



Wave-Based Approach for CPU Latency Optimization in Encoding

Apoorva Reddy Proddutoori

San Diego, CA

Email id–apoorvaproddutoori@gmail.com

ABSTRACT

Low latency is the modest requirement with the rise in the technology. Faster operating times of the I/O latency of CPU is mainly required for OS processing and the operation of the cores accounts for the majority proportion. The following paper proposes reduction of the generated overheads and improving the prediction of data latency using wave-based approach. Further, understanding of the CNN granularity is utilized to granularize the layers and decompose the latency based on layer level. Regardless of the overhead, this paper proposes the work in lowering the overall CPU latency. Based on this concept a novel framework is developed for encoding the latency analysis to build the prediction structure. Furthermore, the low latency design objective is to maintain low penalty on performance. Moreover, DVFS decomposition technique is implemented in parallel to minimize the workload decomposition.

Keywords: Latency, CNN, DVFS, workload, performance monitor.

INTRODUCTION

The evolution of societal requirements led to rise in the demand for low latency high performance CPU. A general framework for encoding latency interleaves a multiprocessor encoder motivating the workload decomposition as independent as possible. Multi-view prediction structure can be generated using the decomposition into customizing encoding architecture. Given the target latency, the multi-view prediction structure models the latency encoder to meet the desired latency, under the assumption that various predictions are unbounded. Hence, a prediction structure with number of processors can figuratively propose low latency delay for the model. The results achieved in this model generated using wave-based are proven to be accurate for multi-view processor encoding the number of processors and attain the minimum desired latency while maintaining the battery power consumption by CPU.

Further, low power design is the vital requirement for high ends CPU systems cooling and packaging costs and low reliability often tied with implications of on-chip power dissipation. The workload of a task is often represented by the number of CPU clock cycles required to complete the task and is given in advance for a hard real-time activity or predicted during the execution of a soft real-time activity. such as multimedia processing. In both cases, the issue of workload distribution according to processor- and memory-related instructions is often not considered. Remember that main memory is asynchronous to the CPU and often has its own clock. Now that the execution time of a task is governed by memory usage time, the processor speed can be slowed down, but this has little effect on the overall execution time of the task. However, this can result in significant CPU power savings.

The study focuses on the unpredictability of data arrival times on CPUs. It aims to reduce fixed costs and bottlenecks associated with CPUs to enable a wider range of applications to benefit from CPU acceleration. Also, introduces important contributions including the DVFS technique for power saving by dynamically distributing workload on- and off-chip, efficient calculation of on-chip computation time ratios, a timing model for access overhead, and the implementation of DVFS policy for energy savings in various applications. Hardware-based energy saving measures are demonstrated on a popular platform.

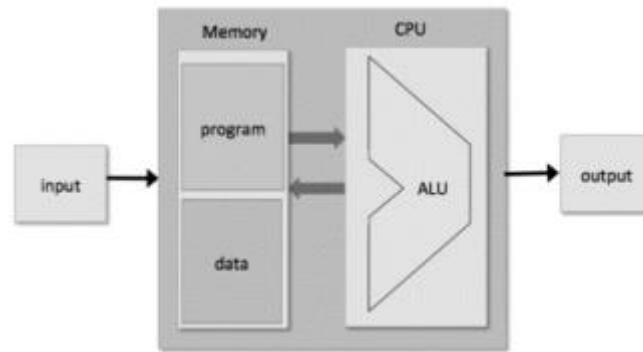


Figure 1: Computer System Basic Flow

RELATED WORK

A. Low Latency Memory

Low latency memory cells are rising to be the demanding storage units with new memory technology. The advancements of the memory cell using NAND stand library memory units supporting higher capacitance meaning higher storage is emerging as the newest development. The sequential read throughput of the low latency NAND memory is remarkably proven to be lower than regular memory cells, mainly highlighting the read latency.

