PHILIP WOJCICKI

filip.wojcicki18@imperial.ac.uk +48 600 800 121
linkedin.com/in/filip-wojcicki filipwojcicki.com

EDUCATION

IMPERIAL COLLEGE LONDON *'22 - '26*
PhD in High Performance Computing:

- Custom Computing research group
- Software Performance Optimisation research group

Research on **FPGAs**, **AI acceleration**, **compilers**

IMPERIAL COLLEGE LONDON *'18 - '22*
MEng Computer Engineering (eq. *summa cum laude*):

- **Distinguished Final Year Project** (scored 88%)
- Y4: **Ranked #1 of 65, 1st Class Hons** (84%)
- Y3: **Ranked #4 of 65, 1st Class Hons** (80%)
- Y2: **Ranked #1 of 74, 1st Class Hons** (83%)
- Y1: **Ranked #1 of 74, 1st Class Hons** (80%)

CZACKI HIGH SCHOOL *'15 – '18*
Finals: **100th percentile** in all advanced exams

RESEARCH EXPERIENCE

I have **presented my work** at:

- FPT'22, CERN FastML workshops, NANDA'23, School of Microelectronics in Tianjin University, FDF'25, ACACES'25, ScalPerf'25, REACH'25,

I have served as a **reviewer** for:

- FPGA, FCCM, CF, IMPACT, PACT

Since 2020, I have been a **teaching assistant**:

- I provide support during lab sessions, mark coursework/exams, offer 1-to-1 tutorials
- I **develop** and maintain labs and **coursework** about creating a C90 to RISC-V compiler in C++
- Awarded prize for the **Top Undergraduate Teaching Assistant** in 2021

SKILLS

PROGRAMMING

- Advanced in **Python, C/C++, SystemVerilog, RISC-V (Vector Ext.) and ARM Assembly**

TOOLS & TECHNOLOGIES

- Machine Learning frameworks (**PyTorch**, **HLS4ML**)
- Quant. Aware Training (**Brevitas**, **QPyTorch**)
- Hardware design/verification (**Vivado**, **Quartus**)
- Compiler infrastructure (**MLIR**, **LLVM**)
- Linux environment, version control (**Git**, **SVN**)
- Unix shell and scripting (**Bash**, **tcsh**)
- Software profiling (**Intel VTune Profiler**, **perf**)

WORK EXPERIENCE

FRACTILE *Nov '24 – Jan '25*
Low-latency Researcher

- Evaluated latency and area trade-offs of block floating-point formats (**Microscaling**) for an analog in-memory-compute AI-accelerator ASIC.
- Implemented and benchmarked low-latency paths as **ASM RISC-V (RVV)** kernels and **SystemVerilog** blocks to guide HW/SW partitioning.

JUMP TRADING *Mar '21 – Sep '21*
Software & Hardware Engineer

- Created a library in **Python** for exploring topology, configuring, connecting and graphing multi-device (ASIC, FPGA, CPU, GPU) hardware systems
- Built an efficient, distributed arbitrage trading **C++** application composed of several parallel processes running on ASICs, FPGAs, and x86 machines
- Configured and benchmarked formal verification using **SystemVerilog** and a novel custom **Python** tool for use in ASIC and FPGA development

JUMP TRADING *Sep '20 – Oct '20*
Software & Hardware Engineer

- Built an ultra low latency ASIC validation test platform for floating point calculations on x86 and RISC-V architectures using **C++**, **C** and **Python**

ARM *Jun '20 – Sep '20*
Systems Architect

- Improved autonomous driving platform's verification in **SystemVerilog**, overhauled its documentation, added support for formal (**JasperGold**)

RESEARCH PROJECTS

LOW-BIT OPS FOR BLOCK FLOATING POINT *May '25 - Nov '25*

- Designed a bit-accurate hardware–software (**C++**/**SystemVerilog**) framework for compensated accumulation with MX-like formats as adder trees
- Built an automated FPGA design-space exploration (DSE) tool to suggest Pareto-optimal MX-like accumulators

TRANSFORMER NEURAL NETWORK ON FPGA *Nov '21 - Dec '22*

- Developed a novel TNN architecture for GPUs (PyTorch) and FPGAs (HLS) for High Energy Physics experiments in collaboration with CERN
- FPGA solution outperformed SoTA models on GPU by ∼**1000 times** thanks to software/hardware-aware optimisations, without accuracy loss
- Experimented with quantisation-aware training (QAT), developed a quick FPGA-friendly post-training quantisation (PTQ) scheme for HLS4ML

QUBIT VISION TRANSFORMER ON FPGA *Jan '23 - Jan '24*

- Designed a low-latency FPGA pipeline (CameraLink) for real-time qubit-state classification in collaboration with experimental quantum physicists
- Optimised the ViT architecture to achieve millisecond-scale end-to-end detection with up to ∼**120**× **lower latency** than a GPU baseline

LIGHTWEIGHT ZERO-SHOT LEARNING *April '23 - Dec '23*

- Developed a lightweight zero-shot learning framework with attribute knowledge graphs, reducing parameters by ∼**100**× while retaining accuracy
- Designed an accelerator on FPGA for CNN feature extraction and attribute recognition, achieving ∼**67**× speedup over a software-only baseline

POLYHEDRAL SYSTOLIC ARRAY COMPILER *Jan '23 - Dec '23*

- Developed an MLIR-based HLS framework that maps C/C++ and PyTorch programs to systolic-array accelerators using polyhedral optimisations
- Achieved SoTA performance while enabling MLIR pass interchangeability