

ExaGraph: Combinatorial Methods for Enabling Exascale Applications




Pacific Northwest
NATIONAL LABORATORY


Proudly Operated by **Battelle** Since 1965

Mahantesh Halappanavar¹, Aydin Buluç², Erik Boman³, Alex Pothen⁴,
Antonino Tumeo¹, Arif Khan¹, Marco Minutoli¹, Nathan Tallent¹, Nitin Gawande¹,
Saliya Ekanayake², Sayan Ghosh¹, Seher Acer³, Siva Rajamanickam³, S. M. Ferdous⁴
¹Pacific Northwest National Laboratory, ²Lawrence Berkeley National Laboratory, ³Sandia National Laboratories,
⁴Purdue University

Overview




Approximation algorithms for b-Edge cover problem




Example: A b-Matching of maximum weight with $b=2$

- Applied improved edge cover algorithms to solve the adaptive anonymity problem
- Enables solution of problems three orders of magnitude larger than previous work, in minutes on parallel machines




cuRipples: Scalable Influence Maximization on multi-GPUs




- Problem: Given a network $G=(V,E,\omega)$ and a model M of a diffusion process, the goal is to find k initial adopters who will lead to most number of adoptions
- 790x speedup over a state-of-the-art serial implementation on 8 GPUs and 80 cores

Scalable graph $\frac{1}{2}$ approximate matching and clustering





- Developed UPCXX versions of $\frac{1}{2}$ approx. matching (joint work with LBL Pagoda team)
- UPCXX-RMA version beats the MPI RMA (+neighbor collective) by 2-4.6x
- UPCXX-RPC version improves productivity (reduced LoC)



- Vite performs graph clustering
- Results show good performance and quality on KNL nodes of ALCF Theta

