

ngVLA Memo #29

An Integrated Receiver Concept for the ngVLA

Matthew Morgan and Stephen Wunduke (on behalf of the IRD Team)

Central Development Lab (CDL), NRAO

E-mail: `matt.morgan@nrao.edu`

Abstract. We present a high-level plan for the integrated warm electronics packages on the next-generation Very Large Array (ngVLA). Based on the advances of an in-house Integrated Receiver Development (IRD) program which has been ongoing for some years at the CDL, these front-end assemblies perform all the necessary functions to convert the broadband, analog outputs of the cryogenic stage to digital, optical data streams suitable for transmission over moderate distances with maximum signal fidelity. Key aspects of this proposed methodology are minimal analog electronics supporting a single-stage, direct-to-baseband downconversion (when needed) [1–3], followed by a very low-power, low-overhead digitization scheme, and an industry-standard fiber optic interface carrying unformatted serial data [4–6].

1. Introduction

The role of the integrated receiver packages which are the subject of this memo is to further amplify the signals provided by the cryogenic stage, downconvert them if necessary, digitize them, and deliver the resultant data streams by optical fiber to a moderately remote collection point where they can be launched onto a conventional network for transmission back to the array central processing facility. Hooks are needed to provide for synchronization of local oscillators (LO's) and sample clocks, power leveling, command and control, health and performance monitoring, and diagnostics for troubleshooting in the event of component failure.

We propose to achieve these goals with compact, fully-integrated, field-replaceable, warm electronics modules. Specific details of these modules is presented in the following sections.

2. Frequency Plan

The cryogenic front-ends of an ngVLA telescope station will cover a frequency range of 1.2–116 GHz in up to six bands (neglecting an opaque window around the 60 GHz Oxygen absorption band). There is a desire to have the capability of digitizing this entire bandwidth simultaneously, and transmitting it all back to the correlator for processing. Of course, to do so generally requires a great deal of complex hardware in the antenna, and is at odds with the requirements for low operating costs and long-term reliability.

One of the novel concepts to have come out of the IRD program is a design for a very low-power, resolution-adjustable sampler with a direct serial interface to the laser driver. Known at the Observatory as the *Serial ADC* [4–6], this device would offer more than an order of




Figure 1. Full-coverage frequency plan for ngVLA downconverters/digitizers

Band	Frequency	Receiver Type	# bits	Sample Rate
1	1.2–3.5 GHz	DS	8	7 GS/s
2	3.5–10.5 GHz	DS	8	7 GS/s
		DS	8	7 GS/s
3	11–18 GHz	2SB	8	7 GS/s
4	18–30 GHz	2SB	4	14 GS/s
		2SB	4	14 GS/s
6	70–116 GHz	2SB	4	14 GS/s
		2SB	4	14 GS/s
		2SB	4	14 GS/s
		2SB	4	14 GS/s

Table 1. Details of full-coverage frequency plan in Figure 1. ‘DS’ refers to direct-sampled receivers, while ‘2SB’ refers to sideband-separating receivers. Note that the Band 1 and 2 frequency ranges are adjusted slightly from the baseline plan (originally 1.2–3.6 GHz and 3.6 GHz–10.8 GHz, respectively).

magnitude improvement in size, weight, and power (SWaP) compared to conventional high-speed, fixed-resolution samplers and serializers, while also providing the flexibility to sample at different resolutions and bandwidths as required. One solution for ngVLA is thus presented in Figure 1. Here, we have assumed that the Serial ADC has a maximum serial output rate of 56 Gbps in either of 8- or 4-bit modes (giving sample rates of 7 and 14 GS/s, respectively). Note that although the Serial ADC is capable of providing dynamically-adjustable resolution, in this design the resolution in each receiver band is fixed to minimize the complexity of the analog electronics. The analog electronics in bands 3 through 8 are single-stage direct-to-baseband downconverters with calibrated, numerical sideband separation [1–3]. Bands 1 and 2 use direct-sampled receivers operating in the first, second, and third Nyquist zones.

