ngVLA Electronics Memo No. 1

Downconversion and Digitization Methodology for the ngVLA

Matt Morgan

January 6, 2020

1 Introduction

A key functional requirement of the Integrated Receivers and Digitizers (IRD) work-package within the Antenna Electronics Integrated Product Team (IPT) of the next-generation Very Large Array (ngVLA) is to downconvert and digitize high-frequency spectral information from the cryogenic front-end for delivery to the central signal processor (CSP). Optimally, the design selected should

- provide the required frequency coverage,
- comprise the fewest integrated modules per antenna (to minimize cost and maintenance),
- yield the greatest signal fidelity (free of spurs, gain ripple/slope, and excess noise),
- enable cost-effective, large-scale manufacturing, and
- be adaptable to changes in scope and future upgrade paths.

The purpose of this memo is to describe and justify the chosen downconversion and digitization/data-transfer schemes for the warm antenna electronics design in light of these goals. Section 2 discusses the general advantages of integrated electronics design and why it was selected for the ngVLA. In Section 3, we explore several common topologies for downconversion and compare them according to the aforementioned strengths and weaknesses of integrated design. Finally, in Section 4, we discuss the overheads of the conventional framework for digital data transfer designed by the data and telecom industries, and explain why a novel, leaner approach is considered more optimal for radio astronomy.

2 Integrated Receiver Design

Even prior to the emergence of the ngVLA as a project, researchers at the Central Development Lab (CDL) had begun to explore reoptimize architectures for radio astronomy receivers that leveraged modern advances in integrated electronics and digital signal processing. These two general technologies are deemed to be mutually complementary, in that numerical digital processing (with appropriate calibration) provides accuracy and flexibility that is far superior to pure-analog designs, while integrated construction of the residual analog hardware delivers uniformity and stability of performance to ensure the greatest precision and longevity of digital calibrations.

The two guiding principles of the integrated receiver development program since the beginning have been

- to digitize signals from the sky as close as possible to the focal plane of the telescope, deferring as much functionality as is feasible into the digital domain, and
- to maximize integration, especially across the traditional boundaries of analog, digital, and photonic components, so that the greatest performance could be achieved in a compact, power-efficient, field-replaceable unit.

Early digitization, compactness, and power efficiency, more than simply conveniences, have significant performance and life-cycle cost advantages, as will be described in the following sections. These considerations are what have driven the design to the architecture which was ultimately selected for the ngVLA [1].

2.1 Strengths and Weaknesses

It is important to understand that integrated electronics, like any technology, has both strengths and weaknesses that differ from those of earlier approaches which drove the architecture of previous generations of instruments. It is a very common mistake to repurpose a preexisting block diagram, optimized around the strengths and weaknesses of older technology, to be implemented with integrated components without taking into account those differences.

In this case, the older generation of technology comprised individually packaged and connectorized components (amplifiers, mixers, etc.) plumbed together with coaxial cable or waveguide in isolated analog, digital, and photonic subassemblies. As an example, a photograph of a traditional, connectorized analog sub-assembly is shown in Figure 1(a). Integrated modules, in contrast, at high frequency would comprise multiple semiconductor die — Monolithic Microwave Integrated Circuits (MMICs) — as shown in Figure 1(b), or at lower frequency, high density circuit boards of surface-mount components. In practice, most integrated modules will incorporate a mix of bare die (for the LO and RF) and integral circuit boards (for baseband, digital, and photonic parts) as required by the application.

The relative strengths and weaknesses of these two approaches is summarized in Table 1. The low replication cost and repeatability of performance offered by integrated techniques is reason enough to select them for a large-N array such as the ngVLA, but the weaknesses, particularly the propensity for cross-talk between elements, must be a driving factor in the architectural design of the integrated modules — in fact, it was pivotal in the decision to use single-stage, direct-to-baseband downconversion for this project.

