

**Conceptual Design and Cost Estimate
for a Phased Array Feed
for the
William E. Gordon Arecibo Telescope**

Final Report

Performed under sub-agreement number 06800-01.01
between Universities Space Research Association (USRA)
and
Cornell University
dated January 22, 2015

November 30, 2015

TABLE OF CONTENTS

Summary	3
1. Introduction	4
1.1. Background	4
1.2. Personnel involved in the study	4
1.3. Scientific Objectives	4
1.4. Phased array feed systems for radio astronomy.....	5
1.5. Arecibo Field-of-View at L-band	5
1.6. The AO19 cryogenically cooled PAF	6
2. Conceptual Design and Specifications of the AO40 PAF	11
2.1. Conceptual design	11
2.1.1 Cryostat	11
2.1.2 Post amplifiers.....	11
2.1.3 Monitor and Control in the Dome	11
2.1.4 Operations building.....	12
2.1.5 The F-engine.....	12
2.1.6 X-engine beamformer and the correlator	12
2.1.7 Data Storage	12
2.2. Specifications.....	13
2.2.1 Operational specifications.....	13
2.2.2 RF specifications	13
2.2.3 Beam-former specifications.....	13
2.2.4 Physical specifications	13
2.2.5 Cryogenic system specifications	14
2.3. Frequency range and RFI	14
3. Physical Constraints	15
3.1. Limits on the size of the dewar	15
3.2. Skirt.....	17
3.3. Weight.....	17
3.4. Rotator.....	17
4. Cryostat, dipole array, LNAs and noise calibration	18
4.1. Cryostat and dipole layout.....	18
4.2. Scaled Dipole Design	21
4.3. Low Noise Amplifiers (LNAs)	22
4.4. Noise calibration	24
4.5. Dipole/LNA assembly	24
4.6. Post cryostat RF stages.....	26

4.7. Monitor and control.....	27
5. Cryostat structural and thermal modeling	28
5.1. Cryostat structural issues.....	28
5.1.1 Dewar size	29
5.1.2 Integrating multiple cryocoolers.....	30
5.1.3 Instrument mount with rotator	31
5.2. Cryostat thermal analysis.....	32
5.2.1 Number of cryocoolers	33
5.2.2 Foam & dipole thermal FEA	33
5.2.3 Stripline vs. coax lines	34
6. Digital receiver	35
6.1. A040 Correlator/Beamformer Design	35
6.2. Functional Description	36
6.3. Digital Receiver Design Parameters and Specifications	39
6.4. Feasibility Performance Analysis for GPU Processing	41
7. Schedule	43
8. Cost estimate	43
8.1 RF sections	44
8.2 Digital Receiver and Data Storage	45
8.3 USRA/Arecibo Observatory installation and integration costs	45
8.4 Total estimated cost	45
Appendix A (Digital Receiver major equipment quotes).....	46

Summary:

In January 2015 a sub-agreement was signed between the University Space Research Association and Cornell University to fund a “Conceptual Design and Cost Estimate for a L-Band Phased Array Feed for the 305 m William E. Gordon Arecibo Telescope”. The work was to be performed by Cornell University and by Brigham Young University (BYU) under a sub-agreement with Cornell. The initial completion date of August 31, 2015 was extended to November 30, 2015 under a No Cost Extension dated August 5, 2015.

As discussed in the report, the suggested design of the 40-beam phased array feed (PAF) system for the Arecibo telescope (AO40) is based on the experience gained by the Cornell University group in building a fully cryogenic 19 dual polarization element prototype system that was tested on the Arecibo telescope in 2013, and on the experience of the BYU group in designing a digital beamformer system for a 19 dual polarization element PAF for the 100 m Green Bank Telescope. Because of this background, and the requirement for a realistic cost estimate, the baseline design presented in the report is much more detailed than a simple conceptual design. The overall concept is for 40 dual polarization formed beams on the sky spaced approximately 2 arc minutes apart filling the approximately 15 x 20 arc minute field of view of the Arecibo telescope at ~1.4 GHz. Based on science and rf interference issues, the overall bandwidth is from 1,280 MHz to 1,720 MHz. This can be modified in the final design. The digital beamformer will be able to process a 300 MHz bandwidth selectable within the overall bandwidth. The implementation is based on a cryogenically cooled array of 80 dual polarization dipoles and 160 low noise amplifiers. The dipole/LNA assemblies are designed to be as easily field replaceable as possible. After further amplification and filtering, the rf signals will be sent directly over fiber optic lines to the operations building. Further amplification, filtering and mixing will select the 300 MHz to be passed to the digital beamformer. In the baseline design the beamformer will utilize CASPER ROACH II boards and HPC servers and GPUs.

As detailed in Section 8 of the report, the estimated capital costs (i.e. costs for parts that will be incorporated in the instrument) will be \$2,435,000. The remaining costs for services, shipping, personnel, etc are very dependent on the institutions that will carry out the work. On the assumption that the instrument will be constructed by Cornell and BYU, the direct costs (i.e. travel, shipping, software licenses, materials & supplies, etc) other than for personnel are estimated to be \$270,000. Based on a four year construction, testing and installation schedule (Section 7), personnel costs including benefits are estimated to be \$1,750,900. USRA/Arecibo Observatory has estimated that their personnel costs to support the installation and testing on the Arecibo telescope would be \$111,500. Indirect costs on the direct costs, including personnel plus benefits costs, for all three institutions would be \$1,121,860. Under the above assumptions the total instrument costs including installation and testing is estimated to be close to \$5,689,000.

1. Introduction:

1.1 Background: On January 27, 2015 Universities Space Research Association (USRA) issued a sub-award to Cornell University for a “Conceptual Design and Cost Estimate for a Phased Array Feed for the William E. Gordon Arecibo Telescope”. The effective date of the award was October 1, 2014. On February 9, 2015 Cornell issued a sub-award to Brigham Young University (BYU) for the “Digital Back End Conceptual Design for the Arecibo AO-40 Phased Array Feed”. Both awards had a completion date of August 31, 2015. A no cost extension with a new completion date of November 30, 2015 was signed by Cornell and USRA on August 5, 2015. The Cornell sub-award with BYU was extended to November 30, 2015.

The primary objective of the project was the conceptual design of a L-band cryogenically cooled phased array feed and associated digital beam former that would provide approximately 40 beams on the sky with the specifications discussed in the next section. A cost estimate in 2015 dollars would also be developed. As discussed in the report, design decisions were made based on what seemed like the sensible, cost effective approach, in 2015. At the time that the phased array feed system (PAF) is constructed alternative approaches may be considered more desirable and/or more cost effective.

1.2 Personnel involved in the study were:

Prof. German Cortes, Cornell University and the University of Antioquia, Colombia
Prof. Donald Campbell, Cornell University (PI)
Mr Ganesan Rajagopalan, Cornell University
Mr Stephen Parshley, Cornell University
Prof. Brian Jeffs, Brigham Young University
Prof. Karl Warnick, Brigham Young University

Arecibo Observatory personnel who have participated in the study are:

Eng. Luis Quintero, Head of Electronics
Eng. Felipe Soberal, Head of Facilities
Mr Phillip Perillat
Eng. Felix Fernandez

1.3 Scientific Objectives: An assessment of the primary objectives for the construction of a 40-beam PAF at L-band for Arecibo was not included as a task under the award. The proposal did encourage the Arecibo Observatory to sponsor a workshop on possible applications of the system. However, discussions were held with Cornell professors James Cordes and Martha Haynes, long time heavy users of the Observatory and the current 7-beam ALFA receiver system, as to high impact science that the PAF system could address. These discussions did inform the specifications for the system. Searches for new pulsars, especially millisecond period ones with very stable periods, are a high priority including for the NANOGrav project for using high precision timing of pulsars to detect gravitational waves. A

related high impact field is the search and study of Fast Radio Transients (FRTs). These are very short, msec or less, bursts of radio emission. To date, only a small number have been detected, primarily by the Australian Parkes telescope but also by Arecibo. As discussed later, their very high dispersion measures ($>1,000$) drive the coarse filter bandwidth requirement in the PAF's digital beam former. The study of atomic hydrogen (HI) emission from external galaxies will continue to be of high interest with one objective being the detection of small haloes – low mass bound clouds of HI and dark matter with or without stars – in the local universe. Current cosmological theories predict large numbers of these low mass haloes, many more than are currently observed and a high sensitivity survey with a PAF on the Arecibo telescope could resolve this issue.

1.4 Phased array feed systems for radio astronomy: L-band PAFs, radio cameras, have been under development for about a decade as they hold the possibility of greatly improving the efficiency with which radio telescopes can be utilized for survey applications. The two groups leading this effort are the Australia Telescope National Facility (ATNF), which is developing a 30 beam PAF system for its 36 antenna ASKAP (Australian SKA Pathfinder) array, and the ASTRON group in the Netherlands for the 14 antenna Westerbork array. Both of these systems are uncooled and have system temperatures of about 50 K and 75 K, respectively, with the high Westerbork temperature being at least partially due to the need to mitigate RFI. The advantage of PAFs for single dish telescopes is that only one unit is required and resources can be used to reduce the system temperature by cooling the LNAs or the entire front-end. The NRAO in conjunction with Brigham Young University (BYU) is developing a 7-beam (19 dual polarization dipoles) for the Green Bank Telescope (GBT) that cools the low noise amplifiers (LNAs) to 70 K. NAIC, while being administered by Cornell University, initiated the development of a prototype fully cryogenically cooled 19 dual polarization dipole element PAF. This system is discussed below in the context of the current project. The Max Planck Institute for Radio Astronomy in Bonn is obtaining an ASKAP system from the ATNF for their 100 m telescope. However, this is an uncooled system and may require modifications because of the high RFI environment at the telescope's site.

1.5 Arecibo Field-of-View at L-band: One consequence of the complex Gregorian optics of the Arecibo telescope (Fig. 1.1) is a restricted field-of-view (FOV). Before any serious development work on an L-band PAF for the telescope the size of the FOV had to be measured to determine the maximum number of Nyquist spaced pixels – the Nyquist spacing is 50% of the half power beamwidth – that a PAF “camera” could support. In 2012 NAIC, while under the management of Cornell University, initiated a study aimed at determining the FOV. A 19-element focal plane array at L-band was provided by the BYU group under Brian Jeffs and Karl Warnick and a special mount developed that allowed the array to be moved radially and vertically with the rotary Gregorian receiver room floor providing motion in the third dimension (Fig. 1.1).




Figure 1.1: Left, the optics of the Arecibo Gregorian system; Right, the positioning system and 19-element array mounted on the receiver room rotary floor.

The BYU group provided a low bandwidth beam former system. Radio source measurements with the system showed that the FOV to the -1 db level of sensitivity was 15.2×19.8 arc minutes with the array in the focal plane. The asymmetry is due to the asymmetric illumination pattern of the Gregorian optics. Assuming a pixel spacing of 2 arc min on a square grid – the Nyquist spacing is close to 1.8 arc min – the FOV could support about 60 pixels. As A040 implies, the number has been set to 40 dual polarization pixels with the exact spacing being an input parameter to the beam former.

1.6 The A019 cryogenically cooled PAF: As a consequence of the FOV measurements, a group at Cornell University under the direction of Dr German Cortes has developed a fully cryogenic prototype Phased Array Feed (PAF) system with 19 dual polarization elements (A019) operating at L-band. The elements are crossed dipoles with the dipoles and the low noise amplifiers (LNAs) connected to the 20 K stage in the dewar. A unique RF lossless foam, Rohacell-71HF, supports an entrance window ~ 75 cm in diameter, capable of holding off a vacuum load of about five metric tons. There are three layers of foam with the dipoles imbedded, but not touching, the lowest layer that contacts the 70 K temperature ground plane (Fig. 1.2). Each crossed dipole, dual LNAs and noise calibration coupler are incorporated in a single plug-in unit that allows relatively simple field replacement by removing the front window and its supporting foam (Fig. 1.2). Development and construction of the system was done under NSF and Cornell University funding with the NSF funding being initially through the National Astronomy and Ionosphere Center and, later, via an ATI grant.




Figure 1.2: *Left:* A cross section of the AO19 dewar showing the window, the “plug in” dual polarization dipole and LNA assemblies, and the overall cryogenic design. Cryogenic design by Stephen Parshley, dipole and plugin assembly by German Cortes-Medellin. *Right:* fully assembled cryo-PAF AO19.

Testing of the cryo-PAF AO19 system was carried out at the Arecibo Observatory in 2013 with support by the BYU group under Brian Jeffs and Karl Warnick, which provided the down converter, digitization and beam former systems plus analysis software. Despite problems primarily related to thermal cycling failures of a number of the low noise amplifiers, the tests demonstrated that the radio transparent vacuum window design was able to successfully support the atmospheric pressure, and that the cryo-PAF is capable of producing beam-formed system noise temperatures of the order of 35K. It showed that the overall AO19 design is viable, and furthermore, that the concept and cryogenic design could form the basis of a larger system. A detailed description of the AO19 PAF and the results of its tests can be found in Cortes et al, “A Fully Cryogenic Phased Array Camera for Radio Astronomy”, *IEEE Trans. Antennas and Propagation*, Vol. 63, No. 6, 2471-2481, June 2015.

As shown in Figure 1.2, in the original cryo-PAF AO19 approach, the cryostat top flange was aligned with the topmost ROHACELL foam layer to ensure that the load on the Kapton film was purely vertical. As a result, there is a metallic thermal shield that extends from the ground plane, just next to the outermost dipoles, almost to the window flange. The proximity of this metallic wall alters the dipoles’ radiated fields. Figure 1.3 shows a 3-D cut of the calculated field radiated by a dipole next to the metallic thermal shield. Figure. 1.4 shows, on the left panel, an inset of the current AO19 numbered dipole layout, and in the right panel, the calculated full embedded radiation patterns corresponding to the dipoles 1, 2 and 8. The calculated dipole patterns close to the metallic thermal shield are highly distorted.




Figure 1.3: E-Field cut of dipole radiation being reflected by the inside cryostat metallic thermal shield.




Figure 1.4: *Left:* Section of the dipoles in a hexagonal layout with 45° arm rotation. *Right:* Calculated embedded radiation patterns for dipoles 1, 2 & 8 inside the cryostat walls.

Subsequent detailed electromagnetic modeling indicated that modifications to the cryogenic design to minimize the effects of the aluminum thermal shield and the dewar walls on the radiation patterns of the dipole elements would improve the performance. A new design was developed and modifications to the AO19 dewar system to implement and test this design was part of the current work under the sub-award with USRA.




Figure 1.5: A cross section of the modified AO19i dewar showing the HDPE top-hat dome and Rohacell foam layer support that replaces the flat vacuum window. Also shown is a new shorter external dewar vacuum shield terminated at the ground plane level. The dipoles were rotated to improve dipole pattern characteristics.