A summary of the frequency ranges, sample rates, and bit resolutions for this scheme is given in Table 1. A total of ten integrated receivers is required (as will be shown later, both Nyquist zones of the Band 2 range can be serviced by a single module). Note that the RF frequency ranges for Bands 1 and 2 have been modified from the baseline plan (originally 1.2–3.6 GHz and 3.6–10.8 GHz) to allow more efficient coverage of the bandwidth with direct-sampled receivers. This widened the gap between Band 2 and 3 from 200 MHz to 500 MHz. Should the project wish to close this gap, we recommend shifting Band 3 down 500 MHz, but keeping the same 7

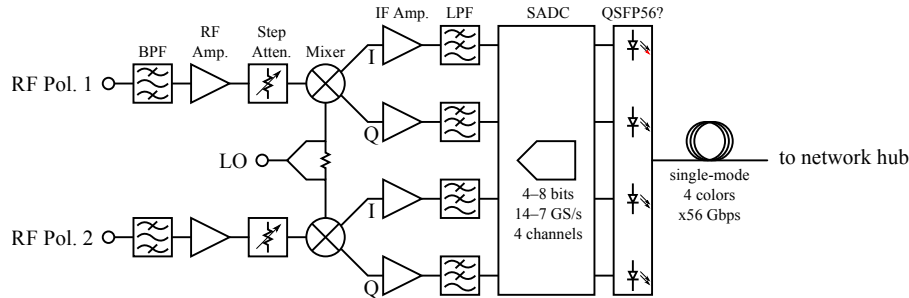


Figure 2. General block diagram of a sideband-separating (2SB) integrated receiver module.

GHz bandwidth, and expanding Band 4 by 500 MHz to fill it in. This would be most efficient for the integrated receivers since the Band 4 downconverter has excess bandwidth anyway.

Note that the use of more than one downconversion/sampling channel per RF band implies that splitters will be needed between the cryogenic front-end and warm-electronics modules. It is not recommended to integrate these splitters and redundant signal paths within the modules as it would be difficult in a multichip block to ensure adequate isolation between the independent LO's, increasing the likelihood of RF spurs. It is also desirable to create field-replaceable units in relatively small functional blocks for ease of manufacture and maintenance, rather than integrating them all into a single, large sub-assembly which is more difficult to handle and to work on, and more expensive to fabricate.

3. Integrated Module Block Diagrams

Although the detailed requirements have yet to be specified, sufficient early prototypes have been developed ranging from L-band to W-band to draw a general block diagram for the two kinds of integrated receiver modules described in Section 2.

3.1. Sideband-Separating (2SB) Integrated Receiver

The first, a sideband-separating receiver module for the higher bands, is shown in Figure 2. It has two inputs, corresponding to the two native polarizations produced by the cryogenic front-end in a given RF band. Each signal path begins with a bandpass filter to limit the signal power and protect the electronics downstream from strong interferers that might lie out of band. An RF post-amplifier increases the signal strength in preparation for the lossy and potentially noisy components which are to follow. After the amplifier, a step attenuator is included for gain adjustment, ensuring that the samplers downstream will always be driven at the optimum level for quantization efficiency. Next, an I/Q mixer downconverts the selected portion of RF spectrum to baseband. This is followed by an additional IF amplification stage. The final analog component is a lowpass, anti-aliasing filter.

Immediately following the filters is the Serial ADC described earlier. We intend to implement this as a four-channel ASIC with direct serial interface to the fiber optics. The lasers and laser drivers will be industry standard transceivers of the Quad-lane Small Form-Factor Pluggable (QSFP) type. Current commercial offerings include the QSFP28 standard, which delivers 28 Gbps on each of four colors. For the ngVLA, we are anticipating that an equivalent 56 Gbps-per-lane unit will be available by the time construction begins. Commercial companies have already begun developing Serializer-Deserializers (SerDes) operating at this speed for both copper and optical interconnects [7]. (Should suitable QSFP modules fail to become available in the next decade, alternatives include the existing new QSFP-DD standard, which has eight lanes at 25

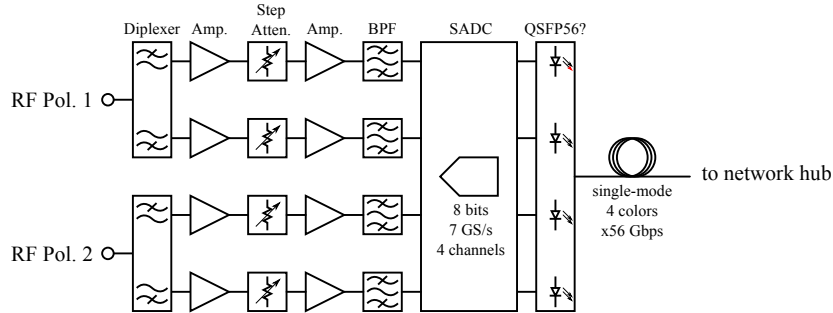


Figure 3. Block diagram of direct-sampled (DS) integrated receiver module.

