

Versatile 20 GHz Wideband RF Digitizer for Test and Measurement

*Nikhil Adnani, Mathieu Lamanque, Tarek Helaly,
Mohammad Farhan, Tim Hember, and Ian Ward*

Wireless device manufacturers are challenged with the test and measurement of ever-evolving complex signals while at the same time contending with shrinking test equipment budgets. For example, modern wireless signals can have bandwidths up to 80 MHz [1] and can vary in operating frequency and amplitude over time. In contrast, affordable spectrum analyzers typically have narrow (up to 10 MHz) instantaneous bandwidth (IBW). This makes analyzing a time-varying wideband signal that is greater than 10 MHz in bandwidth with such narrowband instrumentation require a *sweep* over several frequency bands. This can yield erroneous measurement results.

To effectively measure such sophisticated waveforms, a wideband vector or real-time signal analyzer is required. These can cost as much as US\$50,000 to \$100,000 each depending on frequency range and IBW. Additionally, these wireless devices and their associated signals have to be tested in the field and/or real-world operating environments. However, it is difficult to field-deploy these legacy analyzers due to factors such as size, weight, and power consumption.

The combination of an RF digitizer and application software that enables the processing and visualization of raw digitized data is referred to as a signal analyzer. Products such as real-time and vector signal analyzers bundle this capability into a single unit of test equipment. A different, modular approach to signal analysis provides RF digitizer hardware within a Peripheral Component Interconnect (PCI) or PCI eXtensions for Instrumentation (PXI) form factor. This hardware interfaces with processing software residing on the host PC. Since the software can be modified to suit any measurement, this approach offers greater flexibility.

This paper introduces the WSA5000, a cost-effective, high performance RF digitizer developed by ThinkRF that enables the acquisition, measurement and processing of

RF signals in a variety of operating environments. The WSA5000 has taken an approach similar to the aforementioned software-defined instrumentation with one major difference: the RF digitizer resides in a stand-alone, small enclosure. This form-factor provides the ability to easily integrate it within a number of different test and measurement applications. For instance, it can be embedded within a specialized test module designed to perform third-order intercept point (IP3) measurements. Alternatively, it can also be interfaced with PXI or PCI-based wideband digitizers to enable measurements and record data over even wider bandwidths, those greater than 100 MHz.

The WSA5000 is a small (24cm x 16.5cm x 3cm) RF digitizer that includes a radio receiver front-end, a wideband digitizer, real-time processing hardware, and a network interface. It operates from 100 kHz to 20 GHz and has a wide IBW of 100 MHz. PyRF is an open source application development framework implemented in the Python programming language that is currently used in conjunction with the WSA5000 to enable a wide range of applications including spectrum analysis, signal demodulation and a variety of waveform measurements. The WSA5000 can be used in automated test equipment (ATE) setups for measurements of signal strength, bandwidth, intermodulation distortion, and other parameters. Additionally, it can be used to analyze complex modulated waveforms such as Wi-Fi and LTE in a variety of operating environments.

This paper first discusses key considerations to be taken into account for the design of general-purpose signal analyzers. It briefly explains the three main radio receiver front-end architectures. The paper then describes the architecture of the WSA5000 which is a reconfigurable software-defined radio receiver. Finally, it describes the software architecture of the platform including PyRF, a Python-based open application development framework. An application is also illustrated with a brief example.

“This paper was presented at AUTOTESTCON 2013 (© 2013 IEEE, in *Proc. IEEE AUTOTESTCON 2013*, used with permission), [19].) It has been slightly edited and formatted for the *IEEE I&M Magazine*.”

Key Considerations for the Design of General-Purpose Signal Analyzers

This section provides an overview of key considerations and related specifications for the design of general-purpose spectrum or signal analysis hardware.

Frequency Range

A large percentage of commercial waveforms are limited to 6 GHz, and therefore, this is a reasonable upper limit for most wireless test and measurement applications. There is, however, a number of applications such as microwave backhaul systems operating at higher frequencies. Therefore, providing a means to extend the frequency range of the signal analyzer would be an important consideration.

Instantaneous Bandwidth

The bandwidths of commercial wireless systems have been increasing steadily to accommodate the growing demand for larger data rates from a growing user base. For instance, devices on the market today that use the new 802.11ac standard utilize waveforms that occupy up to 80 MHz in bandwidth. Furthermore, the Third Generation Partnership Project (3GPP) introduced its Long Term Evolution (LTE)–Release 10, namely LTE-Advanced, which outsteps the bandwidth requirements for fourth generation (4G) mobile systems as it aims to utilize bandwidths of up to 100 MHz [2], [3]. The instantaneous bandwidth for an analyzer, therefore, should be at least 100 MHz if not higher to account for the bandwidth requirements of upcoming standards.

Acquisition Memory and Trigger Capability

Modern waveforms such as those associated with the wireless LAN standards utilize packet-based signaling techniques. Therefore, it would be essential to allow for enough memory to enable real-time capture of multiple data packets. In addition, both time and Fast Fourier Transform (FFT) based trigger capabilities to initiate the capture of relevant data would be useful.

Scan Speed

In production it is often required to test an RF product at multiple different frequencies. The speed at which such tests can be conducted is an important consideration. While a wide IBW is important in this regard, the scan or sweep speed of an analyzer is also an important parameter. This speed determines how fast the analyzer can jump from analyzing one set of frequencies separated at most by the IBW to another set.

