Seminar Neuste Trends im Hochleistungsrechnen: Microns Automata Processor

Eine Ausarbeitung von Dennis Struhs

26. März 2016

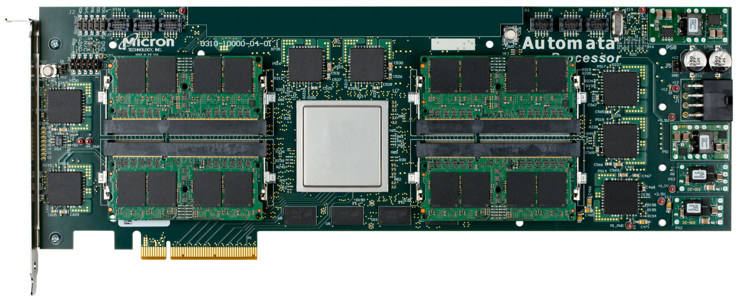

Inhaltsverzeichnis

1 Introduction

1.1 The challange of handling Big Data through Data Mining

"Data Mining is the practice of searching through large amounts of computerized data to find useful patterns or trends"-(Merriam Webster Dictionary)

"Data Mining is the process of identifying valid, novel, potentially useful, and ultimately comprehensible knowledge from databases that is used to make crucial business decisions"-(G. Piatetsky-Shapiro)

During the past years only few other aspects of the information technology branch have grown equally in significance then Big Data. Especially the development of the Internet and the increase of accessability of internet infrastructure throughout the globe cause the amount of data circulating through the web to grow at frightening speeds.

Numerous companies throughout the IT sector have to face the problems this amount of data causes. For example:

- The MEDLINE text database with 14 million published articles
- Google with over 40,000 search requests per secondand more than 3.5 billion search requests per day
- The NASA MODIS satellite Coverage: 250m resolution, 37 bands, whole earth, every day
- Walmart transaction data beeing in the order of 100 million transactions per day

Trends show that the amount of data gathered will raise even more in the future, this phenomenon is known as the: **Data Flood / Data Explosion**

The problem about this situation is that only a small portion (5% - 10%) of the collected data is ever analyzed. This **Data Wasting** causes Data that may be never analyzed continue to be collected at great expenses, doing nothing but bloating the system that is supposed to handle the data not drowning in it.

Fortunately Micron® has come up with a promising new way of approaching this Data Wasteland and may help to increase the amount of actually analyzed data greatly in the future. In the following we'll deal with the features and benefits of Microns® new Automata Processor.

2 The Micron Automata Processor

2.1 Introducing: The Automata Processor

The Micron Automata Processor (AP) is a highly parallel reconfgurable non-von Neumann architecture designed for the emulation of Non-Deterministic Finite Automata (NFA). The Automata Processor is effective in modeling complex regular expressions as well as other more complex automata. The architecture allows for effective solutions for pattern matching problems found in computational biology, cyber security, graphics processing, and more. Its purpose-built design exceeds the capabilities of similar FPGA-based solutions while also signifcantly outperforming them. The Automata Processor connects to a system through the PCI Express and is designed to supplement a CPU as an acceleration device rather than a standalone utility.

Sequential processor implementations of NFA in traditional CPU's

A von Nuemann sequential processor (CPU) may take a direct approach of modeling an NFA by defining a set of states and connections. By maintaining a list of which states are active, given the next token of input data, the next set of states can be computed and updated. Unfortunately, the CPU must make considerations for each of its active states and their connections sequentially, and such state processing time is increased by increased activity across the decision graph. In the worst case, the entire NFA may be active resulting in the entire graph needing to be considered for a given character of input - yielding a processing complexity of $\mathcal{O}(\mathrm{n}^2)$. Although such direct solutions offer a simple storage cost $\mathcal{O}(\mathrm{n})$, although the large processing complexity makes them unrealistic for large problems. While CPUs struggle with the processing of multiple simultaneous states in NFA, Deterministic Finite Automata (DFA) require singularity in state, which means that only one state is active at a time. This property allows for the direct processing time $\mathcal{O}(1)$ of DFA, and encourages a translation from NFA to DFA. For an NFA to be modeled as a DFA, the equivalent DFA must be able to model any combination of simultaneous states possible in the NFA. For this reason while NFA can be reliably converted into equivalent DFA, it comes with the challenge of significantly inceased storage cost.

