



BlackLynx on Dell: Heterogeneous Signal Processing for SDR

Use Case: Accelerated GNU Radio Waveforms

BlackLynx, February 2019, v1

Overview and Problem Statement

Existing high-performance mission-focused waveform engines typically rely on a set of proprietary signal processing architectures often requiring many racks of expensive, difficult to program and hard to maintain equipment based heavily on isolated CPU, GPU, ASIC and/or FPGA technology. Typical *status quo* deployment environments prefer one type of compute methodology over another. Recent advancements in massively multi-core architectures running software defined radio (SDR) and related signal processing technology allow for a new class of low-cost yet highly adaptive waveform strategies, leveraging best-practices from both open-source and proprietary technologies. Unfortunately, SDR implementations leveraging traditional CPU software techniques do not always perform to levels needed for a variety of mission-critical signal processing needs.

One example of this is the implementation of low-density parity check (LDPC) codes. LDPC is a commonly occurring digital signal processing (DSP) operation that implements forward error correction for transmitting a message over a noisy communications channel. It is heavily used in both commercial and military applications. It is, however, a notoriously computationally intensive algorithm that does not fare well on traditional CPU fabric, which can greatly impact the performance of a larger SDR waveform implemented in popular open source toolkits such as GNU Radio.

This whitepaper benchmarks a GNU Radio LDPC decoder block implemented in BlackLynx FPGA fabric, comparing performance against a traditional CPU implementation. A single Dell commodity 1U Dell R640 server outfitted with a single low-power ½-height, ½-width Xilinx VU9P-powered FPGA PCIe expansion board is used in the benchmark. The results discussed in this whitepaper can also be achieved running on FPGA-accelerated cloud infrastructure, such as the AWS F1 or Nimbix Cloud with Xilinx Alveo Accelerator Cards.

We've chosen to highlight LDPC as opposed to a simpler digital signal processing (DSP) block such as a fast Fourier transform (FFT) since LDPC is significantly more computationally challenging. BlackLynx can also accelerate many other types of compute-intensive DSP blocks, including custom blocks leveraging accelerated SDKs. At BlackLynx, we are built from the ground-up to solve the most difficult analytics challenges at scale, with best-of-class performance at the lowest possible total cost, with the smallest possible form factors and lowest possible power consumption.

BlackLynx Recommends Dell Servers

Although BlackLynx technology is deployable on practically any Intel CPU-based Linux platform server, BlackLynx recognizes Dell as a best-of-breed Linux-friendly hardware server provider. As such, BlackLynx recommends the dual socket (supporting a variety of 64-bit Intel processors) 1U Dell R640, dual socket 1U Dell C4140, or dual socket 2U Dell R740 servers. These recently introduced Dell servers have already been heavily deployed to address some of the world's most complex analytics problems. And, with Dell's world-class supply-chain, OEM, and hardware and software partner ecosystems, IT departments simply can't go wrong leveraging Dell hardware.

The BlackLynx software technology referenced in this whitepaper leverages a **single** commercial off-the-shelf (COTS) Dell R640 1U server in our lab at BlackLynx headquarters in Rockville, MD. Our test server is outfitted with ~755GB of usable



RAM across two 20-core (40-thread) Intel(R) Xeon(R) Gold 6148 CPUs @ 2.40GHz. This provides 40 cores and 80 threads of processing power that the BlackLynx technology stack will effortlessly manage, while running on top of the well-regarded Ubuntu 16.04 LTS Linux operating system on our Dell R640 benchmark system. BlackLynx also deploys natively on other modern Linux operating systems including CentOS 7.6 and Scientific Linux 7.6 (a variant of CentOS) for those customers who prefer CentOS distributions.

A single low-power ½-height, ½-width Xilinx VU9P-powered FPGA PCIe expansion board is added to the Dell server and used in the accelerated benchmark.

Although this test system is outfitted with a rather large amount of usable RAM (~755GB), only a small percentage (on the order of about 300MB) will be used by the operations we will describe and benchmark in this whitepaper. That means that a system with significantly less RAM (low GB) would have achieved the exact same performance. Similarly, although our system has 80 Linux-facing CPU threads, we can reasonably leverage only very few of these given the nature of real-time SDR frameworks, and similar accelerated performance would be achieved on a system with many fewer cores, such as a single-socket quad-core system.