Dynamic Range

The dynamic range of a spectrum analyzer depends on a number of factors including the resolution of the analog-to-digital converter (ADC). ADCs with sampling rates in excess of 100 MS/s required for wideband signal processing typically have dynamic ranges between 70 to 80 dB. This is consistent with an effective 12 to 16 bits of resolution. However, RF

measurements for characterizing IP3 generally require a dynamic range of around 100 dB. A means for enabling both wideband measurements as well as high-sensitivity RF measurements is therefore desirable.

Interfaces

While USB remains one of the most ubiquitous and simple to implement interfaces, its main application is to enable a point-to-point link. Fast Ethernet connectivity is more desirable as it enables network connectivity, which is important in ATE setups that need to be controlled remotely, in addition to alternatively supporting a point-to-point link.

Size, Weight, and Power

So that the signal analyzer is field-deployable, easy to integrate within larger setups and suitable for portability in applications such as drive testing, it is important that its dimensions are small. Typical lab spectrum analyzers weigh around 22 kg and consume about 500 W of power, which makes them unsuitable for field applications. To be consistent with the size and power specifications of a typical business laptop, a field-deployable analyzer should weigh less than 2 kg and consume less than 50 W.

Cost

Finally, the ever-increasing demand for low-cost mobile devices is driving the need for efficiencies in the costs of production. This necessitates lowering the costs of test equipment in several stages of the research, development and production cycle. Alternatives to signal analyzers that cost in excess of US\$50,000 are therefore highly desirable.

Receiver architectures

This section provides an overview of three main receiver front-end architectures that are used in modern signal or spectrum analysis hardware: super-heterodyne, homodyne, and direct digitization (Fig. 1).

Super-Heterodyne Receivers

In super-heterodyne receivers, the incoming RF signal is translated down to a lower intermediate frequency (IF) where it is filtered and subsequently digitized. A pre-select filter bank placed right after the RF input is used to reject strong out-of-band signals that can interfere with the signal of interest. This is followed by a low noise amplifier (LNA) to reduce the overall noise figure of the receiver. Right after are filters to reject images prior to mixing the signal with a local oscillator (LO) down to the first IF. A surface acoustic wave (SAW) filter is typically used to filter the IF output from the mixer. The signal is translated down once again to a final lower IF.

The filtered IF signal is subsequently digitized by an ADC and data is transferred to a field programmable gate array (FPGA) and subsequently to a computer for further processing. Digital signal processing (DSP) on the digitized data is done either in the FPGA or on the host computer.

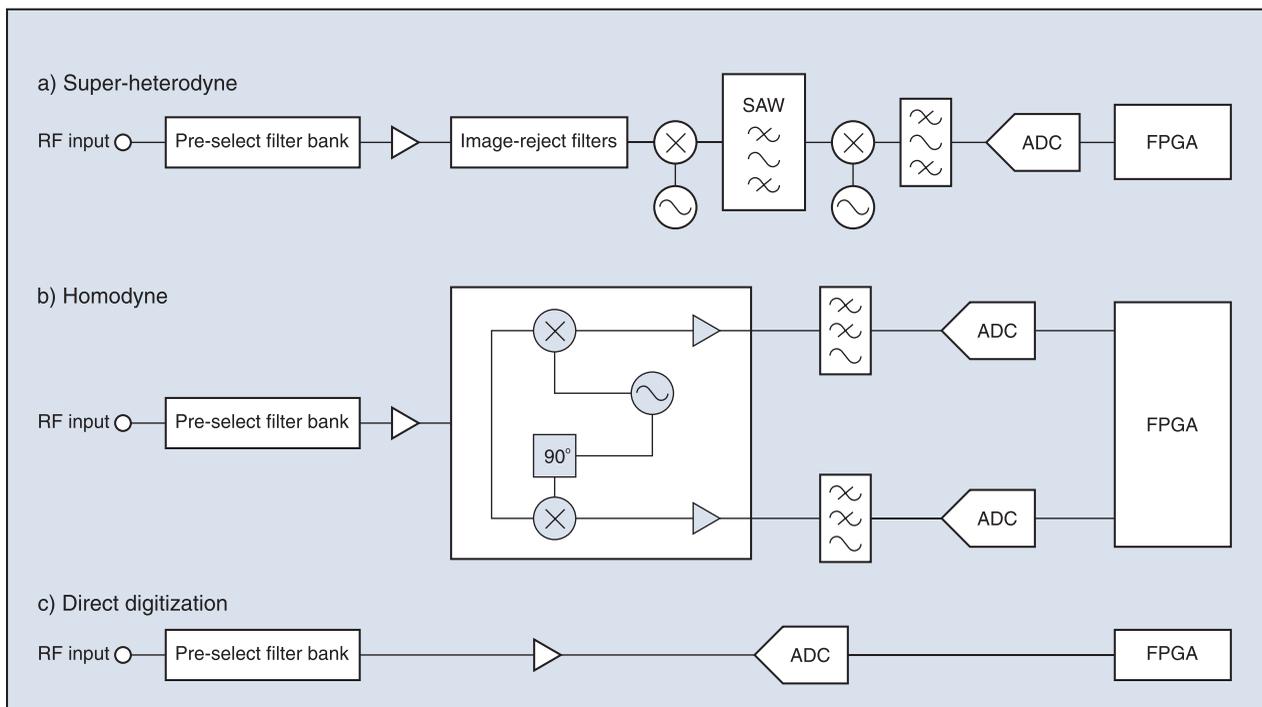


Fig. 1. Receiver architectures.

This basic architecture can be extended to include multiple LNAs and mixing stages to achieve desired system specifications. For instance, a requirement for low receiver noise figure might require multiple cascaded LNAs.

