

Public Abstract

First Name:Marziyeh

Middle Name:

Last Name:Nourian

Adviser's First Name:Michela

Adviser's Last Name:Becchi

Co-Adviser's First Name:

Co-Adviser's Last Name:

Graduation Term:FS 2016

Department:Computer Engineering

Degree:MS

Title:Critical Analysis and Evaluation of Different Automata Processing Accelerators on Large-scale Datasets

Many established and emerging applications perform at their core some form of pattern matching, a computation that maps naturally onto finite automata abstractions. As a consequence, in recent years there has been a substantial amount of work on high-speed automata processing, which has led to a number of implementations targeting a variety of parallel platforms. More recently, Micron has announced its Automata Processor (AP), a programmable silicon device for accelerating non-deterministic finite automata (NFAs). Despite the abundance of work in this domain, the advantages and disadvantages of different automata processing accelerators and the innovation space in this area are still unclear.

In this work we target this problem. In particular, to allow an apples-to-apples comparison, we focus on NFA acceleration on three platforms: GPUs, FPGAs and Micron's AP. We discuss the automata optimizations that are applicable to all three platforms. We perform an evaluation on large-scale datasets: to this end, we propose an NFA partitioning algorithm that minimizes the number of state replications required to maintain functional equivalence with an unpartitioned NFA, and we evaluate the scalability of each implementation to both large NFAs and large numbers of input streams. Our experimental evaluation covers resource utilization, throughput, and preprocessing cost. Based on our experiments, FPGAs outperform GPUs by a 900x factor and Micron AP by a ~10-20x factor in terms of throughput while requiring multiple devices to handle large-scale datasets and suffer from substantial preprocessing time. Despite offering low performance on a single input stream, GPUs can accommodate large datasets on a single device and have limited preprocessing time. AP still require multiple boards on very large datasets, and have a preprocessing time comparable to FPGA, which can be reduced on datasets consisting of a large number of small NFAs of known structure.