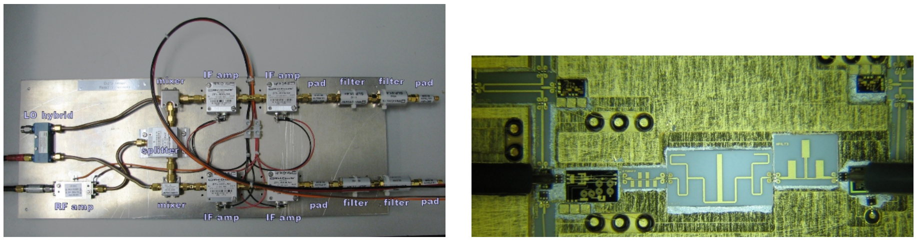

(a) (b)

Figure 1: Examples of (a) a traditional sub-assembly of individually packaged, connectorized components, and (b) an integrated module containing semiconductor chips.

Table 1: Connectorized vs. Integrated Construction

Technology	Strengths	Weaknesses
Connectorized	• ease of assembly and test • hand-tuned performance • high part-to-part isolation • low development cost	• bulky • prone to pass-band ripple & slope • poor repeatability • high replication cost
Integrated	• compact • amplitude & phase stability • uniformity/repeatability • low replication cost	• needs special equipment • higher loss, lower Q • prone to leakage & cross-talk • high development cost

2.2 Advantages of Compactness

One of the aspects of the IRD program which is least understood by those outside the group is the importance that is placed on compactness. While most understand the present trend in radio astronomy that dish diameters are shrinking, even as the number dishes in an array is growing, and therefore the room available near the focal plane of such a dish to support the electronics for an ever-increasing operational bandwidth is limited, once the electronic package is small enough to fit, they would say, there is no advantage to making it any smaller.

In fact, there are significant performance as well operational advantages to having a more compact module. It is critical for high-performance, wideband systems to maintain flat gain and phase structure so that sensitivity and dynamic range can be optimized simultaneously. Tighter integration ensures that the electrical length between components internally is minimized so as to avoid the creation of rapid pass-band ripples due to standing waves. Smaller modules, in turn, can be placed closer to the cold electronics package, reducing the external cabling and loss slopes, while easing the stabilization of temperature.

Operationally, a large, widespread antenna array will be easier and less costly to maintain if a single technician can take a full complement of warm receiver modules with him or her into the field, stopping at several telescopes to swap out new units for defective ones, before returning to a central repair facility to unload and resupply. Conventional receiver packages would have required multiple personnel with equipment to carry and install large, heavy pieces of equipment. These integrated modules, in contrast, fit in one's hand or shirt pocket and can be installed with nothing more than a screw-driver and a wrench.

It is not yet clear in the case of ngVLA whether individual IRD modules are to be replaced at an antenna in the field by a single technician as described above, or if instead a larger portion of the front-end electronics containing the IRD modules as a sub-element will be removed from the antenna and brought back to a repair facility in whole. If the latter turns out to be the case, than the importance of compactness in the IRD modules is not at all diminished, but rather is the enabling factor that makes it feasible to treat such large portions of the antenna electronics as a unified "field-replaceable" element in the first place.

2.3 Power Dissipation Drives Life-Cycle Cost

Another important but underrated aspect of the IRD program is the pursuit of the lowest possible power dissipation. While this reduces the overall power consumption of the antenna (probably negligible compared to other subsystems, say, the cryogenics), it more importantly improves the lifetime of the active components. Considering that the integrated modules incorporate all of the functions of the warm analog electronics, the samplers, and the fiber-optic transceivers into one small package, there is quite a bit of hardware packed into them, so it is essential for thermal management that each component dissipates as little power as possible.

3 Downconversion

A significant portion of the ngVLA's 1.2–116 GHz frequency range is beyond the reach of practical digitizers, so the spectrum must first be downconverted to a lower frequency. A critical performance parameter of analog downconverters is the suppression of leakage from undesired image bands. A number of approaches have been developed over the years to achieve this. These will be summarized and discussed in the following sections.

