


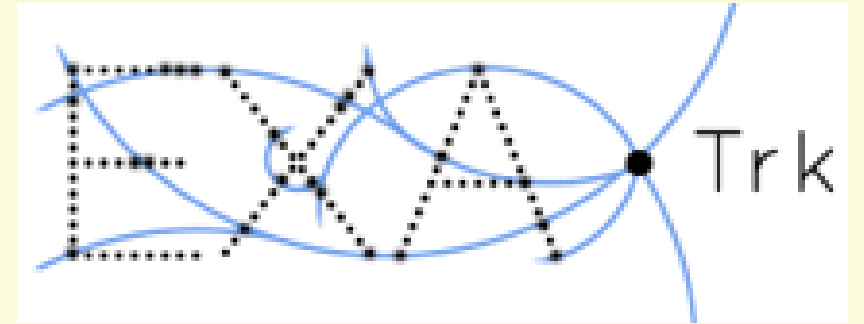
Optimizing the Exa.TrkX Inference Pipeline for Manycore CPUs

Nirajan Acharya¹, Emma Liu², Alberto Lucas³ and Alina Lazar¹
on behalf of the **Exa.TrkX** collaboration

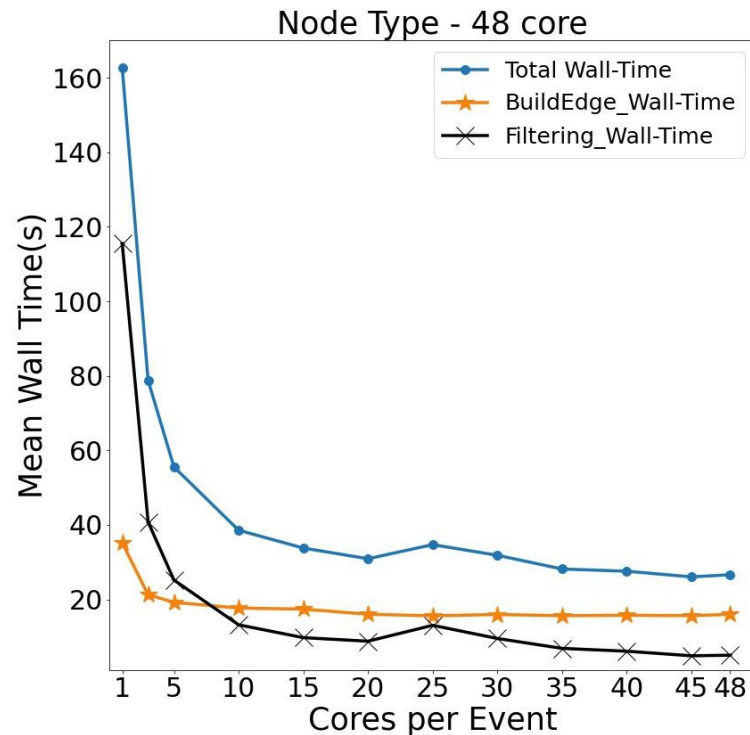

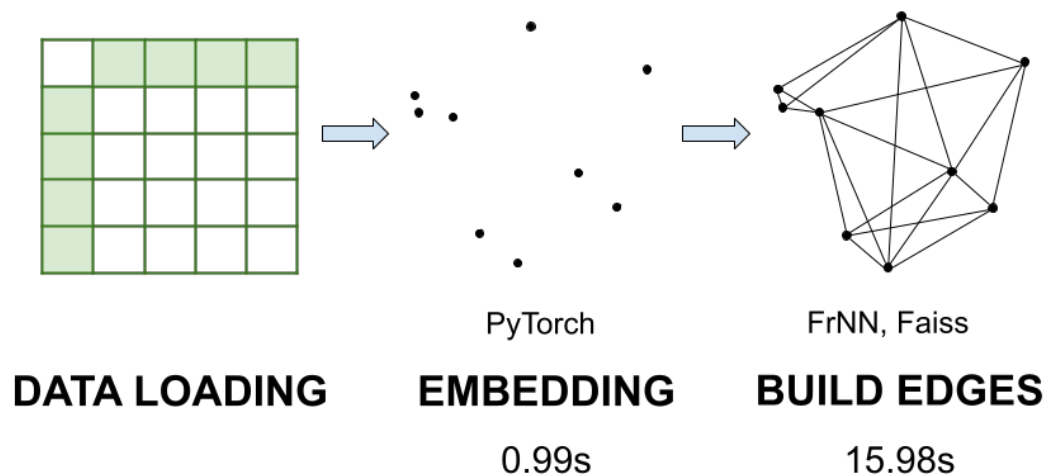
¹Youngstown State University,

²University of California, Los Angeles,

³California State University, Monterey Bay



The Exa.TrkX Inference Pipeline



MPI was used to run events in parallel, using multiple cores.

The most time-consuming steps of the pipeline are Build Edges and Filtering. To speed-up Build Edges we used Faiss with 2 threads and multiprocessing for the Filtering for-loop.

The results indicate that it is best to use between 10 and 15 cores per event, however running it on the GPU is still 38 times faster.