Gbps each, but this would require an approximate doubling of IF signal paths in the ngVLA.)

3.2. Direct-Sample (DS) Integrated Receiver

The second block diagram, in Figure 3, is a direct-sampled receiver for the lower frequency bands. In the case of Band 2, each polarization input is applied to a diplexer to separate the second and third Nyquist zones (the Band 1 receiver operating in the first Nyquist zone needs no diplexer). These then pass through amplification and power leveling stages before a final anti-aliasing filter.

The rest of the module is identical to the previous one, comprising a four-channel Serial ADC and quad-lane 56 Gbps laser transmitter.

4. First-Mile Data Transmission

The unformatted serial link employed by these downconverters was designed to be a low-power, low-RFI, low-overhead companion to the analog front-ends in the tight focal planes of small-diameter interferometer dishes (or in the densely packed focal planes of larger dishes). While the deserialization algorithm itself is compatible with fiber runs of virtually any length, the relatively low-power transceiver modules which we prefer to integrate with our front-ends in the spirit of the low-overhead development do not reach over the transcontinental distance scales needed to connect the outermost stations of the ngVLA. It is further assumed that for cost-effectiveness (if nothing else) the most distant arms of the instrument will make use of existing commercial fiber to bring their data back to the correlator. The unformatted links in this design, then, are thought of as *first-mile* links, serving to stream the data as efficiently as possible to a network junction, where it can be time-tagged and reformatted for the rest of the journey back to the digital signal processing center of the array.

The *first-mile* label should not be taken literally, of course. Current QSFP-style transceivers are available with maximum reach ranging from a few tens of meters to at least 80 km [8]. The shortest-reach units (which are probably also the cheapest and lowest power) could be used, for example, to connect the focal plane to a rack in the base of the telescope for long-distance transmission. This may not be the most efficient solution however. A more economical approach may be to deliver the unformatted samples to a junction, or hub, where the data streams from several telescopes may be aggregated, as shown in Figure 4. These could be beamformed locally, or simply reformatted and packaged for transmission over the rest of the distance to the array central processing facility. It may be advantageous to locate the junction where it will have access to existing fiber infrastructure for this purpose.

It is worth pointing out that many QSFP-style plug-in units are interchangeable. Thus, a single front-end integrated receiver module design could be used in conjunction with 10m-class




Figure 4. Illustration of two-stage data transmission, wherein the unformatted serial streams from a cluster of telescopes is gathered together and repackaged for delivery over longer distance through commercial data communications infrastructure.

or 80 km-class plug-in units, and simply installed with the shorter-range units near the core of the array where the telescopes are closest together, and with the longer-reach plug-in units in the outer parts of the array where the telescopes are more sparse.

5. Clock and LO Distribution

An important consideration in any interferometer design is the distribution of timing references, specifically sample clocks and local oscillators. Although the full details of the timing distribution system for the ngVLA are beyond the scope of this memo, it bears discussing how the integrated modules proposed here would interface with such a system.

Since the output data stream is digital, having a known sample period, it is only necessary to distribute a coherent LO and sample clock *to* the telescopes, ensuring that the data is captured by the samplers at a known instant in time with a predictable delay offset between antennas. The latency in the data stream coming back *from* the telescopes to the correlator need not be tracked precisely in real-time. That data would be read into first-in-first-out (FIFO) buffers, and then drawn out from those buffers in sync with the data streams from the other telescopes. The initial alignment of those buffers will occur naturally as part of achieving phase closure with the array pointed at a calibration source.

Nevertheless, compensating for the time-dependent delay getting these reference signals *to* the telescopes will likely require a closed-loop, round-trip measurement. Since the Band 1 receiver needs only two of the four QSFP lanes (one Nyquist zone from each of two polarizations) one of the others could be used to implement this loop, as shown in Figure 5. (In addition to the two spare output channels, each QSFP will have four high-speed receive channels which are essentially idle in this design.)

Given this time reference, distributed through the Band 1 receivers to each telescope in the array, the sample clocks and LOs for every band could then be derived. This function could be encapsulated along with monitor and control (M&C) in a separate assembly, as shown in Figure 6. It is not the intent of this memo to elucidate the specific details of the timing distribution system, only to illustrate the fact that *some* kind of closed-loop synchronization between all the telescopes will exist, that it results in a local reference in the telescope that is coherent with all the others, and that the local sampler clocks and LOs will all be derived from it.