Automata Processor implemenations of NFA

The Automata Processor allows for a physical implementation of NFA. It has many configurable elements which act as individual states for a modeled NFA. These elements can be dynamically connected to one another and configured to respond to different stimuli. In this way the Automata Processor is capable of literally modeling NFA directly. Because of this the Automata Processor has a direct NFA storage cost of O(n). The physical design of the Automata Processor allows for processing to be done at an element level, and gives the input to all computing elements simultaneously. Because of this, all state processing is done in parallel. This means that regardless of how many states are active in an Automata Processor modeled NFA, the next-state computation is constant. This gives it an O(1) processing cost, and an impressive advantage over generic CPU solutions shown in the Table 2.1 below.

	Processing Complexity	Storage Cost
Sequential Processor NFA	$\mathcal{O}(n^2)$	$\mathcal{O}(n)$
Sequential Processor Equivalent DFA	$\mathcal{O}(1)$	$\mathcal{O}(\Sigma^n)$
Automata Processor NFA	$\mathcal{O}(1)$	$\mathcal{O}(n)$

Table 2.1

2.2 Technical Specifications

- Micron® D480 Automata Processors 32 / 16 Automata Processor chips 1536K STE per board / 768K STE per board Up to 4 ranks of AP chips per board Up to 8 AP chips per rank
- Altera FPGA 270,000 LEs / 101,620 ALMs 17Mb embedded memory
- DDR3 on-board memory 4GB DDR3 SDRAM
- Micron Automata Processor SDK supports
  Centos 6, SUSE 11, Ubuntu 14, Windows 7 (Design tools only)
- Both active and passive cooling solution
- Temperature Operating: 0 °C to +45 °C
- Mechanical/Electrical Full length, full height, double-width PCIe board PCIe-compliant, x8 land PCB connector 300 watt
- RoHS compliant

