Analog Optical Computing for Machine Learning and Optimization

H Kremer[1], H Ballani[1], G Brennan[1], B Canakci[1], J Chu[1], H Clegg[1], D Cletheroe[1], G Gkantsidis[1], J Gladrow[1], K P Kalinin[1], D J Kelly[1], G O'Shea[1], F Parmigiani[1], L Pickup[1], B Rahmani[1], and A Rowstron[1]

[1] *Cloud Systems Futures, Microsoft Research, Cambridge, UK*
Contact Email: t-hkremer@microsoft.com

Digital computing is reaching its fundamental physical limits just as compute-intensive workloads like machine learning are taking off. We propose a next generation hardware accelerator tailored to machine learning and optimization tasks. Our accelerator combines low cost and scalable commodity optics technologies with analog electronics to achieve an estimated 100x increase in throughput per energy compared to current state-of-the-art GPUs. By retaining the information in a feedback loop in the analog domain until convergence to a fixed point, our accelerator is particularly suited for iterative methods and energy-based models. We demonstrate the feasibility of our approach in hardware for real-world optimization problems as well as regression and classification tasks. The large speed-up at low power and low cost highlights the potential of analog optical computing in the post-Moore Law's era.

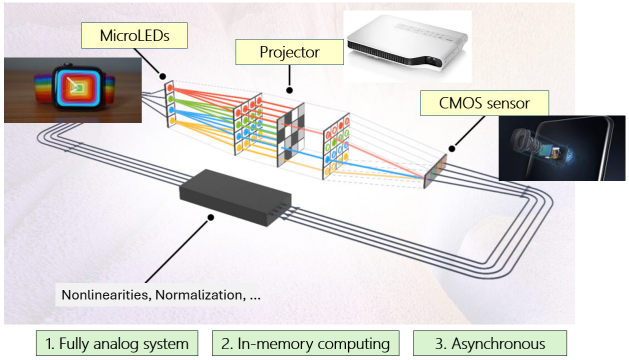

Figure 1: Analog optical computing. By combining commodity optics technologies with analog electronics, we achieve an estimated 100x increase in throughput per energy compared to the latest GPU architectures