Manos Pavlidakis

• **Email** • **LinkedIn** • **Github** • **Scholar** • **+30 6974121338** • **Greece, Crete**

I am a software engineer with a Ph.D. in Computer Science and Engineering from the University of Crete, Greece. My doctoral research on the "*Transparent and Efficient Spatial Sharing of Multiple and Heterogeneous Accelerators*" provided extensive hands-on experience in various accelerator programming models and runtimes, including OpenCL, CUDA, and ROCm. In addition to my research achievements, I have developed a profound familiarity with heterogeneous accelerators, encompassing CPUs, GPUs, and FPGAs. This expertise has been cultivated through active participation in multiple courses, seminars, and workshops, reflecting my commitment to staying abreast of advancements in the field. Moreover, my adeptness extends to low-level knowledge about the functionality and details of popular machine learning frameworks such as PyTorch, Caffe, and TensorFlow. These frameworks have played a pivotal role in the evaluation of the runtime designed and developed during my Ph.D., showcasing my well-rounded proficiency at the intersection of cutting-edge technologies.

TECHNOLOGIES AND LANGUAGES

C/C++, CUDA, OpenCL, ROCm, CUDA PTX virtual assembly, Java, CMake, nvcc compiler, g++
Git, Containers, Bash scripting, Python, Machine Learning: TensorFlow, PyTorch, Caffe
Host and GPU code performance evaluation using NVIDIA Nsight Systems, Flamegraphs, VTune, Perf
NVIDIA GPUs, AMD GPUs, Intel Altera FPGA, Accelerator runtimes

PROFESSIONAL EXPERIENCE

Graduate Research Assistant at CARV, ICS-FORTH • Design and develop *Arax,* G-Safe, and TReM, which are systems enabling efficient sharing of multiple and heterogeneous accelerators	Sep. 2017 – Jan. 2024
Research Engineer at CARV, ICS-FORTH • Design and develop *Vinetalk,* a framework that simplifies FPGA access while enabling FPGA sharing across different applications	Jan. 2017 – Sep. 2017
Intern at CARV, ICS-FORTH • Working on implementing an HDFS library over NoSQL DB	Jun. 2016 – Sep. 2016
Freelancer • Android application and website development	Jun. 2013 – Jul. 2014

EDUCATION and CERTIFICATIONS

Ph.D. at Computer Science Department, University of Crete	Sep. 2017 – Jan. 2024
MSc at High Performance Distributed Computing track, Vrije Universiteit (VU) Amsterdam	Sep. 2014 – Jun. 2016
BSc (integrated master) at the Department of Information and Communication Systems Engineering, University of Aegean, Samos	Sep. 2006 – Jun. 2012
NVIDIA **certification**: Fundamentals of Accelerated Computing with CUDA C/C++	Aug. 2022
NVIDIA **certification**: Scaling CUDA C++ Applications to Multiple Nodes	May 2022

TEACHING EXPERIENCE

Teaching Assistant • Digital design (CS120) • Programming lab (CS255) • Embedded System Lab (CS428) • Theory of Computation (CS280)	Winters 2017 – 2022 Springs 2017 - 2019 Springs 2020 – 2022 Winter 2023

PUBLICATIONS

G-Safe: Safe GPU Sharing in Multi-Tenant Environments (Under Submission) - **Technologies used**: CUDA PTX kernels, Binary instrumentation, GPU allocator, LD_PRELOAD, bitwise AND-OR, Modulo, C++ data structures (vectors, map, unorder map), nvcc, cuobjdump, nsight, rdtsc.	2023
Arax: A Runtime Framework for Decoupling Applications from Heterogeneous Accelerators (SoCC'22) - **Technologies used**: Shared Memory Segment, IPC, pthreads, OpenCL, CUDA, HIP, ROCm, Streams, Command queues, Scheduling, nvidia-smi, nsight, nvprof, CUDA events, LD_PRELOAD, Perf, Flamegraphs.	2022
TReM: A Task Revocation Mechanism for GPUs. (HPCC'20) - **Technologies used**: CUDA dynamic parallelism, CUDA unified memory, Scheduling, Processes, Process pools, multiple CUDA contexts, GPU time-sharing.	
The VINEYARD integrated framework for hardware accelerators in the cloud. (SAMOS'18)	2018
VineTalk: Simplifying Software Access and Sharing of FPGAs in Datacenters. (FPL'17) - **Technologies used**: SDaccel, OpenCL	2017
KVFS: An HDFS library over NoSQL DB. (CLOSER '16)	2016
Introducing Touchstroke: Keystroke-based Authentication System for Smartphones. (Journal: Security and Communication Networks)	2016

VOLUNTEERING

Member of the organizing committee of the 3rd Career Fair: Meeting the companies, CSD, UoC	2023
Head of the organizing committee of the 2nd Career Fair: Meeting the companies, CSD, UoC	2022
Head of the organizing committee of the 1st Career Fair: Meeting the companies, CSD, UoC	2021
Head of the organizing committee of the Career Fair: Prepare for the interview, CSD, UoC	2021
Host virtual sessions at the European Conference on Computer Systems (EuroSys'20)	2020
Head of the organizing committee of the first Graduate Student Conference, CSD, UoC	2019
Head of the Graduate Student Association, Computer Science Department, University of Crete, Greece	2018 - 2020

FELLOWSHIPS and AWARDS

Scholarship from ADMHE for the academic year 2020-2021	2021
Scholarship from ADMHE for the academic year 2019-2020	2020
$1,000 in Google Cloud Platform from Google for Education	2020
Two Intel Altera FPGAs (Arria 10) from the Intel FPGA Academic Program	2019
One NVIDIA Titan V from the NVIDIA GPU Grant Program	2018

CONFERENCES, WORKSHOPS, and other EVENTS

Participate at NVIDIA event hosted at SC23	2023
Participate in the event "NVIDIA AI and Data Science Virtual Summit"	
Participate in the event organized by EuroCC, "NVIDIA Developer Day for EuroCC"	
Participate in the workshop "Seamlessly Migrate CUDA Code to SYCL Code with SYCLomatic"	
Participate in the event "Facebook Systems @Scale Summer"	
Participate in the event "Advanced SYCL Concepts – Graphs and Dependencies"	
Present "Arax" (and participate) at the 13th Symposium on Cloud Computing (SoCC'22)	2022
Present at EuroCC EU event: 1st presentation: Why Accelerators? GPUs and their advantages.	

2nd presentation: Introduction to CUDA	
Reviewer at ACM Transactions on Architecture and Coder Optimization (TACO'22)	
Present "TReM" at Microsoft Research, Cambridge, Virtual Workshop on Next-Generation Cloud Infrastructure	2021
Present "TReM" at International Conference on High-Performance Computing & Communications (HPCC'20)	2020
Participate in the conference Symposium on Operating Systems Design & Implementation (OSDI'20)	
Participate in Intel's webinar 2020, "Migrate Existing CUDA Applications to DPC++ Code"	
Present "Flexy; Elastic Provisioning of Accelerators in Single Node with Multi-GPUs" at ACACES'18	2018
Present "Vinetalk" (and participate) at the Conference on Field-Programmable Logic and Applications (FPL'17)	2017