Once again, it is important to recognize that the latency of the outgoing data streams need not be part of this closed loop — the FIFO's in the backend will take care of that. While it may

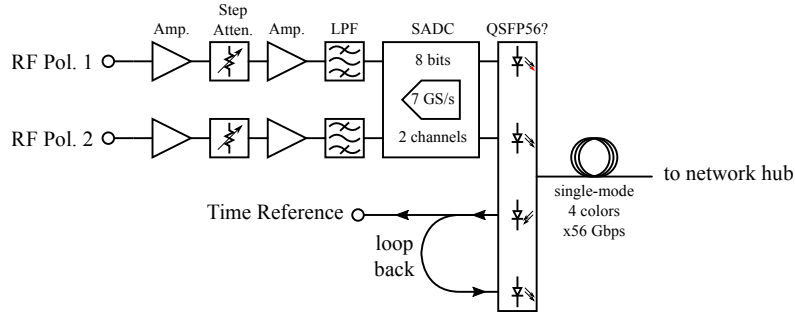


Figure 5. Band 1 integrated receiver with two polarization channels and a loopback channel for distribution of the timing reference.

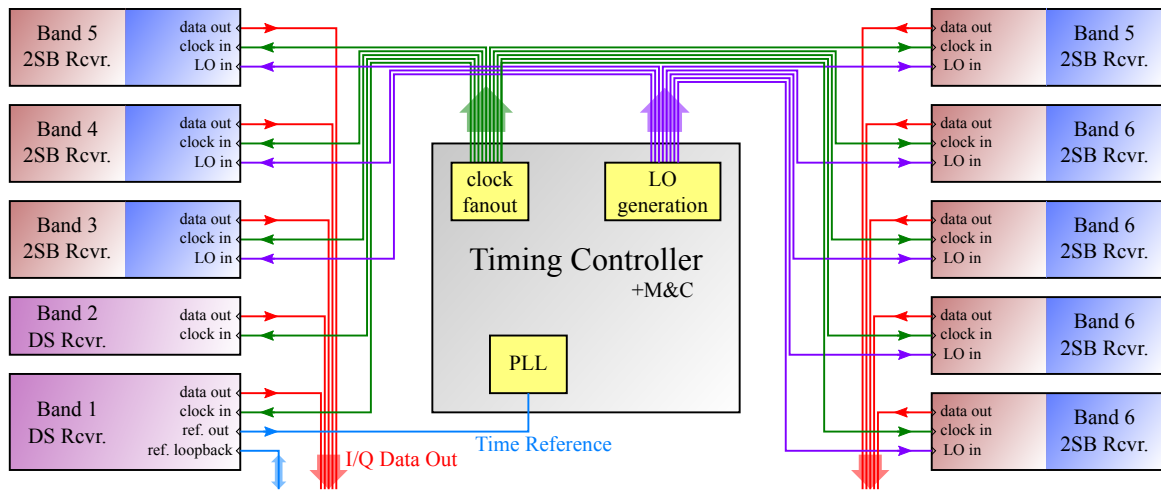


Figure 6. Schematic of clock and LO distribution inside the telescope front end.

be convenient to put timing reference distribution fibers in the same bundle as the return data paths, mutual coherence of their delay fluctuations is not required.

6. Backend Processing Requirements

The front-end modules proposed here are designed to simplify and minimize the hardware installed in the antenna, and to defer as much functionality as possible to the digital domain. It is imperative, then, to disclose exactly what extra processing is required of the backend to support this scheme, aside from the core functions of correlation and spectrometry.

6.1. Deserialization

The premise of the unformatted serial links described above is that the basic characteristics of the raw, sampled data from the receiver meet the fundamental requirements of the link hardware, and that it carries enough information to reliably parse the bit-stream in the backend. For example, DC balance (effectively, an even ratio of 0's and 1's) is required to pass the waveform through AC-coupled components in the laser driver and photodetector. This balance, fortunately, is guaranteed statistically by the symmetry of the sample code ladder if the offset of the digitizer is sufficiently small [5].

Another basic characteristic is transition density, needed for proper clock recovery at the receive end of the link. For binary-modulated waveforms, this transition density is once again ensured statistically when the coding is straight binary (as opposed to two's complement), and when the signal amplitude is at an appropriate level for high quantization efficiency. Once again, experiments have shown that the offset of the digitizer is important to the reliability of this link, and offset adjustment is called out as a critical specification of the Serial ADC.

