Microwaves&RF

KAUE MORCELLES, Independent Technical Writer,
Per Vices Corp., pervices.com

Monitor and Record More Data than Ever with SDRs

Software-defined radio can be used to increase the amount of data that can be monitored and recorded, with minimal effect on latency.

As the number of wireless devices, IoT systems, and 5G networks continues to escalate, the radio-frequency (RF) environment is becoming more challenging to monitor and manage, especially considering the higher frequencies and complex communication protocols involved in modern radio. In this scenario, spectrum monitoring and recording devices play an essential role in the analysis and control of the RF content at a certain location, which is highly beneficial for several applications.

For instance, these devices can help mobile service-providers optimize the use of the spectrum, avoiding channel saturation and improving decision-making. Furthermore, spectrum monitoring is essential in the military industry, being a crucial operation in signal intelligence (SIGINT) and electronic warfare (EW), helping to detect, locate, and identify malicious signals in challenging environments.

To properly characterize the spectrum, the monitoring device must acquire and process a huge amount of data, especially considering the high bandwidths involved in these applications. Thus, spectrum-monitoring systems must constantly evolve to keep up with the rapid development of the technological landscape, and such flexibility can't be achieved with traditional hardware-based approaches.

Software-defined radios (SDRs) are digitally based RF transceivers that can help solve the data throughput limitations in spectrum-monitoring devices. They introduce embedded ultra-low-latency computation units that can process multiple channels simultaneously and powerful

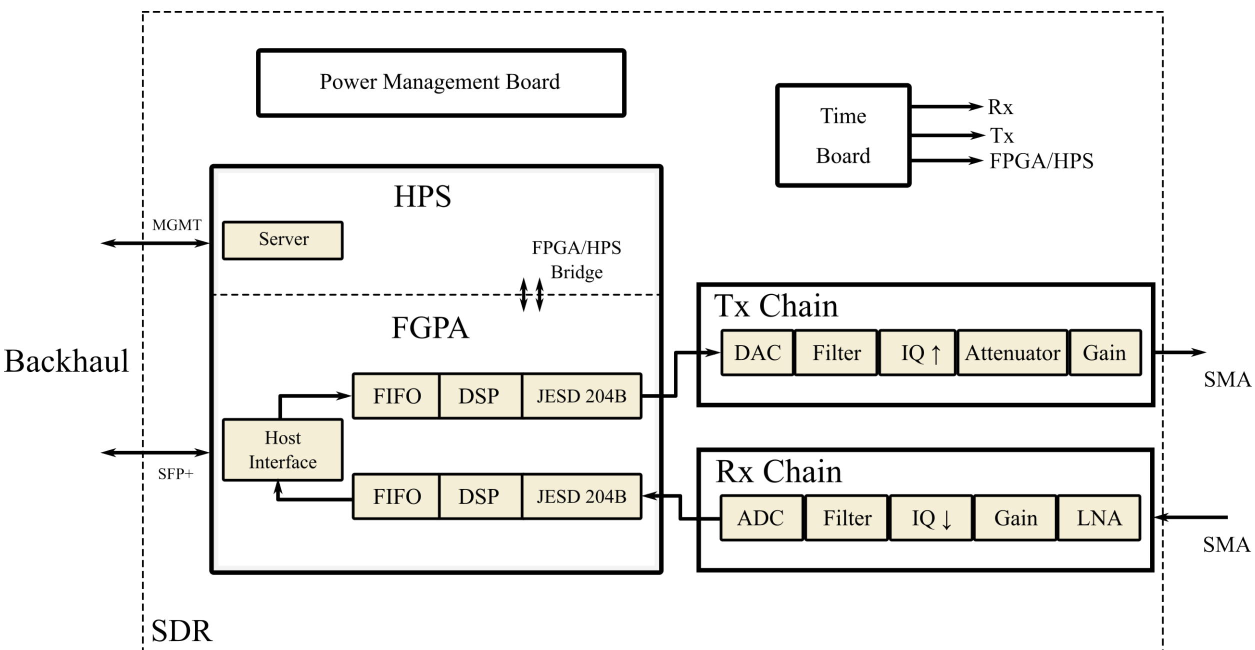

Shown is a high-level diagram of an SDR board.

high-speed host communication systems.

This article discusses how the use of SDRs in spectrum-monitoring systems can significantly increase the amount of data that's monitored and recorded without drastically affecting the latency, which would be expected in a traditional setup based on analog systems. It first presents the basic concepts of SDRs, covering the essential technical aspects of both the digital backend and the radio front-end (RFE).