B. Maintaining the Integrity of the Specifications

The utilization of optimizing compilers will definitely improve the code layout and throughput when provoked by likely/unlikely hints. Sometimes, the performance is highly impacted based on the placement of the encoding latency code while compiling, thus making it dependable on the offset instructions of memory to make potential use of the low latency memory cells.

There are several optimizations in research, such as model architectural search quantization, weight compression, and graph pruning to run CNNs completely on the edge. On the other hand, there are studies on using CNNs with resourceconstraining peripherals in their original form. In order to propose an efficient hardware design for CNN inference on FPGA and CGRA, respectively. However, most edge devices use CPUs and GPUS in off-the-shelf SoCs for CNN inference.

Most current machine ML frameworks use embedded CPUs instead of GPUs, mainly because the performance of previously embedded GPUs was insufficient for edge inference. However, embedded GPUs have seen significant performance improvements since then, but are still nowhere near their non-integrated counterparts. Therefore, there have been several efforts to synergistically use both CPUs and GPUs in SoC to perform high-performance edge inference with CNNs pipeline between asymmetric multi-core CPU clusters to perform CNN inference to improve performance. To propose a CNN inference pipeline between CPU and GPU to improve performance. But a sliding-line design can only improve performance of inference, not latency.

C. Memory Controller

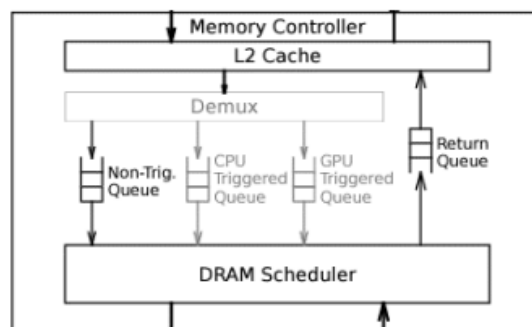


Figure 2: Memory Controller Architecture

Memory requests arrive over the link, pass through the L2 cache, and then enter one of three queues. Requests are held at the head of the specified queue until the trigger condition is met. By default, applications have no triggers, and these applications continue through the memory manager normally. On the other hand, if the specified trigger condition is not met (for example, the CPU is trying to read an empty location), the request will stop at the head of the queue. If the condition is met by a request that passes through another queue (for example, a CPU writes to the same location), the original request continues. Requests are merged into a single channel in the low latency memory scheduler. Currently, if an update function is selected, the F/E bits are updated. Apart from the full/empty bit and the two new intermediate lines, this process is the same as traditional CPU.

Each of the Measure scheduler lines is entirely requested for the purpose of equipment straightforwardness. The Measure scheduler checks the trigger conditions (in case indicated) of the demands at the head of each line; as it were the head of each line is ever watched. On the off chance that the trigger of task isn't fulfilled, the low latency memory scheduler will stall that ask (and thus all other demands queued behind it within the same line), and it'll at that point snoop the demands passing through the other lines for an activity which causes the trigger to gotten to be fulfilled. When this happens, the already blocked ask is at that point discharged. Other lined demands, in case there are any, can at that point continue as well.

