

A Software-Defined Spectral Line Observing System for the DVA1 Telescope

Stephen T. Harrison^{*(1)}, Timothy Robishaw⁽¹⁾, and Gary Hovey⁽¹⁾

(1) Dominion Radio Astrophysical Observatory, National Research Council.

Abstract

A spectral line observing system for the Dish Verification Antenna 1 (DVA1) telescope at the Dominion Radio Astrophysical Observatory (DRAO) was developed using a software-defined radio (SDR) approach, and is used to observe hydroxyl masers for the purpose of calibrating the telescope's polarization response. A description of the FPGA and software processing pipelines are given, focusing on the digital downconverter (DDC) and software-defined spectrometer. The firmware is implemented on the National Research Council's Kermode FPGA platform, using industry standards for form factor and data transport. Requirements for an upgrade of the John A. Galt 26 m telescope at DRAO are also given, with a description of how the DVA1 backend design will be expanded to meet them.

1 Introduction

In this paper we present the spectrometer backend currently in use on the DVA1 telescope at DRAO [1] used to observe highly polarized 1665 MHz hydroxyl (OH) maser emission from Galactic star-forming regions for the purpose of calibrating the polarization response. The spectral line observing component of the backend is realized in software on a general purpose processor (GPP) without the use of graphics processing unit (GPU) acceleration. Although the bandwidth of this design is limited to a few MHz, this approach is sufficient for the purposes of this calibration effort. This software-based approach also allows for easy experimentation with the number of channels, filter shapes, and active RFI mitigation techniques [2]. The flexibility of this approach has also enabled the debug of equipment issues that were obscured in the accumulated spectra but were readily apparent in the time-domain data.

2 Backend Design

The polarimetric characterization of the DVA1 telescope required the observation of the 1665 MHz Galactic OH maser sources W3 and W49. The radiation from these sources is highly polarized due to the Zeeman effect. In order to reliably measure the polarization signature of Zeeman splitting in the emission from these maser sources, a bandpass of 500 kHz and a channel resolution of 500 Hz were required. A secondary goal of these observations was to measure the time stability of the polarized emission from these

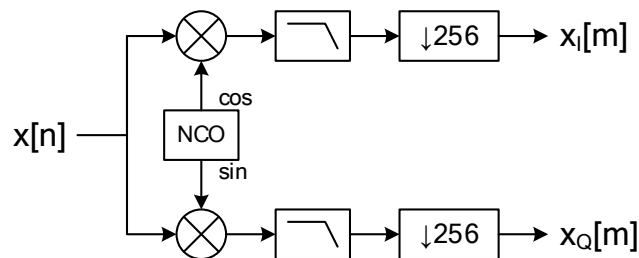


Figure 1. DDC architecture. The real-valued input signal is multiplied by cosine and sine outputs from the NCO to create in-phase and quadrature components. The filtering and downsampling operations are done in three stages to obtain a total sample rate reduction by 256 and a final output bandwidth of $0.71 \frac{f_s}{256}$.

OH maser sources over the period of weeks to months.

2.1 Firmware Architecture

The FPGA portion of the backend is implemented on the Kermode platform [3], a Xilinx Virtex-6 Advanced Telecommunications Computing Architecture (ATCA) processing blade. The ADC card mounted in the Kermode's FMC slot samples at 10-bit resolution and can cover four channels up to 625 MHz of bandwidth in non-interleaved mode, or two channels up to 1250 MHz of bandwidth in interleaved mode [4].