Homodyne Receivers

In homodyne or direct-conversion receivers, the frequency of the LO is the same as that of the incoming RF signal. As a result, unlike the super-heterodyne architecture, the signal is mixed down to an IF of zero. This architecture has several advantages over the super-heterodyne scheme:

- ▶ There is no requirement for image or SAW filters in homodyne receivers.
- ▶ The ADC sampling rate required to process a fixed bandwidth needs to be at least twice as high for the super-heterodyne scheme relative to the homodyne scheme. Post-digitization signal processing circuitry, therefore, needs to operate at higher speeds in super-heterodyne systems.

Because of these reasons, direct-conversion receivers are less expensive. Such receiver architectures have gained in popularity and are commonly used in consumer wireless devices, such as those operating in the 2.4 GHz industrial, scientific, and medical (ISM) band.

The use of such receivers in instrumentation has been limited due to two main drawbacks:

- ▶ Direct Current (DC) offsets: The input signal is mixed with an LO set to the same frequency, and there is always some LO signal that couples to the RF port. This spurious LO signal mixes with itself to result in unwanted signal content at DC. This DC content can mask weaker signals present at the center of the RF band being processed by

the receiver.

- ▶ In-phase and quadrature offsets: The in-phase and quadrature (I/Q) signals at the output of the receiver are never perfectly matched in amplitude and never exactly 90° out of phase. This offset creates signal artifacts in the output spectrum.

There are various techniques to compensate for such offsets that have been documented in open literature [4]–[6]. In general, it is harder to compensate for such offsets over the large frequency ranges that are required for general-purpose instrumentation applications. It is easier to do this over a limited frequency range of operation where the offsets can be characterized and various pilot-based or calibration techniques can be developed and applied.

Direct Digitization Receivers

Direct digitization of the RF input signal yields the simplest and most ideal front-end architecture for software-defined radio systems. A wideband ADC processes RF signals that can vary in frequency up to about half of its sampling rate. All subsequent processing is done either in the FPGA and/or in software. Such ADCs required to process multiple GHz of input signals instantaneously, however, are quite expensive, and their power consumption as well as that of associated circuitry can be high. Direct Digitization architectures have found applications in software-defined HF and VHF receivers.

Architecture of the WSA Family of Receivers

This section provides some background on architectural considerations, available technologies and signal processing requirements for RF digitizers. Our first generation of

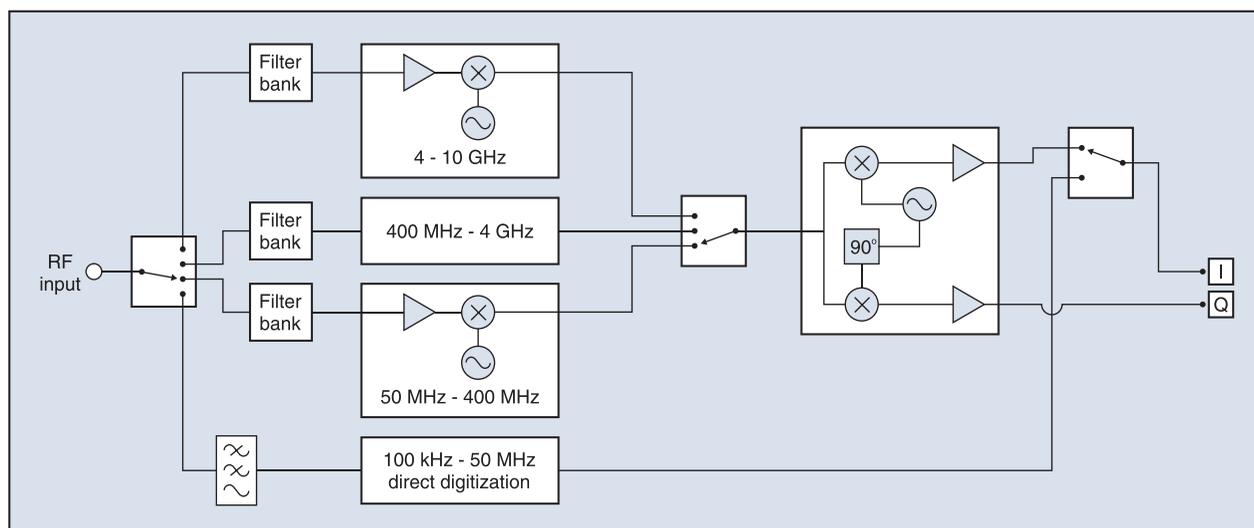


Fig. 2. Architecture of the WSA4000.

products were limited to cellular or ISM bands only and intended for band-specific applications. With the second generation WSA4000 we addressed the problem of monitoring up to 10 GHz of the radio spectrum (Fig. 2). The WSA5000, which represents the third and present generation product, extends the capabilities of the WSA4000 to monitor up to 20 GHz in addition to providing other enhancements for test and measurement applications. This section explains various system considerations in the development of our receiver platforms.

Our goal with the WSA family of receivers has been to build a common platform that could be used for a variety of signal processing applications. The primary design goals were to achieve the following key specifications at an affordable price point:

- ▶ Large operating frequency range: this requirement was primarily driven by the objective to design a general purpose instrument for the analysis of most commercial wireless signals.
- ▶ Wide instantaneous bandwidth of at least 100 MHz: this requirement was defined in order to facilitate the monitoring, detection and analysis of all signals in the popular cellular and ISM bands, and enable fast sweep rates for greater measurement speeds.
- ▶ Form factor: a small device footprint was desirable in order to facilitate integration.