This new design replaced the thermal shields and dewar walls extending above the level of the ground plane (Figure 1.5) with a top-hat shaped, RF transparent, high density polyethylene (HDPE) dome, with 5mm thick walls, and a diameter of 785mm by 230mm in height. Fig. 1.6 left shows a cut of the calculated 3-D fields




Figure 1.6: *Left:* Calculated EM fields radiated in the presence of the HDPE top-hat dome. The 5mm thick walls make the top-hat dome radio transparent. *Right:* Manufactured Top-hat dome made of fused HDPE, shown here filled with ROHACELL foam.

radiated by the center dipole inside of the top-hat dome. The calculations confirm that there are no significant reflections in the dome walls and no variations in the dipoles input match characteristics. A picture of the manufactured HDPE top-hat vacuum dome is shown in right panel filled with the three layers of ROHACELL foam including the lower layer with the dipole cut-outs.

Another modification we made was to align the dipole arms as shown in Figure 1.7, left. The figure's right panel shows for comparison the calculated embedded 3-D radiation patterns of dipoles 1, 2 and 8 respectively, inside the top-hat vacuum dome, indicating a vast improvement in the patterns shape.




Figure 1.7: *Left:* Inset of the dipoles in a hexagonal layout with 0° arm rotation. *Right:* Calculated embedded radiation patterns for dipoles 1, 2 & 8 inside the cryostat walls.

Fig. 1.8 shows, on the left the new shorter height thermal shield and dewar wall, and on the right the full system with the HDPE top-hat dome. Thermal modeling showed that the 1st and 2nd stage temperatures would be maintained although the outer ring of dipoles may see a small temperature increase.



Figure 1.8: *Left:* New lower height metal vacuum shield installed on the AO19 cryo-PAF prototype. The dipoles have the new 0° alignment (see Fig. 1.7). *Right:* AO19i dewar after installation of the HDPE “top hat” dome.

This is confirmed during the cool-down experiment. Figure 1.9 shows a comparison of the cool down rates for the 20K and 70K stages of the original AO19 system and the modified system (AO19i) with the polyethylene top hat. The total times were almost identical at about 1.8 days. The steady-state temperatures for both stages were less than 1 K warmer for AO19i compared to AO19.

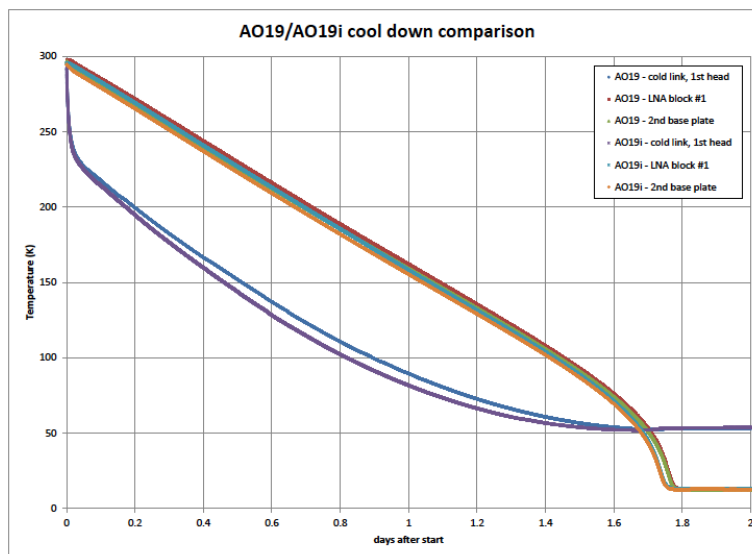


Figure 1.9: Comparison of the cool down rates for AO19 and AO19i. The rates for both the 20 K and 70 K stages are almost identical. The steady-state temperatures for AO19i are less than 1 K warmer compared to AO19, which shows the cryogenics have sufficient cooling capacity to support the increased heat load due to the HDPE top-hat dome.

2. Conceptual Design and Specifications of the AO40 PAF

2.1 Conceptual design: Fig. 2.1 shows a schematic design for the AO40 PAF.

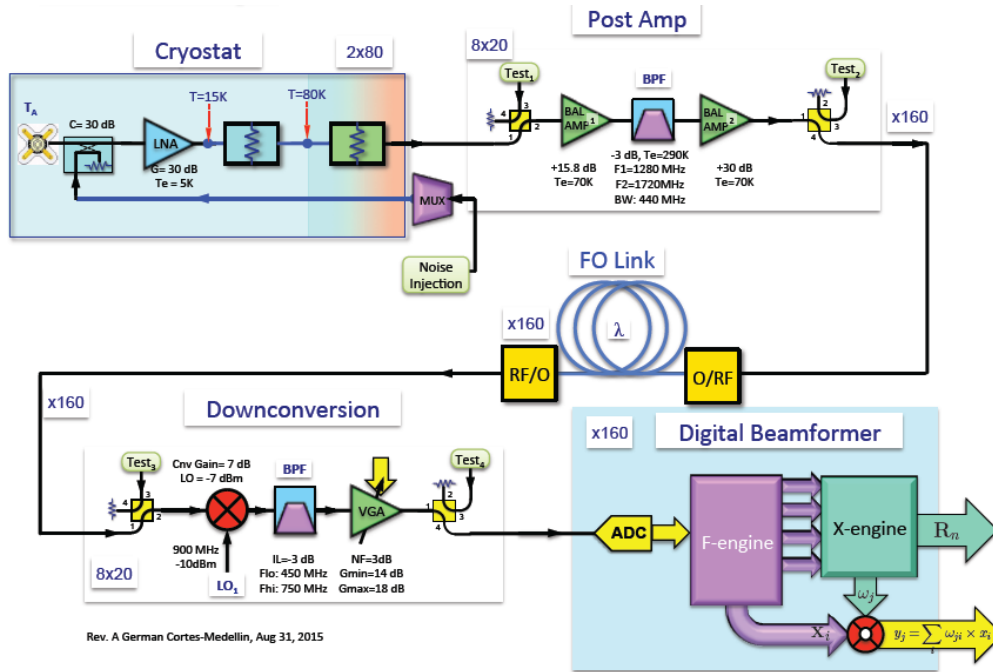


Figure 2.1: The architecture of the proposed AO40 PAF.

2.1.1 Cryostat: 80 dual polarization dipoles feed 160 low noise amplifiers (LNAs) with approximately 30 db of gain. There will be a directional coupler for a calibration noise signal just ahead of each LNA. Four, 20-port multiplexers are used to direct the output of a single noise source one at a time to the two LNAs for each polarization of a single dipole. The cryostat will have two temperature stages, nominally 20 K and 70 K, with the dipoles and LNAs thermally connected to the 20 K stage.

2.1.2 Post amplifiers: The 160 output signals from the cryostat undergo bandpass filtering (center frequency 1500 MHz, BW 1280 to 1720 MHz, this band can be modified in the final design) and amplification stages with about 30 db of total gain. This will be followed by optical fiber transmitters to send the rf signal over fiber about 1,933 ft to the operations building. There will be no mixing stage or digital conversion in the Gregorian dome to minimize both complexity and self-interference in the Dome.

2.1.3 Monitor and Control in the Dome: Cryostat performance, temperatures at several points and vacuum pressure and rf levels will need to be monitored. Control will be required for bias voltages for the LNAs, the operation of the noise source and multiplexers, rotation of the dewar to track the parallactic angle over +/- 15 deg. and monitoring of the rf levels prior to the rf to fiber optical transmitters. This

equipment will be housed in a rf shielded rack mounted, possibly, over the cryostat and its support structure.

2.1.4 Operations building: Optical to rf receivers will be used to convert the 160 signals to rf followed by a bandpass equalization stage and a variable gain amplifier to adjust the signal levels. The signals will then be mixed down to 600 MHz center frequency and filtered with a 300 MHz bandpass filter (450 MHz to 750 MHz) that will give about 70 db rejection of the image signals after 8-bit digitization at a 0.80 GHz rate. The mixer local oscillator frequency will be adjustable to allow the 300 MHz to be selectable within the 1.280 GHz to 1.720 GHz band.

2.1.5 The F-engine: The 160 digitized signals will be polyphase filtered, performed in the baseline design utilizing 25 CASPER ROACH II boards, producing 800 coarse frequency channels per polarization with 390.6 kHz bandwidths spanning the 312.5 MHz bandwidth of the beam-former stage. This will be followed by a “corner turn” stage to convert 160 digital streams each with 800 channels to 800 x 2 streams each with 80 “channels” needed for the input to the X-engine beamformer and the correlator.

2.1.6 X-engine beamformer and the correlator: The baseline design utilizes 25 Windows HPC servers each with 2 GPUs to produce the 40 beams per polarization for each coarse frequency channel. During a calibration operation prior to the observation, the correlator will compute the complex weights needed by the beamformer. In the “pulsar” mode, the output power from the beam formed coarse channels will be sampled at a 64 μ sec interval with 4- or 8-bits. Depending on the number of bits, data rates will be close to 0.5 or 1.0 GBPS, respectively.

The fine spectral resolution mode will produce high resolution spectra over a limited RF bandwidth with the output power sampled at > 0.1 sec intervals. In the case of observations of atomic hydrogen (HI) in external galaxies, the power spectra with spectral resolutions of 12.2 kHz for the formed beams corresponding to 384 coarse channels (~ 150 MHz RF bandwidth) will be obtained and sampled. Assuming 16-bit samples, data rates in this case will be ~ 1 MBPS. The total number of fine spectral channels will be fixed but frequency resolution can be traded for total bandwidth (e.g. for high resolution HI mapping of the galactic plane).

SETI observations would require frequency resolutions of ~ 1 Hz over the maximum possible bandwidth. The SETI project would be responsible for implementing this part of the system.

2.1.7 Data storage: For the “pulsar” mode data storage is a significant issue. An 8-hour observation may require 16 TBytes of storage.

2.2 Specifications:

2.2.1 Operational specifications:

- Number of dual polarization beams: 40
- Rotator: +/- 15 deg; needed for pulsar search integration times of < 15 min.
- System temperature: specification < 35 K, goal < 30 K
- Beam spacing: 2 arcmin, close to Nyquist spacing but selectable within FOV
- Frequency range: 1280 MHz to 1720 MHz (see Section 2.3)
- Total bandwidth: 440 MHz
- RF input BW to beam former: 300 MHz (selectable within the 440 MHz band. The beamformer will process a bandwidth of 312.5 MHz)
- Number of coarse channels per polarization: 800
- Coarse channel BW: 390.6 kHz (dictated by fast transient searches)
- Fine individual channel BW: 12.2 kHz for HI in external galaxies, finer for other observations such as galactic HI.
- Fine channel overall bandwidth: 150 MHz max for fine channel BW of 12.2 kHz, scaled for narrower fine channel BWs.
- Sampling rate for beam formed coarse channels: 64 μ sec
- Sampling rate for fine bandwidth channels: > 0.1 sec

2.2.2 RF specifications:

- Number of dual polarization dipoles: 80
- Low noise amplifiers: 160; < 10K noise temperature, 30 db gain, <20 mW dissipation
- Calibration: ~-30 db noise injection coupler before LNA, polarizations correlated. Multiplexer will feed one dipole polarization pair in each of the 4 dipole array quadrants at a time.
- Second stage amplifiers: 90 K noise temperature, ~30 db total gain

2.2.3 Beam-former specifications

- ADCs: 8-bits specification, 10 or 12-bits goal
- Total beamformer bandwidth: 312.5 MHz
- Coarse channel bandwidth: 390.6 kHz
- Number of coarse channels: 800
- Number of beams per polarization: 40
- Beam angular spacing: Nominally 2 arcmin but selectable within FOV
- Pulsar & FRB mode: 64 μ sec dual polarization, total power sampling rate
- Fine spectral resolution: 12.2 kHz over 150 MHz BW, resolution and bandwidth tradable. Dual polarization, total power sampling rate of > 0.1 sec.

2.2.4 Physical specifications:

- Maximum flange diameter of the cryostat: 1420 mm (set by rotary floor limitations)

- Maximum weight on the rotary floor: ~2,000 lb (set by total weight of ALFA system)
- Maximum diameter of the foam: 1360 mm
- Maximum diameter of the dipole array: 1220 mm, 115 mm dipole spacing
- Second stage amps, filters, etc plus monitor & control in RF shielded enclosure(s) possibly mounted above the dewar
- RF and beamformer rack space in operations building: 5 x 6ft RF shielded standard racks

2.2.5 Cryogenic system specifications:

- Requirement: 80 dipoles, 20 mW per LNA, cold sinks: 70 K @ 1st stage (110 W), and 20 K @ 2nd stage (9 W).
- Number of cold heads: 4
- Refrigerator type CTI 1050
- Compressors: Oxford M700 air cooled – each M700 can support 2 x CTI 1050 refrigerators
- Total power requirements: 26 kW peak, 340-525 VAC, 3 phase

2.3 Frequency range and RFI: The recommended frequency range for the A040 dipoles is 1.280 to 1.720 GHz dictated by RF interference at Arecibo below 1.280 GHz. This frequency range will cover red shifts of up to 0.11 for observations of atomic hydrogen in external galaxies and will allow observations of the four hydroxyl (OH) lines in our own galaxy. The frequency range can be modified at the time of the final design of the system.

Major sources of ground and space based RF interference at Arecibo are:

FAA related radars:

San Juan: 1252.4, 1257.6, 1344.4, 1349.6 MHz

Punta Salinas: 1274.6, 1269.4, 1327.4, 1332.6 MHz

All blanked +/- 2 degrees about the direction to the Arecibo Observatory

Aerostat radar:

Lajas, SW Puerto Rico: 1261 and 1246 MHz

Blanked +/- 20 degrees about the direction to the Arecibo Observatory but still very strong.

Space based transmitters:

GPS L1: 1575 MHz

Iridium: 1615 to 1630 MHz

GLONASS: 1600 – 1608 MHz

INMARSAT: 1530 – 1560 MHz

In May 2015 an A019 dipole was mounted at the Arecibo telescope's Gregorian dome focal point to sample the RF interference environment. The project was organized by Ganesh Rajagopalan with considerable assistance from the Observatory's electronics and computer departments. Fig. 2.1 shows the **peak integrated power** out of the dipole averaged over 300 kHz for the frequency range 300 MHz to 2,500 MHz.

