Accelerating metagenome assembly on GPUs

Muaaz Awan
Application Performance Group

NERSC
BERKELEY LAB
U.S. DEPARTMENT OF ENERGY | Office of Science

Microbiomes

- Microbiota consists of a group of diverse micro organisms.
- These microbial communities form associations with their environments and have profound effect on them.
- For instance, 1-3% of a human body's mass is made up of its microbiome.
- Metagenomic analysis helps characterize the taxonomic diversity in these communities and provide functional annotation of individual organisms and as a group.

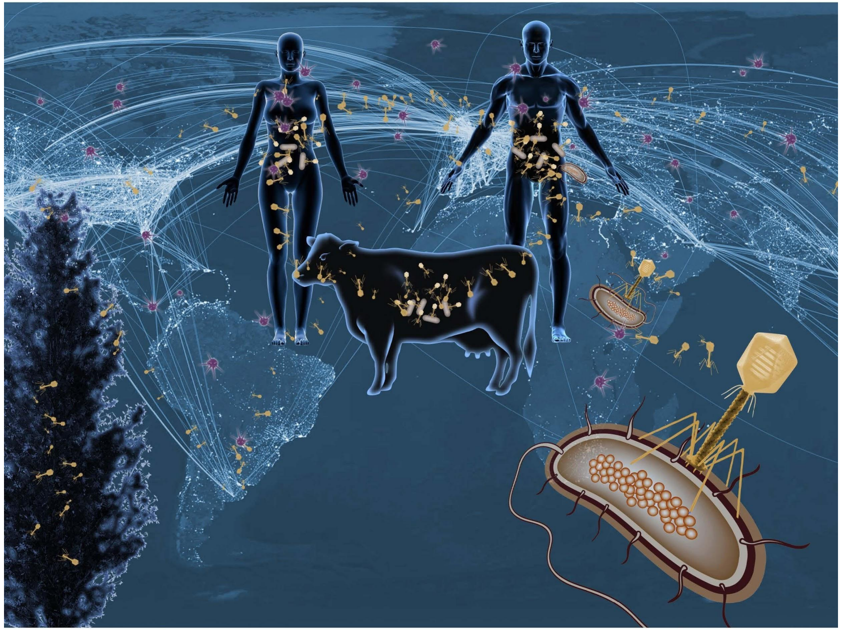

Genome assembly 101

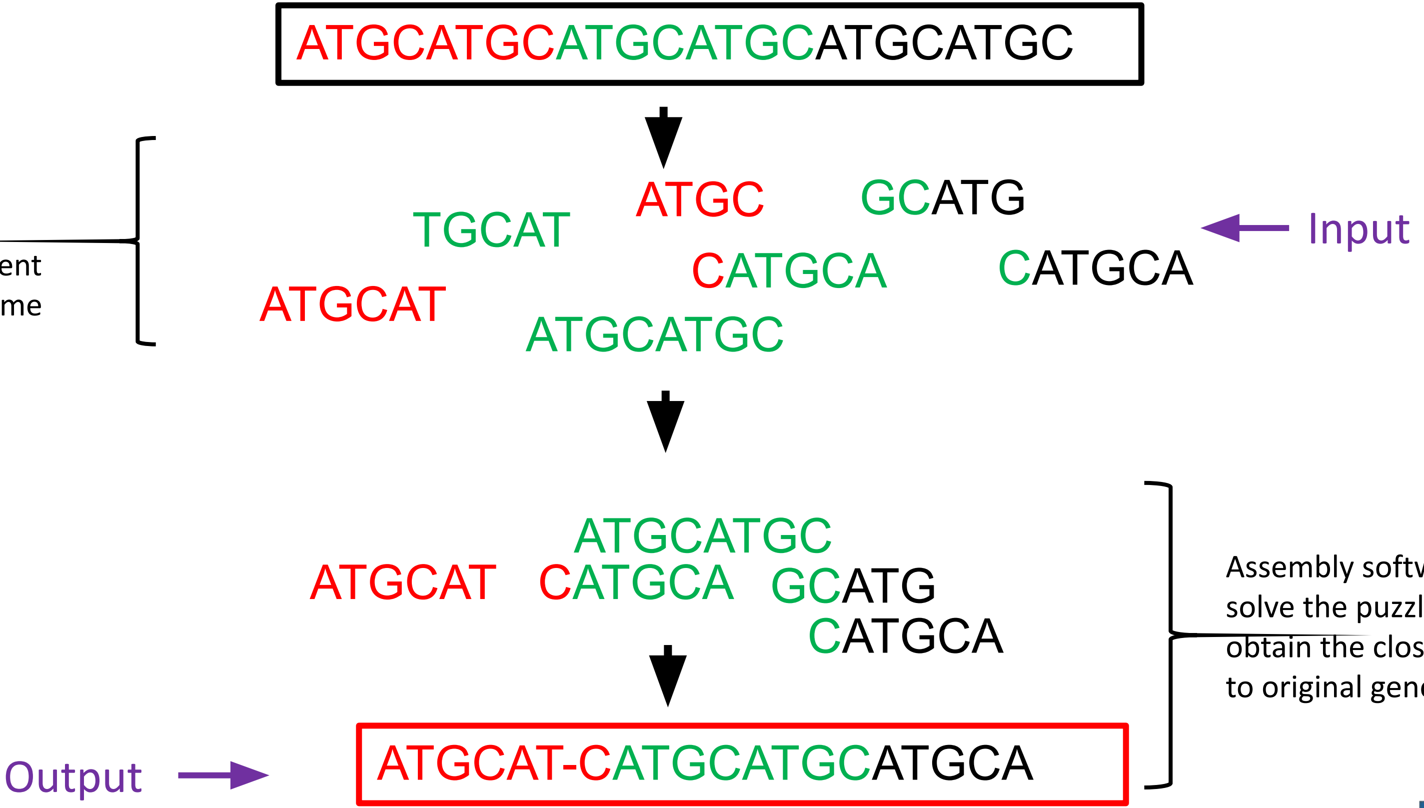

ATGCATGCATGCATGCATGCATGC

TGCAT ATGC GCATG

CATGCA CATGCA

ATGCAT ATGCATGC

Input

Sequencing technologies sample different parts of genome

ATGCATGC

ATGCAT CATGCA GCATG

CATGCA

Assembly software solve the puzzle to obtain the closest to original genome

Output

ATGCAT-CATGCATGCATGCA

Metagenome Assembly

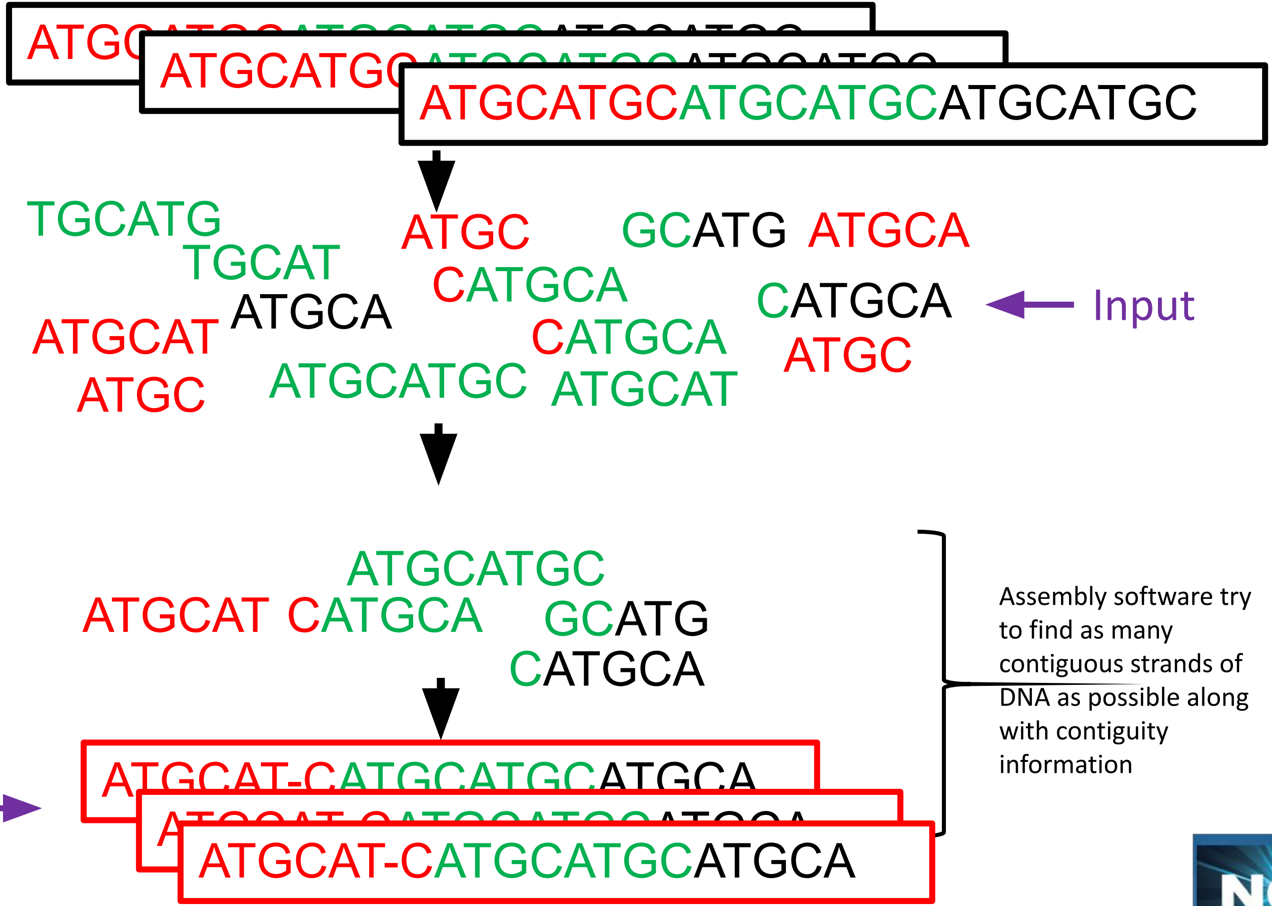

ATGC

ATGCATGC

ATGCATGCATGCATGCATGCATGC

Reads from all the genomes are sequenced together

TGCATG TGCAT ATGCAT ATGCA ATGC ATGC CATGCA ATGCATGC CATGCA ATGCAT GCATG ATGCA CATGCA ATGC

