RF Analysis/Capture using NVMe Storage
Application Note 104

Classifying RF signals for Homeland Security is now a major challenge for law enforcement. While the signal bandwidth might be low, it could appear anywhere in the spectrum, or in multiple bands simultaneously. Reconfigurable FPGAs can examine the spectrum and classify the incoming signals, but when the signal is not recognised, the operator needs to record the data at full incoming rate for offline analysis. The latest NVMe drives offer very high speed storage/replay via a PCIe GEN3 interface which is compatible with the PCIe interfaces included in the latest FPGAs. This article describes how to make use of this capability using a 3rd party IP core to provide the FPGA with full speed recording and replay via the NVMe interface at over 4GB/s.

Requirements and Architecture

Monitoring of RF signals is crucial to many organisations including mobile telephony operators, government RF licensing bodies, intelligence agencies and military personnel. Common signals such as mobile telephony, Wi-Fi, Bluetooth, need to be recognised and excluded so that the remaining signals can be explored. In this paper we propose how standard components can be combined to create a powerful tool for signal classification, recording and offline analysis.

The system is built using standard COTS cards in a 3U VPX platform. A block diagram of the system is shown in Figure 1. The core of the system is a COTS 3U VPX chassis fitted with a standard OpenVPX 9-slot backplane (BKP3-CEN09-15.2.17-3). This backplane has space for 8 payload cards and one PCIe switch card for system control. The chassis uses dual power supplies to ensure sufficient current to drive the backplane when all the cards are fitted, and has powerful forced air cooling. This style of chassis is not particularly rugged but is ideal for use in a lab or stationary vehicle; the same system components could easily be built into a more rugged system for more challenging environments (moving land vehicle, airborne, etc). The RF digitiser comprises an Apissys AV125 card featuring a Xilinx Ultrascale KU115 processor. This board is controlled over PCIe (via a PCIe switch) using a COTS Single Board Computer (SBC) running Linux. In some applications there may be a requirement for an external programmable filter or maybe an RF downconverter – this could be controlled by the SBC using standard USB or GbE
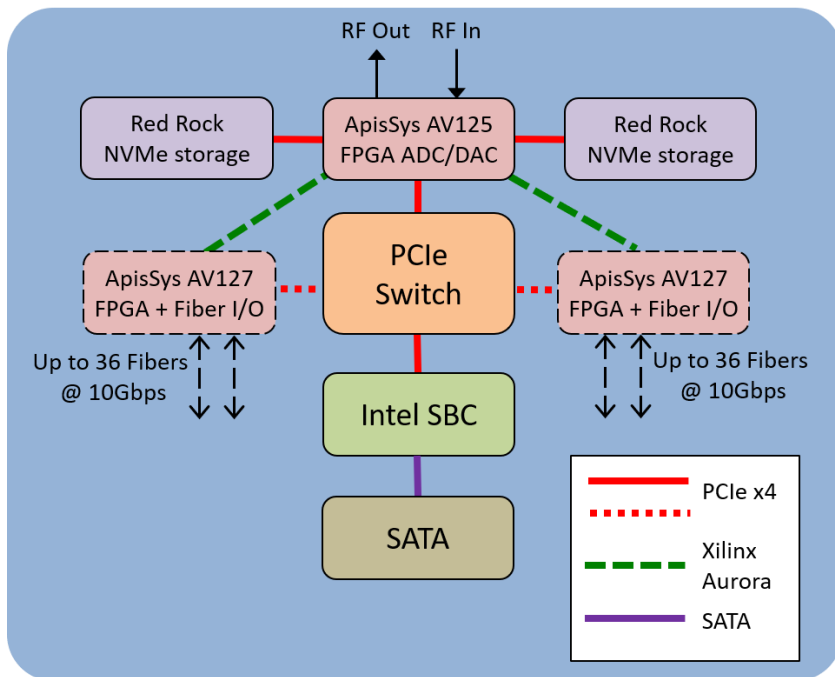

Figure 1 – System Architecture

interfaces. The SBC hosts a small FLASH drive module which is used to hold the OS and also act as a repository for the FPGA firmware images. The FPGA can be loaded with different images, depending on the particular part of the

spectrum being scanned. The SBC can also control additional FPGA cards (E.g. Apissys AV127) in the system which might be installed to add an additional FPGA processing resource. These optional cards could also add extra RF channels or high speed fiber IO for external data streaming. The system uses two NVMe drives from Red Rock Technologies, with removable storage modules. These modules are built using Samsung M2-style NVMe drives which offer the best combination of performance and capacity (up to 1TB per drive). The storage modules can be removed and transferred to a Red Rock Offload Station for data archiving and offline analysis. Each NVMe drive has a direct PCIe x4 GEN3 link to the FPGA card (acting as a PCIe Root Complex), with each link running at around 2GB/s. An off-the-shelf NVMe IP core is used to maximise the speed of data transfer between the FPGA and the NVMe drives.

System Operation

On powerup, the FPGAs are loaded across the PCIe bus with firmware stored on the SBC's SATA drive. Different firmware images can be used depending on the classification requirements. Incoming RF data is sampled at up to 5.4GSPS by the Apissys card and split into two low-latency streams inside the FPGA. One stream is sent directly to the NVMe drives as a permanent record of the RF environment – this raw unprocessed data can be streamed at an aggregate rate of <4GB/s (2GB/s per drive). The other stream is down-converted and then further processed by the FPGA, depending on the firmware image loaded at system start-up and control registers set by the SBC. In applications needing more processing, additional FPGA cards can be installed in the chassis, with direct bidirectional links between the FPGAs. The

Figure 2 – Chassis and cards

FPGA(s) can compare the pre-processed data with a database of known signals to confirm that the signal is recognised. If the processed data matches a known pattern, then this information can be passed back to the SBC and logged. The processed data could also be sent back to the host SBC over the PCIe link. If the data is not recognised then the real time data stream stored on the NVMe drives can be tagged and saved, ready for offline analysis later. Adding a GPS unit to the SBC could allow for time/location stamping on the recorded data. For lower speed applications, an embedded Microblaze core running Petalinux might be sufficient to access the NVMe drive, but for maximum speed this system uses an NMVe IP core inside the Xilinx Ultrascale on the AV125. This core has a standard AXI interfaces and multi-channel DMA which can saturate the NVMe bus to achieve the maximum possible read/write speed.

Conclusion

Combining FPGA processing with high speed NVMe storage creates a powerful tool suitable for many signal processing applications. The NMVe drives offer a large bulk storage system for RF data, recording up to 2GB/s per module. The storage modules (up to 1TB each) can be removed from the capture/processing system and analysed offline. Multiple NVMe drives can be used in parallel to increase performance, and including the COTS NVMe IP core gives the maximum possible NVMe storage bandwidth. Finally, building this system using COTS components, with a COTS chassis and backplane keeps the cost down and simplifies the integration task.