

Commensal Systems at LWA-SV

Ver. 1

J. Dowell* G. B. Taylor†

September 1, 2020

Contents

1	Introduction	2
2	Network Topology	2
3	Capacity for Future Expansion	3
4	Document History	5
A	Fiber Connection from LWA-SV to the UNM SERF Building	7

*University of New Mexico. E-mail: jdowell@unm.edu

†University of New Mexico. E-mail: gbtaylor@unm.edu

1 Introduction

The purpose of this document is to expand upon the commensal system capabilities described in the LWA-SV Station Architecture memo (LWA memo 214 [2]) and provide reference for future commensal systems at the station.

The Advanced Digital Processor subsystem (ADP; [1, 2]) is the digital backend for LWA-SV. ADP supports two styles of commensal systems, one that needs access to the raw F-Engine data and one that need access to the final beamformed (DRX) and transient buffer–narrowband (TBN) data products (Table 1). The F-Engine data currently consists of 12, 3.3 MHz sub-bands. Each sub-band contains dual polarization complex spectra from all 256 antennas in the array. The sub-bands are organized into two tunings, each containing six sub-bands. These two tunings correspond to the two DRX tunings for beamformed data. Within each of the tunings the sub-bands are contiguous in frequency but the two tunings are not necessarily contiguous.

Name	Description	Bandwidth	Data Rate	Address	Port(s)
F-Engine – Complex Spectra					
Sub-band 1	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.11	4015/4016
Sub-band 2	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.12	4015/4016
Sub-band 3	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.13	4015/4016
Sub-band 4	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.14	4015/4016
Sub-band 5	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.15	4015/4016
Sub-band 6	256 stands, dual pol.	2x 3.3 MHz	25 Gb/s	239.168.40.16	4015/4016
Beamformed – Complex Voltages					
Beam 1	2 tunings, dual pol.	19.6 MHz	600 Mb/s	239.168.40.41	10000
Beam 2	2 tunings, dual pol.	19.6 MHz	600 Mb/s	239.168.40.42	10000
Transient Buffer – Complex Voltages					
TBN	256 stands, dual pol.	100 kHz	800 Mb/s	239.168.40.44	10000

Table 1: Available commensal data sources at LWA-SV.

Both of these modes are support through IP multicast¹, a method of sending one data packet from a source to multiple receivers with a single transmission. In particular, ADP uses UDP for the transport layer on top of IP multicast, which is referred to as UDP multicast in LWA documentation. Since the multicast client subscriptions are handled by the switches at the station, additional commensal systems are largely transparent to the station. In principle, all that is needed to add a new commensal system is to configure it to subscribe to the correct multicast address and the switches will take care of sending the data.

2 Network Topology

Commensal systems need access to access to the ADP data network. The topology of the network is shown in Figure 1. Briefly, this network consists of ADP (both the 16 ROACH2 boards and the seven GPU servers), two Mellanox SX1024 switches², and a Dell PowerConnect 8024F switch³. All of this network, with the exception of the second SX1024 switch, is located inside the LWA-SV shelter. The second SX1024 is located in the server room of the UNM Sevilleta Education and Research Facility (SERF) building located on the refuge approximately 500 m north of LWA-SV.

¹See RFC 112 – <https://tools.ietf.org/html/rfc112>, RFC 4604 – <https://tools.ietf.org/html/rfc4604>, and RFC 5771 – <https://tools.ietf.org/html/rfc5771>.

²https://www.mellanox.com/related-docs/prod_eth_switches/PB_SX1024.pdf

³<https://www.dell.com/downloads/global/products/pwcnt/en/switch-powerconnect-8024f-spec.pdf>

The shelter is connected to the SERF building through a ~ 600 m run of multi-strand single mode fiber cable. See A for details of the fiber path.




Figure 1: Network topology of the ADP data network. The arrows indicate the predominate direction of data flow and are labeled by the link speed. Blocks that are blue are located in the server room of the UNM SERF building.

As of September 1, 2020 there are currently seven machines connected to this network: the four data recorder machines inside the LWA-SV shelter and three commensal systems located in the SERF building. The commensal machines are the LWA All-Sky Imager (LASI; [3]), the Wilbur system (which runs on the same machine as Orville; [5, 4]), and the E-Parallel Imaging Correlator machine (EPIC; [6]). The first two are connected to the SX1024 switch in the SERF building and EPIC uses a dedicated fiber pair for its 40G link⁴. The LASI and Wilbur machines receive a combination of DRX and TBN data from ADP and the maximum expected data transit of the 10G link to the SERF building is ≈ 1.4 Gb/s (800 MB/s for TBN at a sample rate of 100 kHz; 600 MB/s for DRX at a sample rate of 19.6 MHz). The EPIC machine receives raw F-Engine data from the ROACH2 boards and has a maximum expected data rate of 25 Gb/s (12.5 GB/s per 3.3 MHz sub-band).

3 Capacity for Future Expansion

As mentioned in LWA Memo 214 [2] the ability to add additional commensal systems is dependent on the available data switch bandwidth, number of free ports on the data switch, and the added power and heat loads. With the current hardware in place it is possible to have:

- Additional machines in the shelter that use either the DRX or TBN data products. The 8024F switch has ≈ 15 SFP+ ports free. The largest constrain on installing machines here is the heat load and it is not likely that additional machines would be allowed.
- Additional machines in the server room that use either the DRX or TBN data products. The SX1024 switch has ≈ 40 SFP+ ports free and the maximum theoretical multicast bandwidth on the 10G link to the shelter is 3.1 Gb/s, assuming TBN and the eventual expansion of ADP to four beams. The largest constraints on installing machines here is the available rack space

⁴There is a problem with the SX1024 switch located in the SERF building that precludes it from using the 40G optical transceivers needed to link the two SX1024's at 40G.

and free IP addresses on the publicly routable network. LWA currently has a single rack in this room and it may be able to hold an addition one or two 4U machines. Similarly, the number of publicly routable address available to LWA is two.

