IBM Research –Zurich

neuroSoC

Master's Thesis – Code Generation for Analog In-Memory Computing SoCs

For decades, conventional computers based on the von Neumann architecture have performed computation by repeatedly transferring data between their processing and their memory units, which are physically separated. As computation becomes increasingly data-centric and as the scalability limits in terms of performance and power are being reached, alternative computing paradigms are searched for in which computation and storage are collocated. A fascinating new approach is that of analog in-memory computing where the physics of nanoscale memory devices are used to perform certain computational tasks within the memory unit in a non-von Neumann manner.

IBM has been spearheading this exciting new paradigm by fabricating chips (see the recent cover of Nature Electronics) that integrate emerging non-volatile memory devices in the backend of advanced CMOS technology nodes [1].

We are inviting applications from students enrolled in the ETH Master's program to conduct their thesis work at IBM Research – Zurich on this exciting new topic. The research focus will be on scheduling optimization strategies and contributing to a code generation infrastructure for executing inference workloads on the state-of-the-art Multi-Processor System on Chip (MPSoC) being developed in 28nm CMOS within the NeuroSoC project (https://neurosoc.eu/). The chip's heterogeneous architecture comprises analog in-memory computing cores leveraging the ultra-dense phase-change memory technology developed by ST Microelectronics along with digital RISC-V cores. A complete software toolchain is being developed to deploy AI Deep Learning algorithms optimized for the chip, in which the work conducted within this thesis will be integrated. NeuroSoC is a large EU-funded project driven by ST Microelectronics and IBM Research, in which many other academic and industrial partners participate including ETH. The ideal candidate should have a strong mathematical aptitude and programming skills.

Note that ETH does not allow MSc students to be paid when the thesis is conducted externally. The project would start on the 1st of September 2024 (this is flexible) and will last for a period of 6 months.

If you are interested in this challenging position on an exciting new topic, please send your most recent curriculum vitae including a transcript of grades by email to: Dr. Irem Boybat (ibo@zurich.ibm.com), Dr. Manuel Le Gallo (anu@zurich.ibm.com) and Dr. Abu Sebastian (ase@zurich.ibm.com).

[1] Le Gallo, M., Khaddam-Aljameh, R., Stanisavljevic, M. et al. A 64-core mixed-signal in-memory compute chip based on phase-change memory for deep neural network inference. *Nat Electron* 6, 680–693 (2023). https://doi.org/10.1038/s41928-023-01010-1