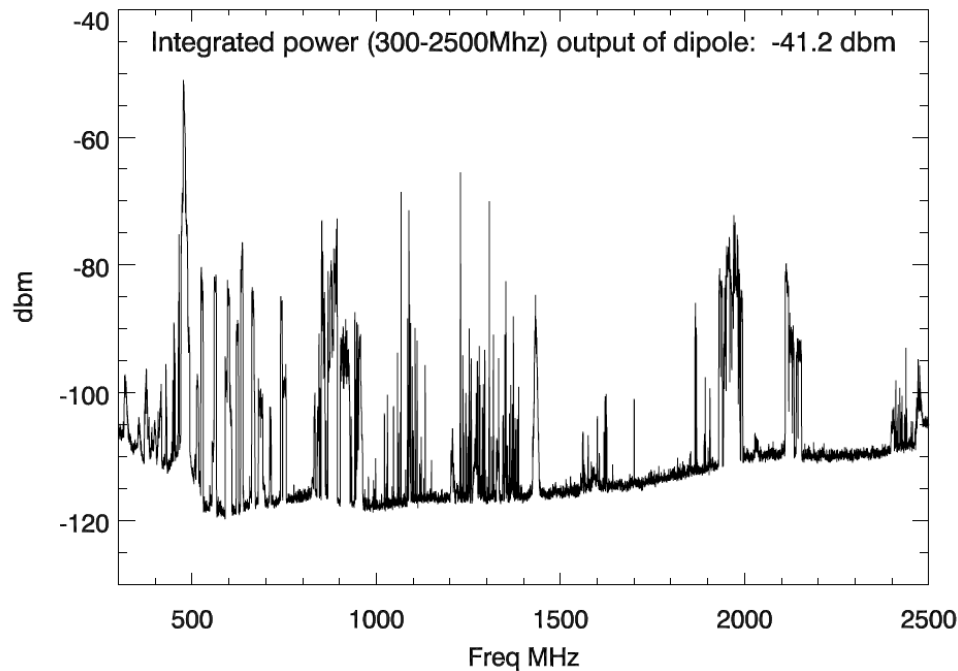


Figure 2.1: Peak dipole output power as a function of frequency. The peak hold was 15 seconds and the power averaged over 300 kHz. The spectrum was scanned in 165 MHz increments in a period of 230 seconds with the zenith angle at 18° and the azimuth scanning. This was repeated 19 times so that each point is the peak power recorded over a ~ 5 minute period. The Aerostat radar was not in operation when these measurements were made. N.B. This is a plot of the peak power recorded, it is not the average power so depicts a worst case situation. Plot courtesy of P. Perillat.

3. Physical Constraints

3.1 Limits on the size of the dewar: The spacing of the radial beams of the rotary floor of the Gregorian dome limit the maximum size of the cryogenic dewar (Fig. 3.1) to 1462 mm. To allow 21 mm clearance on both sides the maximum size of the cryostat including flanges has been set to 1420 mm.




Figure 3.1: The maximum diameter of a circle that can fit within the radial beams of the rotary floor of the Gregorian dome is 1462 mm. Based on rotary floor drawing provided by Arecibo Observatory.

However, there are several structural interferences that will obstruct the installation of the dewar or the rotation of the floor with the dewar in place. Some re-routing of electrical conduit may be required (Fig. 3.2). The most serious obstruction is a bank of switches that limit the maximum floor rotation (Fig. 3.2). Discussions with Observatory personnel have indicated that it should be possible to move the limit switches to a location above the floor rather than below.





Figure 3.2: *Left*: conduit that may interfere with installation of the dewar. *Right*: bank of limit switches that need to be moved to allow the floor to rotate with the 1.420 m diameter AO40 dewar installed.

The sector of the Gregorian receiver room floor that houses the ALFA system has additional beams that were installed to reinforce the aluminum floor plate to

support the weight of the ALFA system. The large diameter of the proposed AO40 dewar means that it will be supported directly by the radial beams of the rotary floor. If the sector of the rotary floor that houses the ALFA system is used for AO40 then the added beams will need to be removed. Discussions with AO personnel have indicated that this should not be a problem although they would wish to consult with the Observatory's engineering firm Ammann & Whitney.

3.2 Skirt: A flared aluminum skirt attached to the metal flange of the dewar has the potential to improve the performance of the AO40 system by reducing spillover of the dipole radiation patterns outside of the tertiary reflector of the Gregorian system. However, the size of the dewar leaves little room for a skirt (Fig. 3.3). A skirt would also potentially simplify the design of a shield that would limit the rf power that the system sees while the S-band or 430 MHz transmitters are operating.




Figure 3.3: Schematic of the 1.420 m diameter cryostat installed on the rotary floor showing the limitation on installing a flared skirt (red dashed line) attached to the metal flange of the dewar can. The limitation is a structural member in the floor that cannot be altered.

3.3 Weight: Based on discussions with Arecibo Observatory personnel, the weight of the AO40 system on the rotary floor of the Gregorian dome will be limited to the current weight of the 7-beam ALFA system, which AO40 will replace. Total weight of the ALFA system from Observatory personnel is 1,981 lb for the receiver and 300 lb for the shutter, limiting the AO40 weight to 2,281 lb maximum including a new shutter (just over 1,000 kg). Removal of the ALFA floor beams and extra thick floor plate may add to the weight allowance.

3.4 Rotator: Based on the very considerable mechanical complications and the associated weight, it was decided not to include a rotator that would track the parallactic angle over the entire feasible observing time. However, long integration times, ~15 minutes, required for pulsar searches require some rotation and +/- 15 deg. has been set as a goal. The option of changing the beamformer weights to track a location on the sky digitally is not feasible due to the time required to determine each new set of weights.

One possible observing mode would have the azimuth arm aligned with the local meridian and observing the sky at a fixed zenith angle as the Earth rotates. This

would keep the beam former weights fixed for an extended period and optimize the use of the observing time. The penalty is that each point on the sky would be observed for about 1 minute – the time to drift through the roughly 15 arcmin diameter of the FOV. The ALFALFA survey of atomic hydrogen in external galaxies was conducted in this mode.

4. Cryostat, dipole array, LNAs and noise calibration

4.1 Cryostat and dipole layout: The AO40 Cryo-PAF camera consists of an array of 80 dual polarized dipoles (160 channels), the maximum that can be accommodated within the physical constraints. The camera will produce 40 simultaneous beams in the sky with a beam spacing of ~ 2 arcmin (selectable within the FOV) and a beam-formed system temperature specification of 35 K with a goal of < 30 K.




Figure 4.1: *Left:* Front view of the Cryo-PAF front-end camera, with relevant dimensions. *Right:* Back view of the front-end camera, indicating the 4 cryo-cooler heads, turbo-pump and the 160 RF channels out.

The Cryo-PAF camera will have a high density polyethylene (HDPE) “top-hat” dome with a flange diameter of 1420 mm and an overall height from the back plate to the dome face of 610 mm (Fig. 4.1). With a maximum foam diameter of 1360 mm, thermal issues dictate that the dipole array be limited to a diameter of 1220 mm. Maximizing the number of elements and the size of the array increases the field-of-view (FOV) and the sensitivity of the outer formed beams. The maximum number of dual polarization dipoles that can be fitted within this diameter is 80 with the outer arms of four of them slightly outside the desired diameter (Figure 4.2). The outer arms of these 4 dipoles may have slightly elevated temperatures.

The Cryo-PAF camera has two temperature stages, the ground plane is at < 70 K and the dipoles and LNA's at < 20 K. The dome will be filled with 3 layers of Rohacell HF foam, with cutouts in the lower layer to provide thermal isolation from the dipoles. The larger thermal load requires the use of 4 CTI 1050 cold heads instead of just the one required for the AO19 prototype. Fig. 4.1 provides an external view of the current Cryo-PAF camera concept.




Figure 4.2: *Left*: partial view of the Cryo-PAF front-end camera, with the top-hat dome and foam layers removed. *Right*: Front view of the square dipole array layout, with element separation of 115mm.

Fig. 4.2 (left panel) shows a partial view of the 80-dipole array without the top-hat dome and foam layers. The right panel shows a front view of the dipole array layout. The decision to use a square grid instead of a hexagonal grid was based mainly on cryo-mechanical considerations. The dipole orientation selected has the dipole arms aligned in a square grid rather than the hexagonal one with the dipole arms at 45° as was used for the AO19 prototype system (Cortes et al, 2015). This maximizes the number of dipoles but at the expense of some co-polar coupling between adjacent dipoles, -10 dB vs -15 dB. The use of a square grid also has very practical advantages as the absence of a central dipole allows division of the cryostat into 4 quadrants, each with exactly 20 dipoles, (see Figure 4.3), greatly facilitating the cryo-mechanical design.

In theory, an hexagonal dipole array could contain more dipoles within a fixed diameter. However, for this layout the dipole spacing only needs to be 0.6λ as opposed to 0.5λ for a rectangular array so the number of dipoles in the two configurations would be roughly the same. Most importantly, the rectangular array does not have a central dipole element allowing the array to be split into four identical quadrants greatly simplifying the cryogenic design (Fig. 4.4). Non-circular options were investigated (Fig. 4.4) but discarded on the grounds that too few dipoles were added given the additional complexity and weight resulting from an

elliptical or distorted-shaped pressure vessel, and some options would prevent rotation of the front end during observations.

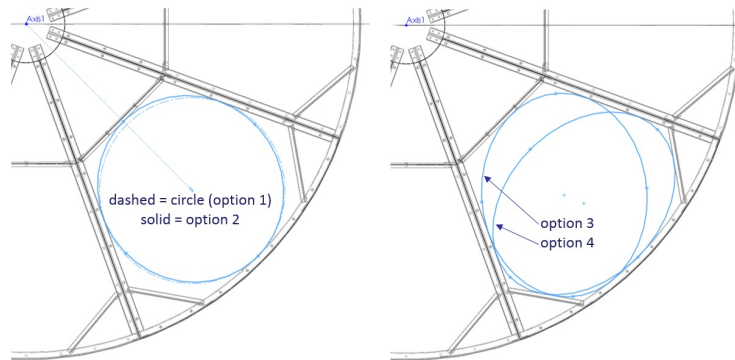


Figure 4.3: Non-circular options for the overall cryostat shape. These options were rejected due to weight, lack of rotator, and small gain in dipole number for the added complexity. Option 4 (wide ellipse) actually has less area than the baseline option 1 (circle).

Fig. 4.4 right panel, shows a single cryostat wedge quadrant. Each quadrant has a single CTI 1050 cold head. The second stage is completely contained inside the first stage for each quadrant and there is no direct thermal or mechanical connection between the temperature stages of different quadrants, which allows for independent thermal contraction of each quadrant to its cold head, simplifying the design.




Figure 4.4: *Left:* Dipole array divided into quadrants. *Right:* Cross-section view of a single quadrant.

The gaps between the first stage quadrant ground planes are bridged with aluminized Mylar to form a continuous ground plane for the dipole array. The first stage structure, along with its insulating standoffs (G10 blades and columns), support and transfer the atmospheric vertical load (~ 4 tons per quadrant), from the front of the cryostat to the back plate.

4.2 Scaled Dipole Design: The dual polarization dipoles are sleeve dipoles with an operating bandwidth of ~ 500 MHz. For the A019 prototype system the dipoles were designed to cover the frequency range from 1.22 GHz to 1.67 GHz. As discussed in Section 2.3, the suggested frequency range for A040 is 1.280 to 1.720 GHz. The A019 dipole design has been scaled to this new frequency range. However, since some components of the dipole are “components off the shelf”, only the dipole height and dipole arm dimensions were scaled (Fig. 4.5) in order to shift the s-parameter responses to the new operating frequency band (Fig. 4.6).


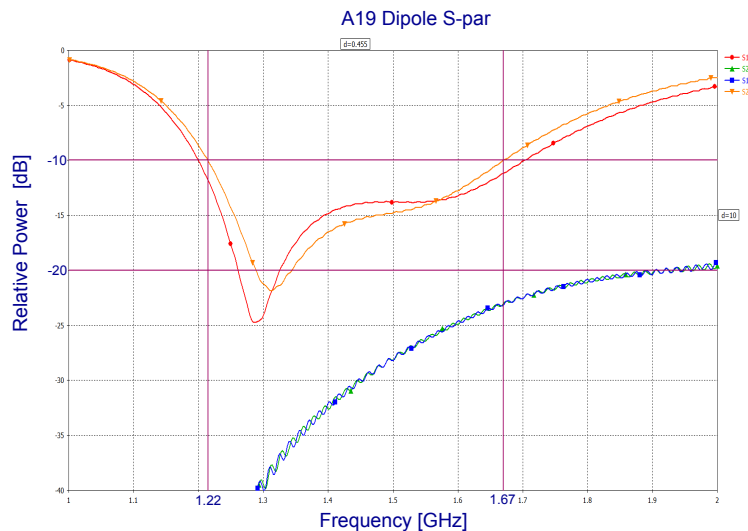


Figure 4.5: *Left:* A019 dipole general dimensions. *Right:* A040 scaled dimensions. Only the height and dipole arm length were modified.

The port-to-port isolation has improved from -23 dB to better than -40 dB (compare Fig. 4.6 top and bottom) but additional tuning will be required prior to fabrication to improve the response over the band.

Figure. 4.7 shows a comparison of the calculated far field patterns between the original A019 and new modified A040 dipoles. The modified A040 far field radiation patterns have slightly more directivity than the original A019 patterns but otherwise are almost indistinguishable.



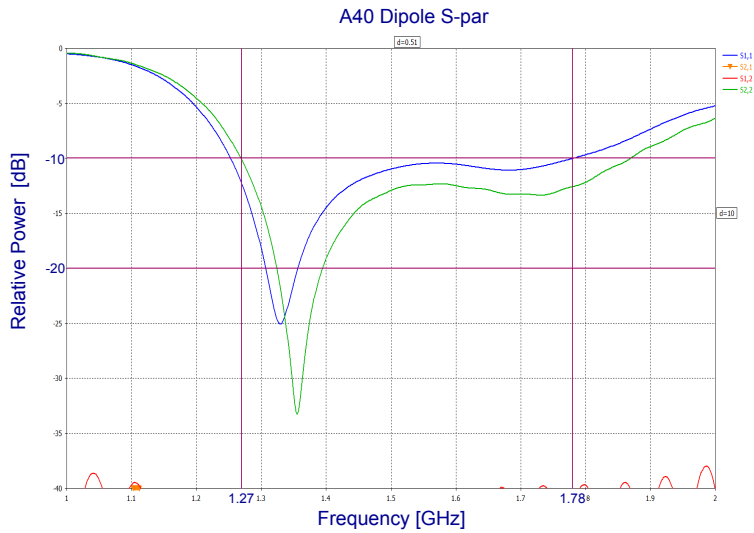


Figure 4.6: *Top:* Calculated S-parameter responses for the original A019 dipole. *Bottom:* modified A040 dipole. The A040 response can be improved further prior to the final design.