Architectural Considerations for the WSA4000

Our second generation receiver, the WSA4000, was based on the direct-conversion radio architecture (Fig. 2). This enabled us to achieve the aforementioned design goals of wideband signal processing over a large frequency range at a compelling price-performance ratio. In designing the WSA4000, we had to overcome the constraints of direct-conversion technology such as I/Q and DC offsets. I/Q offset correction was done in software while the DC offset was corrected partially in hardware.

Other technological issues needed resolution as well. In Fig. 2, the I/Q mixer (also known as a quadrature demodulator) typically restricts the operating frequency range of the receiver as it is difficult to maintain a reasonable level of I/Q phase mismatch over large frequency ranges. Overcoming this limitation to process signals outside of those limited frequency ranges required a frequency conversion stage prior to the I/Q mixer. This was done to shift the frequency of the incoming RF signal to within the limits of the I/Q mixer. As a result, the WSA4000, unlike a conventional homodyne receiver, incorporated two IFs, the first at 1250 MHz and the second at 0 MHz. The WSA4000 uses both up-conversion and down-conversion blocks to operate over the frequency range of 100 kHz to 10 GHz as shown in Fig. 2.

The instantaneous bandwidth of the final system was 100 MHz, which enabled concurrent processing of the entire 2.4 to 2.5 GHz ISM band. This receiver architecture enabled vector signal analysis at an attractive price point. One of the main drawbacks of the direct-conversion architecture for real-time processing tasks such as signal demodulation was the complexity of I/Q and DC offset correction. This processing, when done in software, introduced latencies in signal processing which in turn impacted the ability to stream and demodulate signals in real-time.

As mentioned earlier, the super-heterodyne architecture produces a non-zero IF at its final stage and therefore does not have DC or I/Q offsets. As a result, it is better suited to real-time signal demodulation. The final IBW using the same component technology (ADC sampling rate) would however be lower than when the direct-conversion architecture is used.

Architectural considerations for the WSA5000

With the WSA5000 the goal was to create a versatile signal analysis platform which could offer the best features of both receiver architectures (Fig. 3). More specifically, it was designed to retain the wide 100 MHz instantaneous bandwidth with the flexibility to switch into an 'offset-free' super-heterodyne

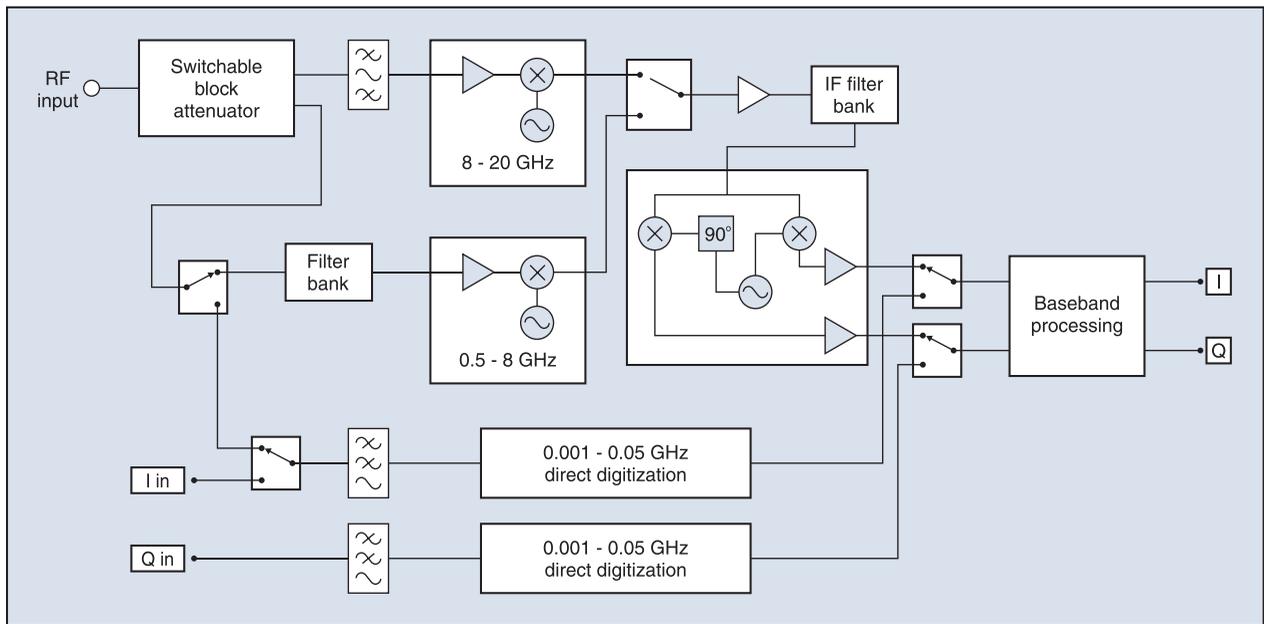


Fig. 3. Architecture of the WSA5000.

mode of operation. Other important enhancements over the previous generation were:

- ▶ Frequency range: The extension of the upper end of the operating frequency range from 10 GHz to 20 GHz was important to keep up with the evolution of new wireless devices that utilize these frequencies.
- ▶ Dynamic range: Enabling a system dynamic range in excess of 100 dB – a key requirement for measuring the IP3 of various semiconductor devices such as LNAs and mixers.
- ▶ Providing the ability to test baseband signals and also to interface with very wideband (> 100 MS/s) digitizers and data recording systems.
- ▶ Spectrum scan rate of 100 GHz/s.
- ▶ Smaller form factor and low power consumption: Fulfilling this requirement was instrumental toward ensuring that the receiver be used within a number of form-factors i.e. in a rack or in a handheld. The dimensions of the WSA5000 are 24cm x 16.5cm x 3cm, and the DC power consumption has been measured to be less than 18 W.