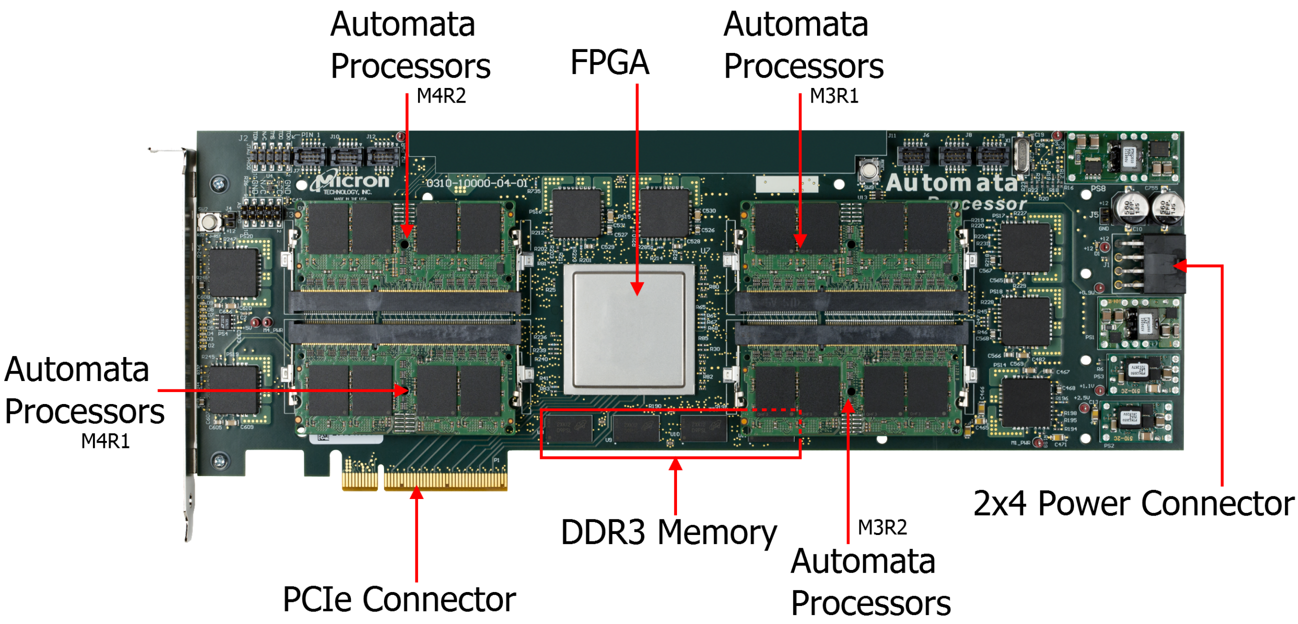

2.3 How it works

The entire Automata Processor exists as a conglomeration of 6 distinct ranks connected to a system through PCI Express. An Automata Processor rank consists of 8 distinct Automata Processor cores on a single chip. Each core consists of two half-cores containing 24K elements each, where no connections can be made between elements of unique half-cores. It is the elements of these half-cores which are used to directly model the states of NFA. With a section of the Automata Processor configured to model a desired NFA, input can be streamed. The modeled NFA views each streamed-in character as the stimulus for a potential state transition. An Automata Processor core processes data at a rate of 1 Gbps. Because one byte characters are the fundamental unit of the Automata Processor this can be better viewed as $128 \cdot 10^6$ characters per second. This means that the state of a modeled NFA can be updated every 7.4510^{-9}s This time is also known as a symbol cycle. Cores can be associated among their rank, in groups of 1, 2, 4, 8 cores. Grouped cores will receive data from the same stream of data, while cores of different groups can concurrently process different streams of data. With all cores associated in one group of 8 the Automata Processor has an effective throughput of 1 Gbps. With all cores grouped individually with their own data streams an effective throughput of 8 Gbps can be achieved.

2.4 Benefits of the Automata Processor: Gaining Parallelism

A common target of NFA solutions is pattern matching. An NFA can be used for pattern matching rather directly. For every character in a target string, a state is created to look for that character. With the states connected in the order in which they appear in the target string, the automata is complete. For the state of the final character to have been entered, it must have passed through all previous states consequtively. This way we know that if the final state has been entered, the pattern has been matched. The Automata Processor is capable of simultaneously processing action for each of its active states for each token of input. The time step in which an input character is processed and the states are updated is known as a symbol cycle. The effective parallelism of the Automata Processor is in part derived from this simultaneous inspection of multiple active states within an NFA. Parallelism for the Automata Processor is further derived from its ability to simultaneously process multiple separate NFA.

2.5 The Automata Processor vs. common Hardware

While NFA are difficult to emulate with a sequential processor, Field Programmable Gate Arrays (FPGAs) are capable of NFA modeling with the use of look up tables. Similarly, GPU based solutions have been devised. Due to its specialization, the Automata Processor is able to show decisive advantage over such FPGA and GPU based solutions. Furthermore, the Automata Processor remains the only such specialized non-FPGA, non-GPU hardware.

2.6 Overview of the SDK features

3 Summary

3.1 Possible future applications

The strength of the Automata Processor is surely its immense processing power considering pattern matching. Especially in a world that is determined by the gaining influence of Big Data and the challenge of handling it.

Therefor any application that requires alot of pattern processing from given datastreams will most likely benefit from using the Automata Processor.

3.2 Is the Automata processor suited for me?

The Automata Processor appears to be an extremely promising new processing architecture for large scale pattern matching applications which outperforms other Hardware thriving to implement similar capabilities. Which concludes the captions question by rephrasing it:

If you need pattern matching on a large scale performing as fast as possible with a maximum of parallelism, then this device would be the right choice for you!

If this does not account for your problem then you might look for a different problem solution as this one would probably not fit your needs.

4 Resources

Automata Processing main page
Micron main Page

Literatur

[Ath14] Désiré Athow. All you need to know about Automata, Micron's revolutionary processor. *Techradar.pro*, (1):1, 3 2014.

[Fis14] Russell Fish. Two Views of the Post PC World - Automata Processor and TOMI Celeste, Part 1-4. *EDN Network*, (1):1–4, 3 2014.

[Mic15] Micron. A Massively Parallel Computing Solution. *http://www.micron.com/about/innovations/automata-processing*, (1):1, 11 2015.