Second, the fundamentals of spectrum monitoring and recording are introduced, including working principle, technical requirements, and some applications. Finally, the article delves into how SDRs can be implemented in each spectrum-monitoring application, highlighting the main advantages obtained by integrating these transceivers in such systems.

What are SDRs?

Software-defined radios (SDRs), versatile and flexible RF transceivers that implement most of the radio and signal processing functionality on the software side, are composed of a radio front-end (RFE) and digital backend. The RFE performs all of the essential analog operations for the receive (Rx) and transmit (Tx) functions, including filtering, mixing, signal amplification, noise suppression, and antenna coupling.

The RFE features can significantly limit or enhance the overall SDR performance, which is why high-end SDRs implement powerful radio Tx and Rx chains. The highest-bandwidth SDR in the market can achieve tuning ranges from 0 to 18 GHz and up to 3 GHz of instantaneous bandwidth per channel.

In multiple-input, multiple-output (MIMO) solutions, each channel is able to monitor a slice of the spectrum, significantly improving the overall coverage. MIMO SDR provides completely parallel operation by implementing independent analog-to-digital and digital-to-analog converters (ADCs/DACs) on each channel *(see figure)*, meaning that the spectrum slices can be monitored simultaneously.

The digital backend is the brain of the SDR. It's typically implemented by a field-programmable gate array (FPGA) with on-board digital-signal-processing (DSP) capabilities, including basic radio operations like modulation, demodulation, up/downconverting, and data packaging, plus more complex algorithms including sophisticated communication protocols and automated sequences.

By implementing most of the DSP functions in the backend, the whole SDR can be completely reprogrammed without any hardware modification, on-the-fly and remotely. Thus, these devices are perfect for modern spectrum-monitoring endeavors, where the system typically has to respond and adapt quickly to detected signals. The FPGA implementation also allows for powerful parallel computation with minimum latency, ensuring that multiple channels can be monitored simultaneously with minimum data lost.

Finally, high-throughput SDRs provide on-board host communication using qSFP+ Ethernet links, which enables very fast data exchange between the SDR and the host or storage system. This further increases the amount of data that can be captured and stored by combining these high-performance SDRs with properly configured storage solutions, which is essential for spectrum-monitoring applications.

Spectrum Monitoring Specifics

Before diving further into the SDR's role in monitoring and recording solutions, let's take a closer look at spectrum monitoring. Essentially, spectrum monitoring consists of capturing and analyzing RF signals across a wide range of frequencies in the electromagnetic spectrum to acquire information and situational awareness.

By capturing, analyzing, and recording the RF data, applications can obtain critical insights into the RF environment and make better decisions in real-time. This is extremely important for SIGINT, EW, and security applications, where knowledge about the RF landscape is critical.

However, to effectively carry out these tasks, the device must meet several hardware and software requirements, including parallel computation, low-latency processing, wide tuning range, and high instantaneous bandwidth to effectively sweep the frequencies and detect signals of interest. The main bottleneck in spectrum-monitoring technology is the data-processing capacity of the RF transceiver. Large slices of the spectrum must be measured and processed simultaneously, and a huge amount of data needs to be transmitted to the host solution, which usually involves a storage solution.

As a result, the combination of high-performance SDRs and data storage and playback systems that can handle large amounts of information are crucial to cover wide spectrum ranges without losing data in the computational process.

How are SDRs Used for Spectrum Monitoring and Recording?

The overall architecture of modern spectrum-monitoring solutions is composed of two main building blocks: the SDR transceiver and the storage/playback solution. The SDR is responsible for receiving and processing the huge amounts of RF data that compose the spectrum being analyzed. It performs all of the basic and complex radio functions and provides a coherent and consistent stream of processed data to the storage/playback system. That system is responsible for recording the measured data and playback the signals to users when requested.

Naturally, it's essential that these solutions provide high-throughput data streaming (ranging from 160 to 400 Gb/s) and high data capacity (up to 100 TB). This section discusses the RF and the DSP requirements of SDRs, the storage/playback system hardware requirements, and the necessary link between these two blocks to optimize performance.

When selecting or designing an SDR for spectrum monitoring and recording applications, one must keep in mind four parameters when it comes to hardware: bandwidth capability, channel count, signal resolution, and backhaul. The bandwidth capability defines the maximum portion of the spectrum that can be monitored instantaneously, where higher bandwidths result in more data being captured instantaneously.