One minor wrinkle with the 56 Gbps SerDes solutions now undergoing commercial development is that the line coding is no longer binary non-return-to-zero (NRZ), but rather four-level pulse-amplitude modulation (PAM4) where the levels correspond to a gray-code scaling of the original binary pairs [9]. This has no effect on the DC balance, but the impact to clock recovery (or how even to quantify the transition density requirements of such a modulation) has yet to be studied in much detail. This will be the subject of a future ngVLA memo. Nevertheless, early estimates are that the PAM4 modulation is actually even better suited for clock recovery than NRZ, and the expectation at this time is that the noisy data streams from radio astronomy front-ends will provide ample transition edges for this to be done reliably, as it is known to be for legacy systems.

So the critical characteristics of the waveform, such as DC balance and transition density, are guaranteed by proper selection of encoding and attention to ADC offsets in the front-end. What is needed in the backend is a deserialization algorithm that can identify the boundaries between samples and align the samples from separate fibers (e.g. I and Q). The algorithms which accomplish this have been under development in the IRD group for several years, and are discussed in detail in the references [5]. What is important to note in this memo is that the computing requirements for implementing the sample-boundary detection are minimal, amounting to little more than an XOR operation on each successive pair of bits and a small bank of counters.

I and Q channel alignment may additionally require inspection of the cross-correlation phase between the two channels. When the two data streams are out of alignment — that is, when the samples withdrawn simultaneously from the individual FIFO's were not originally acquired at the same instant in time — the phase of the cross-correlation between the two channels will have a slope across the IF band equivalent to $n\pi$, where n is the number of samples of misalignment. (This assumes the power spectrum is dominated by the RF signal path, and not by IF or LO contributions, which is the case for a properly functioning radio astronomy receiver). Correction of the misalignment is a simple matter of discarding n samples from one data stream. Fortunately, this operation needs only be done once upon start up, or when something (e.g. a power failure) interrupts the continuity of the data flow, which will be clearly indicated by the sample-boundary detection circuitry.

6.2. Sideband-Separation and Calibration

The backend will also need to derive the upper- and lower-sidebands from the I and Q data. In normal operation this is quite simple. Each sideband is simply a weighted sum of the complex I and Q data after spectral processing through a Fast Fourier Transform (FFT) or Polyphase Filter Bank (PFB). The complex weights, or coefficients, are frequency-dependent and must be derived from a calibration procedure which measures the amplitude and phase imbalances of the I/Q downconversion. This calibration will be unique to each front-end integrated receiver module.

Historically, IRD receivers have been calibrated in the lab by injecting a test tone that is swept across frequency through both sidebands. Integration of the front-ends has ensured that the amplitude and phase imbalances are stable enough to use this initial calibration indefinitely so long as nothing internal changes, such as a bias level or part replacement. Experiments have shown that wide temperature swings and aging have negligible impact on these coefficients. This




Figure 7. S-band integrated receiver prototype. (a) Overview. (b) Digital/photonic circuitry. (c) RF circuitry. An American quarter is shown for scale. (Note: these pictures were taken on two modules of identical design, one gold-plated, the other bare aluminum.)

is especially true of interferometers like the ngVLA, where the level of image rejection needed is much less stringent than on single-dish telescopes.

Nevertheless, it would simplify operations considerably if there was the capability to recalibrate these modules in-situ, on the telescope. On-board injection of a test tone has been provided in the past — e.g. on the Green Bank Phased-Array Feed (PAF) [3, 10] — but this adds complication that we would prefer to avoid on the ngVLA. An attractive alternative is a novel calibration procedure now under development at the CDL [11] which requires only a known, broadband noise input. It is more computationally intensive than tone calibration, but needs only to be performed once (or occasionally) and requires no extra hardware in the front-end.

7. Early Prototypes

To better illustrate the intended design and construction of the integrated receiver modules, we present here an overview of two prototypes that have been produced under the IRD program.

7.1. S-Band Prototype

The first is an S-Band (1.7–2.6 GHz) single-channel unit shown in Figure 7. The basic design of the module is essentially that shown in Figure 2, but only for one polarization.