3.1 Single-Stage Image-Filtered

Perhaps the simplest approach is simply to attenuate the undesired image frequencies prior to downconversion, as illustrated in Figure 2(a). The frequency bands associated with this technique are shown in Figure 2(b), where a lower-sideband downconversion is shown for illustration. When downconverting from

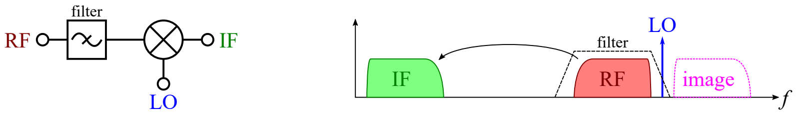

(a) (b)

Figure 2: (a) Single-stage mixer with image-rejection filter. (b) Spectral diagram.

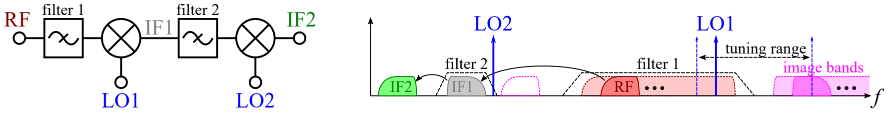

(a) (b)

Figure 3: (a) Two-stage downconversion. (b) Spectral diagram.

a fairly high-frequency, like ngVLA Band 6 (70–116 GHz), this approach demands a very high-selectivity filter to cutoff sharply between the upper end of the RF band and the lower end of the image band. Typically, this would have to be a cavity or dielectric resonator filter, which are large and expensive. To alleviate the selectivity requirement, the IF bandwidth could be pushed higher in frequency, spreading the RF and image bands further apart, but this then demands a very high-speed digitizer, or at least one that has an exceptionally fast track-and-hold stage, driving up the power dissipation.

Further, unless the sampler can digitize the entire RF bandwidth provided by the feed in one shot (rarely the case in radio astronomy), the filter will have to be tunable to cover the image bands which shift laterally as the LO is swept, and will undoubtedly overlap part of the RF input range. Pushing the IF frequency high enough so that there is no overlap between image and feed frequencies is not even remotely feasible.

3.2 Multistage

One way to alleviate the filtering selectivity and tuning requirements is to employ two-stage downconversion, such as that shown in Figure 3. The wide RF feed frequency range is first filtered with a relatively low-selectivity, fixed filter (filter 1 in the figure). A first local oscillator, LO1, is swept through the upper part of the band and above, such that a sliding window of RF bandwidth is downconverted to relatively high intermediate frequency, IF1. The lowest possible image for this downconversion (shown in magenta at the far right of the figure) is still outside of the RF bandwidth as well as the filter1 passband. The high IF1 is then downconverted with a second, fixed local oscillator, LO2, after another fixed filter, filter 2. Although the final intermediate frequency, IF2, is low enough to be digitized effectively, it was downconverted from a more moderate frequency, IF1, than the original RF band, so that filter 2 may still have

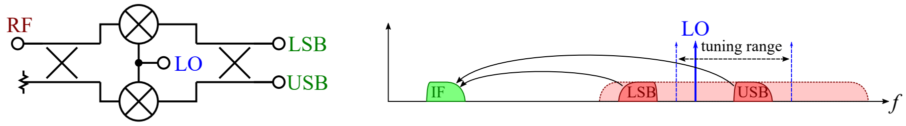

(a) (b)

Figure 4: (a) Sideband-separating mixer. (b) Spectral diagram.

relatively low selectivity. In this way, both filters are fixed and need not have exceptionally high selectivity.

