

# Pavlos Petoumenos

Informatics Forum, 10 Crichton Street, Edinburgh, EH8 9AB, UK

ppetoume@inf.ed.ac.uk

## Esteem indicators

---

Journal/conference papers	15	Papers presented	7
Workshop papers	8	Invited presentations	6
Citations	416	MSc/Dipl students co-supervised	8
H-index	12	PhD students co-supervised	4

## Research highlights

---

- **Deep Learning techniques for compiler optimisations** [CGO'17, PACT'17, ISSTA'18]  
Automatic benchmark generation and heuristic construction using deep neural networks. Learns on real examples how code is structured, then synthesises new, realistic code. Creates a **huge corpus of benchmarks**, orders of magnitude larger than existing collections. For heuristic construction, it lets the neural network extract training features from the code. Effectively learning heuristics directly on code. **Better heuristics with zero expert effort.** **Best paper awards** in CGO'17 and PACT'17, **Distinguished paper award** in ISSTA'18. Contribution: Originated the idea and guided the research.
- **Function merging through sequence alignment** [CGO'19]  
Bioinformatics algorithms for smaller binaries with minimal runtime impact. Identifies optimal alignments between arbitrary functions. Combines well aligned functions into a single function. Aligned blocks shared by all invocations, unaligned ones guarded by function identifier. 25% max, 6% on average size reduction. **Best paper award** in CGO'19. Industrial interest. Contribution: Drove and directed the research.
- **Cache replacement based on reuse distance prediction** [ICCD'07, SAMOS'09, JWAC'10]  
Replacement policy emulating the theoretically **optimal policy**. First hardware mechanism to accurately **predict online data reuse patterns**. **122 citations**, among the most cited cache replacement papers of the 2000s. Contribution: Originated the idea, developed the components of the replacement mechanism.
- **RAEng Research Fellowship** (Deep Learning for easier compiler analysis and optimisation)  
Prestigious five year fellowship, worth **£600,000**. Contribution: Sole author.
- **EPSRC SUMMER project** (EP/P003915/1)  
A **£120,000** successful EPSRC grant application based on IISWC'14. Heterogeneous OS scheduler for energy efficiency. Respects the user's perception of performance. Contribution: Co-author.

## Teaching highlights

---

- **Programming club and programming contests**  
Co-founder and co-runner of the Informatics **programming club**. Teaches students basic programming skills through fun projects and programmable toys. Coach of the University's **programming teams** since 2014. Responsible for selecting team members, training, and securing funding. Typically among the **top 10 teams** in the UK.

- **Teaching assistant**  
Five years of lab exercises, tutorials, projects, and up to half of the lectures for a fifth year course.  
Mostly independent for the first three years, **fully independent** for the last two.  
On my own initiative, I **overhauled the syllabus** and replaced all **teaching material**.
- **Student supervision**  
Co-supervised **four PhD students**.  
Responsible for both their day to day supervision.  
Co-supervised another **five Diploma and three MSc students**.  
Focus on research theses, one of them produced a paper in an international conference.

## Research areas

---

### Active Research

- **Benchmarking methodologies** [IISWC'14, ADAPT'16, CGO'17, ISSTA'18]: Existing benchmarks in many areas of Computer Science are too few and barely representative of real applications. My work aims to develop **automatic techniques** for representative **benchmark generation**, targeting mainly mobile and OpenCL applications. This line of research has produced **two best paper awards** and one **distinguished paper award**.
- **Application tuning on heterogeneous systems** [LCPC'14, ADAPT'16, HLPGPU'16, CGO'17, PACT'17]: Heterogeneous processors make performance tuning complicated. More devices to choose from, each one with different optimal decisions, different for each application. My research uses **Machine and Deep Learning** to make these complex decisions **automatically and with no human supervision**. My work in this area has produced **two best paper awards**.
- **Code compression** [CGO'19]: Most binaries, especially ones relying on multiple libraries and/or templates, contain **similar or identical areas of code**. This work uses algorithms borrowed from bioinformatics to identify such cases, merge them, and **eliminate redundant code**. This work won the **best paper award** in CGO'19.
- **Asymmetric multiprocessor scheduling** []: Existing Operating System schedulers behave poorly when scheduling **mixed workloads** on **asymmetric multiprocessors**. They either ignore the asymmetry, communication bottlenecks, or fairness. My research aims at introducing the first scheduler that can execute **efficiently and fairly** any kind of workload.

### Past Research

- **Cache management** [MoBS'06, SAMOS'06, IISWC'06, ICCD'07, SAMOS'09, JWAC'10, CF'10]: Traditional replacement policies, like LRU, do not work well for modern last level caches (LLC). My work introduced novel replacement and sharing policies for **near optimal utilisation of the cache** by predicting the reuse patterns of the accessed data. My publications in this area have been cited more than 220 times, including ICCD'07 one of the **most cited cache replacement papers of the 2000s**.
- **MLP-aware Instruction Queue resizing** [ARCS'10]: Runtime adaptation of the Instruction Queue size to the amount of available Memory Level Parallelism. **30% energy savings**, with almost **no performance impact**.
- **Efficient power capping** [IPDPS'09, ICPADS'15]: Novel power capping techniques and heuristics to keep the processor's power consumption within acceptable limits.
- **Fine-grained power profiling** [TACO'17]: Measurement and analysis techniques to enable **sub-millisecond resolution** for energy profiling, allowing **fine-grained energy optimisations**.