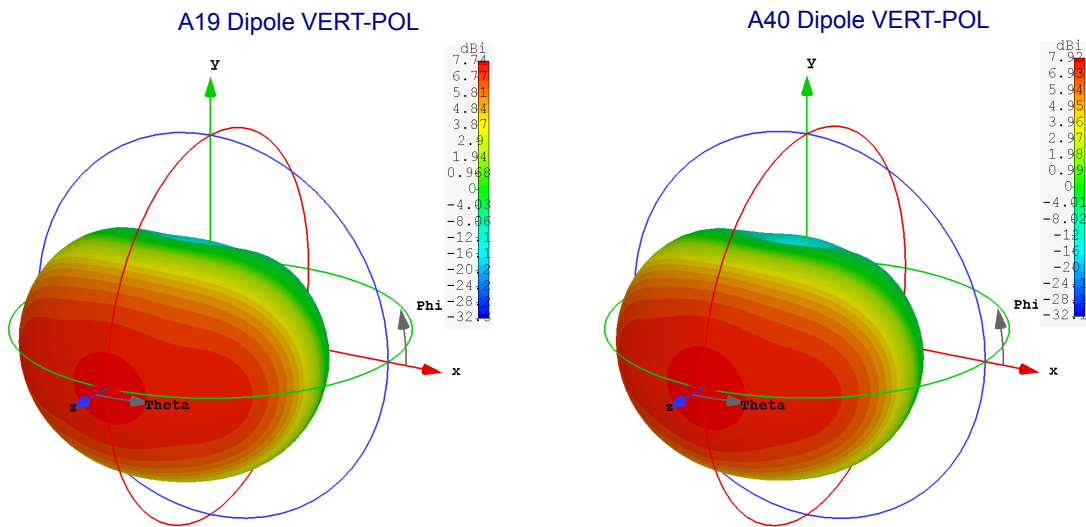


Figure 4.7: Calculated far field's radiation pattern at 1.4 GHz. *Left:* A019 dipole (directivity 7.7 dBi). *Right:* A040 dipole (directivity 7.9 dBi).

4.3 Low Noise Amplifiers (LNAs): The current specifications for the LNAs are as follows:

- Frequency: 1.0 to 2.0 GHz
- Gain: 30 dB
- Noise Temp: 5K
- 1dB Comp: 0 dBm

Power: 20mW

These specifications are based on discussions with the group under Chris Groppy at Arizona State University (ASU) and based on a design by Hamdi Mani. This group is capable and willing to provide the required 160 cryo-LNAs plus spares.

Fig. 4.8 shows the measured gain, input and output match of a cryo-LNA at 10K developed by the ASU group. The gain is 30 dB at the frequency of interest for AO40 with an input match better than -15dB and an output match of -10 dB at 1.4 GHz.

Fig. 4.9 shows the measured noise temperature across the band, indicating a noise temperature better than 5K from 1.0 to 1.6 GHz

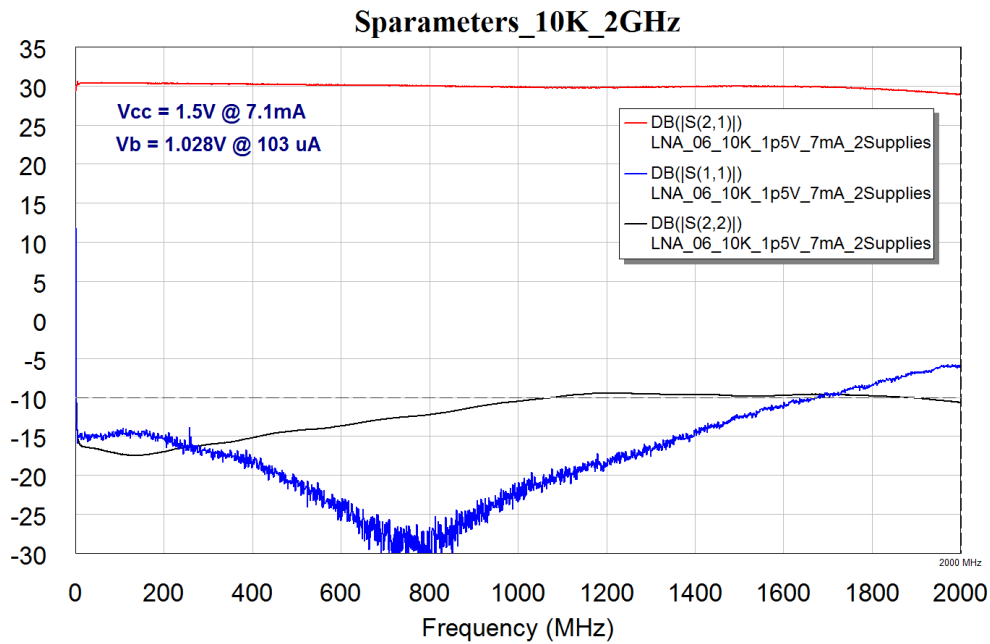


Figure 4.8: Measured S-parameters of a cryogenic ASU LNA, with a gain of 30 dB and input match better than -20 dB at 10K. (Courtesy of Chris Groppy and Hamdi Mani, Arizona State University).

1MHz - 2GHz Cryo LNA #06 Noise Temperature [K] @ 10K

Bias: $V_{cc}=1.5V$ @ $I_c=7.1mA$ / $V_b=1.028V$ @ $I_b=103\mu A$

01/02/2015

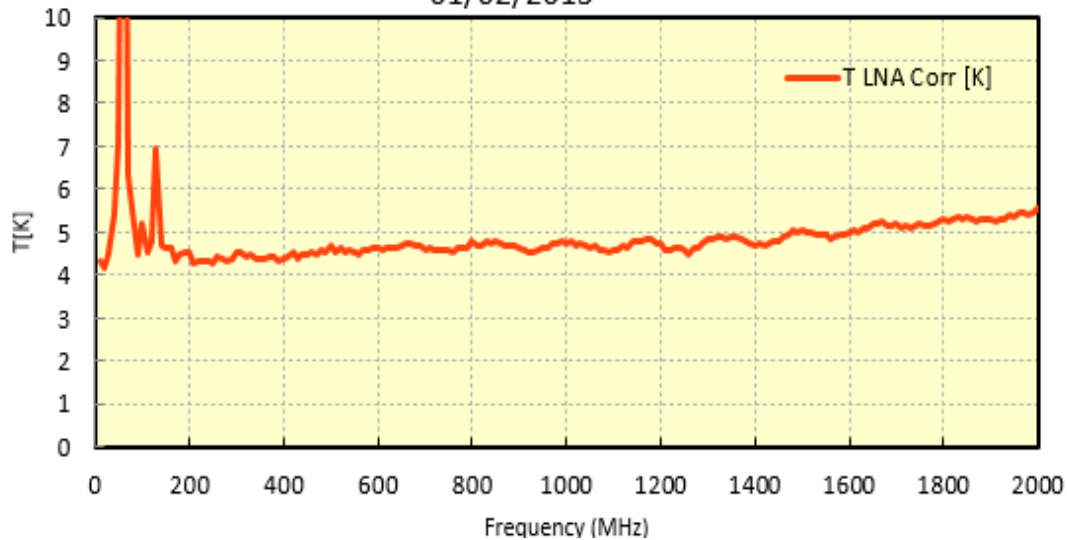


Figure 4.9: Measured noise temperature of the cryogenic ASU LNA, (30dB gain) at 10K. (Courtesy of Chris Groppy and Hamdi Mani, Arizona State University).

4.4 Noise calibration

The ability to inject a calibration noise signal just prior to the LNA will allow fast and frequent monitoring of the health of each of the 160 rf channels. The output of a noise diode can be fed via a switch to four multiplexers each of which has twenty outputs feeding the noise signal via striplines one-by-one to the 20 dipole/LNA assemblies of each quadrant of the dipole array (Fig. 4.4). A power splitter will feed directional couplers just ahead of the LNA for each polarization.

4.5 Dipole/LNA assembly:

The concept of a re-insertable sleeve dipole-LNAs package developed for A019 cryo-PAF is also used in A040 as shown in Fig. 4.10 left. In order to ensure a good thermal contact between the dipole-LNA package and the 15K second stage we also use a cold finger cylinder with a nylon bracing (Fig. 4.10, right). At cryogenic temperatures the nylon brace contracts around the cold-finger cylinder and dipole-LNA cylinder assembly yielding a very strong thermal contact.



Figure 4.10: *Left:* Assembled dipole-LNA package with cylinder cover in place and blind-mate connector plug at the bottom. *Right:* View of cold-finger cylinder fixed on the 2nd stage plate with nylon ring that braces firmly the dipole/package at cryogenic temperatures.

The LNA's will be mounted inside of the dipole-LNA package as seen in Fig. 4.11. The LNA carrier will be split in two in AO40 to allow independent field replacement of an individual polarization LNA.




Figure 4.11 Internal view of dipole-LNA package showing the LNA's carrier.

In the baseline design we use the custom made blind-mate RF connectors developed for AO19 cryo-PAF, for easy insertion and extraction, which allows servicing through the front opening of the cryostat in a plug-and-play style manner.

Fig. 4.12 shown on the left a detail of the blind-mate receptacle mounted on the bulkhead of the 2nd stage of the cryostat. On the right we show the plug connector. The press-fit or blind-mate connectors provide DC bias, lines for voltage and current

monitoring, and RF ports for Pol-A, Pol-B and noise injection.




Figure 4.12 Baseline Blind-mate connector. *Left:* Receptacle mounted on 2nd Stage plate (4 RF coax lines will be used). *Right:* Designed plug connector on Dipole/LNA package.

4.6 Post cryostat RF and IF stages

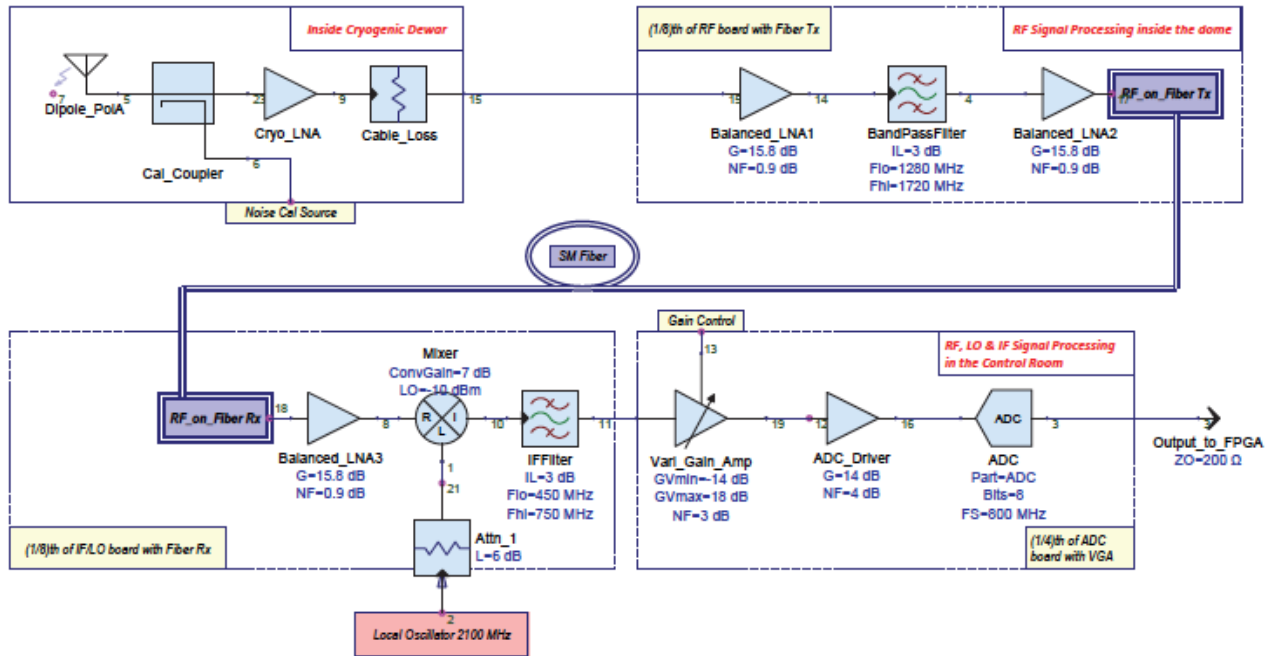


Figure 4.13: Simplified block diagram of the AO40 rf and IF signal chain for a single polarization channel showing specifications for amplifier gains (G), noise figures (NF) and bandpass frequencies. The fiber optic line carrying the 160 rf over fiber signals from the Gregorian dome to the operations building is close to 1,933 ft.

Fig. 4.13 shows the entire rf signal chain for a single channel. The 160 rf outputs from the cryostat are amplified and band pass filtered (1,280 MHz to 1,720 MHz) prior to being sent to the operations building via a rf over optical fiber lines about 1,933 ft in length. In the operation building the signal is first amplified. A mixing and bandpass filtering stage allows selection of the 300 MHz band with a center frequency of 600 MHz that will be processed by the digital beamformer. Controllable variable gain amplifiers will be used to set the final signal levels before they are passed to the 8-bit analogue-to-digital converters (ADCs). Since the ADCs will have sampling bandwidths of 400 MHz (800 MHz sampling rate) the second sampling Nyquist window (400 to 800 MHz) is being used. The rf components in the dome would be housed in RF shielded enclosures.

The two major issues are the design and cost of the fiber optic transmitters and receivers and the installation issues related to running the fiber optic cable from the Arecibo telescope's Gregorian dome to the operations building. Extensive discussions were held with potential suppliers of these devices. Among them, Optical Zonu Corporation and Emcore Corporation supplied us with sample fiber transmitters and receivers on printed circuit board (PCB) modules. Short of developing transmitters and receivers from individual components, we have budgeted for Emcore's 3GHz Fiber Optic Links PCBs at a quoted estimated price of ~\$700 for a transmitter/receiver pair. Total cost would be between \$112k and \$120k depending on the number of spares. Selecting suitable individual components for the transmitters and receivers and integrating them on to the post amplifier/filter board and the down-converter board, respectively, remains an option for the final design.

Arecibo's azimuth rotary joint has a specially designed cable wrap system that can accommodate three fiber optic cables. One of these is associated with the ALFA system and could be replaced by a single cable with about 200 fibers for the A040 system. A potential issue is the diameter of the new cable versus the size of the grooves in the drums of the cable wrap system. At worst, these drums could be replaced with new ones with suitable groove sizes.

4.7 Front End Monitor and Control:

Fig. 4.14 Shows the Monitor and Control (M&C) system for the A040 cryo-PAF camera. There is a separate bias for each individual cryo LNA, with 4 wires inside the cryostat, 2 wires for biasing and 2 wires for sensing. The bias is controlled by a National Instrument controller (NIC) that also will sense the actual voltages and currents applied to the LNA's terminals mounted on the 15K stage. Noise also will be injected to each cryo LNA path for individual noise calibration. There is a single noise line source, controlled by the Observatory computer, and connected to 4, digitally controlled, multiplexer switches (20P1T) to inject noise to each dipole assembly in sequence. A power splitter will feed the noise signal simultaneously to

directional couplers ahead of the LNA for each polarization. Since four multiplexers are used to distribute the noise signal to the dipole assemblies in each quadrant, the noise signal can be sent to four dipole assemblies, one in each quadrant, as selected through the NIC. The NIC will also record the temperatures and vacuum gauge measurements during system operation and allow monitoring of the rf levels prior to the rf to optical fiber transmitters.