There are costs, however, beyond just the increase in complexity (which is itself a factor). For one, the first LO frequency must extend well above the highest operational frequency of the telescope. More importantly, however, the presence of two local oscillators feeding into the same integrated analog signal path increases the population of spurious mixing products. Recall that one of the disadvantages of integrated construction is the relatively poor isolation between components and the propensity for leakage or cross-talk. This makes it difficult to sufficiently suppress higher-order mixing products, and the fact that LO1 must be tunable makes them impossible to avoid in frequency space — they are literally everywhere. While this approach has been used in non-integrated receivers, it is clearly not a good choice for integrated architectures.

3.3 Analog Sideband-Separating

An attractive alternative is the sideband-separating mixer [2], shown in Figure 4. Rather than prefiltering the image bands, both are downconverted simultaneously through a clever arrangement of mixers, splitters, and quadrature hybrids which use phasing to deliver the two sidebands separately to independent ports.

The first quadrature hybrid may appear on either the RF or LO port [3], but a hybrid is always required on the IF output. Analog hybrids of this type cannot be made to work at DC, so there is a limit to how low in frequency the IF range may go. Although hybrids with decade bandwidths have been demonstrated, they are lossy and have significant gain slope, as well as being physically enormous if made for the first Nyquist zone of any practical digitizer. Generally, analog sideband-separating mixers are therefore designed for at least the second Nyquist zone, so that the hybrid is required to cover at most one octave (unless the mixer is followed by other mixing stages prior to digitization, as is the case in ALMA). This of course requires the digitizer to have a wideband track-and-hold stage supporting the second-Nyquist frequency range, at moderate cost to power dissipation, and placing additional constraints on the digitizable bandwidth.

Furthermore, even over an octave bandwidth, it is difficult for the IF quadrature hybrid to maintain accurate amplitude and phase balance. Image rejection ratios of even 15 dB can be a challenge for such systems if wideband operation

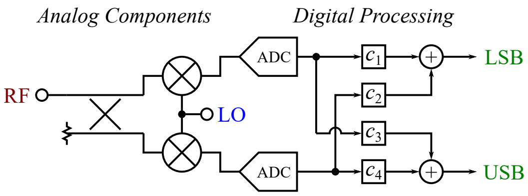

Figure 5: Sideband-separating mixer with a numerically implemented IF quadrature hybrid using calibrated complex-gain coefficients.

is desired.

3.4 Digital Sideband-Separating

To mitigate these issues, the IF hybrid may be implemented numerically, after digitization, as shown in Figure 5. Not only can this "digital hybrid" operate down to arbitrarily low frequency, thus allowing the sampler to work directly at baseband, it is largely immune to the amplitude and phase errors of pure-analog implementations. The complex coefficients (c_1 to c_4 in the figure) can be calibrated to correct for the imbalances of the preceding analog components, thus achieving nearly perfect sideband separation [4]. Performance in this area is limited only by the sensitivity of the calibration procedure, which may be carried out in the lab, and the longevity of those calibrations which is greatly enhanced by the integrated construction of the analog components. For the ngVLA's requirement of 30 dB image suppression, it is likely that the laboratory calibration will be effectively permanent, needing updates only when a part is repaired or replaced.

4 Digitization and Data Transfer

Once downconverted, the analog signal must be digitized for transmission away from the focal plane, and ultimately back to the central signal processor. Although some telescopes transmit analog signals from the focal plane to other parts of the telescope (say, the base or pedestal) prior to digitization, this is not considered a desirable solution for numerous reasons. Coaxial cables over the lengths required would introduce a great deal of loss and gain slope. Analog fiber-optic links, on the other hand, even short ones, are extremely noisy and prone to dynamic range issues, especially for wide analog bandwidth. Either option further exposes the signal path to temperature fluctuations and flexure at the telescope axes which introduce phase variations that ultimately must be tracked and taken out. By digitizing early, we protect the information from any kind of corruption during transport.