For spectral line observing the FPGA firmware provides a digital downconverter (DDC) as typically found in a direct-sampled SDR receiver. As shown in Figure 1, the real signal $x[n]$ is multiplied by cosine and sine outputs from a numerically controlled oscillator (NCO), filtered, and downsampled to produce the complex baseband signal $\tilde{x}[m] = x_I[m] + jx_Q[m]$. The filtering and downsampling operations are performed in three stages, with a final output bandwidth of $BW = 0.71 \frac{f_s}{256}$, and a tuning step of $\Delta_f = \frac{f_s}{4096}$. The complex baseband signals $\tilde{x}[m]$ and $\tilde{y}[m]$ are streamed out of the Kermode's Ethernet interface as UDP/IP packets. The packet payload format follows the ANSI/VITA-49.0 radio transport standard for IF data [5], illustrated in Figure 2.

The DDC was added to an existing firmware design (Figure 3) which also includes an input monitor (IMON), polyphase filterbank (PFB), multiply accumulator (MAC), and channel packer (CHPK). All blocks have independent enables

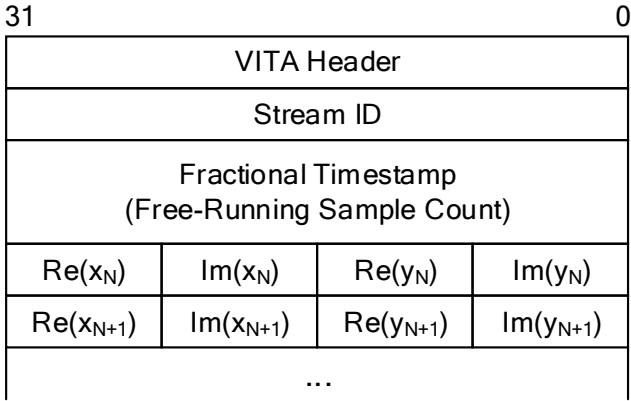


Figure 2. IF data packet payload format showing stream ID, 64-bit timestamp, and 8b+8b complex data samples.

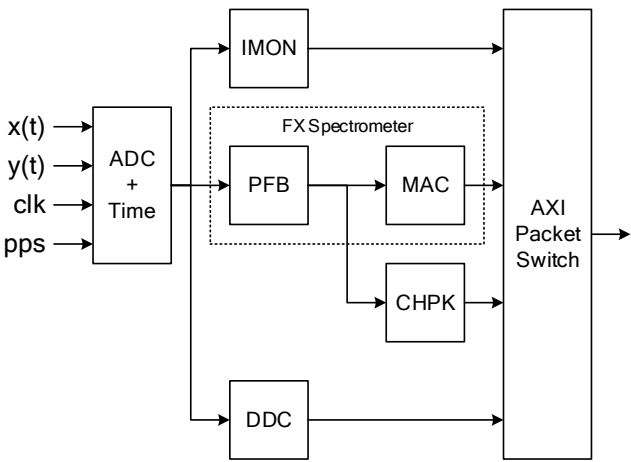


Figure 3. FPGA firmware architecture showing the DDC in context. Data from the ADC is also distributed to an input monitor (IMON) and polyphase filter bank (PFB). The output of the PFB is distributed to a multiply-accumulator (MAC) and channel packer (CHPK). The PFB and MAC work together to create an FX spectrometer.

and can operate concurrently. The PFB and MAC work together to create an FX spectrometer covering the full bandwidth with 1024 channels for continuum measurements. The MAC output is formatted as per the VITA-49.0 Appendix C recommendation for FFT data packets. The CHPK can be used to pack a single channel from the PFB into VITA-49.0 IF data frames, and uses the same packet format as the DDC (Figure 2). The IMON provides a time-domain snapshot capability and a histogram module used to monitor the input distribution and power level. It has also been used to measure the differential and integral non-linearity (DNL/INL) of the ADC [6]. An AXI packet switch aggregates all of the data products into a single packet stream which is sent to the Ethernet interface.

