Road to scalability for efficient graph search on massively parallel neuromorphic hardware

Oskar von Seeler (University Medical Center Göttingen)

Abstract:
Efficient computation of shortest paths in weighted graphs is a fundamental problem with many applications. Neuromorphic hardware platforms promise massively parallel, efficient computation, changing parallelism tradeoffs.
In this work, We introduce NEURO-MAPP (Neuromorphic-based Min-Add Parallel Propagation), a distributed shortest path algorithm designed to use local computation and network communication available in neuromorphic systems. We provide an optimized implementation of the algorithm on the SpiNNaker 2 platform and evaluate its performance on a selection of synthetic and real-world graphs. These results are compared to Dijkstra's algorithm on a modern CPU. We find that the NEURO-MAPP implementation scales favorably in terms of runtime for many graph types while consuming less energy per shortest-path query than the CPU implementation in almost all cases.
These findings highlight the potential of neuromorphic hardware featuring sparse, spike-based communication as a scalable and energy-efficient platform for computation in graph search and related tasks.