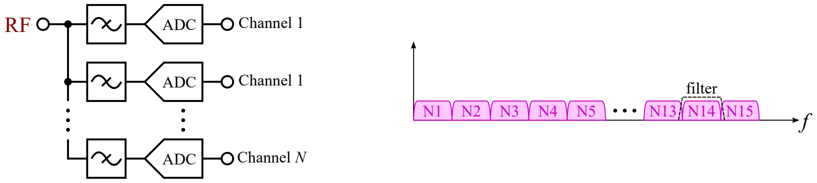

(a) (b)

Figure 6: (a) Direct-sampling architecture. (b) Filtering a high-level Nyquist zone.

4.1 Direct-Sampled

Some designers favor direct-sampling at relatively high-frequencies (in high-number Nyquist zones), as shown in Figure 6, thus bypassing the downconversion altogether. While appealing for its architectural simplicity, it requires an ADC with an extraordinarily high-frequency input stage — inevitably a power hog with compromised performance — as well as a suite of very high-performance filters which are bulky and expensive, and yet still sacrifice some bandwidth at the zone edges. These drawbacks aside, it is highly unlikely that sampling could be achieved as high as ngVLA's Band 6, requiring one of the aforementioned downconversion schemes for at least that band anyway.

4.2 Conventional ADC Boards and Data Formatters

For the bandwidths that we require, conventional ADC's must be interfaced directly to power hungry FPGA's to do the formatting (bit-scrambling, packetizing, etc.) required of industry standard data links. Even the modern JESD204B standard for so-called 'serial-output' ADCs[1] is meant only for short-distance communication, typically from ADC to FPGA on the same circuit board or from a mezzanine card to its motherboard. Long-distance communication between the ADC and FPGA, even tens of meters, is out of the question, due to the need to keep synchronized multiple serial lanes and frame clocks.

In addition to the bulk and power dissipation associated with implementing these schemes, especially for broad-bandwidth data, the localized buildup of high-speed digital electronic components increases the risk of self-interference, both through emission and conduction via supply buses and grounds. While we are assuming the risk of integrating analog and digital electronics in the front-end, it is nevertheless advantageous to keep the digital electronics to a minimum.

---

[1] These standards typically multiplex the sample stream to several independently serial outputs, along with parallel clock channels, making it more accurately a combination of serial and parallel communication.

4.3 Unformatted Links and the Serial ADC

To that end, the IRD group has conducted extensive research in the use of unformatted serial links for transmission of digitized radio astronomy data, without the overhead associated with commercial protocols [5]. While conventional data transmission systems are optimized to work with any kind of data — whether it be compressed streaming video or mostly-empty files having long strings of zeros — we know a great deal a priori about the content of our signals. To a high degree of accuracy, they will essentially comprise Gaussian-distributed white noise (dominated in most cases by the noise of the cryogenic amplifier, if not the source on the sky).

We can use these properties to our advantage by exploiting the natural statistics of our signal as its own "formatting" without the need of high-speed digital electronics in the front-end, aside from the digitizer itself and a bare-bones serializer to drive the laser [6]. There are efficiencies to be gained from this in the front-end, especially if the sampling and serialization functions can be integrated onto the same piece of silicon, as illustrated in Figure 7.

The surface-mount packages for common ADCs are typically on the order of an inch square, not because of the size of the chip inside, but because of the number of pins that must be accommodated around the periphery to support a parallel (and often multiplexed) interface. Even the aforementioned JESD204B devices and those packaged in ball-grid arrays (BGAs) are still among the largest devices a high-speed circuit board will contain, aside perhaps from the monstrous FPGA itself. In our application, the traces from this wide interface are immediately collected again in a serializer which reduces them all once more to a single differential output, connected to a laser driver. By integrating these two parts on a single chip, we eliminate the vast majority of intermediate pins, giving us an order of magnitude reduction in footprint on the circuit board.