Finally, although our test system in this whitepaper was a standard rack mountable 1U Dell R640, the performance that we will demonstrate can also be obtained with a similarly outfitted Dell C4140 1U server, or a 2U Dell R740 server. The advantage of the Dell C4140 is that it has a wider array of PCIe expansion options (which translates to massive amounts of near-compute storage and/or more acceleration card slots), which could even further improve seamless scale-up acceleration for a variety of complex analytics. There is one distinct disadvantage of the extra expandability afforded by the Dell C4140, though, that being that the Dell C4140 is a rather deep 1U server, which means it may not fit standard-sized server rack infrastructure. The 2U Dell R740 also sports exceptional PCIe expansion slot capability, yet it fits standard racks, but its disadvantage is that it takes up an extra 1U of vertical rack space as compared to the Dell R640 or Dell C4140 servers.

BlackLynx Technology Connected to Mission Architecture

Figure 1 below depicts a high-level view of how a BlackLynx-enabled architecture fits into a larger mission environment. An arbitrary external transceiver topology feeds or is fed by our solution using digitized data streams (such as, perhaps, by simple multicast protocols) across a standard 1GbE, 10GbE, 40GbE or 56Gbps FDR InfiniBand network topology. BlackLynx-enabled COTS hardware, such as a variety of popular Dell servers, simultaneously act as sinks or sources as needed and as defined by mission requirements leveraging a variety of BlackLynx-accelerated SDR tools, such as (but certainly not limited to) GNU Radio.




Figure 1: System Architecture of BlackLynx-Enabled Arbitrary Mission Waveform Analysis



Non-accelerated LDPC Benchmarks

We will measure the performance of the computationally-intensive LDPC operation using our Dell R640 test server, running GNU Radio, first leveraging a CPU-only environment. That is, no BlackLynx acceleration will be employed for our initial benchmark.

We start with a stock/traditional LDPC decoder that a knowledgeable DSP designer might create, either directly in C++, Python, possibly with the assistance of capable tools like MATLAB. Figure 2 below shows a partial consecutive codeword performance profile of a stock LDPC block with maximum noise injection (since LDPC is used precisely because of high-noise environments). The codeword length used was 2560 bits, and the code rate used is 2/5 ($n=2560, k=1024$). The single codeword performance was observed to have significant variance and ranges from about 11 ms to about 20 ms per codeword. Those performance levels and the associated variance are problematic for many real world applications.

0.015619	seconds
0.013243	seconds
0.014035	seconds
0.014638	seconds
0.014249	seconds
0.015671	seconds
0.010976	seconds
0.014509	seconds
0.020245	seconds
0.014139	seconds
0.014089	seconds
0.012519	seconds
0.014859	seconds

Figure 2: Standard CPU-Only LDPC Single Codeword Performance

BlackLynx FPGA-Accelerated LDPC

Many DSP engineers at this point would give up on SDR techniques and migrate to very costly pure hardware solutions that could take up to a year or more to develop. Leveraging BlackLynx technology, which seamlessly integrates FPGA acceleration for targeted waveform blocks, this all-or-nothing hardware approach is no longer required. Users are free to leverage the best of what SDR has to offer, without sacrificing performance.

The BlackLynx FPGA-accelerated LDPC processing block is a drop-in block that can integrate with any GNU Radio flowgraph/waveform requiring an LDPC decode function. There is nothing special about it that the user must be concerned with, except for a few minor configuration parameters which a knowledgeable DSP engineer understands. For example, we will configure our block for one codeword per batch and a maximal number of LDPC iterations – the combination of these configurable parameters present effectively a worst-case scenario for the accelerated benchmark, since it won't batch codewords and uses maximum decode iterations. Figure 3 below shows the block as it might appear in isolation in a flowgraph:




Figure 3: BlackLynx FPGA-Accelerated LDPC GNU Radio Block

The BlackLynx FPGA-accelerated LDPC block is a great example that makes use of the Xilinx SDAccel framework. BlackLynx heavily leverages key aspects of SDAccel to implement its FPGA-accelerated algorithms for a variety of use cases including text analytics, artificial intelligence and machine learning, and signal processing, to include our LDPC decoder. The Xilinx SDAccel framework is a hybrid data movement architecture that allows for seamless integration between control plane, data, and algorithms flowing to and from one or more Xilinx FPGAs across a CPU-managed PCIe bus running a variety of supported Linux-based operating system platforms, including variants of Ubuntu, CentOS and Red Hat, both on-premise and in the cloud. The beauty of the Xilinx SDAccel framework is that it allows rapid development and integration using multiple technologies, including VHDL, Verilog, OpenCL, and where appropriate, even higher level languages like C and C++. It also affords the ability to quickly swap out multiple algorithmic “kernels” so that the architecture can support a multitude of simultaneous use cases. BlackLynx leverages all of these SDAccel strengths to support complex analytic patterns that fuse use cases such as signal processing, text analytics, and ML/AI frameworks all in real-time.