Input

ATGCATGC

ATGCAT CATGCA GCATG

CATGCA

Output

ATGCAT-CATGCATGCATGCA

Assembly software try to find as many contiguous strands of DNA as possible along with contiguity information

ExaBiome: Project Overview

Exascale algorithms & systems for previously intractable problems

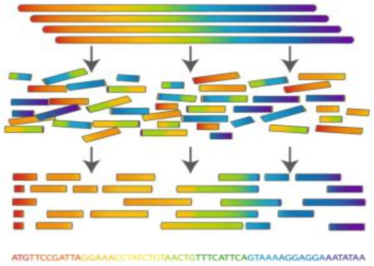

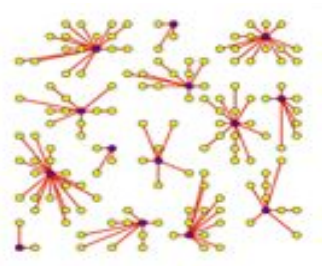

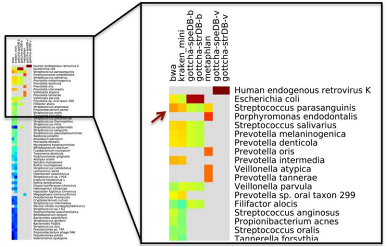

Problem Domain	Metagenome Assembly	Protein Clustering and Annotation	Comparative Analysis
Science need	Find species, genes and relative abundance in microbial communities	Improve understanding of tree of life for microbes; aid in identifying gene function	Track microbiome over time or space, changes in environment, climate, etc.
Computing techniques	hash tables, alignment, k-mer counts, graph walks	Hash tables tables, alignment, k-mer counts, sparse matrices, ML (clustering, GNNs)	hash tables, alignment, k-mer counts, ML (dimensionality reduction)

http://exabiome.org

Motifs of Genomic Data Analysis

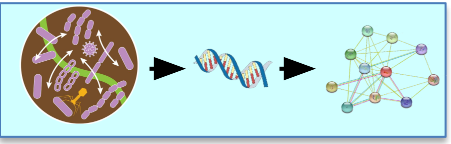

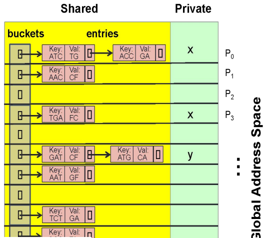

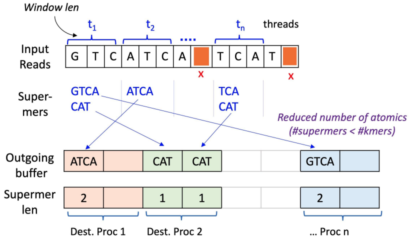

GPU challenges:

<u>Computation</u>

- Random access
- Poor locality
- Strings (little float pt)
- Varying parallelism
- Load imbalance

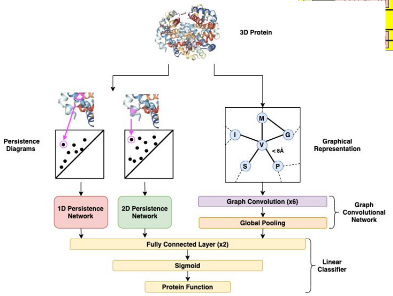

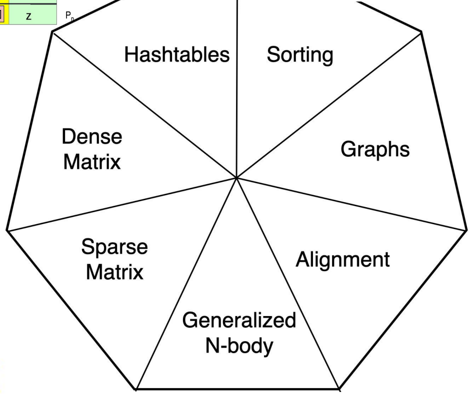

Hashtables

Sorting

Dense Matrix

Graphs

Sparse Matrix

Alignment

Generalized N-body

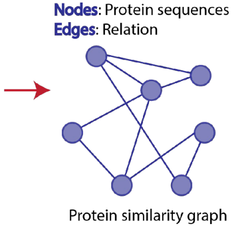

Nodes: Protein sequences

Edges: Relation

Protein similarity graph

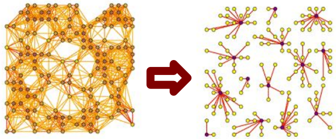

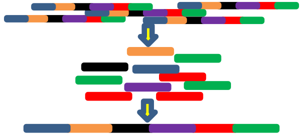

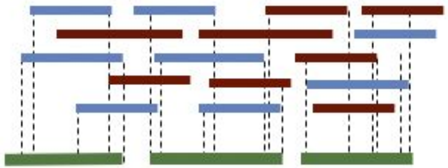

Yelick, et al. “The Parallelism Motifs of Genomic Data Analysis”, Philosophical Transactions A, 2020

Motifs of Genomic Data Analysis

Short Read Assembly:
- Dynamic Programming Algorithms.
- Distributed and local Graph Traversals.
- Distributed and local hash tables.

Long Read Assembly:
- Dynamic Programming Algorithms.
- Sparse matrix multiplication.
- Distributed hash tables.

Protein Similarity and Clustering:
- Dynamic Programming Algorithms.
- Sparse Matrix Multiplication.

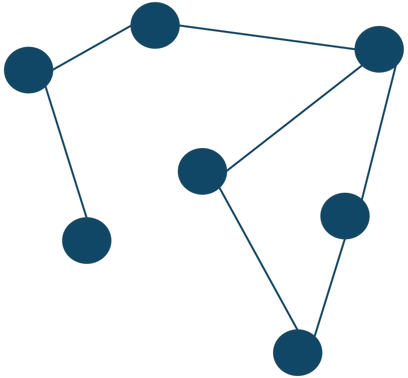

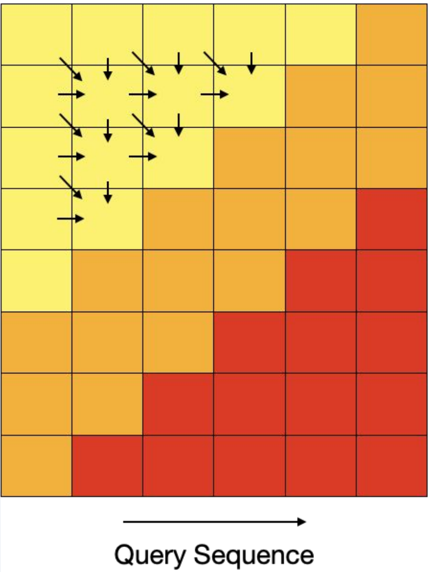

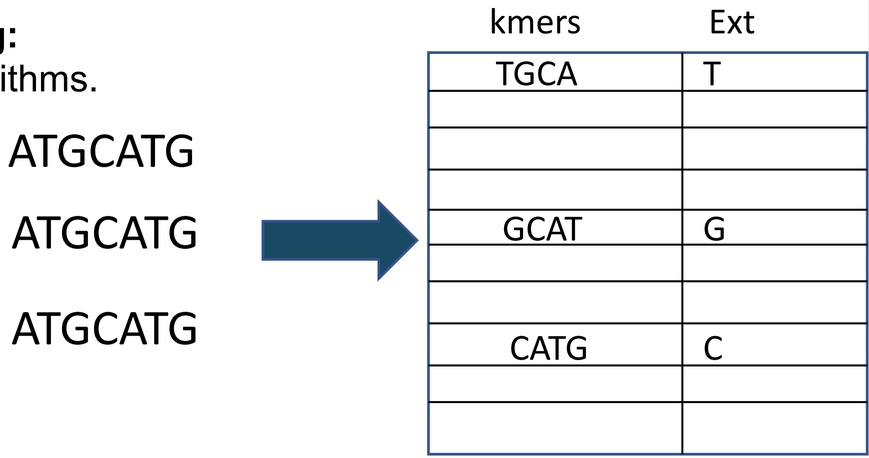

ATGCATG

ATGCATG

ATGCATG

kmers	Ext
TGCA	T
GCAT	G
CATG	C

Yelick, et al. “The Parallelism Motifs of Genomic Data Analysis”, Philosophical Transactions A, 2020

Challenges on GPUs

Ideal for GPUs:

- Localized and predictable memory access pattern.
- Alot of computations per each memory access.
- Equal amount of work can be distributed across threads.

What we have:

- Random or along diagonals memory access pattern.
- Integer only computations bound by memory bandwidth.
- Non deterministic amount of work.
- Varying or limited parallelism (DP and graph algorithms).