This module downconverts the RF input to I and Q basebands of approximately 0–310 MHz. It samples those basebands at 622.08 MS/s with 4-bit resolution. The I/Q channels are serialized and then modulated onto 1310 nm optical carriers for transmission over single-mode fibers with a serial rate of 2.48832 Gbps per lane. The transmitters in this case were not QSFP modules, but dual, low-power Fabry-Perot Lasers with integrated drivers and a 10 km reach.

These unformatted data streams are then received by an FPGA which uses our patented in-house algorithm to locate the sample boundaries and align the I/Q lanes. A screenshot of the time-domain panel from the laboratory graphical user interface (GUI) for this experiment is shown in Figure 8(a). Live data streams from three modules identical to the one shown in Figure 7 are displayed on this screen simultaneously. The middle column shows time-series data for each of the I/Q pairs in blue traces. The right column shows Gaussian histograms of this data with 16 bins corresponding to the levels of the 4-bit samplers. On the left are user controls for testing various aspects of the deserialization algorithm, and green indicator lights showing that each of the six lanes is properly aligned. Beside one set of indicator lights is a confidence number, ranging from 0.9845 to 0.9996. This number represents the density of statistically derived alignment triggers in the data stream and indicates a confidence level *very much* higher than 99%.

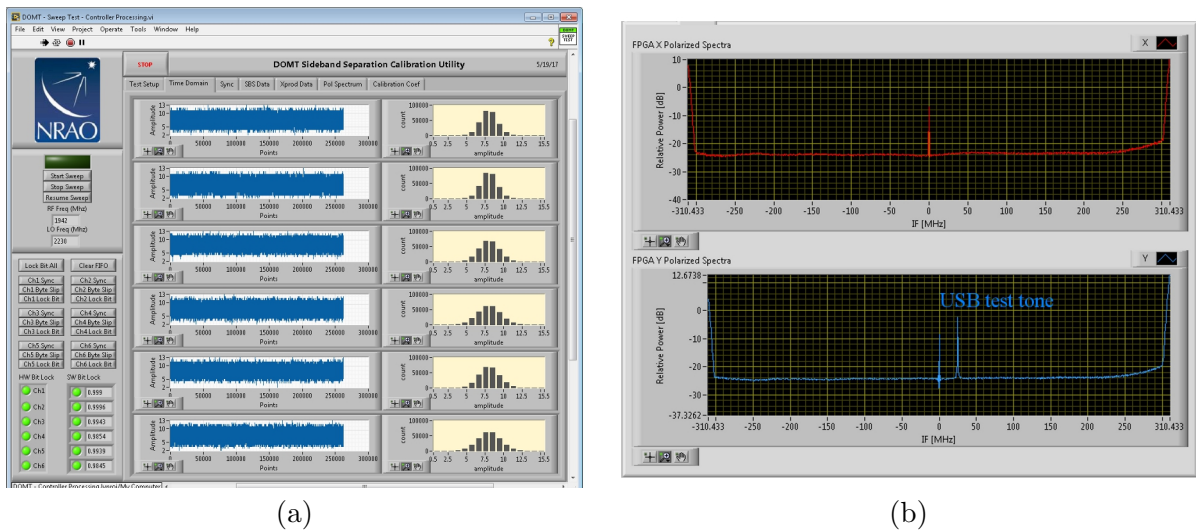


Figure 8. GUI screenshots for S-band prototype modules. (a) Time-domain panel showing time-series (blue traces, middle column) and histograms (grey, right column) for three I/Q pairs. (b) Frequency-domain panel showing sideband- and polarization-processed spectra with polarized test tone input at 25 MHz offset.

A *polyphase filter bank* (PFB) converts the data into the spectral domain, and upper- and lower-sidebands are reconstructed using calibrated coefficients. The experiment in which these three integrated front-ends were used also implemented a unique polarization-synthesis methodology [12]. While this feature is not pertinent to the current ngVLA system design, it is useful to show how the front-ends and the associated algorithms were able to produce fully-processed sideband- and polarization-separated spectra in Figure 8(b). Note that a polarized, single-sideband test tone was injected for this measurement at +25 MHz offset in the Y-polarization. The spikes at IF=0 are DC terms associated with offsets in the samplers. The high noise “wings” beyond ± 310 MHz in both spectra correspond to faulty extrapolation of the sideband and polarization coefficients outside the range of the calibration sweep.

