Power-Efficient Accelerated Genomic Short Read Mapping on Heterogeneous Computing Platforms

Ernst Joachim Houtgast*†, Vlad-Mihai Sima†, Giacomo Marchiori†, Koen Bertels* and Zaid Al-Ars*
*Computer Engineering Lab, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands
†Bluebee, Molengraaffsingel 12-14, 2629 JD Delft, The Netherlands
E-mail: {ernst.houtgast, vlad.sima, giacomo.marchiori}@bluebee.com, {k.l.m.bertels, z.al-ars}@tudelft.nl

***Abstract*—We propose a novel FPGA-accelerated BWA-MEM implementation, a popular tool for genomic data mapping. The performance and power-efficiency of the FPGA implementation on the single Xilinx Virtex-7 Alpha Data add-in card is compared against a software-only baseline system. By offloading the Seed Extension phase onto the FPGA, a two-fold speedup in overall application-level performance is achieved and a 1.6x gain in power-efficiency. To facilitate platform and tool-agnostic comparisons, the base pairs per Joule unit is introduced as a measure of power-efficiency. The FPGA design is able to map up to 34 thousand base pairs per Joule.**

Introduction— The extreme scale of genomics data necessitates the use of high performance and power-efficient solutions. BWA-MEM [1] is the de facto standard for mapping DNA short reads onto a reference genome. Following the Seed-and-Extend paradigm, exactly matching seeds are generated for each read, which are subsequently extended using inexact matching similar to the Smith-Waterman algorithm. Here, we accelerate this Seed Extension phase, which forms a major bottleneck in BWA-MEM requiring 30%-50% of total execution time. Few accelerated BWA-MEM implementations exist: two implementations on the Convey HC-2^{EX} platform with four Xilinx Virtex-6 FPGAs, one accelerating only the Seed Extension phase [2] and achieving a 1.5x speedup, and one accelerating multiple phases [3] for an overall 2.6x speedup; and a GPU implementation [4]. Our work extends [2], but is able to achieve a 2.0x speedup using only a single Xilinx Virtex-7 FPGA, which offers about 23% of the resources as compared to the Convey platform.

Approach— Our implementation offloads the Seed Extension phase that performs the inexact matching onto an FPGA, accelerating the Smith-Waterman-like algorithm using a systolic array. Our design contains six modules that are each able to process the Seed Extension phase. The remainder of the logic is filled with arbitration logic to distribute reads, with the PCI-Express interface and with the memory controller. The design is limited by the amount of LUTs available.

Results— Results were gathered using BWA-MEM v0.7.8, which is highly multi-threaded, with the publicly available 150bp-se-small-indel GCAT data set [5]. To measure the power-efficiency, an emonPi energy monitor [6] was used to track the system-level power consumption as measured at the power plug. Performance and power-efficiency results are summarized in Table I. On a kernel-level, the FPGA is 1.8x faster as compared to software-only execution for the Seed Extension phase. In our implementation, this phase is executed in parallel with the other phases, thus resulting into a two-fold improvement to overall application performance. Note that the accelerated Seed Extension phase requires less than half of the overall application execution time. This shows that the FPGA is not fully utilized and can be used to accelerate a more powerful system. Moreover, the FPGA-accelerated implementation is much more power-efficient: it requires only 35 kJ, an energy efficiency improvement of 60%. The power-efficiency would be even greater for a more balanced system. A conservative estimate indicates that the FPGA-accelerated platform is able to achieve an up to 2.1x improvement in power-efficiency and is able to map up to 44 kbp/J.

References

[1] H. Li, "Aligning Sequence Reads, Clone Sequences and Assembly Contigs with BWA-MEM," *arXiv preprint arXiv:1303.3997*, 2013.
[2] E. Houtgast, V. Sima, K. Bertels, and Z. Al-Ars, "An FPGA-Based Systolic Array to Accelerate the BWA-MEM Genomic Mapping Algorithm," in *International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation*. IEEE, 2015.
[3] N. Ahmed, V. Sima, E. Houtgast, K. Bertels, and Z. Al-Ars, "Heterogeneous Hardware/Software Acceleration of the BWA-MEM DNA Alignment Algorithm," in *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, ser. ICCAD '15. Piscataway, NJ, USA: IEEE Press, 2015, pp. 240–246.
[4] E. Houtgast, V. Sima, K. Bertels, and Z. Al-Ars, "GPU-Accelerated BWA-MEM Genomic Mapping Algorithm Using Adaptive Load Balancing," Accepted for Publication in Architecture of Computing Systems, 2016.
[5] G. Highnam, J. J. Wang, D. Kusler, J. Zook, V. Vijayan, N. Leibovich, and D. Mittelman, "An Analytical Framework for Optimizing Variant Discovery from Personal Genomes," *Nature comm.*, vol. 6, 2015.
[6] The OpenEnergyMonitor Project, "OpenEnergyMonitor Website," http://openenergymonitor.org/, accessed: 2015-12-14.

TABLE I: POWER CONSUMPTION, PERFORMANCE AND ENERGY-EFFICIENCY FOR SOFTWARE-ONLY AND FPGA-ACCELERATED PLATFORMS (TESTED ON: INTEL CORE I7-4790 @ 3.6 GHZ + 16 GB RAM; ALPHA DATA ADM-PCIE-7V3 WITH XILINX VIRTEX-7 XC7VX690T-2 @ 160 MHZ + 16 GB RAM)

	Kernel Performance		Overall Performance		Power Consumption		Energy Efficiency		
Platform	**Time**	**Speedup**	**Time**	**Speedup**	**Idle**	**Load**	**Power**	**Efficiency**	**Improvement**
Software-Only	237 s	-	552 s	-	37 W	105 W	58 kJ	20.7 kbp/J	-
FPGA-Accelerated	129 s	**1.8x**	272 s	**2.0x**	62 W	129 W	35 kJ	34.1 kbp/J	**1.6x**

978-1-5090-2356-1/16 $31.00 © 2016 IEEE
DOI 10.1109/FCCM.2016.17

IEEE computer society