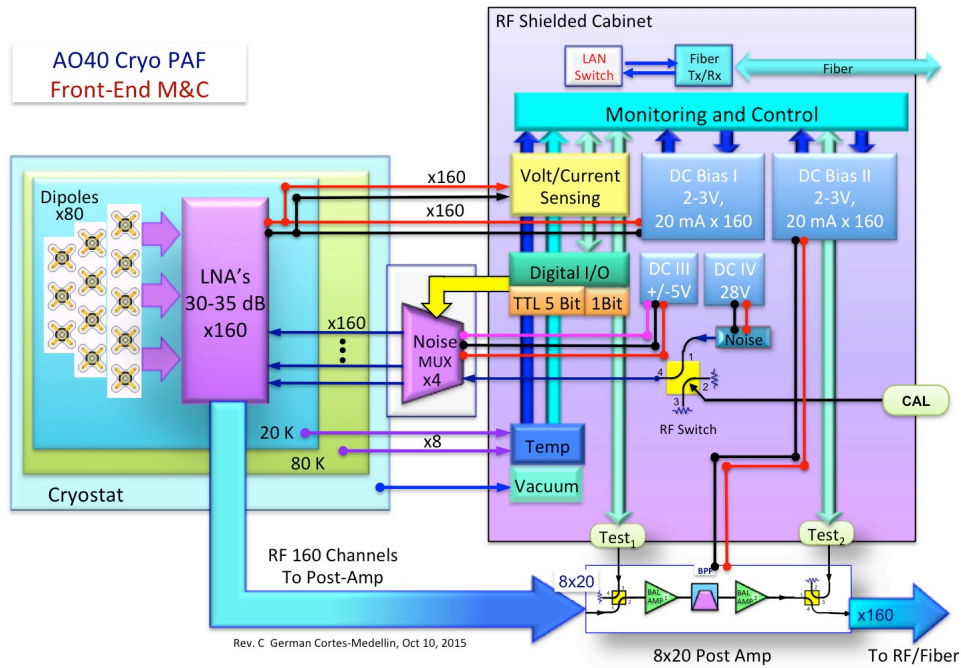


Figure 4.14: AO40 monitor and control system

The post amp section will have an independent bias box with no bias sensing. We anticipate that the post amp boards will have 2 test points with digital RF switches, along each LNA path to verify path normal gain operation.

5. Cryostat Structural and Thermal Modeling

5.1 Cryostat structural issues: The structural design for AO40 is based on the successful design of AO19i, however, there are three key differences between the two instruments. The first is the dewar diameter (~1.4 m vs. ~0.8 m), the square of which is proportional to the atmospheric load while under vacuum. The second difference is the use of four cryocoolers to handle the thermal load (see §5.2), which must utilize some sort of flexible connection to the inner cold stages to prevent damaging loads from developing due to differential thermal contraction. Finally, the third difference is the desire to rotate the cryostat to remove the parallactic angle over long integrations.

5.1.1 Dewar size: Extensive finite element analysis (FEA) was utilized for the design of AO19i to ensure that the HDPE top-hat dome would be adequately supported by the Rohacell foam. Whereas the foam in AO19 only had to support the vertical atmospheric load against the thin Kapton the window, AO19i (and AO40) must also support the compressive radial load from the atmosphere against the HDPE, since it is both too weak and too thin to resist this load on its own. The minimum HDPE top-hat wall thickness is set to ~5 mm for manufacturability, which is also good from an RF standpoint, minimizing distortion of the incoming electromagnetic field. Figure 5.1 shows the stress distribution in the foam layers and the top-hat for AO19i, both are well within the material's respective yield strengths. Since this loading is basically hydrostatic, extrapolating to the larger size of AO40 gives similar results, and thickening these components for strength is not required (e.g. the AO40 HDPE can be the same thickness as for AO19i).




Figure 5.1: FEA results broken out by component for AO19i showing stress distribution under vacuum. *Top:* Rohacell foam layers. *Bottom:* HDPE top-hat dome. Both are well within material yield strengths. Distortions are greatly amplified to illustrate model displacements.

Stress in the cylinder of a thin-walled pressure vessel is proportional to its radius and inversely proportional to its wall thickness, so in this case we can scale up the vacuum shield dimensions of AO19i to accommodate the larger size of AO40. Bolt circle details are notional and would need to be validated during the final design effort. The back plate also has to be scaled up by the same factor (thickness proportional to radius) to support the vacuum load. However, the final design can benefit from optimizing a light-weighted version of the plate given the load points

transferred from the 1st temperature stage and the uniformly distributed atmospheric load on the outside.

5.1.2 Integrating multiple cryocoolers: The number of cryocoolers is set by the heat load (see §5.2). Cold head locations are determined by the dipole grid, located between dipoles to minimize mechanical interference issues with other parts and ease cable routing. The cold head is approximately centered within its quadrant minimizing the maximum thermal path for any individual dipole (Fig. 5.2). Within the vacuum back plate, the cold heads are situated at the corners of a 460 mm virtual square (~650 mm diagonal) centered on the plate. If the cold stage base plates were solid instead of broken into quadrants, the thermal contraction to cryogenic temperatures would force the cold heads together by 1.8 mm (2.6 mm on the diagonal), potentially damaging the cold heads.




Figure 5.2: Top-down line drawing view of one quadrant of A040 (only one dipole is shown – bottom right). The cold head with its 8 (6) thermal links for the 1st (2nd) stage is centrally located within the dipole grid (blue circles) to minimize interference with parts and cabling, and to centrally locate the heat sink for the dipoles.

One solution would be to provide a flexible link between the cold heads and the base plates, but this requires links with three degrees of freedom (DoF), two in-plane and one vertical. Instead, by breaking the stages into quadrants we take advantage of symmetry and simplify the problem to two DoF, allowing us to reuse the thermal link design from A019. The link allows vertical and radial compliance while still maintaining a highly conductive connection. Figure 5.3 shows the mechanical and thermal FEA of the 2nd stage thermal link, and a photo of the fabricated part. These links worked very well in A019, with a measured temperature drop of only 0.9 K, nicely matching our thermal FEA predictions.




Figure 5.3: Thermal link for the 2nd cold stage. *Left:* Temperature distribution from thermal FEA showing $< 1\text{ K}$ temperature rise for a $\frac{1}{4}\text{ W}$ heat load at 14 K base temperature ($\frac{1}{2}$ model due to symmetry). *Center:* Fabricated thermal link from OFHC copper, instrumented with temperature sensors for test cool down. AO40 links will be gold coated to protect their surface. *Right:* Mechanical FEA showing strain in notches for expected displacement ($\sim 1\text{ mm}$) well below material limits.

5.1.3 Instrument mount with rotator: Although originally outside the scope of this conceptual study, an instrument rotator was deemed necessary to satisfy some of the scientific objectives. A rotator adds to the overall mass budget and reduces the area available in the vacuum back plate for feedthroughs due to the required keep out zones for the swept area of the mount, but our preliminary design accommodates these additional constraints. Figure 5.4 shows a concept for mounting the cryostat on the Arecibo rotary floor with an integrated rotator. The mount is a truss structure that provides a stiff, light-weight support for the instrument. The slew ring, an off-the-shelf component, is a self-lubricating, low-friction sliding element capable of handling very high loads in all directions. Two motor-gearbox units are used to take out system backlash and drive the slew ring gear. Cable wrap details and ultimate analysis would be completed during the final design phase.




Figure 5.4: A040 mount concept shown installed on the rotary floor. For proper focus, the cryostat must sit almost entirely below the floor plate surface. A truss structure transfers the instrument weight directly to the rotary floor I-beams. A $\pm 15^\circ$ instrument rotator couples the cryostat to the truss. Drive motors are blue, gearboxes orange, and slew bearing with integrated gear is purple.

5.2 Cryostat thermal analysis: The overall heat load for the instrument is based on a custom model developed and used for over a decade in the design of cryogenic instruments at Cornell, where we account for material properties as a function of temperature. The custom model is supplemented by results from thermal FEA for more advanced loading conditions (e.g. the radiative load from the foam to the dipoles). Results for the heat budget are presented in Table 5.1. With four CTI 1050 cryocoolers, the instrument has sufficient cooling for both the first and second temperature stages. Experience has shown that having a 30% margin, or greater, on the cooling power is appropriate for a detailed conceptual design.

Heat Load in Watts	1 st Cold Stage (63 K)	2 nd Cold Stage (18 K)
Radiation	33	0.2
Support Structure	17	0.3
Wire Conduction	4	1.6
Joule Heating	0.1	3.3
Foam (1 st Cond, 2 nd Rad)	59	3.8
Auxiliary Cooling	-6	0
TOTAL	107	9.2
Cooling Power (4x 1050)	180	12
Margin (%)	68%	30%

Table 5.1: Heat load results grouped by source for the 1st and 2nd temperature stages of A040. The Rohacell foam dominates the 1st stage load, while the foam and LNA waste heat dominate the 2nd stage. Auxiliary cooling on the 1st stage comes from the 2nd stage acting as an alternate heat sink through conduction and radiation.

5.2.1 Number of cryocoolers: The number of cold heads is driven by the second stage heat load, requiring four CTI 1050 cryocoolers. This in turn provides ample cooling for the first stage. Options for reducing the number of cold heads or using different types (i.e. combining 1020's and 1050's) were investigated but ultimately discarded due to added complexity to the cryostat design. Only having one type of cold head is also preferable from a maintenance and serviceability standpoint.

The second stage heat load has three main contributors: wire conduction, Joule heating from the LNAs, and the radiative heat load on the dipoles (coupled to the 2nd stage) from the foam (coupled to the 1st stage). Wire conduction is mainly due to the RF transport lines, and could be reduced by increasing the stripline length (see §5.2.3) between stages, but this is a linear effect and is only reasonable up to a point. Joule heating from the waste heat of 160 LNAs could be reduced by using lower power LNAs, but that would negatively impact the receiver performance. The radiative coupling from the foam to the dipoles is also fixed. Adding thin HDPE “shields” around the dipoles, coupled to the first stage, would destroy the structural integrity of the bottom foam layer. Adding more foam is not an option because the dewar size is already at its maximum limits (see §3.1).

5.2.2 Foam & dipole thermal FEA: In order to accurately design AO19, a finite element model was created to simulate the conductive and radiative thermal effects between the foam, ground plane, and dipoles. This model was adapted and expanded for AO40 (Fig. 5.5). The conduction through the foam dominates the first stage heat load, but given the size of the dewar window, the magnitude of the heat




Figure 5.5: Thermal model for the front section of one quadrant of AO40. *Left:* Isometric view of the temperature distribution resulting from a 313 K outside temperature and a 68 K ground plane. *Right:* Looking up through the ground plane reveals the radial temperature gradient that dipoles in different positions will be exposed to. The difference between the temperature of the foam cavity for the central dipole (top-right hole) and the one closest to the edge is ~30 K.

transferred is relatively low. Although there is no conductive load from the foam to the dipoles, the radiative load is a significant fraction of the total load on the second stage, exceeding even the waste heat from the LNAs by nearly 20%. However, due to its careful design, the dipole temperature does not rise significantly above its sink temperature, ensuring the LNAs function at or below their operating temperature (Fig. 5.6). Preliminary results predict the variation in LNA base temperature will be ± 0.2 K across the entire AO40 instrument.




Figure 5.6: *Left:* Thermal model for a $\frac{1}{2}$ dipole-cell, including ground plane, foam, HDPE top, dipole, dipole cylinder and LNA mounting block. *Right:* The LNA heat load and radiative coupling between the dipole and the foam cavity results in heat flow to the 2nd stage, and a small temperature gradient along the dipole. Boundary conditions are outside 313 K, ground plane 68 K, and dipole sink 18 K. Maximum dipole temperature is equal to 20.4 K.

5.2.3 Stripline vs. coax: An important lesson learned from the development of AO19 is that stainless steel coax, while provided excellent signal isolation and low thermal conduction, is difficult to route between the vacuum, first, and second stage base plates for 57 lines (Fig 5.7, *Top*). To mitigate this difficulty for AO40's 240 lines, we have designed a stripline, 300 mm long, which can run between stages, and even serve as a vacuum feedthrough (Fig. 5.7, *Bottom*). The calculated heat load from the stripline verse the coax is the same for the first cold stage, while being about 0.8 W higher for the second stage. This represents less than 10% of the total heat load and is well justified by our experiences. The final design can investigate ways to reduce this conductive load.




Figure 5.7: *Top*: Routing the 57 stainless steel coax lines from the vacuum feedthroughs to the dipoles was a challenge in A019, and would simply not be practical for 240 lines in A040. Striplines with sufficient signal isolation can be fabricated from 18 μm copper sheet bonded to 100 μm dielectric. *Bottom-left*: They can be potted to a vacuum fitting to form a feedthrough. *Bottom-right*: Zoomed in section of stripline with upper copper and dielectric layers hidden showing individual traces 118 μm wide on a 1 mm pitch.

6. Digital Receiver

6.1 A040 Correlator / Beamformer Design Philosophy

The A040 back end is a digital receiver system that performs the primary functions of real-time array correlation and real-time beamforming. This is followed by data product storage of integrated beamformer spectra and array covariance matrix estimate outputs to a large, very high data transfer rate Lustre file storage system. Signals for all dual polarized antenna ports of the A040 PAF are included in these continuous (*not* burst mode or batch processed) real-time calculations, with an instantaneous signal bandwidth of 312 MHz.

The correlator function is used primarily to estimate array covariance matrices in response to bright, point-like calibration sky sources which are positioned (by pointing the telescope in a grid pattern around the calibrator) at every angle, relative to dish bore sight, which corresponds to a desired beam steering direction. Calibration covariance matrices are used to compute the optimal beamforming weights, which are in turn used in the real-time beamformer to shape the on-sky directional response pattern of each beam. Ideally, this calibration process is repeated every few weeks to account for system electrical and mechanical drifts

over time. Given the non-standard nature of the Arecibo telescope, the calibration procedure should be performed at short intervals over several weeks to determine the time scale of any drifts with time or telescope motion.

The beamformer linearly combines signals from all 80 dual pol array elements, each antenna weighted by its corresponding beamformer weight, to form 40 dual polarization beams. Prior to beamforming, the full 312 MHz bandwidth is broken down into 800, 390 kHz wide, frequency channels. Separate beamformer weights are computed for every beam (2x40) and every channel (800).

The design philosophy for this digital receiver is to use off-the-shelf components for all major hardware systems, and to base the FPGA firmware, GPU code, and CPU code on existing open source programs from related instruments in the radio astronomical community. Though of necessity a significant amount of new programming and firmware development will be required to meet the application-specific design parameters for this instrument, in many cases we can drastically reduce development time by extending capabilities of existing codes. Though there is yet much development work to accomplish, this approach will dramatically reduce program risks.