Fig. 3 shows a block diagram of the receiver front end within the WSA5000. The architecture consists of a super-heterodyne front-end with a backend that utilizes an I/Q mixer similar to that in a direct-conversion receiver.

Depending on the frequency of the signals being analyzed, one of three receiver signal processing paths is selected. Signals in the frequency range 100 kHz to 50 MHz are directly digitized, while all other signals are translated to the frequencies of the first IF block via one of the other two signal processing paths. The IF block consists of a bank of multiple surface acoustic wave (SAW) filters. Any one SAW filter in this bank is selected depending on the frequency of the input signal. The output of the SAW filter feeds the I/Q mixer. Depending on the mode of operation, i.e. super-heterodyne or

homodyne, either one or both outputs are utilized to process either 40 MHz or 100 MHz instantaneously.

The IF analog outputs labeled I and Q in Fig. 3 are digitized within the WSA5000 using one of two ADCs with the following specifications:

- ▶ 125 MS/s sampling rate with a dynamic range of between 60 and 70 dB.
- ▶ 300 kS/s with a dynamic range in excess of 100 dB.

Either of the above two ADCs can be selected depending on the desired dynamic response, which in turn depends on the nature of the RF measurement being conducted. Alternatively, I and Q outputs in Fig. 3 can be switched into an external higher performance digitizer, for example one that would be able to sample at 400 MS/s [7].

WSA5000 Signal Processing Hardware

After the ADC, the signal is processed in an FPGA. The different processing components are shown in Fig. 4. Embedded within the FPGA are data processing logic blocks that include a MicroBlaze microprocessor [8], Gigabit Ethernet (Gig-E) interface, a digital signal processor, and logic for data capture control. Implementation within a large FPGA enables the capability of remote device maintenance, updates, and customization while maximizing integration and minimizing size, weight, power, and cost.

Digitized data from the RF front end can be processed in either the frequency or time domain. This is done prior to storage and subsequent transfer to the host across the Gig-E interface. The different processing blocks are described in the remainder of this section.

When analyzing a signal that has bandwidth smaller than the IBW, the embedded digital down-converter (DDC) [9] processing blocks can be used. The DDC first down-converts the signal of interest and then decimates the sampling rate to one

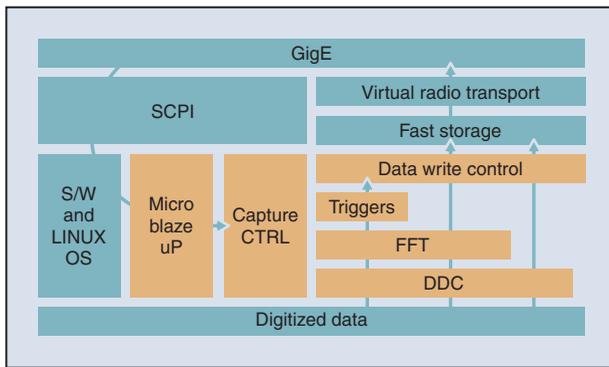


Fig. 4. Data processing in the FPGA.

that is lower and consistent with the bandwidth of the signal of interest. The lower sampling rate enables lower rate processing within the FPGA and also reduces the throughput requirements across the network interface. The DDC has two major elements: A direct digital synthesizer (DDS) that generates a complex sinusoid which is mixed with the I/Q input using a complex multiplier to shift or offset the signal spectrum from the selected carrier frequency. The complex multiplication is followed by a cascaded integrator-comb (CIC) filter. The CIC filter has a *droop* associated with it in the passband. To compensate for this droop, the CIC filter is followed by a compensating finite impulse response (FIR) filter. The filter block is followed by a decimator. This process effectively reduces the sample rate and filters the signal to remove adjacent channels, minimize aliasing, and maximize the received signal-to-noise ratio (SNR). Digitized time domain data can be then either stored in memory on the WSA5000 and/or transferred across the Gig-E interface.

An embedded FFT included in the FPGA fabric can be used to provide frequency domain output data. This alleviates host-side FFT processing. Additionally, performing a logarithm on the output of the FFT within the FPGA itself can compress the data further and enables more efficient data transfer across the Gig-E interface. The WSA5000 has an embedded data capture engine that enables the user to define and execute sophisticated triggers, traces and sweeps. A trace and a sweep are defined as a single capture and a series of captures, respectively. Traces and sweeps are controlled by the capture controller, as illustrated in Fig. 4. The capture controller configures the RF front-end, DDC, triggers, time stamping, and data output for traces and sweeps.

Hardware-based triggering provides a means of qualifying the storage of captured time domain I/Q data samples based on external, periodic or frequency domain events, and thereby selecting signals of interest for the purposes of subsequent visualization and/or analysis. The qualification of triggers is performed at the rate of 125 MS/s, thereby ensuring that events are reliably captured to within an 8 ns window of uncertainty in the time domain. In the frequency domain, using a resolution bandwidth of 122 kHz, a signal has to persist for at least 8.2 μ s for 100% probability of detection.

The WSA5000 has up to 512 MB of fast dedicated data storage. This provides the ability to buffer, store and forward up to one second of contiguous time domain I/Q data sampled at 125 MS/s. Lower sampling rates consistent with signals having narrower bandwidth enable longer time duration signal recordings.