MetaHipMer

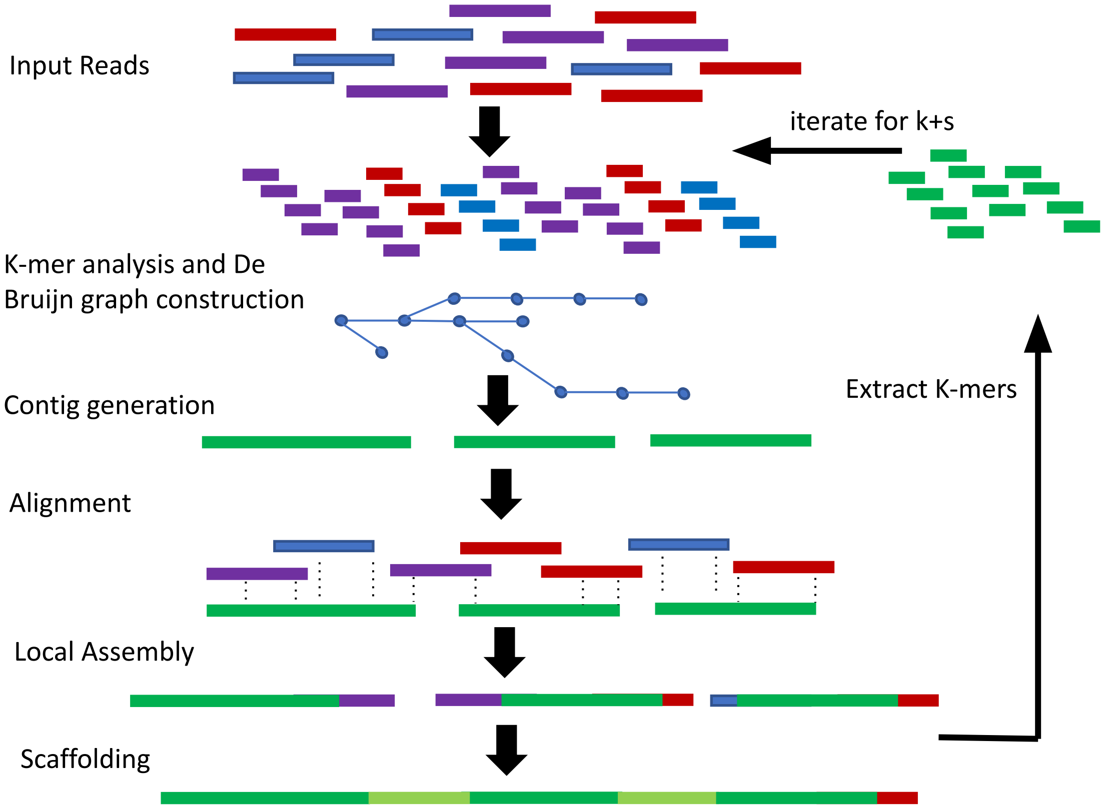

Input Reads

K-mer analysis and De Bruijn graph construction

Contig generation

Alignment

Local Assembly

Scaffolding

iterate for k+s

Extract K-mers

*ADEPT: A GPU Accelerated Sequence Aligner

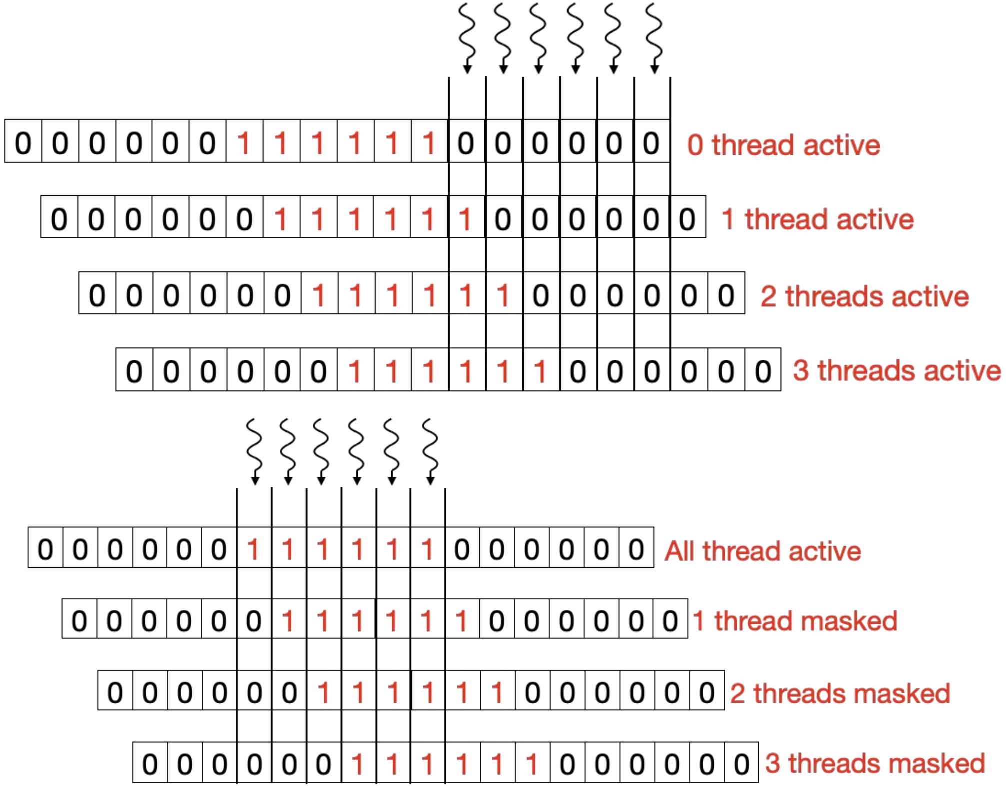

0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 0 thread active
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 1 thread active
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 2 threads active
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 3 threads active

0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 All thread active
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 1 thread masked
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 2 threads masked
0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 3 threads masked

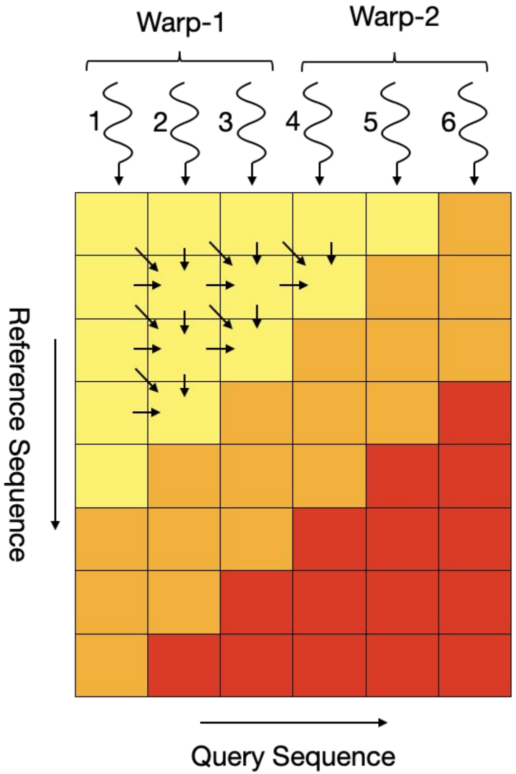

*Awan, Muaaz G., et al. "ADEPT: a domain independent sequence alignment strategy for gpu architectures." BMC bioinformatics 21.1 (2020): 1-29.

