

FPGA Acceleration of GWAS Permutation Testing

Friday, 3 December 2021 12:45 (30 minutes)

The large sample sizes of modern genetic datasets has necessitated the development of high-throughput accelerators in order to allow bioinformatics research to be performed in a reasonable amount of time. Although the inherently parallel nature of FPGAs makes them well suited to accelerating high-throughput workloads, they are not commonly employed as bioinformatics accelerators (in lieu of CPUs and/or GPUs) due to their high cost and the fact that developing FPGA-accelerated algorithms is a more complex and time-consuming process than the development of software for CPUs or GPUs. The availability of cloud-based FPGA instances, however, has made powerful FPGAs accessible to bioinformatics labs and the continuous improvement of FPGA design tools has reduced much of the complexity of FPGA development.

This work determines the efficacy of FPGAs when applied to the acceleration of GWAS permutation testing - a computationally expensive bioinformatics algorithm that involves the repeated multiplication of a constant matrix with a changing vector - by presenting the design and evaluation of an FPGA-based accelerator designed to run on an AWS EC2 FPGA instance.

This work shows that the FPGA accelerator is orders of magnitude faster than a popular CPU-based GWAS tool without an apparent loss of accuracy. Furthermore, this work demonstrates that FPGA acceleration enables the handling of workloads which are almost unfeasible for current CPU-based methods. This work, therefore, proves that FPGAs can effectively accelerate high-throughput bioinformatics workloads at relatively low cost.

Student?

Yes

Supervisor name

Professor Scott Hazelhurst

Supervisor email

scott.hazelhurst@wits.ac.za

Primary author: SWIEL, Yaniv (Wits)

Presenter: SWIEL, Yaniv (Wits)

Session Classification: HPC Technology

Track Classification: HPC Technology