FPGA Based High-Throughput Real-Time Feature Extraction for Modulation Classification

Joshua Mack and Ali Akoglu
Electrical and Computer Engineering Department
University of Arizona
Tucson, Arizona 85721–0104
Email: {jmack2545, akoglu}@email.arizona.edu

***Abstract*—The spectral correlation density (SCD) function is a feature extraction method used in signal classification systems. Due to its computational complexity, SCD has not been a desirable method for systems under power and real-time constraints. In this study, we present results for a hardware implementation of key kernels of the SCD function on a Field Programmable Gate Array (FPGA). By analyzing profiling results for a state of the art GPU implementation, we developed a preliminary architecture that is able to accelerate the most computationally demanding aspects of the SCD algorithm. We find that this FPGA architecture is able to achieve a 2.03X speedup relative to state of the art GPU-based SCD implementations by coupling SCD's large-scale data-parallel nature with an architecture well suited for fine-grained control flow and data access patterns.**

I. Introduction

Signal classification is the task of determining the modulation of an arbitrary signal, and it has numerous applications in both commercial and military domains. In an increasingly competitive RF environment where numerous entities try to coexist in a limited spectrum, solutions such as software defined radios promise higher spectral efficiency by enabling devices to change transmission characteristics based on environmental conditions. The ability of software defined radios to dynamically change their signal parameters such as transmission frequency or modulation scheme makes the task of classifying and identifying signals they produce more difficult.

Numerous techniques exist for modulation classification, but we choose to focus on the spectral correlation density (SCD) method for its effective operation in low SNR conditions [1]. In particular, we focus on a flavor of SCD known as the FFT Accumulation Method (FAM) [2] as it is considered more computationally efficient than competing algorithms. Real time classification using the SCD is computationally challenging, though, with the most complex stage of the algorithm requiring calculation of 65,536 distinct 32-Pt FFTs. The state of the art implementation by Marshall et al. [3] addresses this computational barrier by mapping SCD onto an Nvidia Tesla K40 GPU and achieving a throughput of 1569 signals/second.

In this work, we map key kernels of the SCD algorithm to a Xilinx Virtex-7 VC707 FPGA and present promising results regarding suitability of FPGAs for this algorithm that will be explored in future work.

II. Results & Future Work

As shown by direct profiling in [4], the most compute-critical sections of the SCD algorithm involve generating a 2D matrix with a maximum dimension of 4096 elements in each coordinate based on an iterative computation of 65,536 distinct 32-Pt FFTs. This matrix formation step is followed by control-flow intensive phase where maximum element from each row is selected to form the 4096 element *alpha profile* feature vector. Marshall et al. [3] present a GPU implementation that accelerates these critical kernels, but also shows that their implementation is still unable to fully utilize the GPU. Targeting this inefficiency, we present preliminary results for mapping these two critical kernels to a high throughput FPGA design that is suited for maximizing hardware occupancy, particularly with regards to the fine grained operations that are required in merging in alpha profile vector updates.

In our FPGA design, we exploit a fully pipelined 32-pt FFT design and couple it with carefully crafted memory layout techniques and state machines. Emphasis is put into optimizing the handoff between the 32-Pt FFT computations and the alpha profile computation for streaming. With this, we are able to run all of these critical kernels together with an end-to-end execution time of 0.311 ms. As such, the design here provides an approximate 2.03X speedup over a runtime of 0.633 ms in the state of the art GPU version, with no degradation in output accuracy. With this, we are confident that, given a well crafted preprocessing datapath and careful data transfer strategies, this performance gain can be extended to the full SCD algorithm, which will be pursued in future work.

References

[1] D. S. Yoo, J. Lim, and M. H. Kang, "ATSC digital television signal detection with spectral correlation density," *Journal of Communications and Networks*, vol. 16, no. 6, pp. 600–612, 2014.

[2] S. R. Schnur, "Identification and classification of OFDM based signals using preamble correlation and cyclostationary feature extraction," tech. rep., DTIC Document, 2009.

[3] S. Marshall, G. Vanhoy, A. Akoglu, T. Bose, and B. Ryu, "GPGPU based parallel implementation of spectral correlation density function," *Journal of Signal Processing Systems*, Apr 2019.

[4] N. Bidyanta, G. Vannhoy, M. Hirzallah, A. Akoglu, B. Ryu, and T. Bose, "GPU and FPGA based architecture design for real-time signal classification," in *Proceedings of the 2015 Wireless Innovation Forum Conference on Wireless Communications Technologies and Software Defined Radio (WInnComm'15)*, pp. p. 70–79, March 2015.

2576-2621/20/$31.00 ©2020 IEEE
DOI 10.1109/FCCM48280.2020.00073