7.2. W-Band Prototype

The second prototype we will discuss is a W-band front-end which is still under construction. An exploded mechanical drawing of the module is shown in Figure 9. Covering an RF frequency range of 75–110 GHz, this two-channel module is very close to that shown in Figure 2. It downconverts the RF spectrum to I and Q basebands of approximately 0–500 MHz, where they are sampled at 1 GS/s with 8-bit resolution. Since the Serial ADC, which is an integral part of the ngVLA plan, has not yet been developed (due to lack of funding), this module instead uses off-the-shelf samplers and a Kintex 7 FPGA to implement the high-speed serializer (as well as other M&C functions). As a consequence, this module is very much larger and requires more power than it would have been had the Serial ADC been available. The FPGA is so under-utilized in this design as to be nearly idle, present only because it is the only practical off-the-self device fast enough to serialize the data streams from the ADCs.

The high-frequency analog electronics in this block are implemented using a combination of commercial and custom-built Monolithic Millimeter-wave Integrated Circuit (MMIC) chips and ceramic substrates, coupled to the input ports via WR-10 waveguides machined into the split-block housing. Once downconverted, the baseband signal paths reside on a circuit board. At the end of each IF signal path are test points brought out as SMP coaxial connectors, allowing

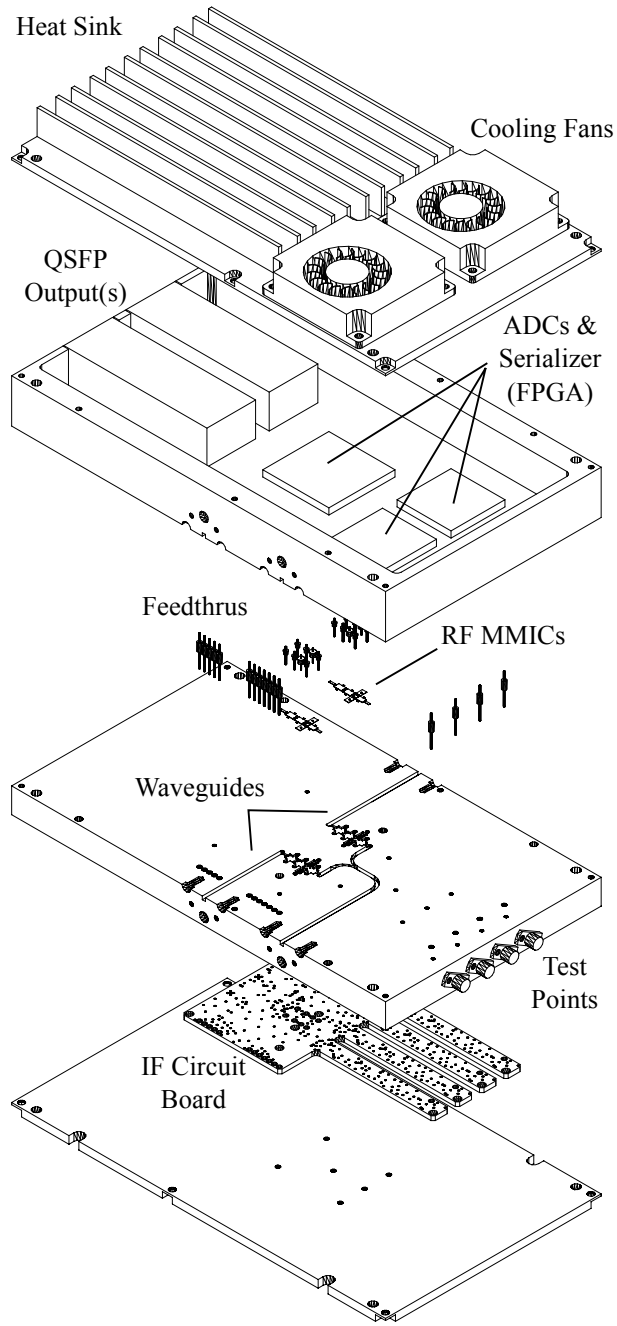


Figure 9. Exploded mechanical view of W-band integrated receiver prototype.

one to monitor the IF spectra just before digitization for troubleshooting purposes.

The IF signals are then passed to a digital/photonic circuit board where the ADCs and serializing FPGA are located. Although one QSFP transmitter is sufficient to output all four I/Q channels from this module, the FPGA has capacity to drive two, so a second QSFP receptacle was included. The FPGA could be programmed to deliver duplicate data streams on this output, or a test pattern, or even processed data should that be deemed useful.