Reorganizing data for better locality

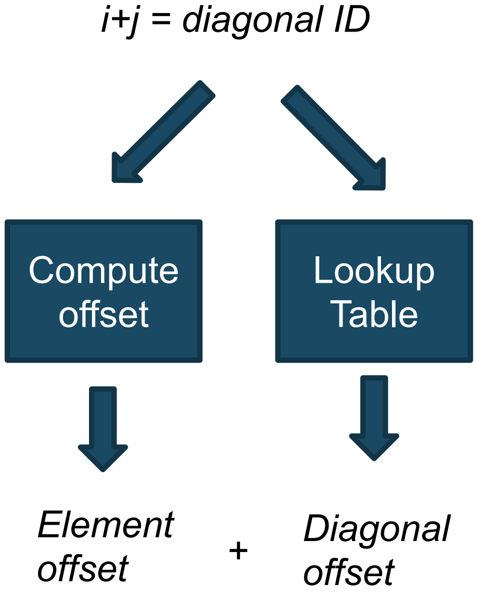

i+j = diagonal ID

Compute offset

Lookup Table

Element offset + *Diagonal offset*

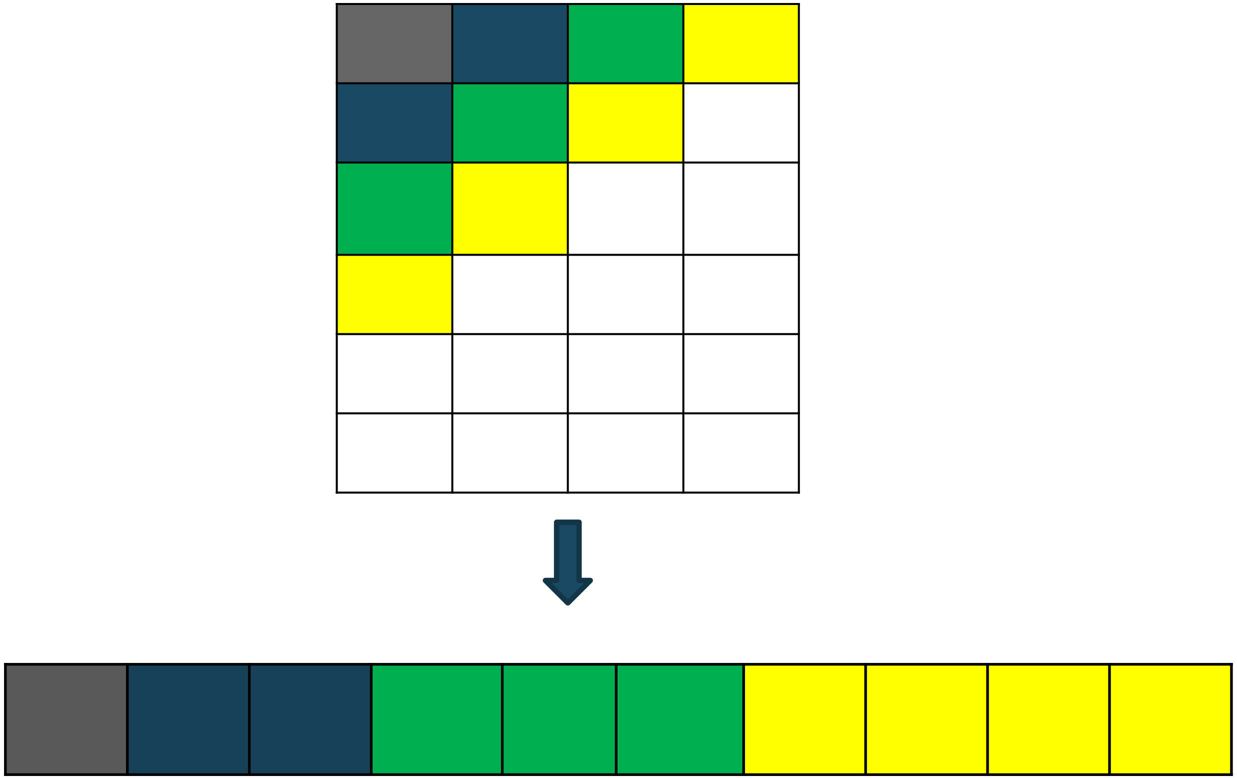

Diagonal Major Indexing

ADEPT’s performance on DNA sequences

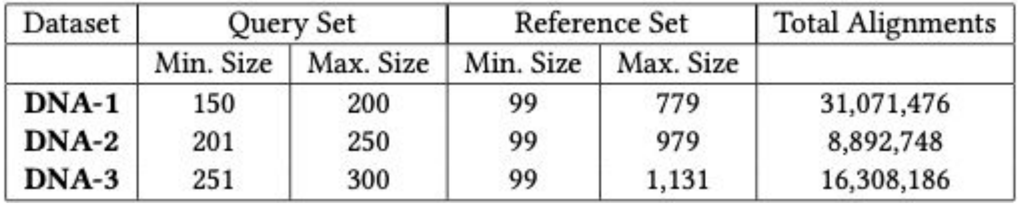

Dataset	Query Set		Reference Set		Total Alignments
	Min. Size	Max. Size	Min. Size	Max. Size	
DNA-1	150	200	99	779	31,071,476
DNA-2	201	250	99	979	8,892,748
DNA-3	251	300	99	1,131	16,308,186

