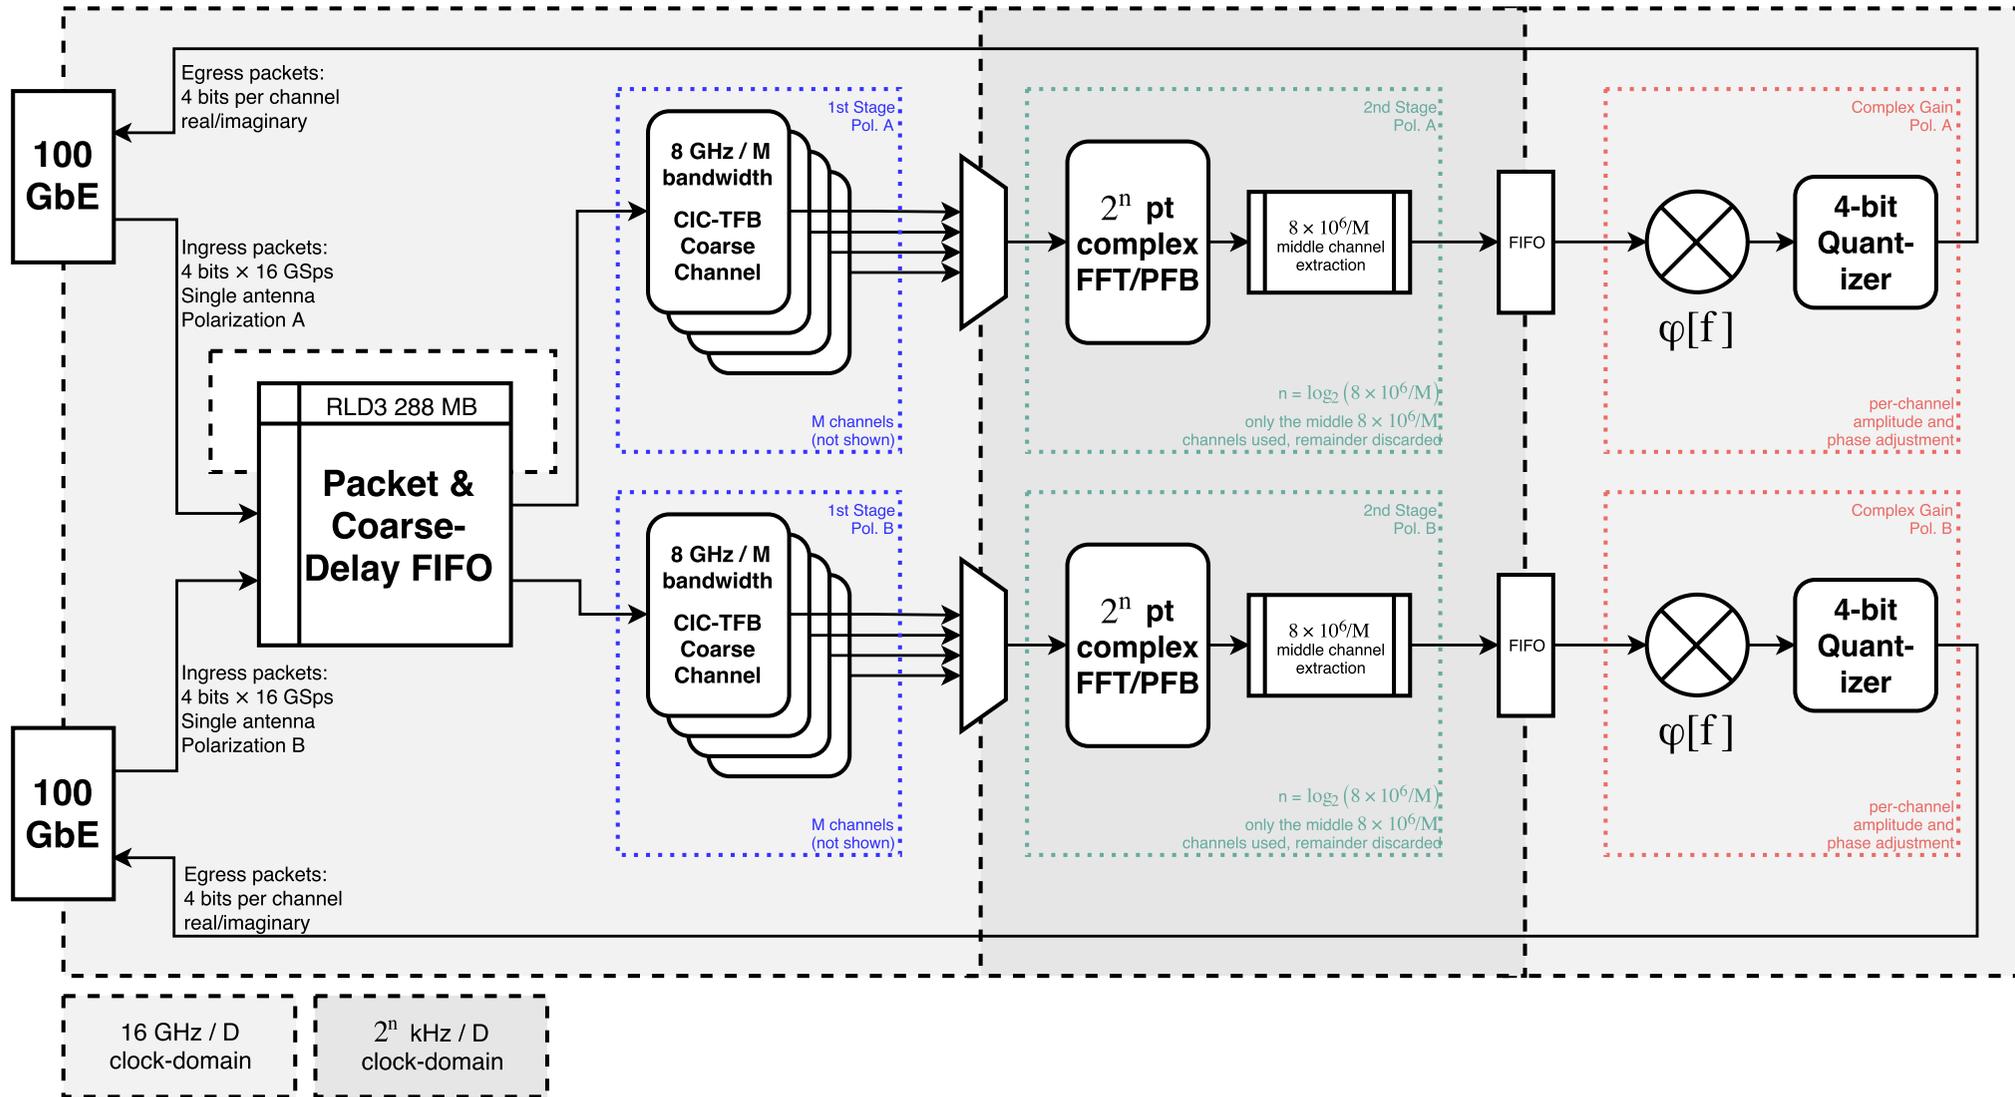
- 7 tap decision feedback equalizer (DFE) vs 5-tap for GTX
- Rx reflection cancellation
- In the Tx, the "Phase Interpolator PPM Controller" which allows fine-grain adjustment of the Tx phase