Command and control of the module is available through either a Serial Peripheral Interface (SPI) bus, or an included Ethernet port. Through these interfaces, the module reports health monitors corresponding to the IF output power in each channel (measured via log-amps), temperature sensors integral to the QSFP modules, and others having to do with the state of the serial data transceivers on the FPGA.

8. Conclusion

We have presented in this memo a proposed plan for warm, integrated, front-end modules for the ngVLA, capable of digitizing the entire frequency range of the instrument from 1.2 to 116 GHz (minus the Oxygen absorption band around 60 GHz). Sampling is at 8-bit resolution up to 18 GHz, and 4-bit resolution beyond.

This plan requires ten front-end integrated receiver modules (Figure 1) of six distinct designs (Table 1) conforming to three basic configurations (Figures 2, 3, and 5). Each module is compact enough to be carried in one hand and installed on the telescope with only a wrench (for the coax connectors) and a ball driver (for waveguide flange and mounting screws), greatly easing the maintenance operations of the ngVLA.

The novel concepts of this design include numerically calibrated sideband-separation at baseband and low-overhead, unformatted serial fiber-optic links, all of which have been demonstrated in laboratory experiments. Some advancement of commercial technology (e.g. faster QSFP-style laser transmitters) is anticipated, while others (such as the Serial ADC) would need to be developed by the Observatory in order to realize the full potential benefits (SWaP) of the technological approach outlined in this memo.

References

- [1] J. R. Fisher and M. A. Morgan, "Analysis of a single-conversion, analog/digital sideband separating mixer prototype," *NRAO Electronics Division Internal Report*, no. 320, June 2008.
- [2] M. A. Morgan and J. R. Fisher, "Experiments with digital sideband-separating downconversion," *Publications of the Astronomical Society of the Pacific*, vol. 122, no. 889, pp. 326–335, March 2010.
- [3] J. J. Castro and M. A. Morgan, "Digital sideband separating downconversion for the green bank telescope phased array feed," in *URSI National Radio Science Meeting*, January 2016.
- [4] M. A. Morgan and J. R. Fisher, "Word-boundary detection in a serialized, gaussian-distributed, white-noise data stream," *NRAO Electronics Division Technical Note*, no. 213, October 2009.
- [5] M. A. Morgan, J. R. Fisher, and J. Castro, "Unformatted digital fiber-optic data transmission for radio astronomy front ends," *Publications of the Astronomical Society of the Pacific*, vol. 125, no. 928, pp. 695–704, June 2013.
- [6] M. A. Morgan and J. R. Fisher, "Statistical word boundary detection in serialized data streams," U.S. Patent 8,688,617, April 1, 2014.
- [7] (2015) Avago demonstrates industry-leading 56Gbps PAM4 SerDes. [Online]. Available: <https://www.broadcom.com/news/press-release/demonstrates-industry-leading-56Gbps-PAM4>
- [8] (2016) Inphi offers 100G PAM4 QSFP28 for 80-km data center interconnect. [Online]. Available: <http://www.lightwaveonline.com/articles/2016/03/inphi-offers-100g-wavelength-pam4-qsfp28-for-80-km-data-center-interconnect.html>
- [9] H. Zhang, B. Jiao, Y. Liao, and G. Zhang, "PAM4 signaling for 56G serial link applications – a tutorial," in *DesignCon 2016*, January 2016.
- [10] D. Roshi, W. Shillue, J. R. Fisher, M. A. Morgan, J. Castro, W. Groves, T. A. Boyd, R. Simon, L. Hawkins, V. van Tonder, J. Nelson, J. Ray, T. Chamberlain, S. White, R. Black, K. Warnick, B. Jeffs, and R. Prestage, "A highly-sensitive cryogenic phased array feed for the green bank telescope," in *32nd URSI General Assembly and Scientific Symposium*, August 2017.

- [11] O. Y. Ojeda, “to be published,” *NGVLA Memo*, 2017.
- [12] M. A. Morgan, J. R. Fisher, and T. A. Boyd, “Compact orthomode transducers using digital polarization synthesis,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 58, no. 12, pp. 3666–3676, December 2010.

The National Radio Astronomy Observatory and Green Bank Observatory are facilities of the U.S. National Science Foundation operated under cooperative agreement by Associated Universities, Inc. This work was supported by awards AST-2034328 (MSIP Prototype Antenna) and AST-2334267 (ngVLA Design Activities); NRAO related activities are funded under award AST-1647378 (NRAO Operations/Development).