For example, all proposed major hardware components are readily available from vendors, and price quotes are included in Appendix A. The FPGA processor units are from the UC Berkeley CASPER collaboration for astronomy signal processing. We have been using these units at BYU for years, and have access to the firmware from a number of groups developing related astronomical array processing systems. We will base our array digitizer and frequency channelizer work on existing CASPER firmware, and on our current development of the FLAG PAF array digital receiver for the Green Bank Telescope.

We have chosen to use a graphical processor unit (GPU) architecture for the correlator and beamformer functions. This also follows our current FLAG development path, and has been shown to reduce firmware/software development time significantly as compared the full FPGA design as was followed by ASKAP. When building 36 identical digital beamformer systems as in ASKAP, the economy of scale argues for the more compact architecture using FPGA processing throughout, and many FPGA chips mounted on each custom board. For one-off projects like A040, it is more important to reduce development costs, so we will design no custom processing boards, and will utilize the simpler programming environment of GPUs for the correlator. This hybrid FPGA – GPU architecture has been successfully adopted by the LIDA, PAPER, and CHIME correlators, the VEGAS spectrometer, the GUPPI pulsar detector, and our nearing-completion FLAG correlator beamformer.

6.2 Functional Description

The digital receiver consists of three main physical subsystems : the sampler/frequency channelizer (or F-engine), the array correlator/beamformer (or XB-

engine), and the mass storage system, (i.e. the Lustre file store). The first two of these are illustrated in the hardware architecture diagram of Figure 6.1. The F and XB-engine configurations are distinguished from each other by the primary mode of parallel independence: the F-engine operates identically and independently per antenna input across all frequency channels; the XB- Engine operates identically and independently per frequency channel. Data paths in each engine are ordered accordingly. The F and XB-engines interconnect via a large “corner-turn” functional block to re-order data packets from antenna-major parallelization, to channel-major parallelization. Five separate 10 GbE non-blocking internet switches, each with 36 10 GbE ports, are used to implement this corner turn and forward data packets from the F-engine to the XB-engine. The F-engine samples IF antenna signals with 8-bit ADCs, and uses Xilinx Virtex 6 field programmable gate arrays (FPGAs) to implement frequency channelization as a large 1024 point polyphase filter bank (PFB). Output data from the PFB are calculated with 18 bit precision, then re-quantized with a sliding window bit-slice function to 8-bits real + 8 bits imaginary for transport to the XB engine. This proposed design adopts the UC Berkeley CASPER group’s ROACH 2 FPGA boards and functional libraries for the F-engine.

Figure 6.2 illustrates the mathematical operations of the XB-engine, which performs either a real-time integrating correlator (array covariance matrix estimate) for each frequency channel, or a 40 dual pol real-time beamformer. Post beamforming, an optional fine PFB function may be enabled to permit high resolution 12 kHz wide frequency channelization over about 150 MHz of the 312 MHz bandwidth. In either coarse or fine PFB mode, the outputs are squared and integrated to form beam-by-beam spectrometer estimates. In coarse PFB mode it is possible to dump these integrated spectra out to the Lustre file store at a fast 64 μ s rate, which is suitable for dispersion measure estimation in pulsar and fast radio burst searches. In fine PFB mode, much longer integration times are used for spectral line (e.g. HI) searches, so the fastest permitted integration dump interval is 100 ms.

All of the operations in Figure 6.2 are calculated using 50 highly parallel graphical processing units (GPUs) hosted, 2 each, in 25 high performance rack mount server PCs (HPCs). The proposed Nvidia G-Force GTX-980 TI boards offer a good mix of high memory I/O bandwidth, large numbers of parallel cores, and an attractive price point because they do not include the (unnecessary for this application) double precision floating point capabilities of the Tesla series. The HPCs have a unique architecture design with two independent memory busses, two PCIe busses, and two 10 GbE internet nics, each with dedicated paths to their corresponding Xeon CPU chips. This permits the two GPUs to operate independently without blockage when accessing shared memory or a 10 GbE network switch.

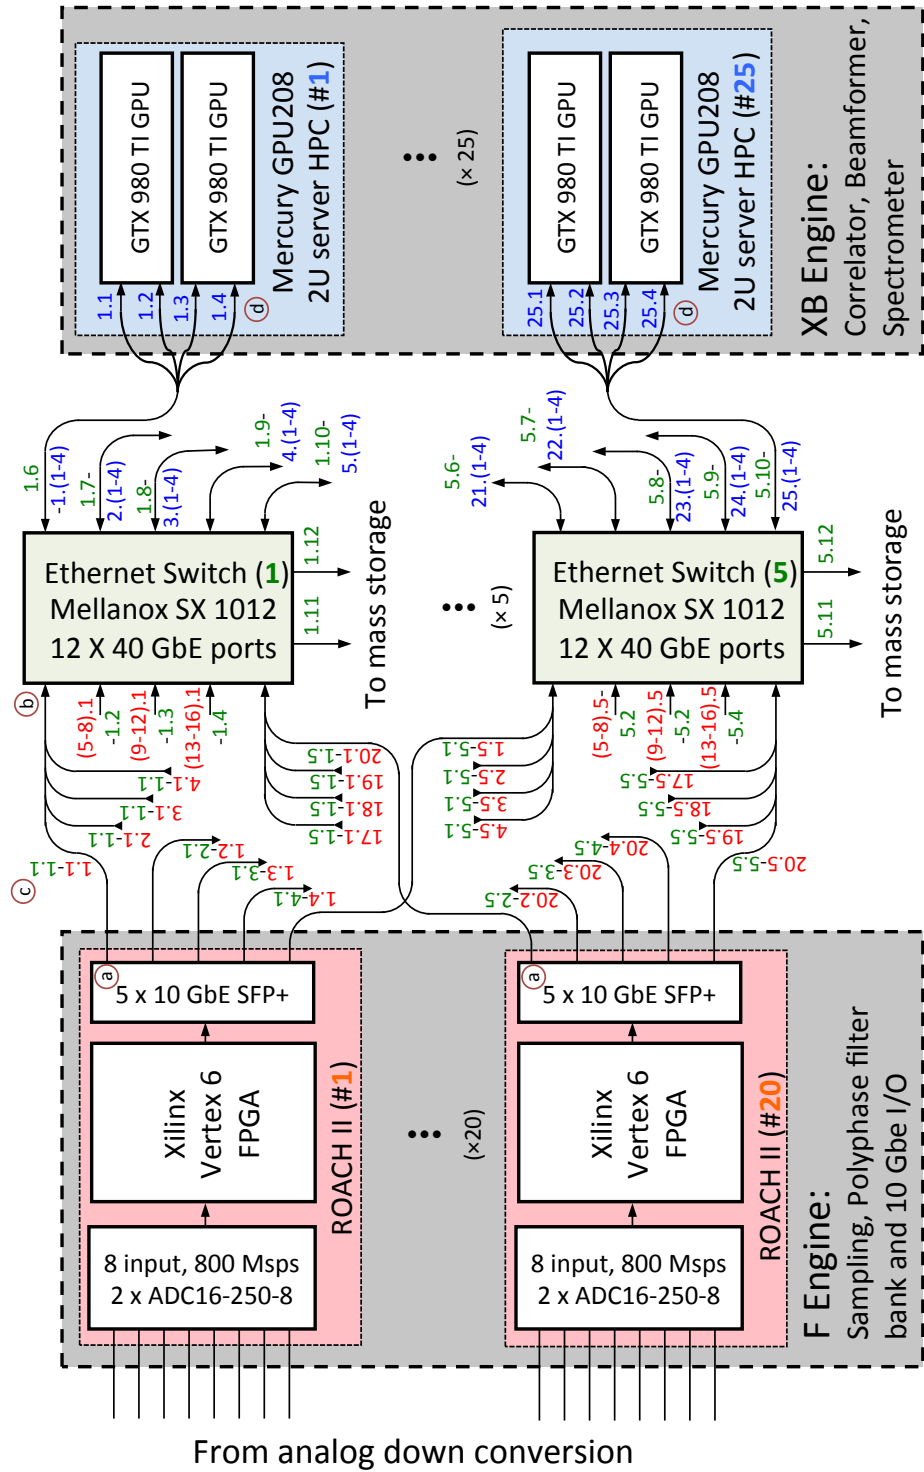


Figure 6.1: AO40 digital backend hardware configuration

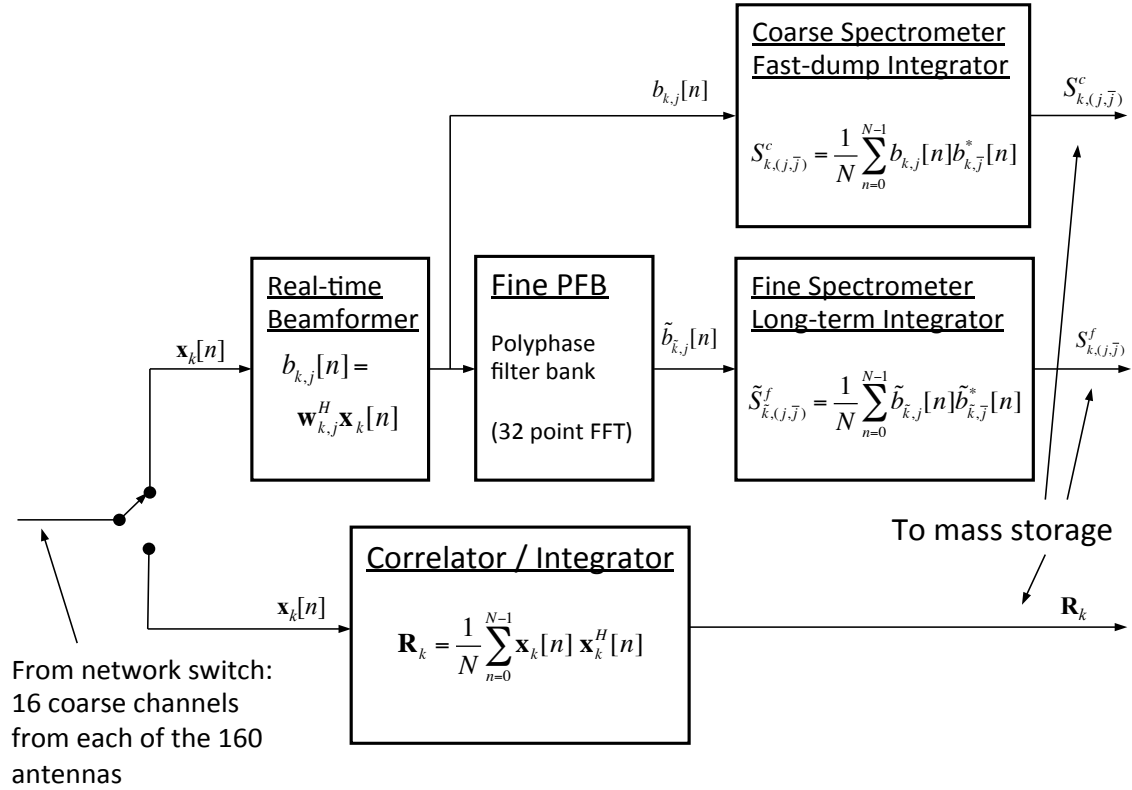


Figure 6.2: XB Engine Correlator / Beamformer Functional Blocks per GPU

Notation for Figure 6.2 is as follows:

1. Bold font denotes vectors or matrices, non-bolds are scalar values.
2. $\mathbf{x}_k[n]$ is the 160 element-long (one per PAF antenna) vector of basebanded array samples for the k th frequency channel at decimated time sample n .
3. $\mathbf{w}_{k,j}$ is the vector of weights to form the j th beam in channel k , and $b_{k,j}[n]$ is the corresponding beamformer output sample.
4. Tilde overbars denote quantities from fine filterbank processing.
5. \mathbf{R}_k is the integrated array covariance matrix estimate for the k th channel.
6. $S_{k,(j,\bar{j})}$ denotes the k th frequency channel integrated cross spectrometer output for beams j and \bar{j} . If $j = \bar{j}$ this is a conventional spectrum; if not, it is a cross spectrum, i.e. for cross polarization observations between the beams pointed in the same direction, but with orthogonal polarizations.

6.3 Digital Receiver Design Parameters and Specifications

Table 6.1 lists the major operational parameters and design specifications which determine the digital receiver system architecture, and firmware and software development requirements. Figure 6.3 defines the internet data packet structure for communications between the F-engine and XB-engine. Each packet contains data from 8 antenna elements, 32 channels, 16 time samples, and one 64-bit header

word. For a single ROACH 2 processor, 25 successive packets are transmitted to cover all 800 coarse channels (32 per packet) from the 8 antenna inputs sampled by that ROACH. In other words, each ROACH produces data for 8 antennas, and all 800 channels. The five network switches reroute (corner turn) the packets so that each of the 50 GPUs receives data for just 16 coarse channels, but from all 160 antennas. The packet structure of Table 6.1 insures proper routing and identification of these data.