## Research grants

---

### Royal Academy of Engineering Research Fellowship £600,000

This fellowship extends the work started with [CGO'17] and [PACT'17]. It aims at using deep learning to produce new powerful optimisation and analysis tools with, ideally, no expert guidance.

2018-2023

## EPSRC SUMMER project (EP/P003915/1) £120,000

A heterogeneous OS scheduler for minimal energy consumption which respects the user's perception of performance. Continues the work started with the [IISWC'14].

2016-2018

## Employment history

---

### University of Edinburgh

Edinburgh, UK

SENIOR RESEARCHER, EPSRC ALEA AND SUMMER PROJECTS, RAENG FELLOWSHIP

December 2013 - Present

- PI for my Fellowship, led most of the research effort for the previous projects.
- One successful Royal Academy of Engineering Research Fellowship application and one successful EPSRC grant proposal.
- One journal, seven conference, and five workshop papers.
- Co-supervised **four PhD students** and working with them on a daily basis.
- Chief editor for **Compucast**, a podcast for computer science, co-founder of the **programming club** for undergraduate students, coach of the University's **programming teams**, and co-host for the **Compiler Optimisation Youtube lectures**.

## Education

---

### University of Patras

Patras, Greece

PHD IN ELECTRICAL AND COMPUTER ENGINEERING

October 2005 - July 2011

- Supervisor: Associate Professor Stefanos Kaxiras (now Uppsala University).
- Thesis: Managing Shared Resources in Multi-core Systems.
- Seven conference and three workshop papers.
- Co-supervised the Diploma theses of five students.
- Responsible for the tutorials, projects, lab exercises, and part of the lectures of a fifth year course.

### University of Patras

Patras, Greece

DIPLOMA IN ELECTRICAL AND COMPUTER ENGINEERING (5 YEAR DEGREE)

October 2000 - October 2005

- 7.9/10 (top 3% of my year)

## Invited Talks and Outreach

---

4th Computing Systems Research Day, Athens	January 2019
Glasgow Systems Section Seminar, Glasgow	September 2018
Compiler Optimisations Youtube lectures - Co-host	January 2017 - Present
Compucast - Chief Editor	March 2016 - Present
Lancaster University, Lancaster	October 2016
Glasgow Parallelism Group Seminars, Glasgow	May 2016
International Workshop on Architectural and Micro-Architectural Support for Dynamic Optimization (AMAS-DO), Barcelona	March 2016
PPar Lunch Lectures, Edinburgh	December 2014

## Teaching and mentoring

---

Lectures for the Compiler Optimisations course, Edinburgh	2016 - Present
Programming club for undergraduate students, Edinburgh	2015 - Present

Co-supervision of four PhD and three MSc students, Edinburgh	2014 - Present
Coach for the University of Edinburgh's programming teams	2014 - Present
Co-supervision of five Diploma projects, Patras	2006 - 2011
Lectures, tutorials, and lab sessions for a fifth year course, Advanced Microprocessors, Patras	2006 - 2011
Teaching assistant for two courses, Microprocessors I and II, Patras	2006 - 2008

## Professional activities

---

General Chair - COSMIC Workshop	2014 - Present
Chair - HiPEAC CSW Session on Big Data Applications	2016
Program Committee - CADO Workshop	2016
Symposium Chair - ParCo Conference	2015

## Awards and Scholarships

---

Best Paper Award CGO'19 ( <i>IEEE/ACM conference with 31% acceptance rate</i> )	2019
Royal Academy of Engineering Research Fellowship ( <i>highly competitive, 24 awarded in 2018 out of hundreds of applications</i> )	2018
Distinguished Paper Award ISSTA'18 ( <i>ACM conference with 23% acceptance rate</i> )	2018
Best Paper Award PACT'17 ( <i>IEEE/ACM conference with 23% acceptance rate</i> )	2017
Best Paper Award CGO'17 ( <i>IEEE/ACM conference with 25% acceptance rate</i> )	2017
Best Paper Award IISWC'14 ( <i>IEEE conference. Top venue for benchmarking methodologies and characterisation</i> )	2014
HiPEAC collaboration grant ( <i>highly competitive with hundreds of applications every year from PhD students from all around the EU</i> )	2007
Greek State Scholarships Foundation Award for Distinguished Undergraduate Students ( <i>awarded to the top 3 students of each year</i> )	2003

## List of peer-reviewed journal and conference publications

---

Research in Computer Science focuses on conference proceedings, while journal publications are mostly used for archival purposes. Conference and Workshop papers that I have presented are marked with an asterisk (\*).