power spectrum of 3-bit sampled noise



FPGA selected as F-engine Platform

(after comparative analysis including GPU and ASIC tech)



One of five possible FPGA top level algorithms, to process 8 GHz to 1 kHz resolution

GPU selected as X-engine platform

(same comparative technology candidates)

xGPU

[Clark et al, 2012]

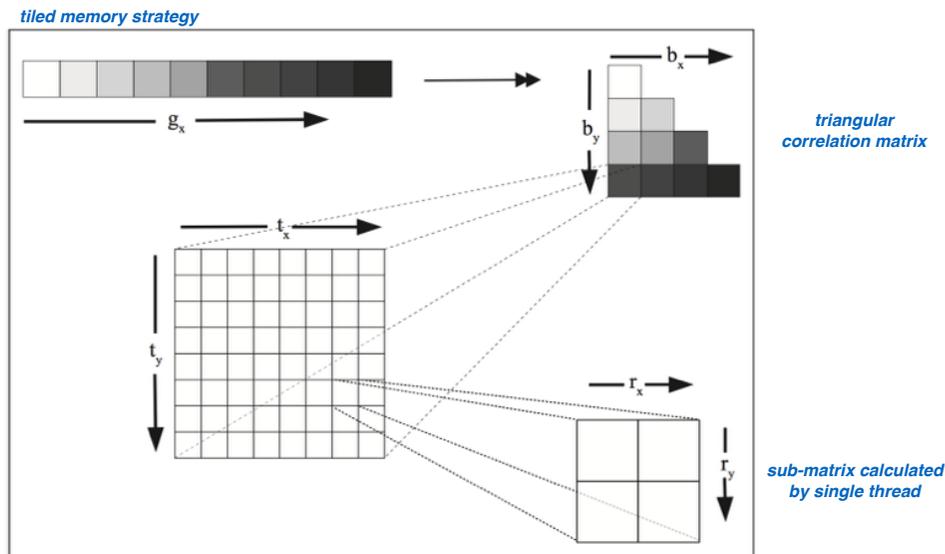


Figure 3: Schematic description of the how threads are mapped to the correlation matrix. The linear grid index g_x is mapped to the triangular block index (b_x, b_y) . Each thread (t_x, t_y) within the thread block is then responsible for calculating an $R_x \times R_y$ tile of the correlation matrix (indexed by (r_x, r_y)). The grid index g_y maps trivially to the frequency dimension (not shown).

may be possible in future with cuBLAS matrix outer product

100 Gbps Ethernet as corner-turn or transpose



- A uniform COTS philosophy: Leveraging industry road maps is the way to go!
- Network switch gives access to antenna, visibility, phased array data, and thus facilitates flexible expansion, and user-instruments
- Architecture is flexibility reconfigurable to, for example, support multi-pixel feeds, perhaps traded against some other feature



NORTH AMERICAN ARC
ALMA Regional Center

North American
ALMA Science
Center



Development Upgrades of the Atacama Large Millimeter/submillimeter Array (ALMA)