2.2 Software Architecture

The software processing pipeline is implemented in Python with the `multiprocessing` library using a flowgraph ap-

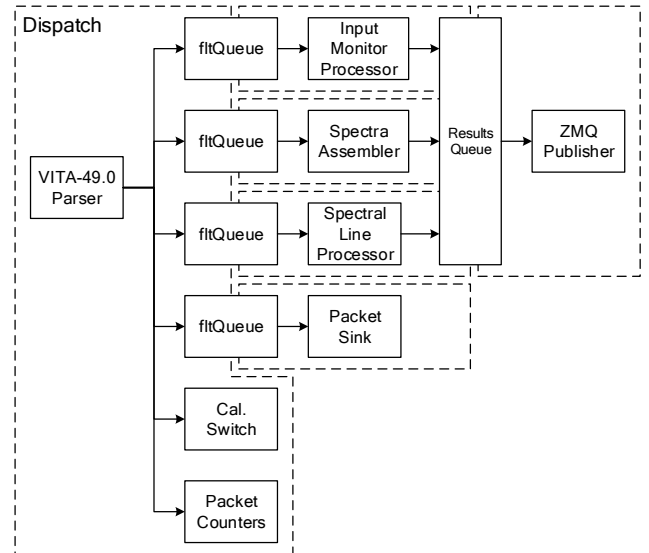


Figure 4. Software processing pipeline. Each dotted line box represents a separate software process. Calibration source switching is handled within the dispatch process to minimize switching latency. All other tasks are handled by separate processes.

proach. A block diagram of the software architecture is shown in Figure 4. In the dispatch process, packets are received, parsed, and sorted into queues based on the VITA-49.0 Stream ID field. Each queue is serviced by a separate process and data products are output to a common results queue. Results are rebroadcast by a ZeroMQ (ZMQ) publisher socket as separate *subscription topics*. As a debug feature, raw packets can be captured to disk for offline processing. In reference to Figure 4, the Input Monitor Processor and Spectra Assembler processes exist simply to unpack and format data products from the IMON and MAC, respectively. Since the DDC and CHPK output a common packet format, the Spectral Line Processor can be used to process data from either one.

Profiling the software reveals that pushing data into the queues is the primary bottleneck, caused by waiting to acquire the lock. For high-throughput queues (such as complex baseband data) the `fltQueue` object aggregates multiple packets to reduce this effect to an acceptable level. A secondary bottleneck is caused by object serialization in `multiprocessing.Queue`.

When acting as the master for calibration source switching, the FPGA emits a timestamped *state change* packet. The software pipeline captures the state change packet and performs a Telnet command to switch the calibration source in the receiver to the appropriate state. This is done within the dispatch process to minimize the switching latency.

The Spectral Line Processor has a process to reassemble the packets into a continuous stream (Figure 5) that can be consumed by the software-defined spectrometer. The samples are converted to complex floating point. The packet times-




Figure 5. Stream assembly process. The packet timestamps are unrolled and the calibration state information is aligned to the timestamps.

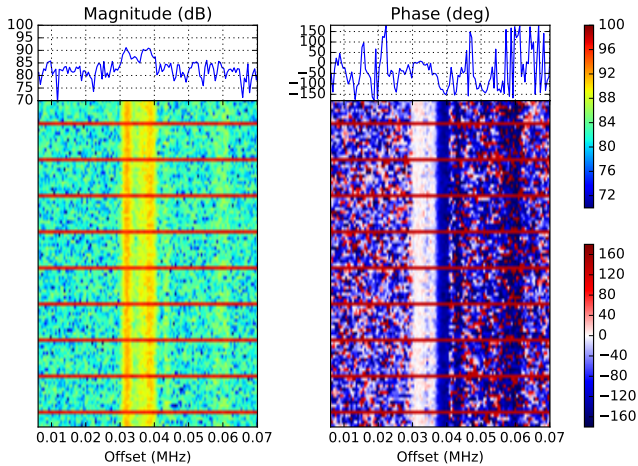


Figure 6. An example of a live waterfall display showing cross-correlation output from the software-defined spectrometer toward molecular cloud W3 at 1665.6 MHz. Polarization features are visible in the phase display. Strong horizontal lines represent the 10:1 duty cycle of the switching correlated noise diode calibration source. The window title bar (not shown) displays the UTC timestamp of the most recent data.