[1] ***Function Merging by Sequence Alignment***

R. Rocha, P. Petoumenos, Z. Wang, M. Cole, and H. Leather. CGO 2019. IEEE CS / ACM SIGPLAN sponsored conference.

**Best paper award**

[2] ***Compiler Fuzzing through Deep Learning***

C. Cummins, P. Petoumenos, A. Murray, and H. Leather. ACM SIGSOFT ISSTA 2018.

**Distinguished paper award**

[3] ***End-to-end Deep Learning of Optimization Heuristics***

C. Cummins, P. Petoumenos, Z. Wang, and H. Leather. PACT 2017. IEEE CS / ACM SIGARCH sponsored conference

**Best paper award**

- [4] ***Synthesizing Benchmarks for Predictive Modeling***  
C. Cummins, P. Petoumenos, Z. Wang, and H. Leather. CGO 2017. IEEE CS / ACM SIGPLAN sponsored conference.  
**Best paper award**
- [5] ***Minimizing the cost of iterative compilation with active learning*** \*  
W. Ogilvie, P. Petoumenos, Z. Wang, and H. Leather. CGO 2017. IEEE CS / ACM SIGPLAN sponsored conference.
- [6] ***ALEA: A Fine-grained Energy Profiling Tool***  
L. Mukhanov, P. Petoumenos, Z. Wang, D. Nikolopoulos, B. de Supinski, and H. Leather. ACM TACO 2017.
- [7] ***Power Capping: What Works, What Does Not*** \*  
P. Petoumenos, L. Mukhanov, Z. Wang, H. Leather, and D. Nikolopoulos. IEEE ICPADS 2015.
- [8] ***Measuring QoE of Interactive Workloads and Characterising Frequency Governors on Mobile Devices***  
V. Seeker, P. Petoumenos, H. Leather, and B. Franke. IEEE IISWC 2014.  
**Best paper award**
- [9] ***Where replacement algorithms fail: a thorough analysis***  
G. Keramidas, P. Petoumenos, and S. Kaxiras. ACM CF 2010.
- [10] ***MLP-aware Instruction Queue Resizing: The Key to Power-Efficient Performance*** \*  
P. Petoumenos, G. Psychou, S. Kaxiras, J. M. Cebrian Gonzalez, and J. L. Aragon. ARCS 2010.
- [11] ***Instruction-based reuse-distance prediction for effective cache management*** \*  
P. Petoumenos, G. Keramidas, and S. Kaxiras. SAMOS 2009.
- [12] ***Efficient microarchitecture policies for accurately adapting to power constraints***  
J. M. Cebrian Gonzalez, J. L. Aragon, J. M. Garcia, P. Petoumenos, and S. Kaxiras. IEEE IPDPS 2009.
- [13] ***Cache Replacement Based on Reuse-Distance Prediction*** \*  
G. Keramidas, P. Petoumenos, and S. Kaxiras. IEEE ICCD 2007.  
**Most cited paper with 128 citations**
- [14] ***Using Value Locality to Reduce Memory Encryption Overhead in Embedded Processors***  
G. Keramidas, P. Petoumenos, S. Kaxiras, A. Antonopoulos, and D. Serpanos. IEEE ETF 2007.
- [15] ***Modelling Cache Sharing on Chip Multiprocessor Architectures*** \*  
P. Petoumenos, G. Keramidas, H. Zeffer, S. Kaxiras, and E. Hagersten. IEEE IISWC 2006.  
**Second most cited paper with 35 citations**

## List of peer-reviewed workshop papers

---

- [W.1] ***Towards Collaborative Performance Tuning of Algorithmic Skeletons***  
C. Cummins, P. Petoumenos, M. Steuwer, and H. Leather. HLPGPU 2016.
- [W.2] ***Iterative Compilation on Mobile Devices***  
P. Mpeis, P. Petoumenos, and H. Leather. ADAPT 2016.
- [W.3] ***Autotuning OpenCL Workgroup Size for Stencil Patterns***  
C. Cummins, P. Petoumenos, M. Steuwer, and H. Leather. ADAPT 2016.
- [W.4] ***Intelligent Heuristic Construction with Active Learning***  
W. Ogilvie, P. Petoumenos, Z. Wang, and H. Leather. CPC 2015.
- [W.5] ***Fast automatic heuristic construction using active learning***  
W. Ogilvie, P. Petoumenos, Z. Wang, and H. Leather. LCPC 2014.
- [W.6] ***Instruction-based Reuse Distance Prediction Replacement Policy*** \*  
P. Petoumenos, G. Keramidas, and S. Kaxiras. JWAC 2010.  
**8th best cache replacement policy out of 25 submissions**
- [W.7] ***Preventing Denial-of-Service Attacks in Shared CMP Caches***  
G. Keramidas, P. Petoumenos, S. Kaxiras, A. Antonopoulos, and D. Serpanos. SAMOS 2006.
- [W.8] ***Statshare: A statistical model for managing cache sharing via decay***  
P. Petoumenos, G. Keramidas, H. Zeffer, S. Kaxiras, and E. Hagersten. MoBS 2006.