High-speed TCP/IP data connectivity is provided over the Gig-E interface. Network connectivity enables remote operation of the WSA5000 in a test rack or over a simple point-to-point link. The Gig-E link has a maximum theoretical throughput of 1000 Mb/s, although practical throughputs have been reported to be limited to 600 Mb/s [10]. The sampling rate within the WSA5000 is 125 MS/s where each sample is comprised of 32 bits. It is impossible to sustain this throughput rate of $125 \times 32 = 4000$ Mb/s over a Gig-E link without data loss.

Depending on the application, there are four different techniques, any combination or all of which can be used to alleviate this throughput bottleneck:

- ▶ DDC, used to reduce the effective sampling rate to stream a narrower band of signal content
- ▶ Hardware triggering in either the time or frequency domain to store relevant portions of the digitized data, which then results in a lower throughput requirement
- ▶ Capturing blocks of data asynchronously for signal measurement applications that do not have a time-varying aspect to them
- ▶ Using a digitizer external to the WSA5000 that has a faster interface such as PCI Express (PCIe) into the host processor

The WSA5000 has an embedded Linux OS and embedded application to enable remote control and management of the WSA5000. A console is accessible via both a local USB console port and remotely over the TCP/IP Ethernet interface. The console provides remote management, updates and customization whereby new versions of the entire FPGA and embedded software system can be remotely uploaded and installed to the WSA5000.

WSA5000 External Ports

The WSA5000 provides a general-purpose input/output (GPIO) port for both trigger sources and hardware expansion to exterior modules such as antenna switches, down converters, GPS, and power-over-Ethernet (POE) sources.

For the purposes of synchronization, a 10 MHz input is provided. A 10 MHz output is also provided from the WSA5000 to enable multi-unit synchronization. The benefit in providing the latter is that the user does not have to create or source external clock distribution circuits or devices.

As shown in Fig. 3, the front end includes one RF input to process signals between 100 kHz and 20 GHz. Two identical baseband inputs labeled I_{in} and Q_{in} are also provided for processing and testing the complex baseband signal output from a radio transmitter. Finally, I and Q output ports are included that can be used to mate the RF front end with an external, high speed digitizer.

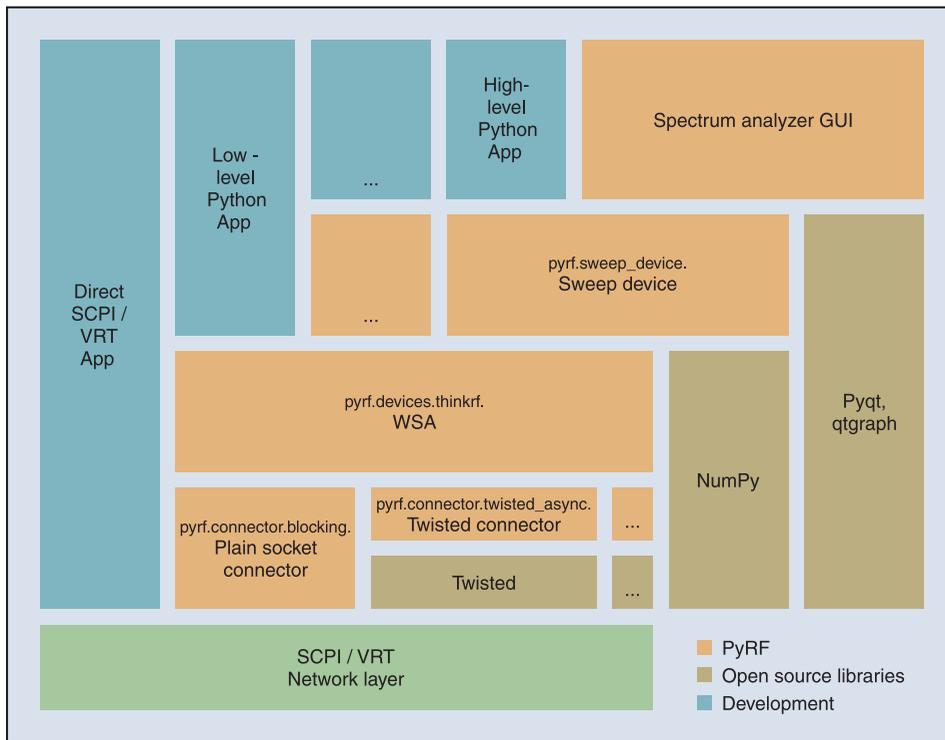


Fig. 5. Software architecture.

Software

This section describes the host-side software architecture that interfaces with the WSA5000 and includes control and communication and application software.

Control and Communication

Communication protocols for the purposes of control and acquiring data from the WSA5000 across the network interface were designed to conform to widely adopted standards. Compliance with standard protocols is critical to the end-user because it provides both device independence and interoperability. One such standard is Standard Commands for Programmable Instruments (SCPI) [11], a language that has been adopted by most test and measurement equipment vendors for controlling programmable instruments. SCPI commands are ASCII textual strings, which are sent to the instrument over the Gig-E interface. Commands can either set a parameter such as the frequency to which the receiver is to tune or query a parameter, e.g., front-end attenuator setting.

Data is returned from the WSA5000 using the VITA-49 Radio Transport (VRT) protocol [12] for digitized data. This is an interconnect standard for transporting digitized data in a link-agnostic format. VRT was primarily intended for communication signal processing applications such as beam-forming, direction-finding and signals intelligence (SIGINT). It includes support for streaming data, time-stamping and other similar features [13] to support tasks such as signal demodulation and fast scans. The protocol is therefore ideally suited for use with the WSA5000 hardware.