Property Description:	Value:
No. of inputs (antennas)	$M = 160$
No. of ADC inputs per Roach	$M_{ADC} = 8$
No. of dual-pol beams	$J = 40$
Sample frequency (real samples)	$f_s = 800 \text{ MHz}$
Bits per ADC sample	$b_{ADC} = 8$
F-engine Coarse PFB complex FFT size	$N_c = 2048$
Total F-engine output bandwidth	$B_{total} = 312.5 \text{ MHz}$
No. of coarse channels passed to X-engine	$K_c = 800$
Coarse frequency channel bandwidth	$B_c = 390.625 \text{ kHz}$
Fine PFB complex FFT size	$N_f = 32$
Total no. of fine PFB channels	$K_f = 25,600$
Fine frequency channel BW	$B_c = 12.207 \text{ kHz}$
F engine output word size, Int. Re + Im	$b_F = 16$
X-engine output word size, complex float	$b_X = 64$
Beamformed spectrometer word size, float	$b_s = 32$
Min. coarse channel integration dump time	$T_c = 64 \mu\text{s}$
Min. fine channel integration dump time	$T_f = 100 \text{ ms}$
Min. correlator integration dump time	$T_X = 10 \text{ ms}$
No. of ROACH 2s	$N_R = 20$
No. of CPUs	$N_{CPU} = 25$
No. of GPUs	$N_{GPU} = 50$

Table 6.1. Design specifications for the digital receiver

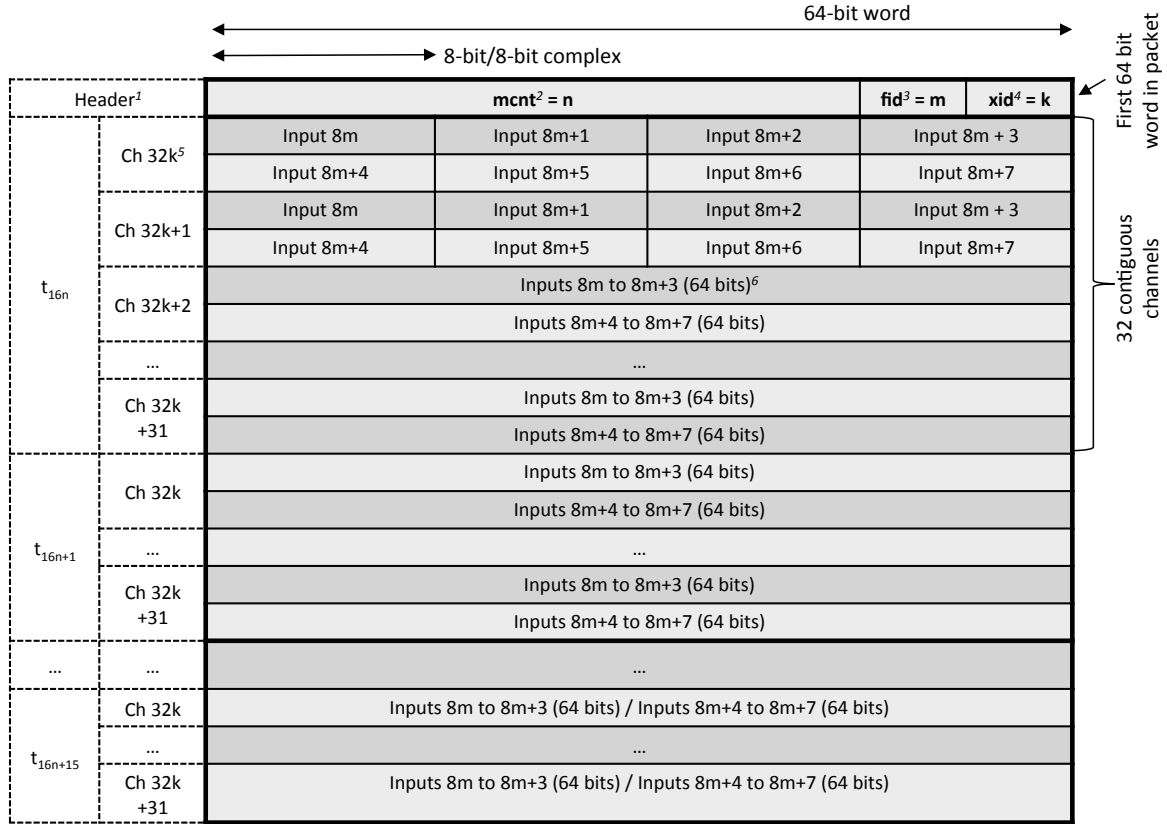


Figure 6.3: ROACH to HPC packet structure for UDP Jumbo packets.

Footnotes:

- ¹ Left-most 2 columns are labels that are not part of the actual packet payload.
- ² The **mcnt** field indicates the packet number for the current observation. This should be armed to 0 by the 1pps signal at the start of every scan. A “scan,” in this context, can consist of multiple pointings such as in the case of a PAF calibration grid. This will effectively align the mcnt field across every F-engine. In this example, mcnt = n.
- ³ The **fid** is an 8-bit value identifying from which ROACH-II board the packet originates.
- ⁴ The **xid** is an 8-bit value identifying to which GPU the packet is destined.
- ⁵ The channel indices are between 0 and 799. Inputs are ordered as shown in t_{16n}, Ch. 32k. Each packet consists of the following: 8 Inputs/Antenna Elements, 32 Frequency Channels, 16 Time Samples, and one 64-bit header word. This yields a packet size of 1025 64-bit words. The maximum UDP jumbo packet payload supported is 1088 64-bit words.

6.4 Feasibility Performance Analysis for GPU Processing

The largest processing task in the GPUs, and a potential limiting factor, is the real-time correlator. This function is performed in the AO40 digital receiver using the CASPER library “XGPU” correlator code for Nvidia GPUs. We performed real GPU processing tests (not simulations) over a range of configuration parameters to see if the proposed number of GPUs was adequate to support the most computationally demanding operational mode. Our tests evaluated whether the actual GPU processing, and the hashpipe routing code for multi-threaded high data rate

processing, shared memory reads and writes, could keep up with the designed input data rate. In these experiments there were no other significant processes running. Also, rather than using real 10 GbE network input of real F-engine data packets, the target input data buffer in shared memory was “looped” to repeatedly cycle through the same pre-stored data. These two provisions insured that the test evaluated only limitations of the GPU processing and hashpipe data routing functions, without being constrained by other peripheral data bottlenecks. These XGPU and hashpipe functions are known to be the biggest demand on processing resources.

Test results are presented in Figure 6.4. The "AO40 Input Data Rate" dotted curve (see legend) is provided as a reference point, and corresponds to the actual AO40 designed input data rate, per HPC, which would get split up among hashpipe instances of XGPU code. Curves above this dotted line represent scenarios where the GPU is able to keep up with that designed input data rate. The two “2 GPU” curves correspond to the proposed XB engine architecture with 2 GPUs per HPC chassis. The horizontal axis shows performance variation as a function of number of input antenna signals processed (160 antennas for the proposed AO40 design.) The GPUs processed the full allotment of 160 frequency channels per HPC for the proposed design.

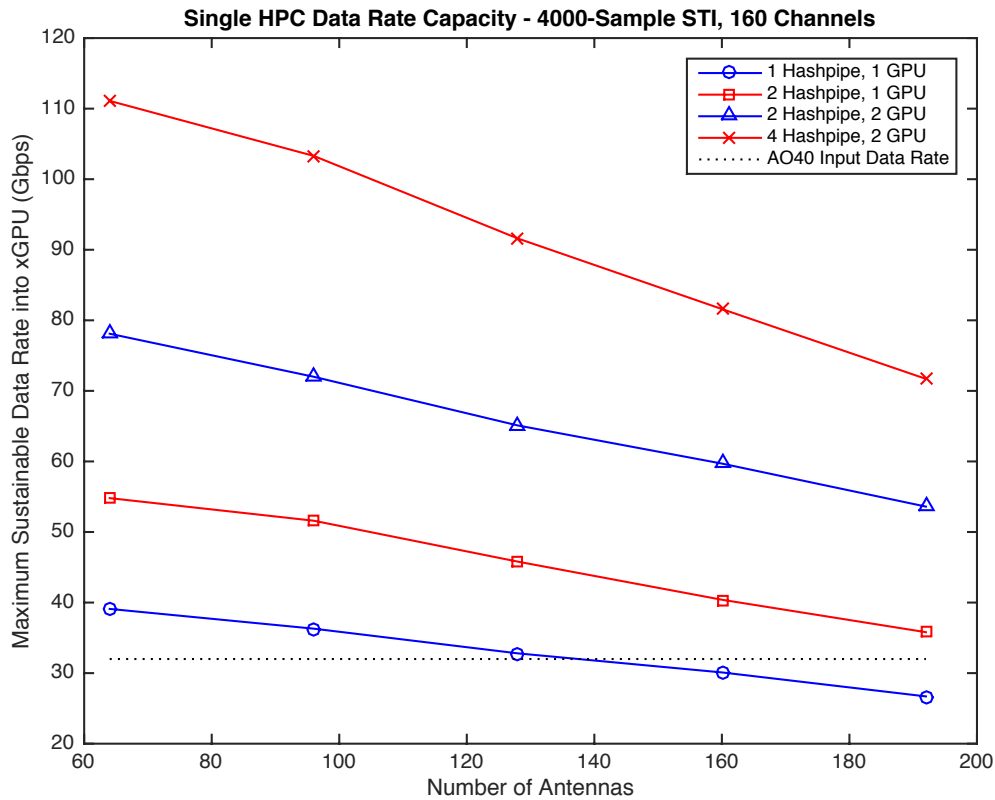


Figure 6.4: GPU throughput experiment for correlation processing.

These test results support several important conclusions: a) Better performance is obtained with two instances of the XGPU correlator per GPU (i.e. four per HPC), rather than one instance per GPU. b) With $M = 160$ antenna inputs (the AO40 specification), four Hashpipe implementations of XGPU, and two GPUs per HPC, this stripped down, XGPU-only system can support 2.5 times the design data input rate of 32 Gbps per HPC. We are comfortable that this leaves more than adequate “headroom” to accommodate all other GPU and HPC haspipe data transfer functions.

7. Schedule

Fig. 7.1 shows a possible schedule for the construction, testing and integration of AO40. This schedule is based on Brigham Young University’s and Cornell University’s time estimates assuming that they were responsible for the project.




Figure 7.1: Suggested schedule for the design, construction, testing and integration of AO40.

8. Cost estimate

An extensive costing exercise was carried out covering all aspects of the AO40 system based on the design presented in this report. The equipment costs are based on quotations where possible, catalog prices, experience based on the AO19 system, and, if necessary, best estimates. They are in 2015 dollars and can be regarded as a reasonably accurate estimate of the equipment and ancillary costs for construction of the AO40 system outlined in this report. In general, no contingency has been applied. Ancillary costs such as software licenses, shipping and travel are based on the assumption that the development of the cryogenic system and dipole assemblies would take place at Cornell University and the remainder of the tasks including final integration and testing at Brigham Young University.

An attempt was also made to estimate the management and labor costs including the requirements and costs of the supporting services by Arecibo Observatory staff. Again, these costs are based on the assumption of a Cornell/BYU collaboration in the production of the system. They may change significantly if other institutions develop the system.

8.1 RF sections – Dipole assemblies, cryostat & post cryostat rf system (Cornell University estimate)

Capital equipment:

Dipole assemblies (includes Dipoles, LNAs and noise injection system):

Dipoles, cylinders & press fit connectors:	350,000
Dipole pattern tests:	15,000
Coaxial and strip lines:	150,000
Low noise amplifiers:	150,000
Noise injection system:	32,000
Small dewar for LNA tests:	20,000
Sub-Total:	\$717,000

Cryogenic system (w/o dipole assemblies and LNAs):

Cryostat:	320,000
Refrigerators and vacuum equipment:	122,600
Monitor and control system with rack:	55,700
Mount in dome, cart, etc:	40,000
Sub-total:	\$538,300

Post cryostat rf system in Dome and operations building:

RF system in the Dome:	84,200
Fiber transmitters & receivers plus cable:	133,600
IF/LO system in ops building:	124,800
Local Osc. Source & distribution:	12,000
Monitor & Control:	20,000
Prototyping costs:	8,800
Sub-total:	\$383,400

Direct costs estimated on the basis of Cornell University's costs:

Personnel plus benefits:	722,700
Travel:	38,000
Software licenses:	37,500
Computer services:	10,000
Materials, supplies, publications, etc:	53,000
Shipping:	12,000
Sub-total:	\$873,200

Indirect costs:

Sub-total:	\$502,400
------------	-----------

Total RF sections costs: \$3,014,300

8.2 Digital Receiver and Data Storage (BYU budget estimates)

Capital equipment:

F-engine, incl. ROACH 2s, clock distribution:	331,418
XB-engine, incl. HPCs and GPUs :	220,250
Corner turn network, 10 GbE switches:	36,380
Shielded racks and UPS power supplies:	33,588
Lustre file store system:	124,550
Miscellaneous:	50,000
(note: does not include analog receiver, down conversion, optical fiber link, amplification, filtering)	
Sub-total:	\$796,186

Non-Capital cost estimates based on BYU constructing the Digital Receiver

Direct Costs:

Personnel and benefits:	1,028,214
Travel, supplies, pubs, tuition:	119,533
Sub-total:	\$1,147,747

Indirect Costs:

Sub-total:	\$570,858
------------	-----------

Total Digital Receiver Estimated Costs: \$2,514,791

8.3 USRA/Arecibo Observatory installation and integration costs:

Direct costs:

Personnel costs:	108,300
Materials & supplies:	3,200
Sub-total:	\$111,500

Indirect costs:

Sub-total:	\$48,600
------------	----------

Total USRA costs: \$160,100

8.4 Total estimated cost: \$5,689,191

Appendix A

The following pages contain vendor price quotes for much of the major equipment and services used in the A040 Digital Receiver back end design. These are the basis for the project's component and subsystem hardware price estimates.

ROACH 2 FPGA assemblies with ADC and 10 GbE nics, less Xilinx FPGA chips, for F-engine. 21 assemblies required, including 1 spare:



EMS Solutions
ISO-9001:2008 ISO 13485:2003

7799 Pardee Lane
Oakland, Ca. 94621
Phone: 510-639-7003
Fax : 510-639-7090

August 7, 2015
QUOTE # 3962

TO : Brian D. Jeffs
Professor
Department of Electrical and Computer Engineering
459 CB,
Brigham Young University
Provo, UT 84602
(801) 422-3062
bjeffs@byu.edu

We are pleased to submit the following quote regarding your inquiry.

Quantity 21, Roach2 board assembly, unit price \$3,850.00, Ext amount; \$80,850.0
Quantity 21, DRAM MODULES, no charge
Quantity 21, Xilinx chip P/N XC6V SX475T model, unit cost \$8,376.0, ext. amount; \$175,896.0
Quantity 42, ADC16x250-8 coax boards, unit cost \$1,300.0, ext. amount; \$54,600.0
Quantity 42, SFP+ Card Mezzanine, unit cost \$450.0, Ext. amount; \$18,900.0
Quantity 21, 1U Enclosures, unit cost \$550.00, Ext. amount; \$11,550.0
Complete chassis includes, power supply, fans, led board, power switch, wired to configuration.

A volume discount amount equal to 0.075 (7.5%) of total value will apply to this order.


Shipping, freight; Not included above
TERMS : Net 30 Days
EXPECTED DELIVERY: 18 Weeks ARO (Delivery time is estimate)

If there are any question, please let us know.
Thank you for your inquiry,
Regards,
M.OHADY
MFG. MGR.
Email: mo@digicom.org

www.digicom.org

Xilinx Vertex 6 FPGA chip, single unit price (21 required for the ROACH 2s, incl. 1 spare), for F-engine:


Part Details
Technical Information



Images are for reference only. See product specifications.

Manufacturer Part #: XC6VSX475T-2FFG1759C
Xilinx
 FPGA Virtex®-6 SXT Family 476160 Cells 40nm (CMOS) Technology
 1V 1759-Pin FCBGA

Avnet Part #: XC6VSX475T-2FFG1759C
ECCN: 3A001.A.7.A
Schedule B: 8542.39.0000
HTSN: 8542.39.0000
Version: v13.0601



Download Datasheet

Compare

Region	Qty Available/ Ships In	Packaging	Price (USD)	Qty	
Americas	Part # : XC6VSX475T-2FFG1759C MFR Xilinx Stock: No Stock Average Factory Lead Time: 13 Week Factory Lead Time	Trays	1-\$10,366.5304 2+-\$10,315.3320	<input type="text" value="1"/> Min: 1 Mult: 1	<p>Add to Cart</p> <p>Add to BOM</p>

Part Details


Description	Value
Package	1759FCBGA1759FCBGA
Family Name	Virtex-6 SXT/Virtex-6 SXT
Device Logic Units	476160476160
Typical Operating Supply Voltage	1 V1 V
Maximum Number of User I/Os	840840
RAM Bits	3922329639223296
Operating Temperature	0 to 85 °C0 to 85 °C
Speed Grade	22
Category	FPGA/FPGA
Manufacturer	Xilinx/Xilinx

Find Similar Parts

-
-
-
-
-
-
-
-
-
-

Search

Mellanox 40 GbE / 10 Gbe 40 port network switch. Interconnect between F-engine and XB-engine. 5 required:



Order By Item # | Quantity Discounts


SHOP ALL STORES | Search All

Home > Networking > Switches & Hubs > Switches > Mellanox Technologies > Item#: 9B12K-00NZ-00023


Mellanox SwitchX SX1012 Layer 3 Switch


Item #: 9B-12K-00NZ-00023 | Mfr. Part #: MSX1012B-2BRS | UNSPSC: 43222612

Be the first to review this product...







