

NetFPGA - Rapid Prototyping of Networking Devices in Open Source

Noa Zilberman Yury Audzevich Georgina Kalogeridou
Neelakandan Manihatty Bojan Jingyun Zhang Andrew W. Moore

Computer Laboratory, University of Cambridge
name.surname@cl.cam.ac.uk

ABSTRACT

The demand-led growth of datacenter networks has meant that many constituent technologies are beyond the budget of the wider community. In order to make and validate timely and relevant new contributions, the wider community requires accessible evaluation, experimentation and demonstration environments with specification comparable to the subsystems of the most massive datacenter networks. We demonstrate NetFPGA, an open-source platform for rapid prototyping of networking devices with I/O capabilities up to 100Gbps. NetFPGA offers an integrated environment that enables networking research by users from a wide range of disciplines: from hardware-centric research to formal methods.

CCS Concepts

• **Networks** → **Intermediate nodes**; *Data center networks*;
• **Hardware** → **Reconfigurable logic and FPGAs**;

1. THE NETFPGA PROJECT

The NetFPGA project (www.netfpga.org) provides software, hardware and community as a basic infrastructure to simplify design, simulation and testing, all around an open-source high-speed networking platform. Beyond the hardware and software, the NetFPGA project is backed by community resources that include online forums, tutorials, summer camp events and developer workshops, all supported by the NetFPGA project team. All the projects developed under the NetFPGA project are open-sourced. By reusing building blocks across projects users can compare design utilization and performance. Past experience has shown that both reference and contributed NetFPGA projects are regularly enhanced by community members, redistributed and encourage a virtuous circle.

The most prominent NetFPGA success is OpenFlow, which in turn has reignited the Software Defined Networking movement. NetFPGA enabled OpenFlow by providing a widely

available open-source development platform capable of line-rate operation and was, until its commercial uptake, the reference platform for OpenFlow.

The NetFPGA project currently supports three platforms:

- NetFPGA SUME, a new platform with I/O capabilities for 100Gbps operation as a networking device, stand alone computing unit or for test and measurement.
- NetFPGA-10G, a platform with multiple 10Gbps interfaces introduced in 2010, and supporting a large number of open source community projects (e.g. OSNT [1] and BlueSwitch [2])
- NetFPGA-1G-CML, a recent platform replacing the original NetFPGA-1G board, designed for low-bandwidth (e.g. gigabit) applications, and especially suited for network-security applications.

2. NetFPGA SUME

The NetFPGA SUME [3] design aims to create a low-cost, PCIe host adapter card able to support 40Gb/s and 100Gb/s applications. NetFPGA SUME, shown in Figure 1, uses a large FPGA, supporting high-speed serial interfaces of 10Gb/s or more presented both in standard interfaces (SFP+) and in a format that permits easy user-expansion, a large and extensible high-speed DRAM, alongside high-throughput SRAM. All this is constrained by a desire for low cost, enabling access by the wider research and academic communities.

At the core of the NetFPGA SUME board is a Xilinx Virtex-7 690T FPGA device, supporting highly complex reconfigurable designs. There are several peripheral subsystems complementing the FPGA and providing an ideal platform for networking devices prototyping; A high-speed serial interfaces subsystem, composed of 30 serial links running at up to 13.1Gb/s, enables 10Gb/s, 40Gb/s and 100Gb/s applications. Users can use these serial links as standard interfaces (e.g. SFP+, PCI-Express Gen.3) or use proprietary interfaces.

The memory subsystem combines both SRAM (QDRII+, running at 500MHz) and DRAM (DDR3 SoDIMM, running at 1866MT/s) devices. These memory devices can be used for different purposes: from flow tables and off-chip packet buffering to serving as RAM for soft-core processor designs. The Storage subsystem of the design can host both a MicroSD card and external disks through two SATA interfaces,

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGCOMM '15 August 17-21, 2015, London, United Kingdom

© 2015 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-3542-3/15/08.

DOI: <http://dx.doi.org/10.1145/2785956.2790029>




Figure 1: NetFPGA SUME Board

thus enabling a complete standalone operation of the board. Additional NetFPGA SUME features support debug, extension and synchronization of the board. Special attention was paid to power instrumentation [3]. The board is implemented as a dual-slot, full-size PCIe adapter, that can operate as a standalone unit outside of a PCIe host. It is manufactured by Digilent Inc. (<http://www.digilentinc.com>).