Application Software

Traditionally, commercial numerical computing and/or lab equipment control software have been used to process raw data acquired from hardware instrumentation. One of the benefits in using such software is the availability of comprehensive signal processing libraries. However, such software is very expensive with single-user licenses costing thousands of dollars.

If an application developer were to distribute a standalone executable, then typically more software is required at additional cost. In addition, from a usability standpoint, the final outcome is not as flexible as an application developed

from scratch in a general-purpose programming language such as C.

When deciding on an application development environment for the WSA5000, it was decided to adopt a language that:

- is easy to use and develop with,
- has a rich environment of diverse and comprehensive libraries, and
- is widely used by the scientific community and affordable.

The Python programming language was chosen for the application development environment as it satisfied all of the above requirements. Python is a high-level programming language, which makes it easy to develop with. Its design philosophy places an emphasis on code readability and programmers can exploit its syntax to express concepts more efficiently than would be possible in languages such as C. As well, libraries such as NumPy and matplotlib have made Python easy to learn for users of other numerical computing languages. In general, Python has gained widespread acceptance, and its adoption within the scientific community continues to increase. Finally, Python is available as open source software *PyRF framework*.

The source code for the Python-based application framework we call PyRF is released under a flexible open-source BSD license [14], [15]. Implementation examples and the source code are made available at the PyRF.org website.

Software Applications

This section explains the application architecture and provides an example to illustrate the use of PyRF with the WSA5000 for

purposes of signal processing. Fig. 5 shows a block diagram with various components representing the software architecture associated with the host-side processing of data acquired from hardware. Software applications can interface directly with the WSA5000 using SCPI with raw data retrieved in VRT format. Using this technique, control software written for any spectrum analyzer that utilizes SCPI can be ported easily. In this case, connections to the WSA5000 device are made using TCP across the Gig-E interface.

Insert Fig. 5

Users can write either low-level or high-level Python applications. User applications that need to exploit the full features of the device can use the *WSA* class. Within this class, commands to control and retrieve data from the device are implemented as Python methods. An example of a low-level application within PyRF is an application that captures a stream of 30 million data samples over a time duration of 250 ms at a center frequency of 1.9 GHz. Alternatively, users can author high-level applications. An example of this would be one that calculates peak power over the frequency range 100 kHz to 20 GHz using a resolution bandwidth of 50 kHz. Details of the sweep are taken care of by the *SweepDevice* Python class. Additional specialized *devices* can be authored for various applications. One such application is frequency threshold triggering where FFT data is recorded into memory every time the amplitude of any signal across two user-specified frequency limits exceeds an amplitude threshold.

The *WSA* Python class interfaces with SCPI through a connector. Two device connectors have been authored at the time of writing this paper, and users can write their own. *PlainSocketConnector* is the default connector for the *WSA5000* class. The term socket refers to an endpoint of a two-way communication link across a network. This connector provides the simplest form of connectivity to the device, and any *method calls* using *PlainSocketConnector* have to wait for a response from the device before returning to the calling function. In other words, no other processing can be performed in the same thread until a response is received.

An asynchronous connector called *TwistedConnector* has also been written based on the Twisted [16] framework. Using this connector, *WSA5000 method calls* return deferred objects instead of blocking as is the case in the default connector. New connector classes can be authored by the user to support other asynchronous Python frameworks such as Tornado [17] or Tulip [18].

The following lines of sample Python code are used to determine the strongest signal across a specified frequency range. It is an example of a high-level Python application that works with *SweepDevice*.

```
# connect to a WSA5000 with a static IP address
dut = WSA()
dut.connect('10.126.110.112')
# perform a sweep and collect amplitude data across the frequency range 2.1 - 2.9 GHz with a resolution bandwidth of 100 kHz
```

```
sweep = SweepDevice(dut)
fstart, fstop, bins = ( sweep.capture_power_spectrum(2.1e9,
2.9e9, 100e3))
# show the peak amplitude
print max(bins)
```

A sample spectrum analyzer application that has been written using PyRF is available as an open-source code example on the PyRF.org website. Released under the open-source BSD license, users may also customize and/or commercialize the spectrum analyzer application for their own purposes.

The spectrum analyzer supports a number of features such as real-time frequency threshold triggering, waveform record and playback and other processing functions. It continues to evolve as new signal processing functionality is added.

Conclusions

Signal analyzers on the market today with instantaneous bandwidths in the range of 100 MHz are very expensive, and there is a growing requirement for versatile, cost-effective platforms. This paper first described key requirements for signal analysis hardware used for processing wideband signals. Various receiver front end architectures used in signal analyzers and the merits and shortcomings of each were then discussed. The system architecture and various design decisions that led to the evolution of the *WSA* family of receivers were described. The architecture of the third generation of receiver hardware, namely the *WSA5000*, was then explained. This RF digitizer presents a novel, cost-effective platform for fulfilling the requirements of testing modern, complex wireless signals. The *WSA5000* can be used in a number of deployment scenarios, and its small form-factor allows it to be embedded within larger ATE systems. Finally, PyRF, the open source Python-based application framework for application development, was described. Using this framework, developers can build and commercialize a variety of signal processing applications for ATE. The combined hardware and software signal analyzer platform satisfies key performance requirements for testing modern wideband signals at a compelling price point.

Acknowledgment

The authors would like to thank Steve Rosinski for reviewing the paper and providing valuable feedback.