Furthermore, the interface between chips using conventional parts comprises resistively-terminated transmission lines — most commonly a protocol known as low-voltage differential signalling (LVDS). Clearly wasteful of power, they are nevertheless required due to the (a priori) unknowable parasitics of the traces on the board. On-chip, the connections between ADC and serializer can be wired directly, gate-to-gate, without any loss of power except when switching.

In early IRD prototypes incorporating all the warm-electronics of a radio astronomy front-end — post-amplification, power leveling, downconversion, filtering, digitization, and fiber-optic transmission — about 75% of the power was dissipated in the digital electronics. Our first-generation Serial ADC, supporting two channels operating at up to 10 Gbps serial rate per channel, will do the same job with one-tenth the power (about 333 mW) in a tiny, 5 mm QFN package. Similar savings are expected from our next-generation product supporting 56 Gbps per lane, as demanded by the ngVLA.

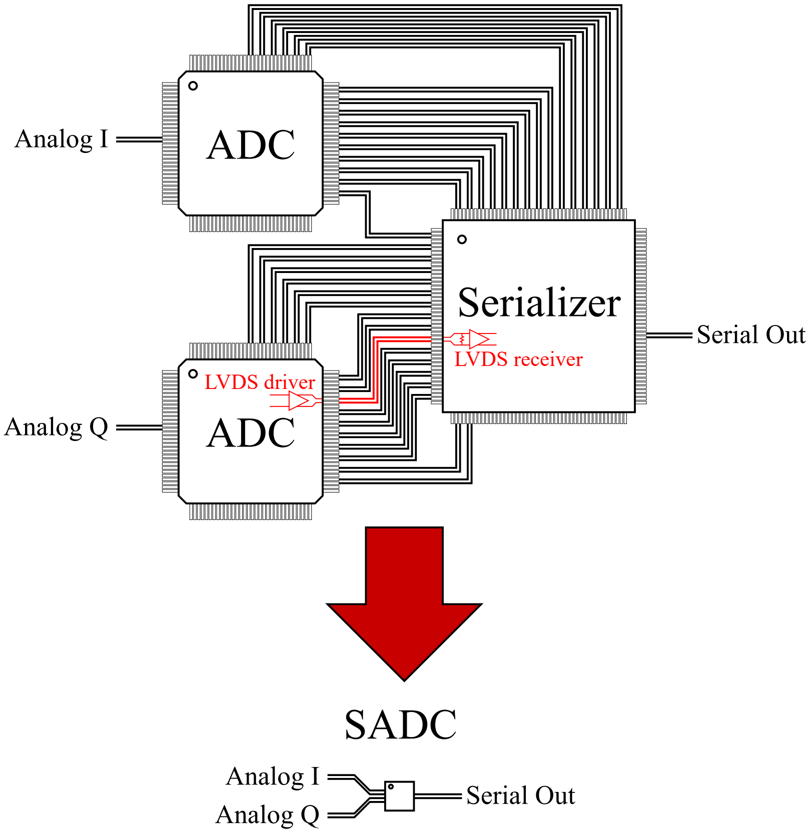

Figure 7: Integration of the sampling and serialization functions in a Serial ADC (SADC) for an order of magnitude improvement in footprint, while saving power in the interface between them.

5 Conclusion

For the reasons that have been described, we have selected for the ngVLA's conceptual design a single-stage direct-to-baseband I/Q downconversion with calibrated numerical sideband-separation, followed by a low-overhead serial-ADC driving an unformatted fiber-optic link. This combination is believed to provide the best performance for cost and manufacturability, in accordance with the system engineering requirements and the general guidelines laid out in Section 1.

References

[1] M. A. Morgan, "An integrated receiver concept for the ngVLA," *ngVLA Memo*, no. 29, November 2017.

[2] S. A. Maas, *Microwave Mixers*, 2nd ed. Boston: Artech House, 1993.

[3] M. A. Morgan, *Principles of RF and Microwave Design.* Norwood, MA: Artech House, 2019.

[4] 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.

[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.