

Towards brain-scale simulation with structural plasticity: A GPU port of RELeARN

Authors:

Marvin Kaster, Fabian Czappa, Felix Wolf

Brain simulations are essential to understanding the brain, but with approximately 10^{11} neurons and 10^{14} synapses, simulating that enormous scale requires extremely large clusters and highly efficient implementations. This challenge intensifies when synapses are allowed to be dynamically created and deleted during simulation. This process is called structural plasticity, and growing evidence shows that it plays an essential role even in the adult brain and is required to understand diseases and learning. Our simulator, RELeARN, supports efficient structural plasticity using a Barnes-Hut approximation, but was previously restricted to CPU execution. Porting it to GPU architectures is challenging because it requires a minimal memory footprint, support for a dynamically changing network graph, and efficient GPU implementations of both the Barnes-Hut approximation and spike transmission. We address these challenges, enabling GPU support that massively accelerates our simulations, bringing us closer to simulating the entire brain with structural plasticity.