References

- [1] "802.11ac: The Fifth Generation of Wi-Fi," Cisco, San Jose, California, Technical White Paper, Aug. 2012. [Online]. Available: http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps11983/white_paper_c11-713103.pdf.
- [2] S. Parkvall, A. Furuskär, and E. Dahlman "Evolution of LTE toward IMT-Advanced," *IEEE Commun. Mag.*, vol. 49, no. 2, pp. 84–91, Feb. 2011. DOI: 10.1109/MCOM.2011.5706315.
- [3] G. Yuan, X. Zhang, W. Wang and Y. Yang, "Carrier aggregation for LTE-advanced mobile communication systems," *IEEE Commun. Mag.*, vol. 48, no. 2, pp. 88–93, Feb. 2010. DOI: 10.1109/MCOM.2010.5402669.

- [4] J. Tubbx, B. Come, L. Van der Perre, L. Deneire, S. Donnay and M. Engels, "Compensation of IQ imbalance in OFDM systems," in *Proc. IEEE International Conference on Communications*, vol. 5, pp. 3403-3407, May 2003.
- [5] J. Yan, Y. Zheng and Y. Xu, "A novel DC-offset cancelling circuit for DCR," in *Proc. IEEE International Symp. Circuits and Syst.*, vol. 1, pp. 396-399, May 2005.
- [6] S. Sampei and K. Feher, "Adaptive DC-offset compensation algorithm for burst mode operated direct conversion receivers," in *Proc. IEEE Veh. Technol. Conf.*, vol. 42, pp. 581-588, Nov. 1993.
- [7] Signatec PX14400D product information sheet. [Online]. Available: <http://www.signatec.com/products/pdf/Signatec-PCIe-Digitizer-PX14400D-Data-Sheet.pdf>.
- [8] Microblaze soft processor core, Xilinx All Programmable. [Online]. Available: <http://www.xilinx.com/tools/microblaze.htm>.
- [9] Richard Lyons, *Understanding Digital Signal Processing*, Third Edition, New York, NY: Prentice Hall, 2011.
- [10] Jeff Atwood, "Gigabit ethernet and back of the envelope calculations," Coding Horror. [Online]. Available: <http://www.codinghorror.com/blog/2005/07/gigabit-ethernet-and-back-of-the-envelope-calculations.html>.
- [11] SCPI specifications. [Online]. Available: <http://www.ivifoundation.org/specifications/default.aspx>.
- [12] VITA Radio Transport (VRT) Draft Standard VITA-49.0-2007. Draft 0.21, 31 Oct. 2007.
- [13] Stephen M. Pereira, "Standardizing digital IF data transfer with VITA-49," *RTC Magazine*, January 2006.
- [14] BSD license overview. [Online]. Available: http://en.wikipedia.org/wiki/BSD_licenses.
- [15] BSD 3-clause license. [Online]. Available: <http://opensource.org/licenses/BSD-3-Clause>.
- [16] Twisted Matrix Labs. [Online]. Available: <http://twistedmatrix.com/trac/wiki/Documentation>.
- [17] Tornado documentation. [Online]. Available: <http://www.tornadoweb.org/en/stable/documentation.html>.
- [18] Tulip documentation. [Online]. Available: <http://www.python.org/dev/peps/pep-3156/>.
- [19] N. Adnani, M. Lamanque, T. Helaly, M. Farhan, T. Hember, and I. Ward, "Versatile 20 GHz wideband RF digitizer for test and measurement," in *Proc. IEEE AUTOTESTCON 2013*, pp. 1-10, 2013. DOI: 10.1109/AUTEST.2013.6645044.

Nikhil Adnani (info@thinkrf.com) is currently CEO of ThinkRF Corp. Prior to this, he held engineering positions at Nortel and Communications Research Centre (CRC). While at CRC, he worked on wideband wireless technologies and

frequency agile transceivers. Over the past seven years at ThinkRF, he has worked on wideband radio receiver technologies for spectrum and signal analysis applications. Nikhil has a BSc and an MSc from the University of Manitoba and a PhD from Carleton University, in Electrical Engineering.

Mathieu Lamanque is a graduate of the Advertising and Marketing Communications Management program at Algonquin College and is a product marketer at ThinkRF. He has helped shape communications and branding for the company and has designed the visual assets used to describe ThinkRF's innovative devices. His interests include cutting edge wireless technologies and new computing platforms.

Tarek Helaly received BSc and MSc degrees in electrical engineering from the MTC, Cairo, Egypt in 1996 and 2001, respectively. In 2010, he received the PhD in electrical engineering from Carleton University, Ottawa, Canada. Since 1996, he has held several research, academia, and engineering posts in several entities, mainly in R&D. Currently, he is a senior algorithmic engineer at ThinkRF in Ottawa, Canada. His research interests are in wireless communications, signal analysis, multiple access and multi-carrier systems, and LTE.

Mohammad Farhan received his BEng degree in Biomedical and Electrical Engineering from Carleton University, Ottawa, Canada in 2012. Since then, he has been an Applications Engineer at ThinkRF. His research interests include medical signal processing and scientific data analysis in Python.

Tim Hember was prior VP and Strategy Director for Alcatel Canada, developing global corporate market strategies, and prior CEO and founder of TimeStep, a network security equipment market leader acquired by Newbridge/Alcatel. Tim has thirty years of experience including Newbridge Networks, Gandalf, Calmos Semiconductor, Genesis Microchip, Pure-Data Research, and Nortel Networks. Tim has a BSc in EE from University Waterloo and an MA in Economics from University Ottawa.

Ian Ward has 10+ years of Python and Open Source development experience. He works on project architecture, visualizations, and technical documentation. Outside ThinkRF he is a core developer on the CKAN data portal project and helped launch the Government of Canada's open data portal. He is a co-founder of Boxkite Inc. and his project updates are posted on excess.org.