Project Proposal

Building the Next Generation Digital Correlator and Phased Array for ALMA

PRINCIPAL INVESTIGATOR: JONATHAN WEINTROUB
INSTITUTION: SMITHSONIAN ASTROPHYSICAL OBSERVATORY
ADDRESS: 60 GARDEN STREET MS78, CAMBRIDGE, MA 20138

(This Project not funded; source for buildout estimates shown next)

Equipment cost and power consumption estimates (full deployment)

Quantity	Item description	P.U. Cost	Extended Cost
576	Xilinx VCU118 FPGA Eval. Board	\$6,995	\$4.03M
144	AMTF Server SYS-6028TP	\$8,595	\$1.23M
48	Trenton Systems Tegra GPU Server	\$18,200	\$0.873M
54	Arista DCS-7060CX-32S-F	\$14,998	\$0.809M
1728	Network Cables CAB-Q-Q-100G-5m	\$450	\$0.778M

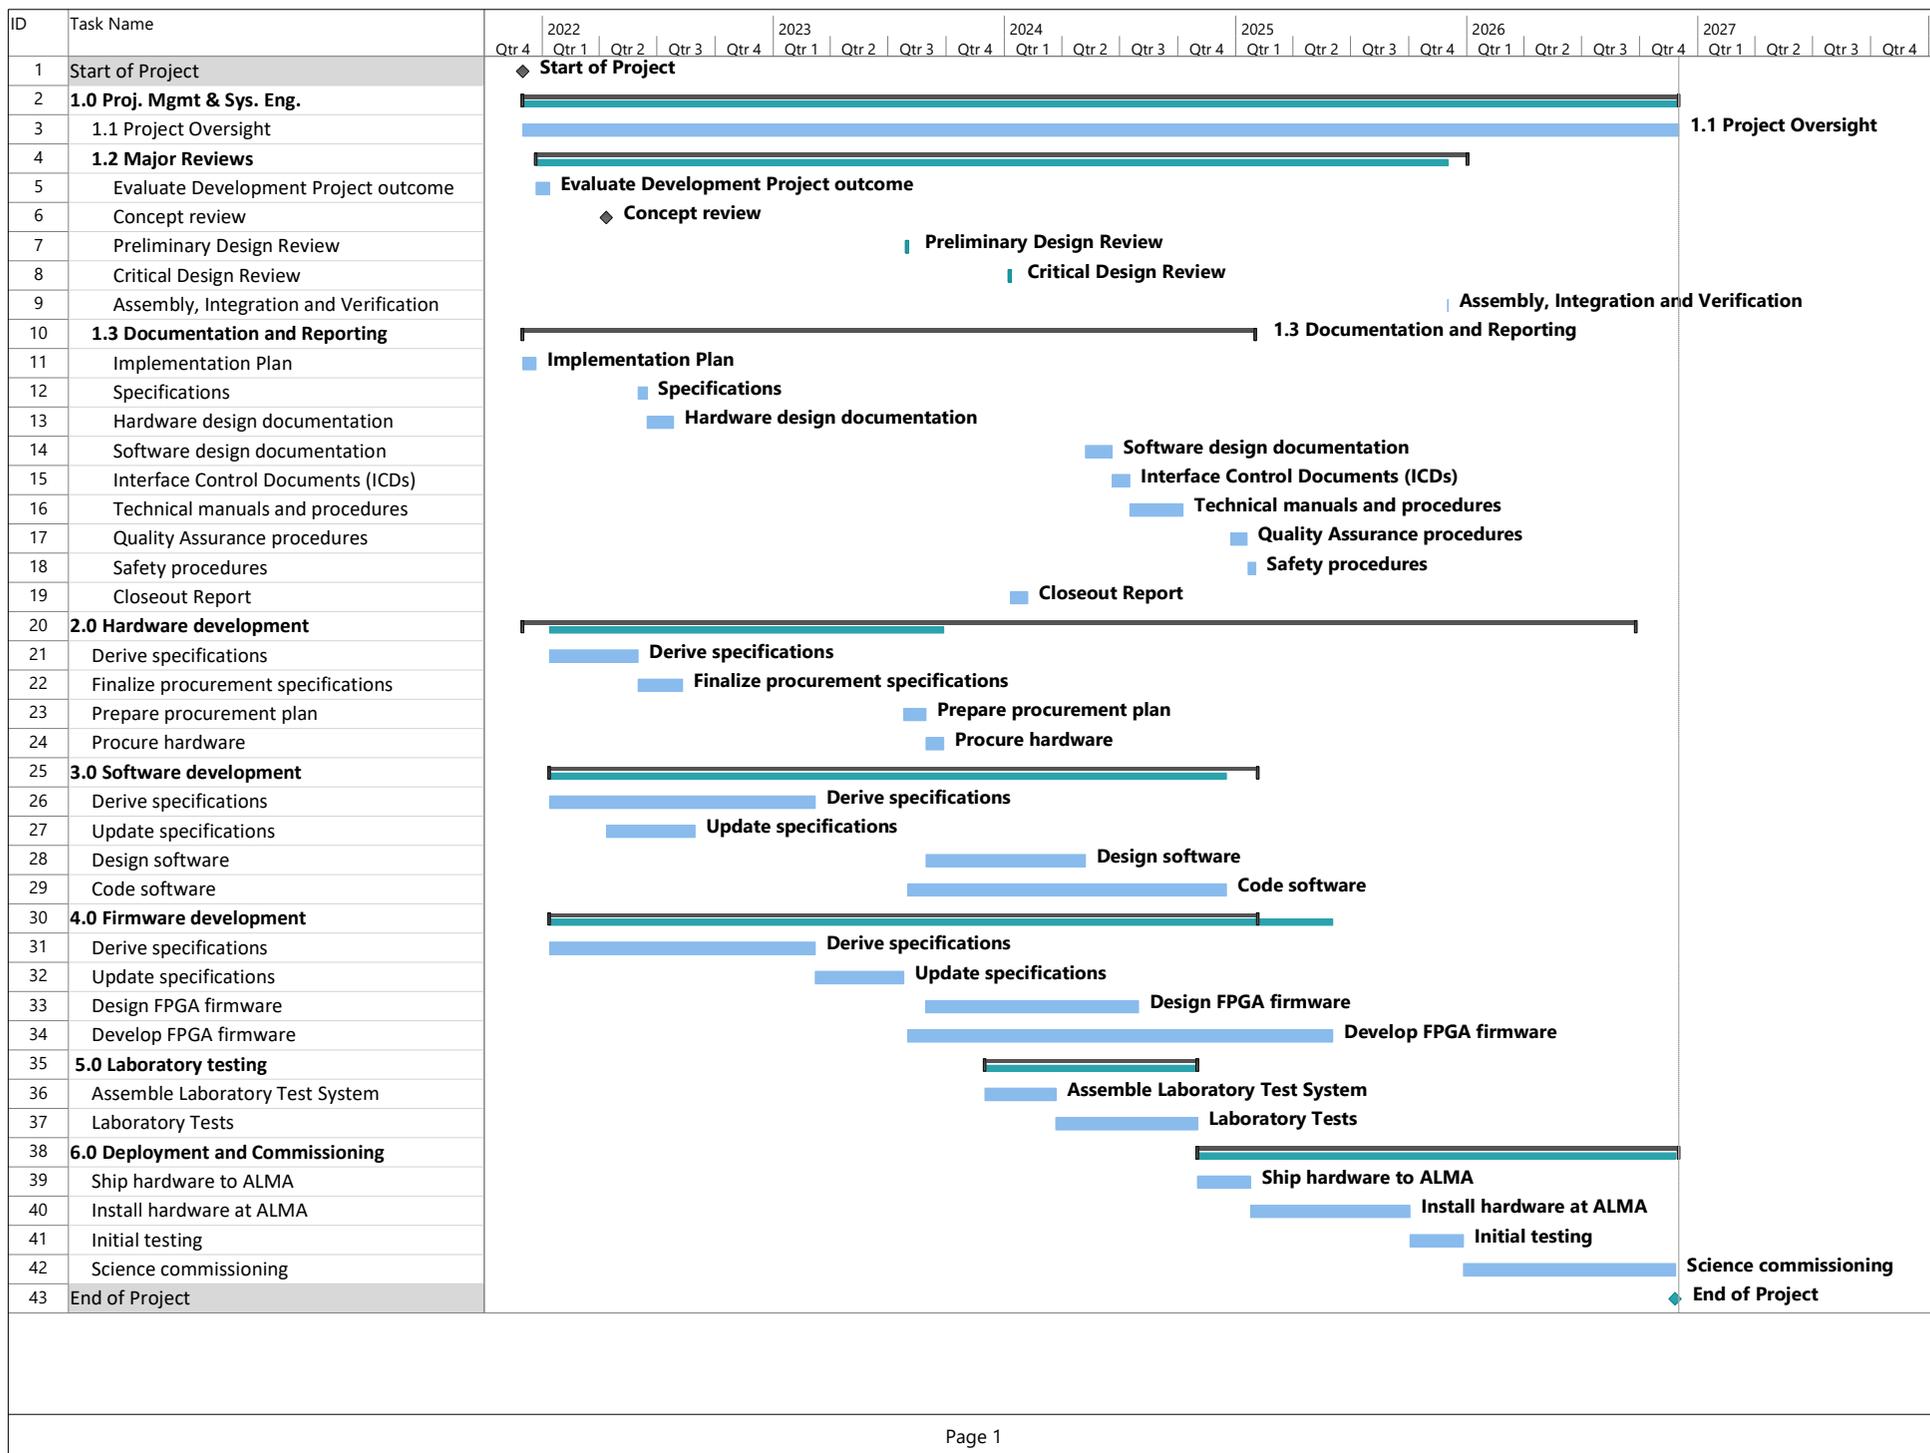
Major components total cost ~\$8M

Quantity	Item description	power p.u. (kW)	total power (kW)
576	Xilinx VCU118 FPGA Eval. Board	0.12	69
144	AMTF Server SYS-6028TP	0.1	14.4
48	Trenton Systems Tegra GPU Server	0.375	18
54	Arista DCS-7060CX-32S-F	0.15	8

Equipment power consumption ~110 kW (present correlator ~140kW)
(however this machine enables 4x BW, ultra-fine-resolution, and other features)

- Cost estimate of major equipment components only, labor, shipping and travel, documentation, other costs excluded
- Power estimate excludes cooling, we have looked into siting the machine at OSF

Phase III: Tentative build-out schedule (prereq.: funded Project)



Questions?