tamp, indicating the time of the first sample in the packet, is *unrolled* over the remaining samples such that each sample is associated with a unique timestamp. The state change information, which arrives separately and is not necessarily synchronous with the start of a packet, is aligned to the buffer and the state is held until a new state change packet arrives. If packets are dropped, the time field is filled in based on the next received packet, the data fields are filled with zeros, and the valid field is set false. The software-defined spectrometer runs slightly behind the stream assembly process, consuming the complex baseband streams and making decisions based on the state transitions.

Various network clients can connect to the ZMQ publisher, such as FITS writers and live displays for any of the results. Most of the live displays take the form of `matplotlib` animations, which are configured to update when the ZMQ socket receives new data. The `matplotlib` library provides basic facilities to pan/zoom and save the plot in various file formats. Figure 6 shows an example of such a display, which can be connected from anywhere on the local network.

2.3 Experimental Setup

To meet the science requirements for this effort, we have reduced the ADC sample rate to 262.144 MHz. The DDC output sample rate is then $\frac{262.144\text{MHz}}{256} = 1.024\text{MHz}$ with a tuning step of $\frac{262.144\text{MHz}}{4096} = 64\text{kHz}$. A 2048-channel software-defined spectrometer gives a channel resolution of 500 Hz. The signal of interest is in the 13th Nyquist zone (1572.864–1703.936 MHz), and a matched pair of analog filters is used to select the hydroxyl band from 1650 to 1680 MHz.

3 Future Work: John A. Galt 26 m Upgrade

An infrastructure upgrade of the John A. Galt 26 m telescope is currently underway. The signal path electronics, control software, and backend were installed many decades ago. The motivation for this upgrade is a forthcoming astronomical survey of the northern sky to map the magnetic field in the interstellar medium. The same state-of-the-art receiver used on the DVA1 telescope, covering the frequencies 900–1800 MHz, will be installed at the prime focus of the telescope.

3.1 Science Requirements

The Zeeman effect leaves a spectral fingerprint on polarized emission from certain interstellar atoms and molecules when they radiate in the presence of a magnetic field. This fingerprint allows astronomers to directly measure the magnetic field strength and direction at the source of radiation.

There are many radio emission lines of astrophysical importance across the frequency band 800–1800 MHz and dozens of them are susceptible to the Zeeman effect. Among them are: the 1420 MHz transition of hydrogen (a.k.a. “the 21 cm line”); the 1612, 1665, 1667, and 1720 MHz transitions of OH; and more than 100 radio recombination lines from electrons recombining with ionized hydrogen, helium, and carbon. The 21 cm line is a particularly attractive target because of the great abundance of interstellar hydrogen in our Galaxy: the spectral line is seen towards every single line of sight in the sky. The upgrade of the John A. Galt 26 m will allow for an ambitious plan to map the Zeeman effect in the 21 cm line while also collecting data from all of the previously mentioned Zeeman-susceptible radio lines in the band.

When an atom or molecule that is susceptible to the Zeeman effect emits spectral line radiation in the presence of a magnetic field, the right- and left-circularly polarized components of the spectral line are shifted in frequency by equal amounts from the line center—when the magnetic field is pointing towards the observer, the IEEE RCP component will be shifted to a higher frequency and the IEEE LCP component to a lower frequency. The splitting for the 21 cm hydrogen transition is $2.8\text{ Hz } \mu\text{G}^{-1}$. Since the average Galactic interstellar magnetic field has a strength of roughly

1 μG and narrowest measured Galactic hydrogen line has a FWHM of about 1 km s^{-1} , the Galactic 21 cm line will never be completely split but only slightly broadened. One measures this broadening by inspecting the difference between IEEE RCP and LCP, otherwise known as the Stokes V spectrum. In order to reliably detect Zeeman broadening, we need at least 6 channels across the FWHM of the narrowest line; we therefore need a spectral resolution of 790 Hz in the 1420 MHz window. Including high-velocity clouds of hydrogen in the outskirts of the Milky Way, the extent of 21 cm emission covers the range 500 km s^{-1} to $+500 \text{ km s}^{-1}$.