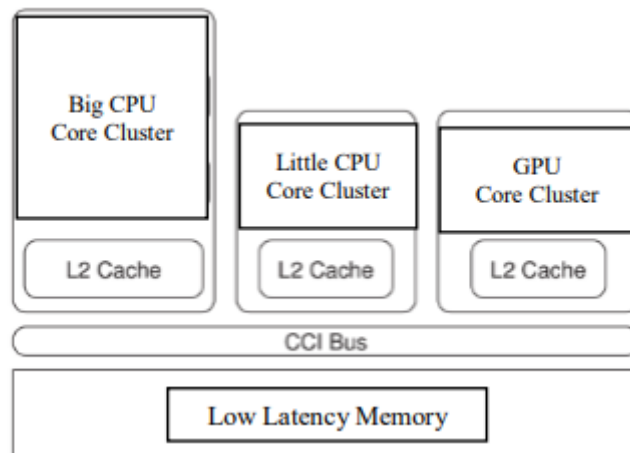


Figure 3: Big-Little CPU Overview

FREQUENCY FOR LOW LATENCY ENCODING

The operating system processing time accounts for a large portion of the total I/O latency for low-latency SSDs. Since the operating system is processed inside the CPU core and its operating speed strongly depends on its operating frequency, we first examine the correlation between core frequency and I/O latency using a low latency SSD as a summary. We measure the average latency of 4K random numbers with or without I/O polling by running on a single core with the synchronous direct I/O option. Figure 1 shows that the average latency decreases almost linearly as the core frequency increases, with or without I/O polling. So it is necessary to maximize core frequency to minimize I/O latency. Since I/O polling greatly reduces I/O latency, we use it in the experiments of this article unless otherwise noted. Of Programmers write CPU programs so that blocks of threads are loaded on the CPU. Each thread block contains a set of computer threads arranged in a one-dimensional, two-dimensional, or three-dimensional grid.

When a compute device becomes available, the CPU hardware scheduler blocks each thread block to an available core. It is hard or difficult for threads in different blocks to communicate with each other correctly during execution since the execution order of the threads and their inputs is unknown. Before starting each call or data transfer, the data required for its use must be ready. This synchronization occurs at different levels. The most detailed synchronization occurs between instructions in a queue or stream. Each statement in the flow must be completely completed before the next one can be executed. This allows synchronization across the core but is slow because it takes several CPU cycles to complete API calls. Other chip barriers can also be used to synchronize power (but not data) within wire blocks.

There are diverse DVFS approaches that make utilize of the asynchrony of memory get to to the CPU clock amid assignment execution. In compiler-assisted DVFS methods were proposed, in which recurrence is brought down in memory-bound locale of a program with small execution debasement. DVFS approaches that depend on micro-architecture or inserted equipment without any help from a compiler or a test system have too been detailed. In a microarchitecture driven DVFS method was proposed in which cache miss drives the voltage scaling. In IPC (instruction per cycle) rate of a program execution was utilized to direct the voltage scaling. Reference displayed an arrangement to select the ideal CPU clock recurrence beneath a settled execution debasement imperative (of say 10%) based on energetic program behavior such as the number of executed informational and memory get to checks amid the complete execution time by employing an execution checking unit (PMU). In a DVFS procedure which empowers more exact energy-performance tradeoff utilizing a PMU was displayed in which the ideal CPU clock recurrence and the comparing least voltage level are chosen based on the proportion of the on-chip computation time to the off-chip get to time.




Figure 4: Traditional Data Ingestion Architecture

FUTURE SCOPE

There are variegate ways to approach CPU latency optimizations to achieve low latency with high performance. PMU has been key importance to validate the throughput of the low latency architecture. Few other methodologies involve code optimizations, core parallelization, improving the performance of the compiler, reducing the IPC count and furthermore.

REFERENCES

- [1]. Satoshi Imamura, Eiji Yoshida, “Reducing CPU Power Consumption for Low-Latency SSDs”, IEEE Non-Volatile Memory Systems and Applications Symposium, June 2018
- [2]. Ehsan Aghapour, Dolly Sapra, Andy Pimentel, Anuj Pathania, “CPU_GPU Layer-Switched Low Latency CNN Inference”, 25th Euromicro Conference on Digital System Design, August 2022
- [3]. Pablo Carballeira, Julian Cabrera, Antonio Ortega, Fernando Jaureguizar, Narciso Garcia, “A Graph Based Approach for Latency Modeling and Optimization in Multiview Video Encoding”, IEEE Xplore, May 2011
- [4]. Daniel Lustig, Margaret Martonosi, “Reducing GPU Offload Latency via Fine Grained CPU-GPU Synchronization”, Princeton Edu, 2020