Mohamed Matar

+1(604)788-9175 • mohamed.omran4192@gmail.com

maomran.github.io/momatar

Experienced ASIC/FPGA Design Engineer with 7+ year experience, currently working in AMD Graphics team as a part of the cache and memory management (CMM) subsystem for the next generation of AMD GPUs. Prior to that I was a grad student at the SoC lab in the University of British Columbia, working mainly on:

1. Hadware accelerators for 5G systems.
2. Hadware accelerators for Machine Learning workloads.

Employment

Sr. Silicon Design Engineer **Vancouver**

AMD, Canada *April 2021 – now*

Part of the cache and memory management (CMM) subsytem for AMD next Gen GPUs. Responsible for 3 blocks inside the the CMM subystem including the Graphics L1 cache, high performance L1-Cache, and Mutlimedia on-chip memory buffer. Responsibilities include:

- Leading a small team of 3 to develop a UVM based TB and environment to test GPU cache blocks.
- Developing performance models to verify blocks throughput and power targets.
- Working with architecture team to model blocks functionality in AMD official in house GPU arch model.

Research Assistant **Vancouver**

University of British Columbia, Canada *Jan 2018 – April 2021*

***MSc Thesis*: Design exploration of Faster than Nyquist Equalizer. (Co-supervised by Mieszko Lis/Lutz Lampe)**

In this work, the target was to explore the design space of the next-generation Faster-than-Nyquist communications hardware accelerator. This accelerator is composed mainly of MAP equalizer engines that works in parallel on different parts of the codeblock, followed by a QC-based LDPC decoder engine, then fed to a symbol mapper. This system works in iteration feedback system to equalize the data.

My thesis describes the architecture trade-offs, and proposes a highly-configurable hardware architecture that vaies based on the design needs in terms of performance/throughput or energy efficiency.

To study the system design parameters, we designed a MATLAB based bit-accurate system model, then adapt the system into a high-level-synthesis (HLS) FPGA design that is then evaluated on an Ultrascale+ FPGA for area/power/latency evaluation.

Part of this work was published in **FCCM 2020** conference (*check publication section*).

Graduate Student **Vancouver**

University of British Columbia, Canada *Jan 2018 – April 2021*

Analysis of DNN Accelerator in SoC simulator. In this work, we integrated a DNN simulator with gem5 system simulator to measure the system performance with DNN workloads. DNN engine is connected on the system bus through memory mapped interface. We integrated NN APIs to run C++ based DNN workloads along with host code. In this *report*, we analyzed DNN workloads performance by running different CNN networks and measure kernel launch latency as well as DNN latency.

DropBack: Continous Pruning During Training GPU implementation.

Dropback introduces an algorithm that aims at reducing training energy by streaming less data from memory and regenerating random weights each iteration on the GPU.

To evaluate this work, the goal of this project was to implement this algorithm on a GPU kernel and measure its energy compared to a baseline trained model running in pytorch.

To implement the algorithm, we used an open source kernel cutlass (Nvidia's open kernels for MatMul) that can be modified for algorithm adaptation, and compare it with CuDNN (Cuda DNN engine used in pytorch).

Design Verification Engineer **Penang**

Intel Corporation, Malaysia *Feb 2016 – Jan 2018*

Verification of MIPI-Audiolink IP

Part of mixed signal group, responsible for MIPI-audiolink IP verification including OVM test development, regression runs, and GLS simulations. This IP was part of different Intel processor chips such as Cannon-lake, Ice-lake and Tiger-lake architectures.

Design Verification Engineer **Selangor**

Symmid Corporation, Malaysia *Feb 2016 – Jan 2017*

Non Volatile Flash Memory with AHB-APB Interface Verification

In this project, I was working with another colleague to develop a UVM based enviroment to test an ECC flash memory interface. This memory uses AHB-ABP interfaces to be integrated with ARM based SoC. This included contributing to the test plan, developing UVM components and test-cases, running RTL regressions, GLS debugging, and measuring functional and code coverage.

Research Assistant **Cairo**

American University in Cairo, Egypt *Jan 2015 – Feb 2016*

Modifying Virtual Memory Architecture to Enhance STT-RAM Performance STT-RAM is a magnetic based RAM that's considered as a potential replacement for DRAM. Due to its high write power, we introduce an operating system solution to reduce the writing energy. Our method introduces a page replacement policy that can reduce the number of bit write operations and consequently reduce the overall power consumption. This is done by using a simple hashing algorithm to select the target page with the nearest distance. We reflected this method on linux kernel, Gem5 full system simulator and a modified NVMain main memory simulator to evaluate the proposed solution. (*check publication section)*

- **BSc. Internship** **Cairo**

  Si-Ware Systems, Egypt *Oct 2013 – June 2014*

  Implementation & Verification of TMS320C25 DSP Core

  My graduation project was an internship to implement an ASIC DSP core for a Texas Instrument open architecture. I was part of the verification team developing assembly based test cases, running RTL tests and GLS debugging, and STA analysis.

Education

- **University of British Columbia**

  MSc in Computer Engineering , Grade: 88% *2018–2020*

  1) CPEN511 Advanced Computer Architecture.
  2) EECE579 Advanced VLSI Design.
  3) CPSC532 Machine Learning and Data Mining.
  4) CPEN411 Computer Architecture.
  5) EECE571 Compute Accelerator Architecture.
  6) CPSC411 Introduction to Compilers.

- **Alexandria University**

  BSc in Electronics and Communication Engineering , Grade: 78% *2010–2014*

Publications

1. Mohamed Matar. Design exploration of faster than nyquist equalizer system. Master's thesis, University of British Columbia, 2020
2. M. O. Matar, M. Jana, J. Mitra, L. Lampe, and M. Lis. A turbo maximum-a-posteriori equalizer for faster-than-nyquist applications. In *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 167–171, 2020
3. M. A. Ewais, M. A. Omran, A. Raafat, and Y. Alkabani. A virtual memory architecture to enhance stt-ram performance as main memory. In *2016 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–6, 2016

Teaching Experience

- (CPEN 311) Digital Systems Design, University of British Columbia.(T2'18,T1'19, T2'19,T1'20)
- (CSCE 231) Computer Organization, American University of Cairo.(Fall'15)

Technical skills

- **Programming Languages:** C/C++, Python, MATLAB, Bash.
- **System Simulators:** Intel Pin, SimpleScalar, Gem5.
- **Frameworks:** CUDA, Tensorflow, pyTorch

- **Digital Design Skills:** VHDL , Verilog , System Verilog, UVM.
- **EDA:** Xilinx Vivado , Altera Quartus, Cadence IES, Synopsis VCS/Verdi, Modelsim .

References

References are available upon request.