Dwells of many hours will be required to observe the Zeeman effect in the Galactic 21 cm emission at each position in the sky. We can also simultaneously attempt to measure Zeeman splitting in any of the four OH emission lines towards lines of sight containing molecular clouds, star-forming regions, or OH/IR AGB stars. In addition, we will simultaneously observe dozens of radio recombination lines being emitted from hydrogen, helium, and carbon in the ionized interstellar medium of our Galaxy. These transitions occur when recombining electrons jump between adjacent quantum levels. Even though the emission might be weak towards any given line of sight, the emission across the frequency band 800–1800 MHz from the dozens of lines from a given atom can be stacked to amplify the signal.

3.2 Backend Design

The design of a new backend for the John A. Galt 26 m telescope at DRAO builds on the success of the DVA1 backend design, with the general architecture largely unchanged from Figure 3. The ADC will sample at 1800 MSPS, covering the receiver’s frequency range of 900–1800 MHz in the second Nyquist zone. The number of DDCs is expanded from one to four for simultaneous observation of spectral lines at 1420 MHz, 1612 MHz, 1665/1667 MHz, and 1720 MHz. The bandwidth of these DDCs is expanded to 7 MHz for OH lines and 14 MHz for hydrogen, in order to accommodate the full velocity range without changing the observing setup.

The CHPK capabilities will also be scaled up from one channel to one-eighth of the total channels. The specific channel numbers will be run-time configurable in order to select the narrow, unevenly-spaced bands where radio recombination lines are known to occur. These will be fine-channelized and accumulated in GPU-based processing nodes.

The expanded capabilities of the design requires the use of 10 GbE for data transport. In the interest of adopting newer standards, the MAC will also be updated to the newer VITA-49.2 format for frequency data [7]. A control system based on Tango Controls is planned to supersede the Python-based software used on DVA1.

4 Conclusion

In this paper we have presented an overview of the software-defined spectral line observing system currently in use on the DVA1 telescope. The design makes use of industry standards such as the ATCA form factor and the ANSI/VITA-49.0 standard for radio transport. We have also given the science requirements for the John A. Galt 26 m upgrade project, and described how the DVA1 backend design will be expanded to meet them.

References

- [1] L. B. Knee *et al.*, “System performance testing of the DVA1 radio telescope,” in *Ground-based and Airborne Telescopes VI*, vol. 9906. International Society for Optics and Photonics, 2016, p. 99063T.
- [2] C. Dumez-Viou, R. Weber, and P. Ravier, “Multi-Level Pre-Correlation RFI Flagging for Real-Time Implementation on UniBoard,” *Journal of Astronomical Instrumentation*, vol. 5, no. 04, p. 1641019, 2016.
- [3] G. J. Hovey *et al.*, “An FPGA based computing platform for adaptive optics control,” in *1st AO4ELT conference-Adaptive Optics for Extremely Large Telescopes*. EDP Sciences, 2010, p. 07006.
- [4] *FMC122/125/126 User Manual*, Abaco Systems, 2016.
- [5] *VITA Radio Transport (VRT)*, ANSI/VITA Std. 49.0, 2015.
- [6] Histogram Testing Determines DNL and INL Errors. Maxim Integrated Products Inc. [Online]. Available: <https://www.maximintegrated.com/en/app-notes/index.mvp/id/2085>
- [7] *VRT Standard for Electromagnetic Spectrum: Signals and Applications*, ANSI/VITA Std. 49.2, 2017.