3. RAPID PROTOTYPING

Networking research attracts interest from a wide range of disciplines. Researchers have different sets of skills and interests and often find it hard to innovate over hardware-based platforms; most platforms require deep understanding of every aspect of the board and an implementation of a complete solution from scratch. This means that a person researching multi-table OpenFlow switch designs is required to know aspects of high-speed serial links, whereas researchers studying power-efficiency in the MAC and physical layers are required to understand aspects of the driver and the DMA engine. Alternately, researchers opt for a partial implementation that is not comparable to real networking devices. A different class of researchers are interested in test and measurement, and do not wish to develop new devices or know implementation details. These researchers often fail to get a hold on commercial devices due to their high cost.

The NetFPGA platform offers a solution to these researchers. First, NetFPGA offers ready-made reference and contributed projects, providing full implementation and an executable application. The user can run these projects, with no further development or modification required. For example, a domain specific language researcher may use the NetFPGA platform for the study of application-level middleboxes.

Second, NetFPGA projects are modular. The hardware and software components are provided as flexible building blocks, that can be modified and replaced without affecting other parts of the design. On one end of the scale, users can opt to modify or develop at the software level alone. For example, an SDN researcher interested in the control plane and lacking any hardware knowledge, can use the BlueSwitch OpenFlow switch project [2] as its data plane, and choose to write a control plane software application to run on top of it. On the other end of the scale, a researcher may choose to explore aspects of hardware-based scheduling, and thus add a new scheduling module to the existing reference router design. A researcher interested in power efficiency can opt to modify only a single interface module in the design, with no need to handle any software aspects or the data path. In all these cases, the users can focus only on research aspects

that interest them, with the overheads being handled by the NetFPGA platform.

The NetFPGA development environment provides an integrated solution that enables rapid prototyping. Every NetFPGA release contains several reference projects that are used as basic building blocks for new projects: a reference NIC project, a reference switch project, a reference IPv4 router project and a project that exercises all the I/O interfaces. Each project consists of hardware, software, testing and documentation components. The hardware portion of a project contains the source code for all the modules used in the design, as well as a large set of scripts that generate the platform's configuration file. The software portion contains embedded code (for a soft-core processor), a driver and relevant applications (e.g. router management). The test environment provides unified tests for simulation and hardware test, allowing simple validation of designs. Each project is also well documented, with the documentation covering every element or module used in the project. In addition to the modules used in the reference project, a large library of modules is provided, as well as contributed projects developed by members of the NetFPGA community.

In the demonstration, we provide a development experience over the NetFPGA platform. We introduce the NetFPGA SUME board, and show the various components of the development platform. Visitors are then encouraged to ad-hoc develop a new project, and experience the meaning of rapid prototyping over the NetFPGA SUME platform. To fully demonstrate the NetFPGA SUME capabilities, we also show a high performance device, and its full implementation, in action.

Acknowledgements

We thank the many people who have contributed to the NetFPGA project. The Xilinx XUP program has been a long-standing supporter of NetFPGA and the NetFPGA SUME project is only possible with their generous support. We thank the people at Digilent Inc. We thank Micron and Cypress Semiconductor for their generous part donations.

This work was jointly supported by EPSRC INTERNET Project EP/H040536/1, National Science Foundation under Grant No. CNS-0855268, and Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory (AFRL), under contract FA8750-11-C-0249. The views, opinions, and/or findings contained in this report are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the National Science Foundation, Defense Advanced Research Projects Agency or the Department of Defense.

4. REFERENCES

- [1] G. Antichi, M. Shahbaz, Y. Geng, N. Zilberman, A. Covington, M. Bruyere, N. Feamster, N. McKeown, B. Felderman, M. Blott, A. Moore, and P. Owezarski. OSNT: Open Source Network Tester. *IEEE Network*, 28(5):6–12, 2014.
- [2] J. Han, P. Mundkur, C. Rotsos, G. Antichi, N. Dave, A. Moore, and P. Neumann. Blueswitch: Enabling provably consistent configuration of network switches. In *ANCS*, 2015.
- [3] N. Zilberman, Y. Audzevich, G. A. Covington, and A. W. Moore. NetFPGA SUME: Toward 100 Gbps as Research Commodity. *IEEE Micro*, (5):32–41, 2014.