Channel count is important for similar reasons, as the total spectrum coverage is also proportional to the number of channels, where each channel is responsible for a slice of the spectrum. Furthermore, additional channels can be used for different functions that can be triggered by detecting signals in the spectrum-monitoring section of the SDR, which is useful for EW countermeasures. This requires MIMO SDRs, as they provide the most powerful solutions for spectrum-monitoring systems, with the best models in the market able to support up to 16 independent RF chains.

However, not only is it important to capture as much data as possible, but also to adjust the bandwidth of each radio chain to obtain the optimal sensitivity and dynamic range for a certain portion of the spectrum. Therefore, the device must provide high resolution for both amplitude and frequency and controlled amplification per channel, to ensure enough sensitivity to weak signals and avoid losing data for saturation or spectral leakage.

To achieve that goal, SDRs must implement Rx chains with low-noise-figure and high-dynamic-range amplification chains with variable attenuators, and high-sampling-rate and high-resolution ADCs to interface with the digital backend. Putting all three elements together drives the digital throughput (or backhaul) capacity, as it limits the amount of data that can be transmitted to the host solution.

For example, with eight channels sampling data at 1 Gsample/s in IQ format, each with 16-bit resolution, results in 8 *1 GSPS * 16 bits * 2 = 256 Gb/s of data that needs to be transferred from the SDR to the host/storage solution. This illustrates the need to select SDRs with powerful backhauls. They should have up to 160 to 400 Gb/s of streaming capacity through optical links to enable fast and efficient data transfer between the SDR and the storage/playback solution.

The Software Side

In terms of software, the implementation of the digital backend using an FPGA with on-board DSP capabilities is crucial to achieve the level of flexibility, programmability, and parallel computation power required for spectrum-monitoring solutions. Different from ASIC solutions, the FGPA can be easily reprogrammed on-the-fly and remotely, without any hardware modification. This makes it possible to easily adapt commercial off-the-shelf (COTS) SDRs to fit any specific requirement of the environment or application.

One of the most commonly used DSP techniques in the spectrum-monitoring field is the fast Fourier transform (FFT), which is applied to transform time-domain signals into the frequency domain. This algorithm is crucial for capturing the frequency-dependent properties of a signal, including phase and magnitude, which reduce the overall data being sent to the storage and playback system—in this case, only the information about the harmonics is sent.

In addition, the digital backend must enable the setting of thresholds or limits for "sweep and lock" signal detection. Several open-source and proprietary toolkits are based on the FFT approach, which can significantly reduce the design and deployment time of the backend firmware.

The SDR alone can only make certain that data is being measured and streamed to a host. Therefore, the complete solution architecture must implement some sort of storage and playback system to prevent data loss and ensure enough data throughput to record and transmit the data. To achieve that, it's important to implement a system that can handle the extremely high data-throughput involved in this application.

In those types of scenarios, FPGA-based network interface controllers (NICs) are crucial to provide high-speed communication between the SDR and the data storage and playback system, ensuring lossless data transfer between the units. The storage portion of the system must be able to keep up with the data transfer, meaning NVMe drives become essential because they offer write speeds much higher than SATA drives. The data storage and playback system should also be scalable and have sufficient capacity to store all of the captured data. Furthermore, having redundant storage and backup is critical to make sure that data isn't lost due to hardware failure or other issues.

SDRs: The Optimal Choice for Monitoring-System Transceivers

With all of these requirements in mind, it becomes clear that the proper selection and design of the radio transceiver has a huge impact on the performance of spectrum monitoring and recording solutions. The SDR is the only transceiver architecture in the market that's capable of properly addressing all challenges and limitations involved in these systems, especially in terms of data capture over a wide frequency range.

As mentioned previously, the data bottleneck in spectrum monitoring involves several steps in the signal chain, including measurement, processing, transmission, and storing. Thus, all of these aspects must be properly designed and optimized to improve overall performance. SDRs can be easily integrated with high-end storage/playback solutions, while providing a powerful backhaul that can be linked to NICs. And the data throughput can be easily optimized by means of proper selection of hardware components and proper programming of the FPGA units.

Therefore, SDR-based spectrum-monitoring systems can significantly increase the amount of data that can be captured and processed, allowing for more effective analysis of the electromagnetic environment and, thus, making a huge impact in key RF industries, including mobile communications, military intelligence, and security. Some companies offer both the SDR and host/storage solution, eliminating the need to try and design your own.

Kaue Morcelles is an electrical engineer, with emphasis on electronic design and instrumentation. He currently works with biomedical research, developing instrumentation devices for tissue engineering. Learning and writing about cutting-edge technologies is one of his passions, along with coffee, musical instruments, and weird animals.