\$6,035.⁹⁹ ~~\$7,899.99~~
Save: \$1,864.00 (24%)

Earn BizPoints as a  member. [Learn how.](#)

Free Shipping Available 

In stock.

Share:    

Sold and Shipped by: Newegg

Overview | Details

Main Features
Manufacturer/Supplier: Mellanox Technologies Ltd
Manufacturer Part Number: MSX1012B-2BRS
Manufacturer Website Address:
Brand Name: Mellanox
Product Line: SwitchX
Product Model: SX1012
Product Name: SwitchX SX1012 Layer 3 Switch
Marketing Information: The SX1012 switch system provides the highest-performing top-of-rack (ToR) solution in a half-width 1U form factor by delivering up to 1.3Tb/s of non-blocking throughput with 10GbE or 40/56GbE form factor. The SX1012 is the ideal top-of-rack (ToR) switch for small pods connecting a small number of servers with 40GbE. Each QSFP port can be used with a passive breakout cable to create four discrete SFP+ interfaces to be used as 10GbE ports, for a total of 48-ports on the smallest form factor 10GbE ToR in the industry.
Product Type: Layer 3 Switch
Port/Expansion Slot Details: 12 x 40 Gigabit Ethernet Expansion Slot
Ethernet Technology: 40 Gigabit Ethernet
Number of Total Expansion Slots: 12
Expansion Slot Type: QSFP
Layer Supported: 3
Manageable: Yes
PoE (RJ-45) Port: No
Input Voltage: 110 V AC, 220 V AC
Form Factor: Rack-mountable
Compatible Rack Unit: 1U

HPCs for XB Engine. 25 required:



QUOTE

Date	Quote #
07/02/15	AH5Q3278

Quote To:
 Brigham Young University
 Brian Jeffs

Phone: (801)422-3062

Ryan Emmerich
Account Representative
Phone: (858) 268-9600
Fax: (858) 268-9605
E-mail: ryan.emmerich@advancedhpc.com

Terms	FOB
NET 30	ORIGIN

Thank you for the opportunity to provide the following proposal.

Qty	Description	Part Number	Unit Price	Ext. Price
25	<p>Mercury GPU208 2U Server to Include:</p> <p>Mercury GPU208 2U Server Dual E5-2620 v3 2.40 GHz Six-Core 85W Processors 64GB DDR4 ECC Registered Memory Two 2TB SATA 3.5" 7.2K RPM ENT Disk Drives Eight 3.5" Hot-Swap HDD Bays One EVGA GeForce GTX980TI PCIE 6GB Graphics Card, (CPU1) One Mellanox ConnectX-3 EN Network Interface Card, 40GigE, Single-Port QSFP, (CPU1) One EVGA GeForce GTX980TI PCIE 6GB Graphics Card, (CPU2) One Mellanox ConnectX-3 EN Network Interface Card, 40GigE, Single-Port QSFP, (CPU2) Two Intel I210AT + One Mgmt LAN Aspeed AST2400 with 32MB VRAM Graphics Quick/Quick Rail Kit 1620W 80PLUS Platinum Certified 1+1 Redundant Power Supply Two C13 to C14 Power Cable-3FT. Pre-configured and Stress Tested Before Leaving AHPC Facility Three-Year Standard Warranty on Parts and Labor. Return-to-Depot Parts Replacement. Toll-Free Phone Support Help Desk Available Mon-Fri 8am - 5pm PST Excludes Battery Backup Modules and Solid State Drives which are Covered by the Manufacturer's Warranty. Lifetime Technical Support</p> <p>Optional Support Warranty</p>	AH-GPU208-AX02	\$6,300.00	\$157,500.00
25	<p>Three-Year Advance Replacement Warranty to Include:</p> <p>Toll-Free Phone Support Help Desk Available Mon-Fri 9am - 5pm PST 48-Hour Advance Replacement of All User Replaceable Parts: Disk Drives, Cooling Fan Modules, Power Supply Modules, System Memory and Software. Free System Firmware Updates.</p>	AH-SRVC/ARWB3YR	\$350.00	\$8,750.00
Total				\$166,250.00



Advanced HPC, Inc.
Corporate Headquarters
4887 Ronson Court, Suite A
San Diego CA 92111-1812

Phone: (858) 268-9600
Fax: (858) 268-9605
URL: www.advancedhpc.com

Nvidia Graphical Processor Unit (GPU). 2 per HPC required, 50 units total:



**EVGA 06G-P4-4995-KR
GeForce GTX 980 Ti 6GB 384-
Bit GDDR5 PCI Express 3.0
SLI Support SC+ w/ACX BP
Video Card**

**Free Metal Gear Solid V: The Phantom
Pain game w/ purchase, limited offer**

★★★★ (78) | [Write a Review](#)

In stock. Limit 5 per customer.

In stock. Limit 5 per customer.

- 6GB 384-Bit GDDR5
- Core Clock 1102MHz
- Boost Clock 1190MHz
- 1 x DVI-I DVI 1 x HDMI HDMI 3 x DisplayPort DisplayPort
- 2816 CUDA Cores
- PCI Express 3.0

[Ask Or Answer A Question](#)

See 16 questions | 42 answers

Special Offers



**FREE GIFT: NVIDIA Bullets or Blades
Bundle - Choose Tom Clancy's Rainbow
Six Siege or Assassin's Creed
Syndicate**

Value:\$0.99

[Tweet](#) 1 [Pin it](#) [G+](#) 2
[Tweet](#) 1 [Email](#)

QTY: 1 **PREMIER** **\$679.99**

\$4.99 Shipping (restrictions apply)

Sold and Shipped by:
Newegg

Option: GTX980 Ti SC+ ACX2.0+ w/BP

[GTX980 Ti CLASSIFIED ACX 2.0+](#)

[GTX980 Ti Hybrid](#) | [GTX980 Ti SC](#)

[GTX980 Ti SC+ ACX2.0+ w/BP](#)

Option: GTX980 Ti SC+ ACX2.0+ w/BP

[GTX980 Ti CLASSIFIED ACX 2.0+](#)

[GTX980 Ti Hybrid](#) | [GTX980 Ti SC](#)

[GTX980 Ti SC+ ACX2.0+ w/BP](#)

Protect Your Investment!

2 Year Extended Repair Coverage **\$59.00**
(more options)

Send It As A Gift!

Make this item a gift

ADD TO CART ▶

[Add to Wish List](#)

[Price Alert](#)

SHOPRUNNER free 2-day shipping
[learn more](#) | [sign in](#)

Receive a USD \$20 prepaid card by mail from
EVGA! Offer expires on 10/31/15. Limit 2 per
household [Print Rebate Form](#)

Lustre high speed file storage system. This capture XB-engine output data (2 pages):



QUOTE

Date	Quote #
08/10/15	AH5Q3439

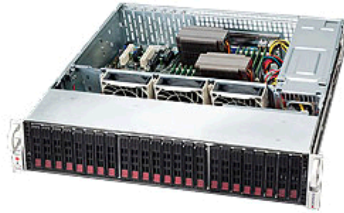

Quote To:
 Brigham Young University
 Brian Jeffs

Phone: (801)422-3062

Ryan Emmerich
Account Representative
Phone: (858) 268-9600
Fax: (858) 268-9605
E-mail: ryan.emmerich@advancedhpc.com


Terms	FOB
NET 30	ORIGIN

Thank you for the opportunity to provide the following proposal.

Qty	Description	Part Number	Unit Price	Ext. Price
Lustre OSS/OST Servers				
2	Lustre OSS/OST Servers to Include:		\$62,000.00	\$124,000.00
	Mercury RM224 2U Rackmount Server Dual E5-2680 v3 2.50 GHz Twelve-Core 120W Processors Intel C612 Chipset 128GB DDR4 2133 ECC Registered Memory Two Intel 240GB DC S3510 Series SATA 6Gb/s 2.5 Inch MLC Solid State Drives for Mirrored OS, (RAID 1) Twenty-four 1.6TB DC S3510 Series SATA 6Gb/s 2.5 Inch MLC Solid State Drives for Data, (Using Four Drive RAID 0 Sets) Twenty-Four 2.5 Inch Hot-Swap SAS/SATA HDD/SSD Bays and Rear 2.5 Inch Hot-Swap SAS/SATA HDD/SSD Bays for Mirrored OS SAS3 12Gb/s Direct Attached Storage HDD Backplane Integrated Ten SATA3 (6Gbps) with RAID 0, 1 and 10 Support Three LSI MegaRAID 9361-8i SAS 8-Port SATA/SAS 12Gb/s RAID Controllers with CacheVault One Mellanox ConnectX-3 VPI Adapter Card, Dual-Port QSFP FDR IB (56Gb/s) and 40GigE with Three Year Bronze Support One Mellanox ConnectX-3 Pro Dual Port 40Gigabit Ethernet QSFP+ Network Adapter Card Intel i350 Dual Port GbE LAN Integrated IPMI 2.0 and KVM with Dedicated LAN Quick/Quick Rail Kit 920W Redundant High-Efficiency Platinum Level Power Supplies RHEL OS Licenses One Intel IEEL for Lustre License with Three-Year L3 Support Pre-configured using RHEL and IEEL and Stress Tested Before Leaving AHPC Production Facility Three-Year Advance Replacement Warranty to Include: Toll-Free Phone Support Help Desk Available Mon-Fri 9am - 5pm PST 48-Hour Advance Replacement of All User Replaceable Parts: Disk Drives, Cooling Fan Modules, Power Supply Modules, System Memory and Software Free System Firmware Updates Lifetime Technical Support	AH-RM224-SB10		
				
				

Advanced HPC, Inc.
Corporate Headquarters
4887 Ronson Court, Suite A
San Diego CA 92111-1812

Phone: (858) 268-9600
Fax: (858) 268-9605
URL: www.advancedhpc.com

Qty	Description	Part Number	Unit Price	Ext. Price
Lustre MDS/MDT Server				
1	Lustre MDS/MDT Server to Include: Mercury RM224 2U Rackmountable Server Dual E5-2680 v3 2.50 GHz Twelve-Core 120W Processors Intel C612 Chipset 256GB DDR4 2133 ECC Registered CL15 Memory Two Intel 300GB DC S3500 Series SATA 6Gb/s 2.5" MLC Solid State Drives for Mirrored OS, (RAID 1) Twenty-four 300GB SAS 6Gb/s 2.5 Inch 10K RPM Disk Drives, (RAID 10) Twenty-four 2.5" On Front and Two 2.5" in Rear, (for OS), Hot-Swap SAS/SATA HDD/SSD Bays SAS3 12Gb/s Expander Storage HDD Backplane One LSI MegaRAID 9361-8i SATA/SAS 12Gb/s RAID Controller with CacheVault One Mellanox ConnectX-3 VPI Adapter Card, Dual-Port QSFP FDR IB (56Gb/s) and 40GigE with Three Year Bronze Support One Mellanox ConnectX-3 Pro Dual Port 40Gigabit Ethernet QSFP+ Network Adapter Card Intel i350 Dual Port GbE LAN Integrated IPMI 2.0 and KVM with Dedicated LAN Quick/Quick Rail Kit 920W Redundant High-Efficiency Platinum Level Power Supplies Pre-configured and Stress Tested Before Leaving AHPC Production Facility Three-Year Advance Replacement Warranty to Include: Toll-Free Phone Support Help Desk Available Mon-Fri 9am - 5pm PST 48-Hour Advance Replacement of All User Replaceable Parts: Disk Drives, Cooling Fan Modules, Power Supply Modules, System Memory and Software Free System Firmware Updates Lifetime Technical Support	AH-RM224-SB14	\$19,500.00	\$19,500.00
				
Luster Cluster FDR IB Network				
1	SwitchX@-2 Based 12-Port QSFP FDR 1U Managed InfiniBand Switch System with a Non-Blocking Switching Capacity of 1.3Tb/s, 1PS, Short Depth, Forward Airflow, RoHS-6	AH-MSX6012F-1BFS	\$5,500.00	\$5,500.00
1	Mellanox Two-Year Extended Warranty for a Total of Three Years Bronze for SX6005 and 6012 Series Switch	AH-EXW-SX600B-3B	\$350.00	\$350.00
Total				\$149,350.00

The total price listed includes system integration, support package software pre-installation, quality testing, and burn-in. (if applicable). Applicable taxes and shipping may be added to your order total. Shipping charges including insurance are prepaid and added unless otherwise specified in writing.

Net terms are subject to credit verification and approval.

This quote is valid for 30 days unless otherwise specified or modified in writing.

Advanced HPC, Inc.
Corporate Headquarters
4887 Ronson Court, Suite A
San Diego CA 92111-1812

Phone: (858) 268-9600
Fax: (858) 268-9605
URL: www.advancedhpc.com

Page 2

Uninterruptable power supply system (UPS):



QUOTE

Date	Quote #
07/16/15	AH5Q3330

Quote To:
 Brigham Young University
 Brian Jeffs

Phone: (801)422-3062

Ryan Emmerich
Account Representative
Phone: (858) 268-9600
Fax: (858) 268-9605
E-mail: ryan.emmerich@advancedhpc.com

Terms	FOB
NET 30	ORIGIN

Thank you for the opportunity to provide the following proposal.

Qty	Description	Part Number	Unit Price	Ext. Price
2	BladeUPS 208V Parallel 12kW with PowerXpert Gateway Card	AH-BLADEUPS/ZC121P0681	\$7,300.00	\$14,600.00
Total				\$14,600.00

The total price listed includes system integration, support package software pre-installation, quality testing, and burn-in. (if applicable).
 Applicable taxes and shipping may be added to your order total. Shipping charges including insurance are prepaid and added unless otherwise specified in writing.
 Net terms are subject to credit verification and approval.
 This quote is valid for 30 days unless otherwise specified or modified in writing.

Advanced HPC, Inc.
Corporate Headquarters
4887 Ronson Court, Suite A
San Diego CA 92111-1812

Phone: (858) 268-9600
Fax: (858) 268-9605
URL: www.advancedhpc.com