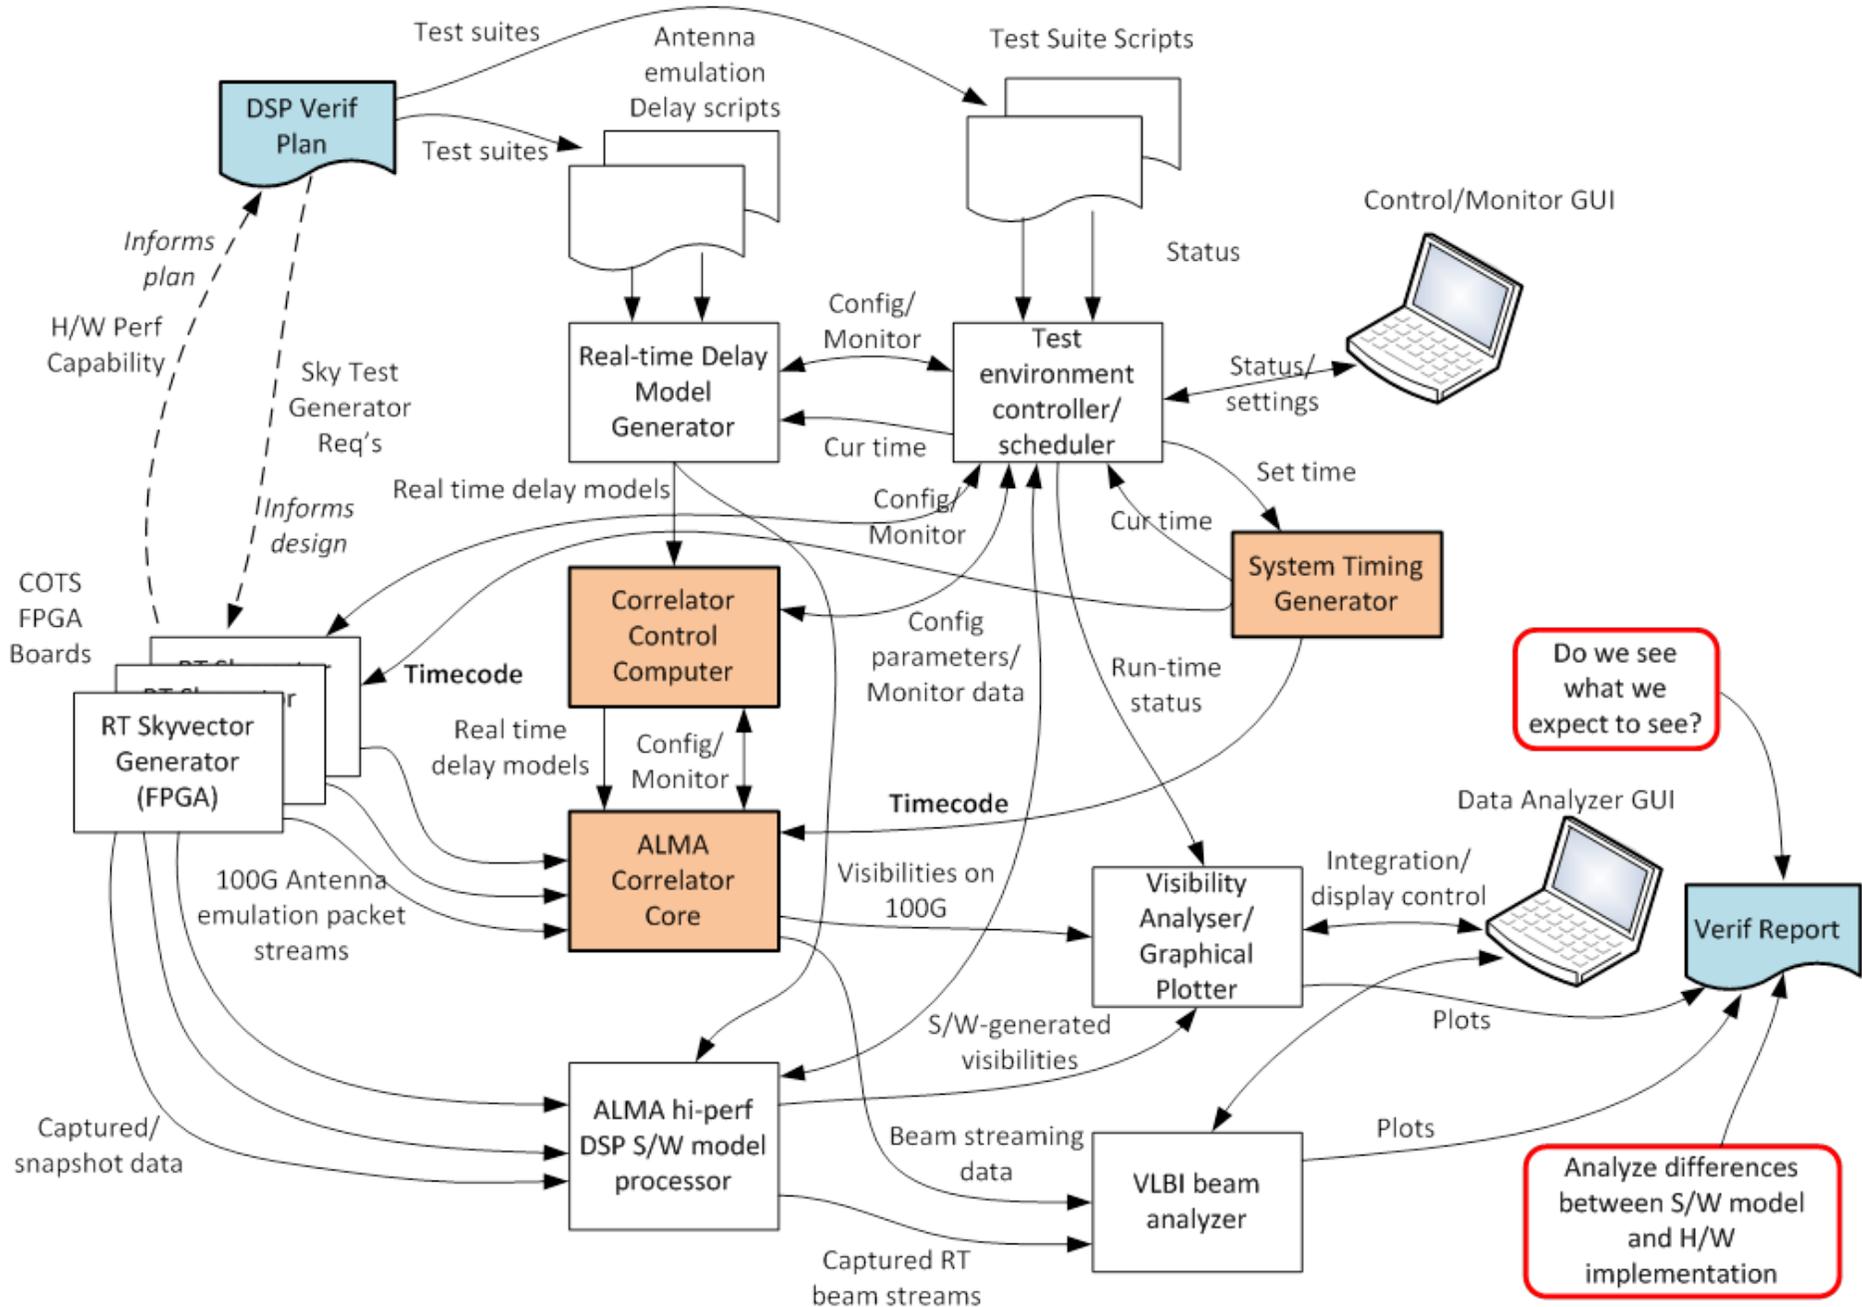
- | | |
|-------------------------------------|--|
| THE VERY LARGE TELESCOPE | <input checked="" type="checkbox"/> |
| THE EXTREMELY LARGE TELESCOPE | <input checked="" type="checkbox"/> |
| THE OVERWHELMINGLY LARGE TELESCOPE | <input checked="" type="checkbox"/> (CANCELED) |
| THE OPPRESSIVELY COLOSSAL TELESCOPE | <input type="checkbox"/> |
| THE MIND-NUMBINGLY VAST TELESCOPE | <input type="checkbox"/> |
| THE DESPAIR TELESCOPE | <input type="checkbox"/> |
| THE CATAclySMIC TELESCOPE | <input type="checkbox"/> |
| THE TELESCOPE OF DEVASTATION | <input type="checkbox"/> |
| THE NIGHTMARE SCOPE | <input type="checkbox"/> |
| THE INFINITE TELESCOPE | <input type="checkbox"/> |
| THE FINAL TELESCOPE | <input type="checkbox"/> |