- An additional machine in the shelter that uses the raw F-Engine data. There is a single 40G port free on the SX1024 switch in the shelter but the additional heat load would likely rule this location out.

It is not currently possible to support additional machines that utilize the raw F-Engine data in the server room since the current 40G link only connects to the EPIC machine. This limitation could be removed with the purchase of new switch that could support the 40G optical transceivers needed to establish the link with the shelter. However, the 40G link only has the capacity to handle a single sub-band pair and the EPIC and additional machine(s) would be required to use the same sub-band at all times. This could also be remedied with two new switches (one in the shelter and another in the server) and, possibly, repairing the damaged fibers (See Appendix A). The two switches could be linked via multiple 40G links combined into link aggregation groups (LAGs⁵) to increase the bandwidth available. A LAG with two links should support two commensal machines that each connect to a single sub-band independently. Increasing the number of links in the LAG to four should allow for all of the F-Engine data to be transmitted to the server room.

⁵https://en.wikipedia.org/wiki/Link_aggregation

4 Document History

- Version 1 (Sep 1, 2020): First version.

References

- [1] M. Cranmer et al., “Bifrost: a Python/C++ Framework for High-Throughput Stream Processing in Astronomy”, 2017, JAI, 650007.
- [2] J. Dowell & G. B. Taylor, “LWA-SV Station Architecture”, Ver. 1, Long Wavelength Array Memo 215, Aug 28, 2020. [online] <http://www.phys.unm.edu/~lwa/memos/index.html>.
- [3] K. Obenberger, et al., “Monitoring the Sky with the Prototype All-Sky Imager on the LWA1”, 2015, JAI, 4, 450004.
- [4] J. Dowell, S. Varghese, & G. B. Taylor “The Orville Wideband Imager,” Ver. 1, Long Wavelength Array Memo 215, Aug 28, 2020. [online] <http://www.phys.unm.edu/~lwa/memos/index.html>.
- [5] J. Dowell & G. B. Taylor, “The Wilbur Single Dispersed Pulse Search System,” Ver. 1, Long Wavelength Array Memo 216, Aug 28, 2020. [online] <http://www.phys.unm.edu/~lwa/memos/index.html>.
- [6] J. Kent et al., “A Real-Time, All-Sky, High Time Resolution, Direct Imager for the Long Wavelength Array”, 2019, MNRAS, 486, 5052

A Fiber Connection from LWA-SV to the UNM SERF Building

The fiber connections from the LWA-SV shelter to the SERF building server room are shown in Figure 2. All of the fiber shown here is single mode fiber and it is lit using LR type transceivers. Although the path has the potential to have five pairs in use between the shelter and the SERF building, several of the fibers were damaged during installation, particularly for the ~600 m run between the two outdoor termination boxes. This damage was due to trying to pull fiber cable with the FC connectors already installed. The current state of the fiber as of September 1, 2020 is:

- All five pairs are operable from the shelter to the termination box nearest it. Three pairs are currently in use.
- Four of the six pairs are operable between the two termination boxes. Three pairs are currently in use.
- Four of the six pairs are operable between the SERF telecom room and the termination box nearest it. Three pairs are currently in use.
- All three pairs from the telecom room to the server room are operable. Two pairs are currently in use.

It should be noted that it is likely possible to return part of the broken fibers to a working condition by re-connectorizing them. This, however, requires FC fiber pigtailed, a fiber cutter, a splicing machine, and expertise to perform the work.

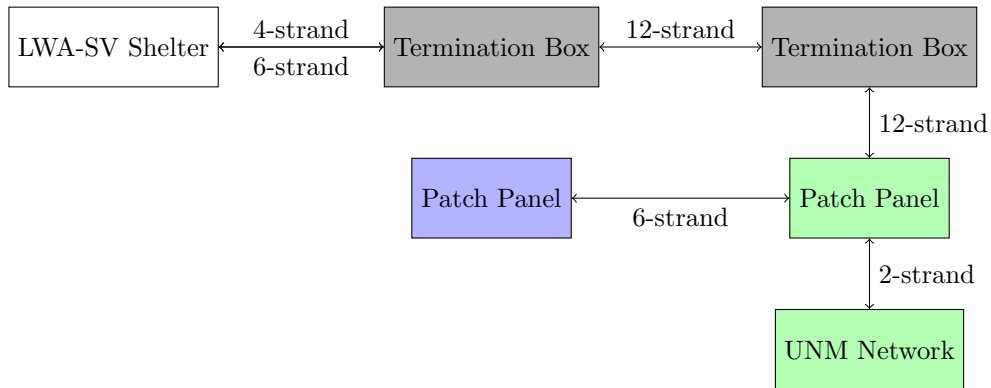


Figure 2: Fiber connections from LWA-SV to the server room of the SERF building. The boxes are color coded by location: white is in the LWA-SV shelter, gray is outdoors, green is the telecom room in the SERF building, and blue is the server room in the SERF building. The links are labeled by the total number of fibers installed.