Scalable P4 Deparser for Speeds Over 100 Gbps

Jakub Cabal
CESNET a.l.e.
Zikova 4, Prague 160 00
Czech Republic
Email: cabal@cesnet.cz

Pavel Benáček
CESNET a.l.e.
Zikova 4, Prague 160 00
Czech Republic
Email: benacek@cesnet.cz

Jana Foltová
CESNET a.l.e.
Zikova 4, Prague 160 00
Czech Republic
Email: jana.foltova@cesnet.cz

Juraj Holub
Faculty of Information Technology
Brno University of Technology
Božetěchova 2, Brno, Czech Republic
Email: xholub40@stud.fit.vutbr.cz

***Abstract*—The P4 language is a language suitable for the description of packet processing inside a network device. The typical P4 device consists of three main building blocks: Parser, Match+Action Tables and Deparser. The deparsing is the most challenging block because the main task of this block is to assemble the output packet based on changes in Match+Action Tables. This operation can be quite complicated in the case of high-speed networks. In this work, we present the scalable architecture (in term of the throughput) of a deparsing circuit which is suitable for implementation in FPGAs.**

I. Introduction and Related work

Research about P4 [1] compilers for FPGAs is an attractive area. In our knowledge, there are two most known implementations of P4 compiler for FPGAs [2], [3]. Unfortunately, such architectures are not scalable on higher data rates.

II. Deparser design

We introduce MFB Deparser architecture that requires only the necessary logic in the dependence on a specific P4 application. The high throughput is achieved by the processing of multiple packets per one clock cycle, the data bus used in MFB Deparser is described here [4]. Usual architectures support only one packet per clock cycle, which makes it impossible to always use the whole data word. MFB Deparser consists of two main blocks. The first and simplest block is a configurable packet **Editor** that modifies selected bytes. The second block is complex **Spacer**. The Spacer block is used to add or remove space (some bytes) in packets.

III. Results

In this section, we provide results of MFB Deparser and FL Deparser [3]. We chose four different P4 applications for experimental measurements. All values provided in the results are after the implementation for the **Xilinx UltraScale+ VU3P** FPGA using the Vivado 2017.2 tool. As the target frequency, we chose 250 MHz. The timming has always been met.

The figure 1 shows the amount of CLB blocks per 1 Gbps of throughput in the worst case. This comparison combines the resource utilization and performance of implemented deprasers. Our Depraser achieves significantly better results for simple applications (L2/L3 Switch, Port Switch/Filter). The significant worsening of results in the case of FL Deparser for the 2048 bit wide data bus (the light blue bar) is caused by the inefficient use of available data bus capacity.

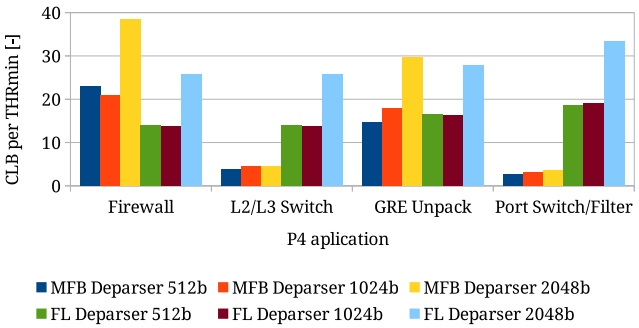

Fig. 1. The relation between resource utilization per 1 Gbps of worse case throughput and P4 application for both Deparsers with multiple data widths.

IV. Conclusion

We introduced a scalable architecture of deparsing pipeline suitable for FPGAs (Intel or Xilinx). The biggest advantage of our architecture is the easy scalability to higher data bus widths where the nearest comparable solution isn't able to reach the quality of ours. We are also able to sustain the throughput of deparsing process independently on packet length. MFB Deparser allows efficient implementation in the pure editing use case where the nearest comparable solution requires the full implementation of the deparsing pipeline. However, the FL Deparser is slightly better in terms of consumed resources for use cases where the bus width is up to 512 bits and the output packet is strongly rebuilt.

Acknowledgment

This research has been supported by the Technology Agency of the Czech Republic project TH02010214.

References

[1] P4 Language Consortium, "P4." [Online]. Available: http://p4.org/

[2] H. Wang, R. Soulé, H. T. Dang, K. S. Lee, V. Shrivastav, N. Foster, and H. Weatherspoon, "P4FPGA: A Rapid Prototyping Framework for P4," in *Proceedings of the Symposium on SDN Research*, ser. SOSR '17. New York, NY, USA: ACM, 2017, pp. 122–135. [Online]. Available: http://doi.acm.org/10.1145/3050220.3050234

[3] P. Benáček, V. Puš, H. Kubátová, and T. Čejka, "P4-To-VHDL: Automatic generation of high-speed input and output network blocks," *Microprocessors and Microsystems*, vol. 56, pp. 22 – 33, 2018. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0141933117304787

[4] J. Cabal, P. Benáček, L. Kekely, M. Kekely, V. Puš, and J. Kořenek, "Configurable FPGA packet parser for terabit networks with guaranteed wire-speed throughput," ser. FPGA '18. ACM, 2018, pp. 249–258. [Online]. Available: http://doi.acm.org/10.1145/3174243.3174250