The performance when running the BlackLynx FPGA-accelerated block for a single codeword and maximum iterations, with the same input as the previous unaccelerated benchmark, is nothing short of astounding at *less than half a millisecond* per codeword:

```
0.000403 seconds
0.000361 seconds
0.000412 seconds
0.000360 seconds
0.000369 seconds
0.000377 seconds
0.000384 seconds
0.000380 seconds
0.000364 seconds
0.000383 seconds
0.000478 seconds
0.000444 seconds
0.000456 seconds
```

Figure 4: BlackLynx FPGA-Accelerated LDPC Single Codeword Performance

Comparing Benchmark Performance and Performance-per-Watt

When running our benchmarks for CPU-only and BlackLynx FPGA-accelerated scenarios, we also measured the power consumption of our benchmark Dell R640 server. We measured the input power to the device at the power plug, to ensure that all power consumption was accounted for, to include the inefficiencies of various on-server power supplies. A breakdown of performance and power consumption appears in Table 1 below:



	Stock LDPC Algorithm	BlackLynx FPGA-Accelerated LDPC
Decode duration (lower is better)	~11 to ~20 ms	~0.5 ms
Codewords per second (higher is better)	~100	~2,000
Computation-specific power consumption (lower is better)	~24W	9W
Performance per Watt advantage	1x	53x

Table 1: Benchmark Performance per Watt Comparison

At 53x greater performance per Watt, we are well over an order of magnitude better. The BlackLynx FPGA-accelerated operation provides the best possible performance with the lowest possible power consumption differential, running in the same form-factor server.

Aside from achieving an astounding performance and performance per Watt, the obvious cost savings in terms of power, number of servers, IT management, multi-machine software licensing, server recapitalization plans, as well as other tangible and intangible costs, although they certainly will vary somewhat from use case to use case, are impossible to ignore. The result is that return on investment numbers are guaranteed to be extraordinarily good by any measure when leveraging BlackLynx-accelerated solutions.

Cloud Ready: BlackLynx Deploys Seamlessly to AWS Instances

It is interesting to note that the Xilinx VU9P FPGA that was employed in this benchmark is the same Xilinx FPGA variant that is used in the Amazon Web Services F1 cloud-based instance type. BlackLynx solutions deploy identically, and seamlessly, even into the AWS cloud, with or without FPGA acceleration.

This means that DSP engineers are no longer constrained by on-premise-only or cloud-only deployment strategies, but instead can leverage hybrid deployment strategies, intelligently enabling businesses to choose when and where to solve problems using the identical BlackLynx software stack in all such cases. Since the same SDAccel-capable FPGA family used in these benchmarks in a Dell server is used in the AWS F1, the resulting performance will be practically identical. The same holds true when running on the latest generation SDAccel-capable Xilinx Alveo Accelerator Cards in the Nimble Cloud.

Summary

BlackLynx technology seamlessly deploys to a wide variety of COTS servers including the Dell R640, Dell C4140, and Dell R740. BlackLynx's industry-leading software control framework, to include native integration into popular SDR toolkits such as GNU Radio, allows for accelerated signal processing at tremendous scale, enabling end users of existing and emerging infrastructure to achieve unparalleled performance at extremely low power and at lower cost that contemporary techniques typically allow.

Utilizing a single Dell R640 server with a single low-power Xilinx SDAccel-capable FPGA PCIe board installed, and leveraging BlackLynx software technology natively in a GNU Radio framework, we demonstrated FPGA-accelerated LDPC signal processing decode operations, achieving up to a remarkable 53x performance per Watt advantage over stock SDR solutions, without disturbing the rest of the signal processing waveform chain. Similar performance is obtained when



leveraging BlackLynx acceleration using other similarly outfitted COTS servers such as the Dell C4140 or Dell R740, or in the cloud on FPGA-accelerated infrastructure such as the AWS F1 instance or in Nimbix Cloud leveraging Xilinx Alveo Accelerator Cards.

[Learn More](#)

To learn more about BlackLynx's COTS-deployable accelerated solutions and how they can help your business thrive in the era of always-on analytics, please visit our website at www.blacklynx.tech.