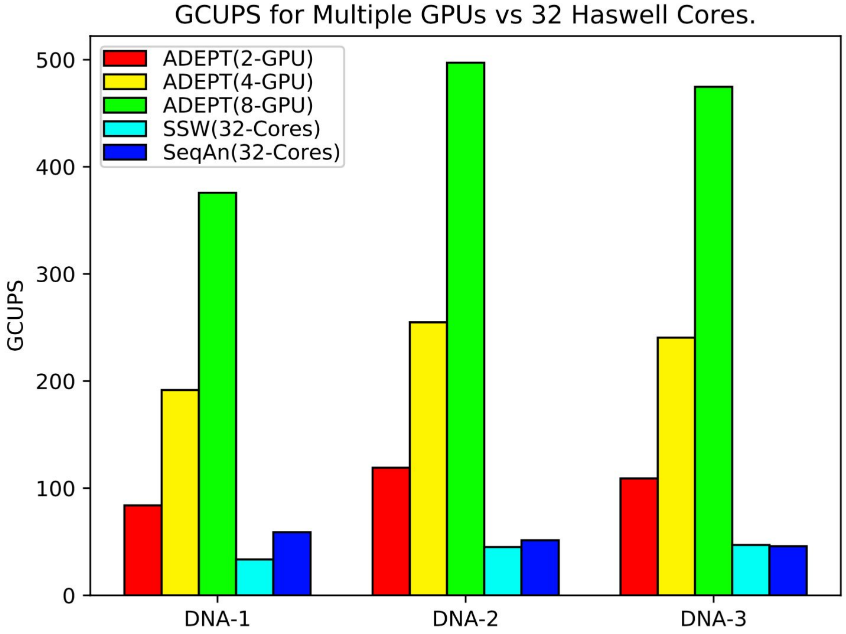

*Hofmeyr, Steven, et al. "Terabase-scale metagenome coassembly with MetaHipMer." *Scientific reports* 10.1 (2020): 1-11.

ADEPT’s performance on Protein sequences

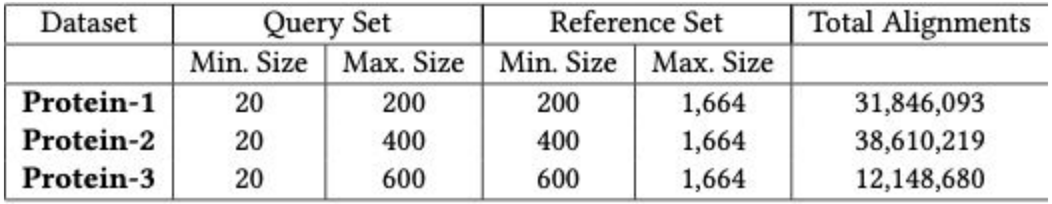

Dataset	Query Set		Reference Set		Total Alignments
	Min. Size	Max. Size	Min. Size	Max. Size	
Protein-1	20	200	200	1,664	31,846,093
Protein-2	20	400	400	1,664	38,610,219
Protein-3	20	600	600	1,664	12,148,680

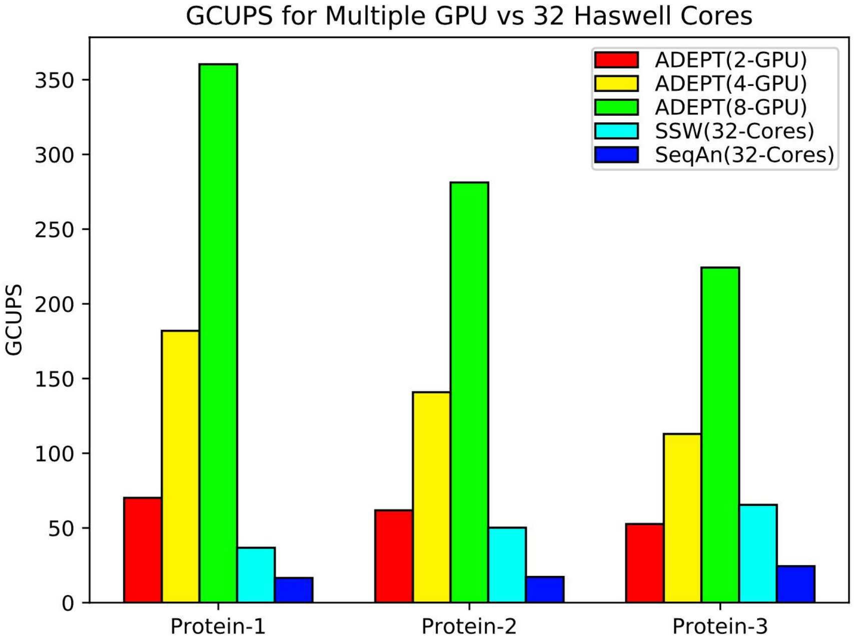

*Fox, Naomi K., Steven E. Brenner, and John-Marc Chandonia. "SCOPe: Structural Classification of Proteins—extended, integrating SCOP and ASTRAL data and classification of new structures." *Nucleic acids research* 42.D1 (2014): D304-D309.

