

Accelerated Large Graph Creation with CUDA: Empowering Graph Neural Networks in Particle Physics

Graph neural networks have emerged as a powerful tool in various physics studies, particularly in the analysis of sparse and heterogeneous data. However, as the field of particle physics advances towards utilizing graphs in high-luminosity scenarios, a new challenge has emerged: efficient graph creation. While GNN inference is highly optimized, graph creation has not received the same level of attention. This becomes a significant bottleneck when working with graphs that have a large number of nodes such as those arising from HL-LHC collisions.

In this talk, we will delve into the specifics of this issue and explore the impact of graph creation on performance. We will highlight that even for graphs with a specific structure in mind, the process of creating them can be several orders of magnitude slower on a CPU compared to running inference with a network. To overcome this limitation, we will investigate how hardware-based acceleration, specifically utilizing CUDA on GPUs, can address these challenges.

By leveraging CUDA acceleration, we can unlock the full potential of graph neural networks without being hindered by inefficient graph creation. This talk aims to highlight the importance of fast and efficient graph creation in the context of large-scale graphs in particle physics and how leveraging hardware acceleration can propel research in this exciting field.

Co-authors: SHIELDS, Edward; KAKATI, Nilotpai (Weizmann Institute of Science (IL))

Presenter: KAKATI, Nilotpai (Weizmann Institute of Science (IL))