SPHYNX: Spectral Partitioning on HYbrid and aXelerator-enabled systems

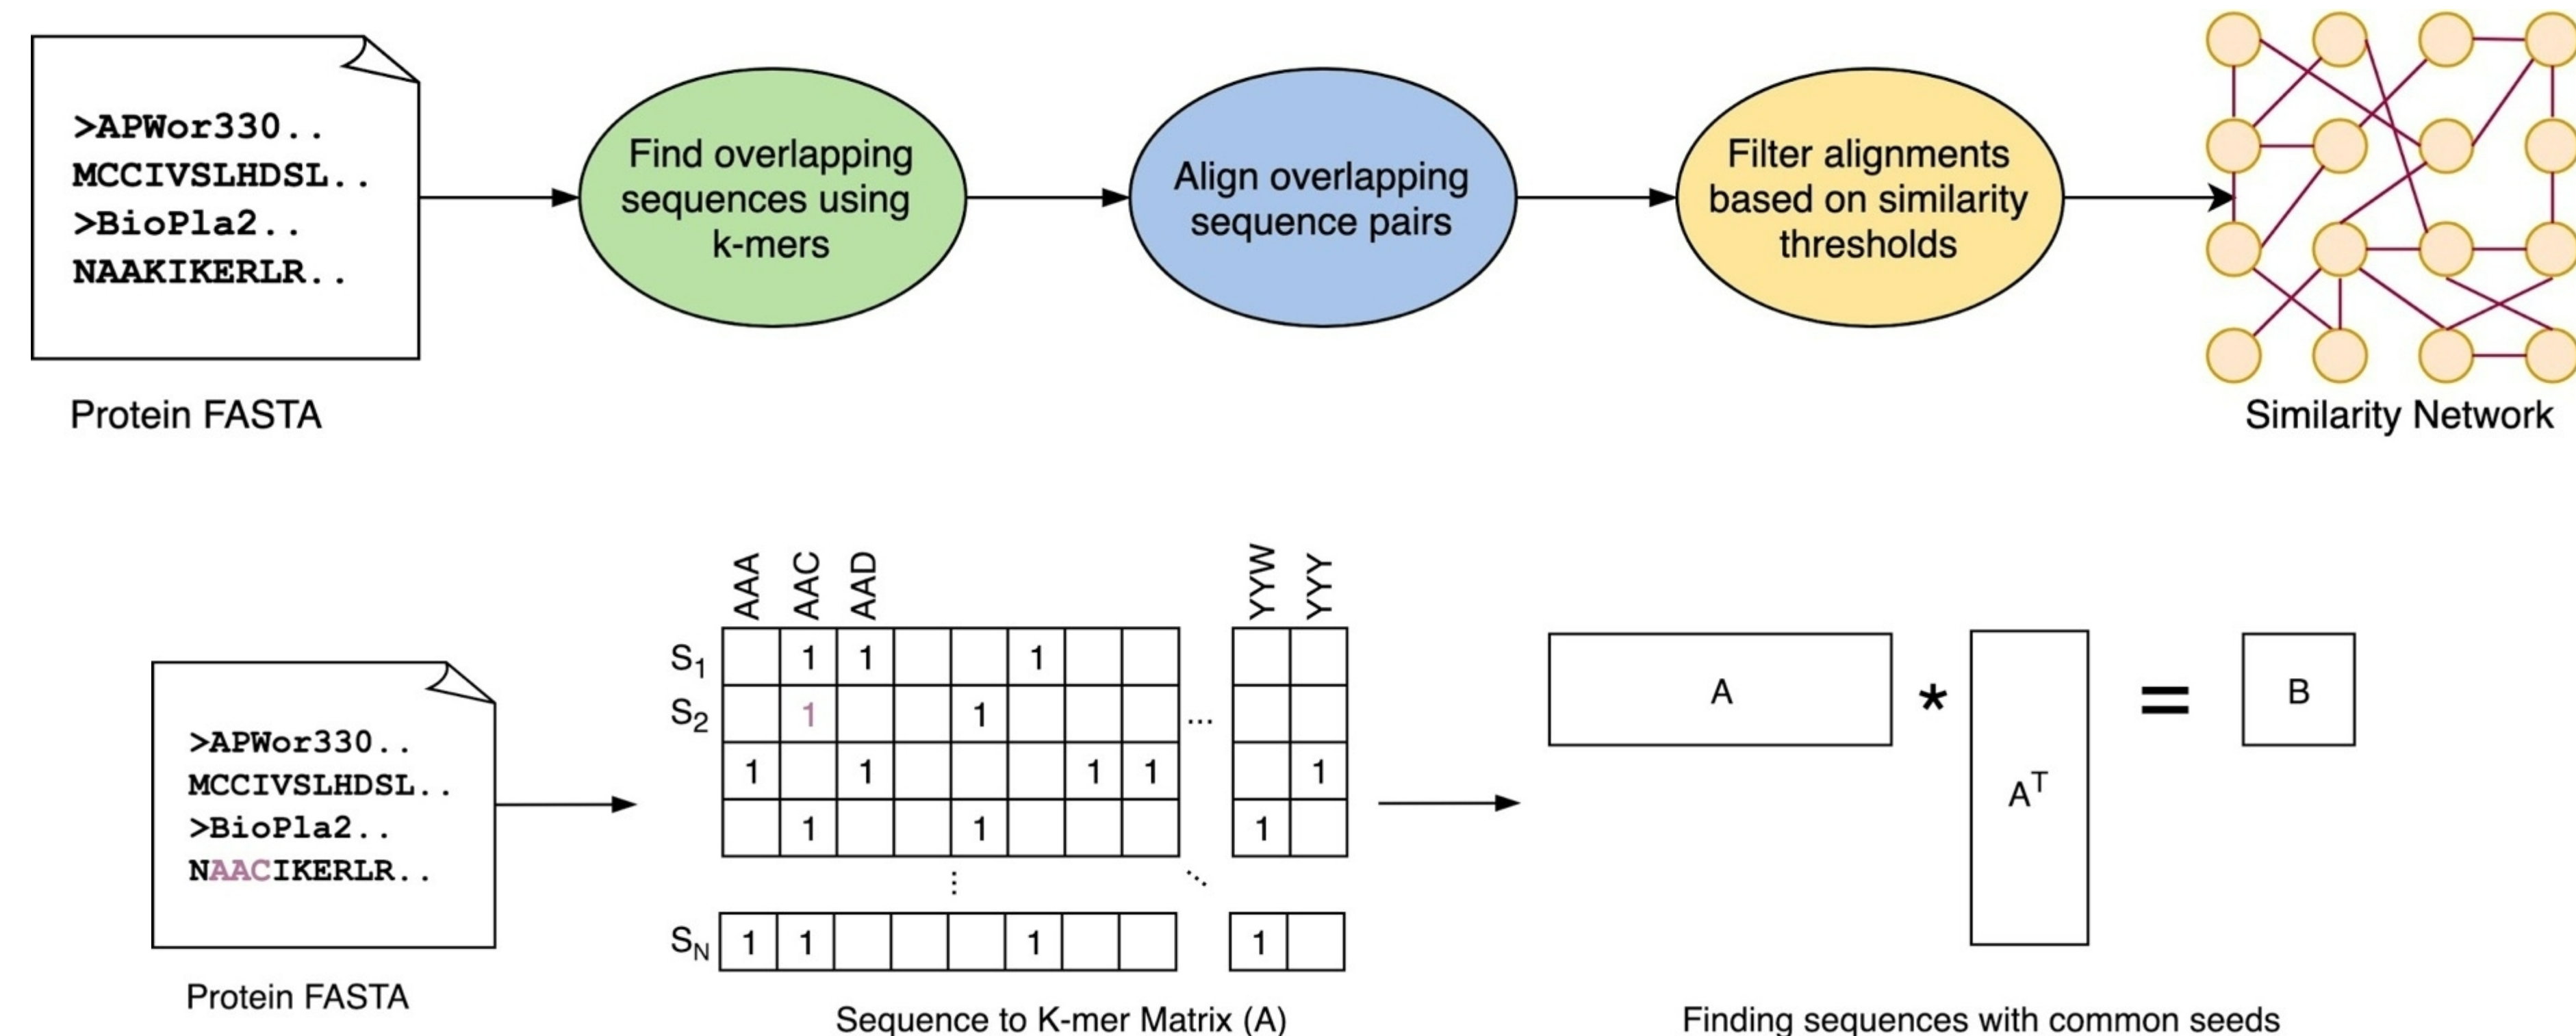
- First graph partitioning tool for GPU-enabled distributed-memory systems
- Based on Trilinos, uses Kokkos & KokkosKernels for portability across architectures



Comparison on 24 GPUs against ParMETIS (24 MPI ranks) on Summit:

- On ASIC_680k, it is 54x faster with 1.66x increase in cutsizes
- On com-Orkut, it is 27x faster with 1.95x increase in cutsizes
- On cit-Patents, it is 2.3x faster with 3.79x increase in cutsizes

Fully-parallel protein similarity graph computation



- Similarity graph generation is the bottleneck for protein clustering application in ExaBiome
- Both the k-mer index creation and querying are fully parallelized
- Relies on sparse matrix computations, followed by pairwise alignment.
- Generalizable to more sensitive indices through inexact k-mer matching