ADEPT in MetaHipMer

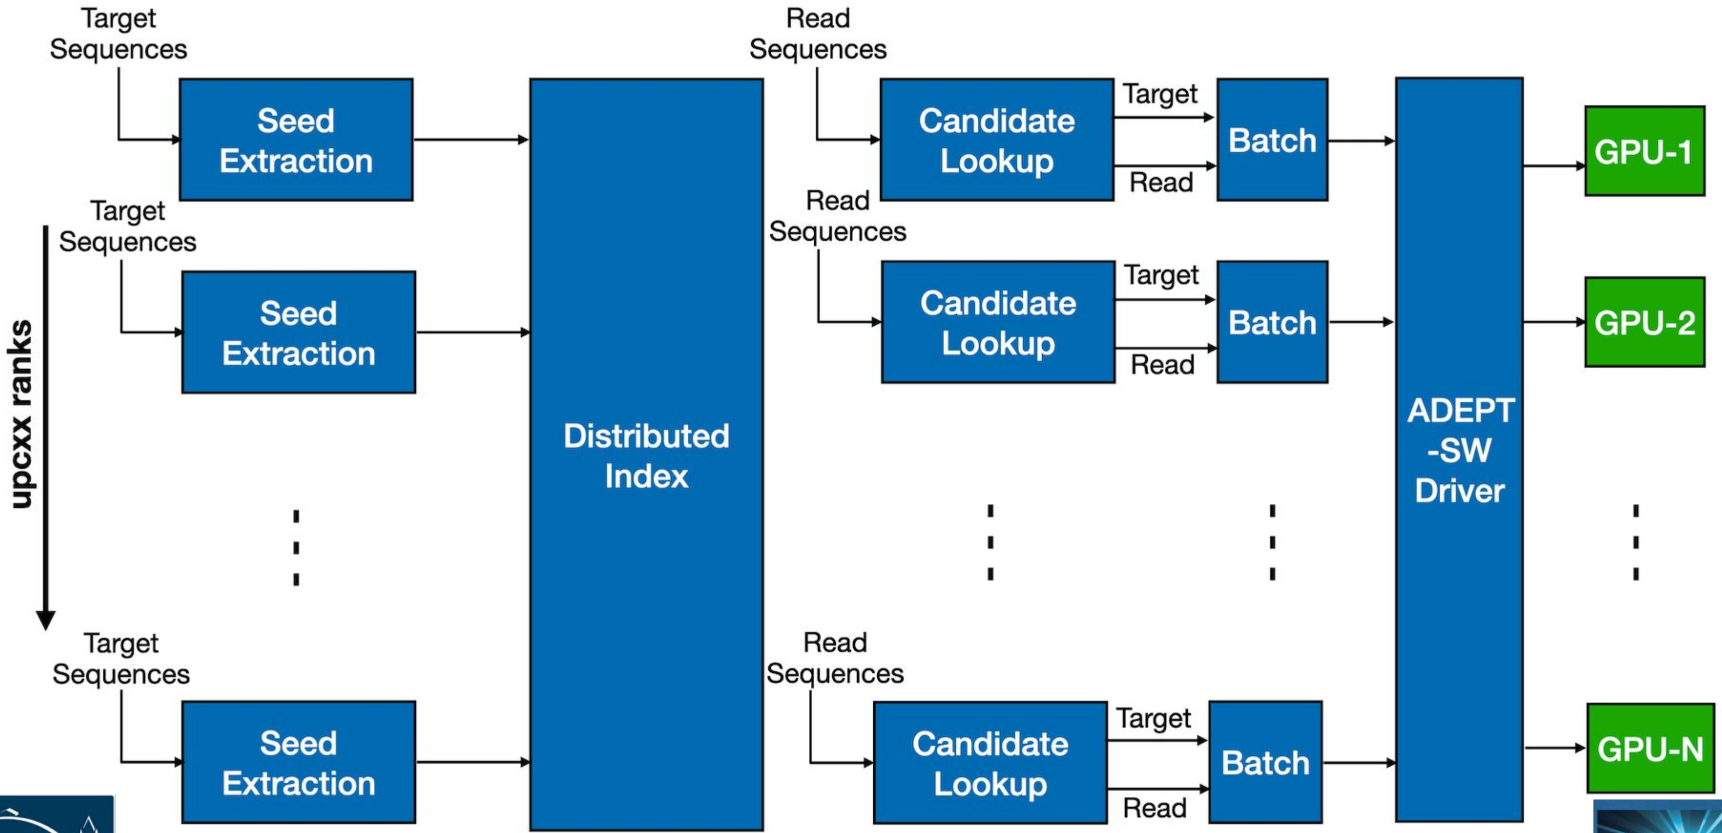

ADEPT in MetaHipMer

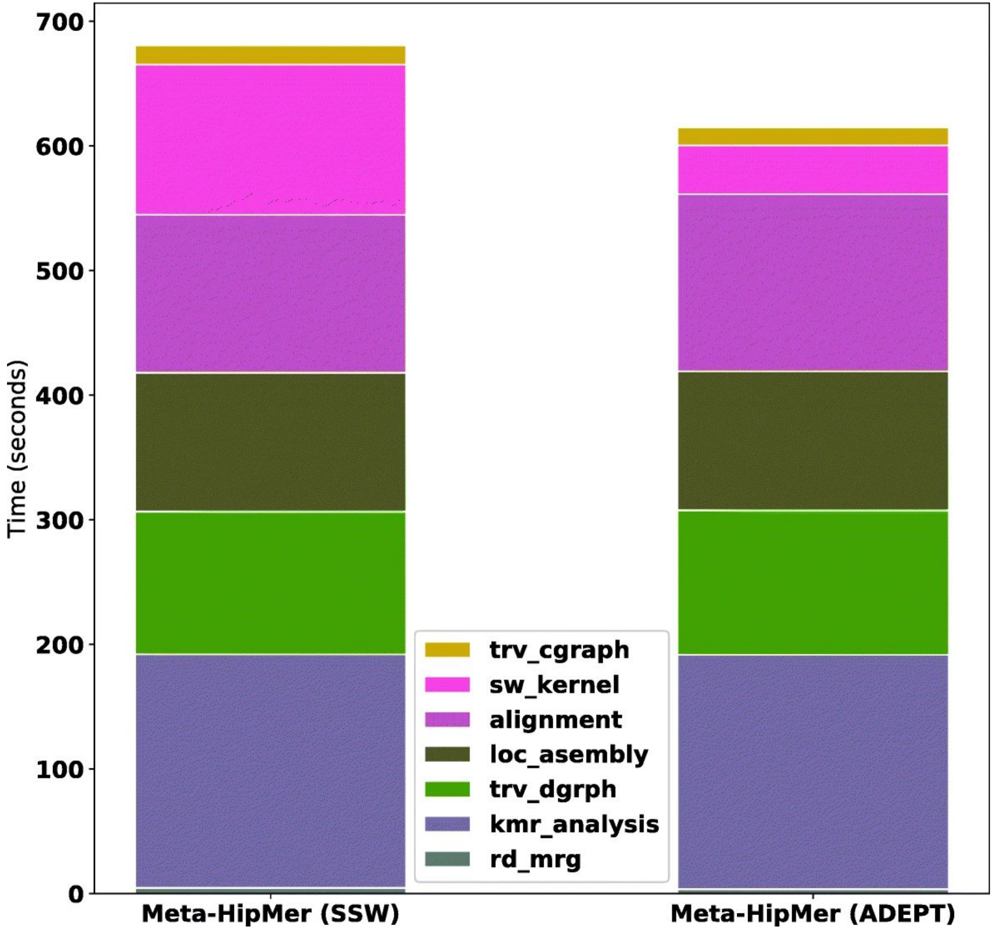

ADEPT in PASTIS

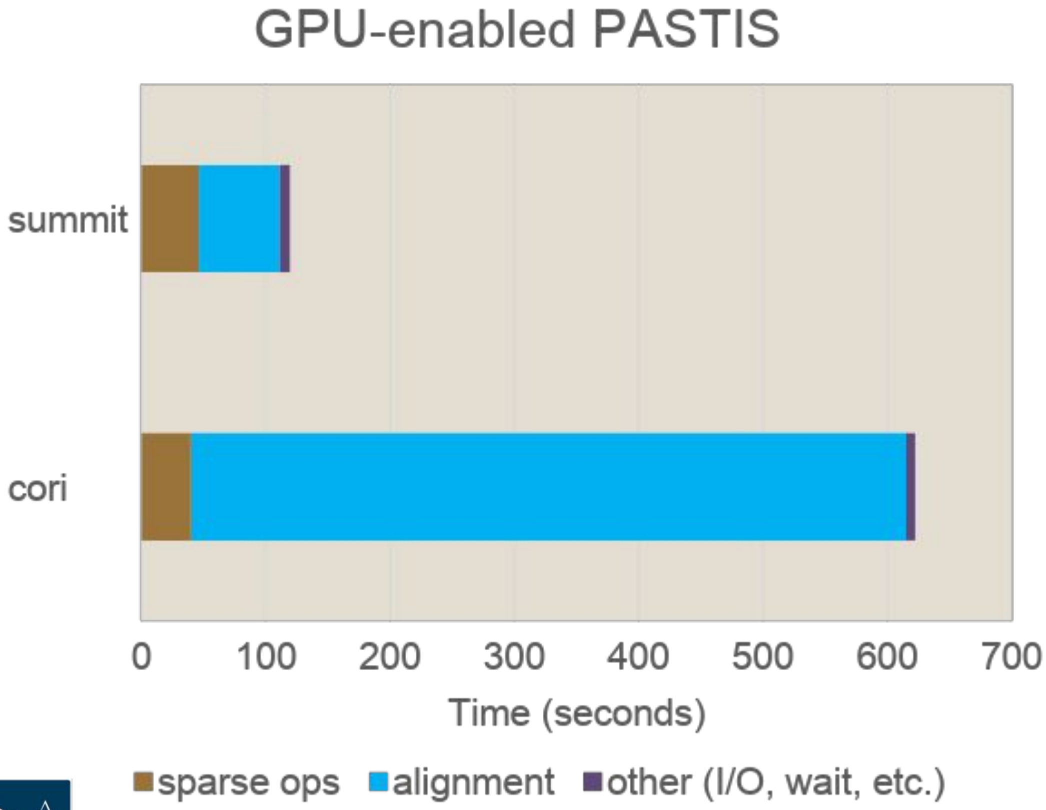

Dataset size: 5 million protein sequences

9.9 billion candidate alignments

~853 Million alignments

Protein similarity graph

- 5 Million nodes
- 64.6 million edges

100 nodes of NERSC Cori and Summit

- ADEPT
- SeqAn

Alignment: 5.2x speedup

Slide from Oguz Selvitopi

ADEPT in PASTIS

- PASTIS was used to analyze one of the largest available protein dataset with a search space of 313 * 313 * 10^{12} proteins.
- Using about 12000 GPUs on Summit a rate of 320 alignments per second was achieved.