Supplementary material

Some next steps: We want to proceed with prototyping DSP platforms (currently in progress for wSMA and EHT)
We to continue discussions with ALMA and study group growing to address system engineering challenges for ALMA
Since the correlator/phased array touches all ALMA systems, it is the thing to get right and make the most versatile.

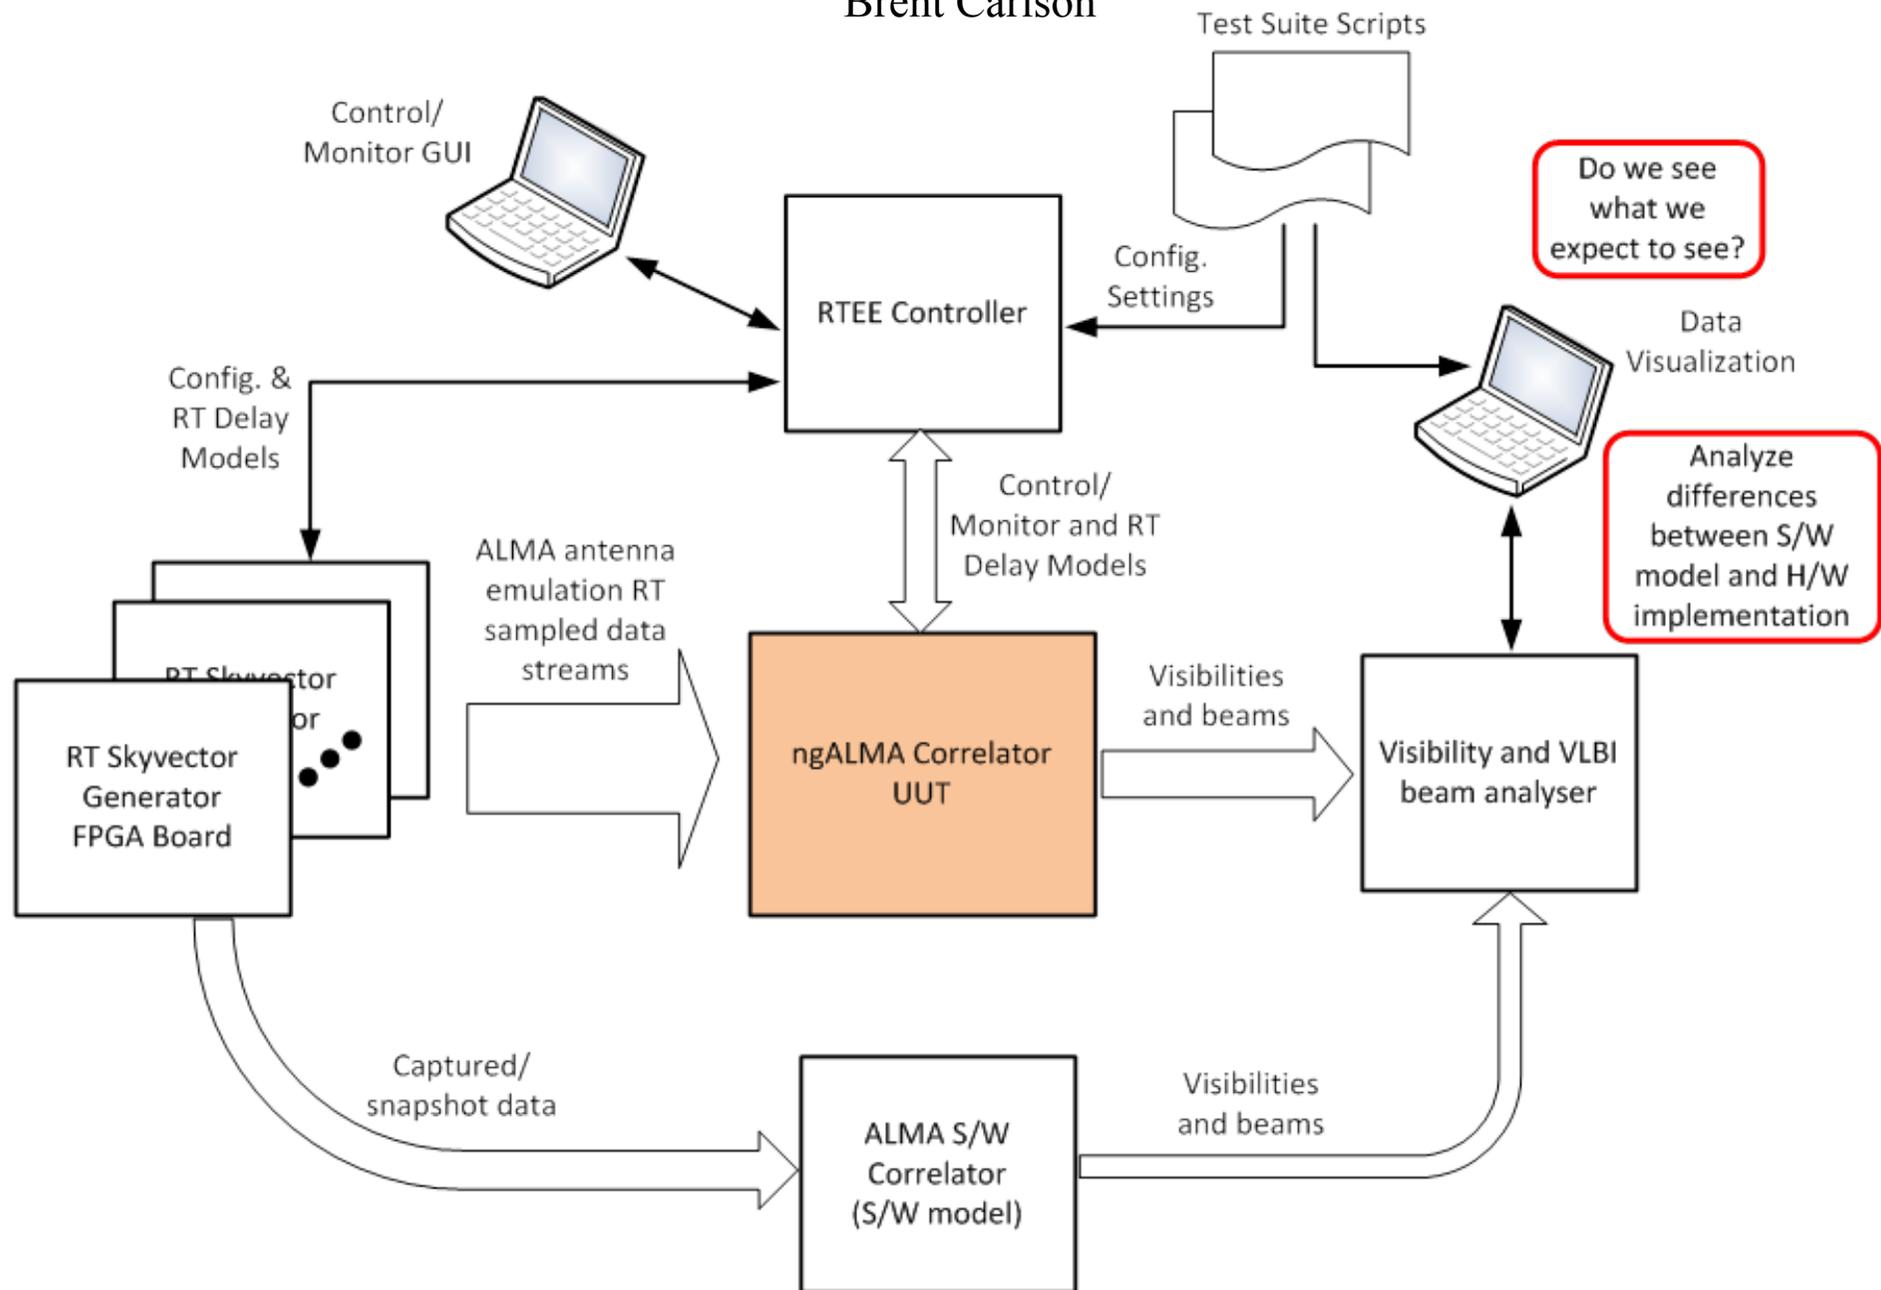
ALMA antenna emulator

Brent Carlson



ALMA antenna emulator

Brent Carlson



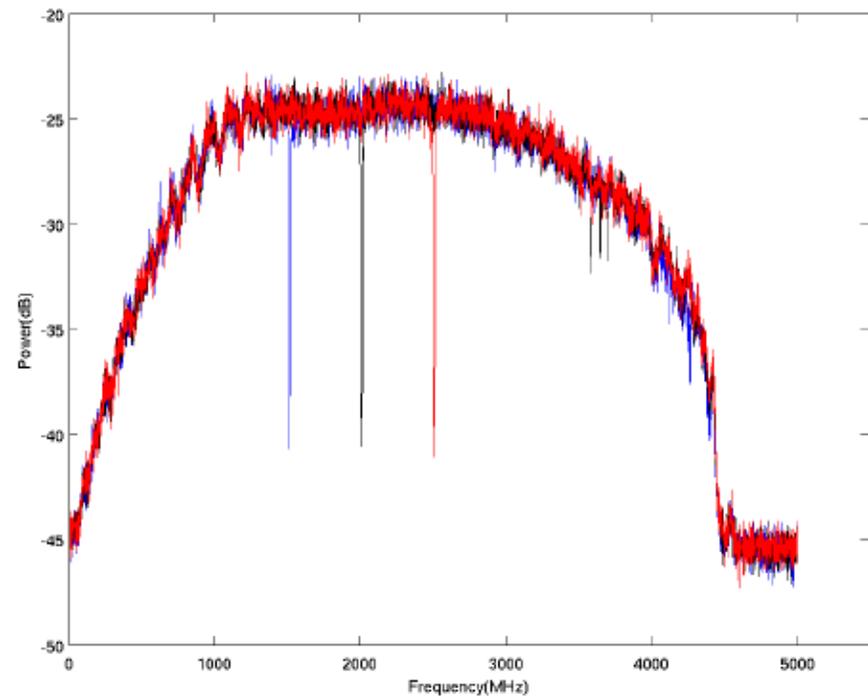
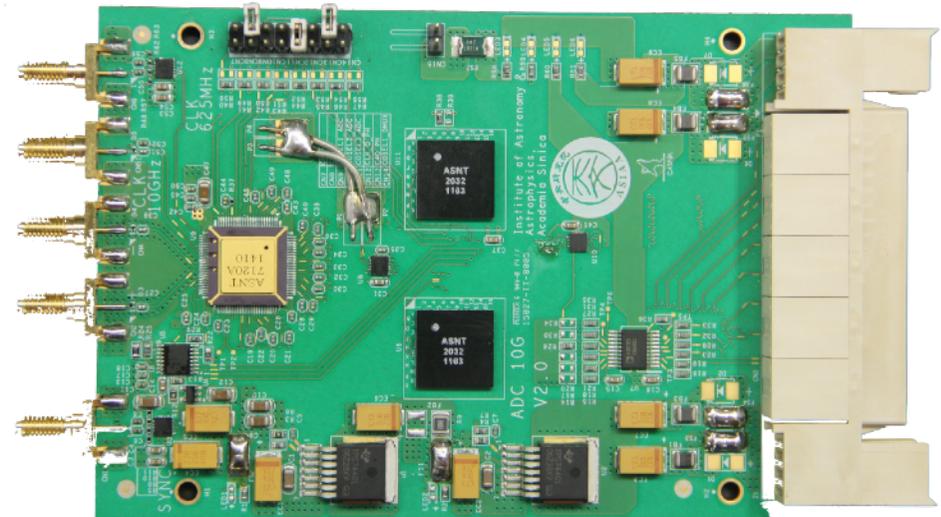
A 10 GSa/s single core CASPER ADC from ASIAA

based on Adantec ANST7120A-KMA

Jiang, Yu & Guzzino (2016)



Homin Jiang, ASIAA, with 10 Gsps ADC, yesterday



Analog frequency response 0 to 5 GHz