- An unprecedented 143.9 TCUPs Performance was demonstrated by PASTIS.
- Gordon Bell submission for this SC22.

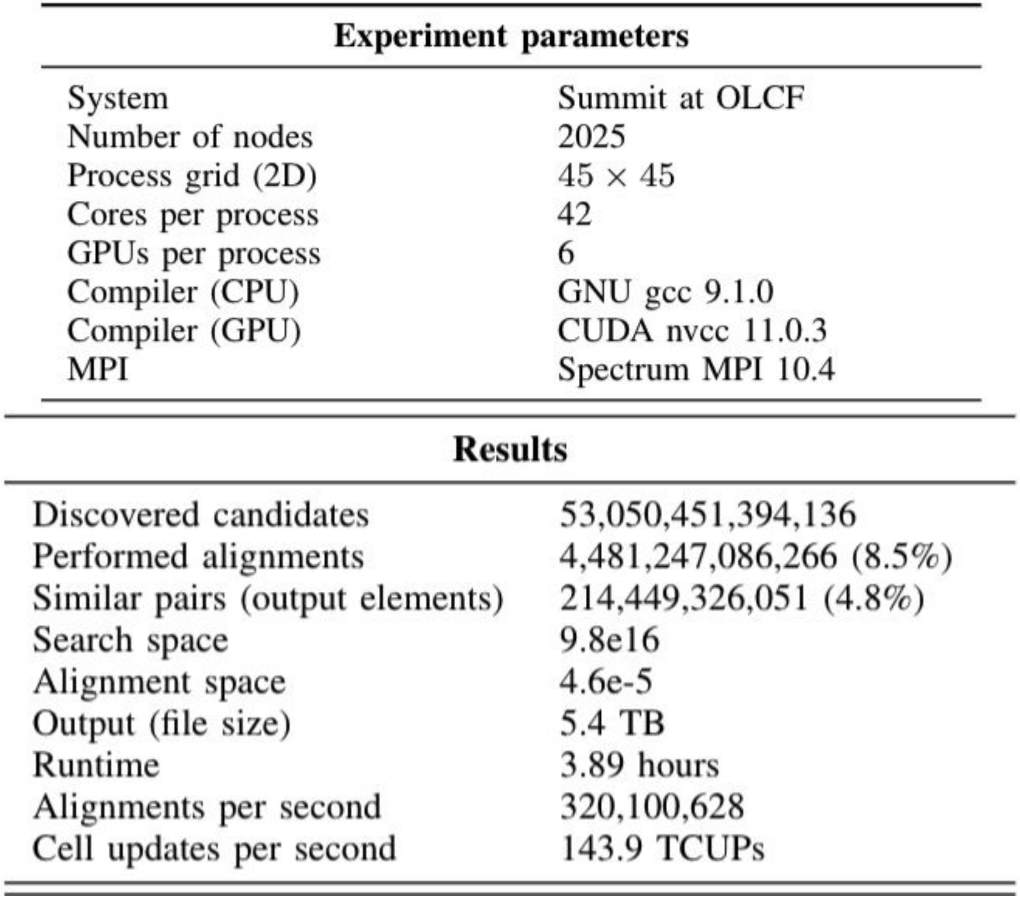

Experiment parameters	
System	Summit at OLCF
Number of nodes	2025
Process grid (2D)	45 × 45
Cores per process	42
GPUs per process	6
Compiler (CPU)	GNU gcc 9.1.0
Compiler (GPU)	CUDA nvcc 11.0.3
MPI	Spectrum MPI 10.4

Results	
Discovered candidates	53,050,451,394,136
Performed alignments	4,481,247,086,266 (8.5%)
Similar pairs (output elements)	214,449,326,051 (4.8%)
Search space	9.8e16
Alignment space	4.6e-5
Output (file size)	5.4 TB
Runtime	3.89 hours
Alignments per second	320,100,628
Cell updates per second	143.9 TCUPs

Local Assembly

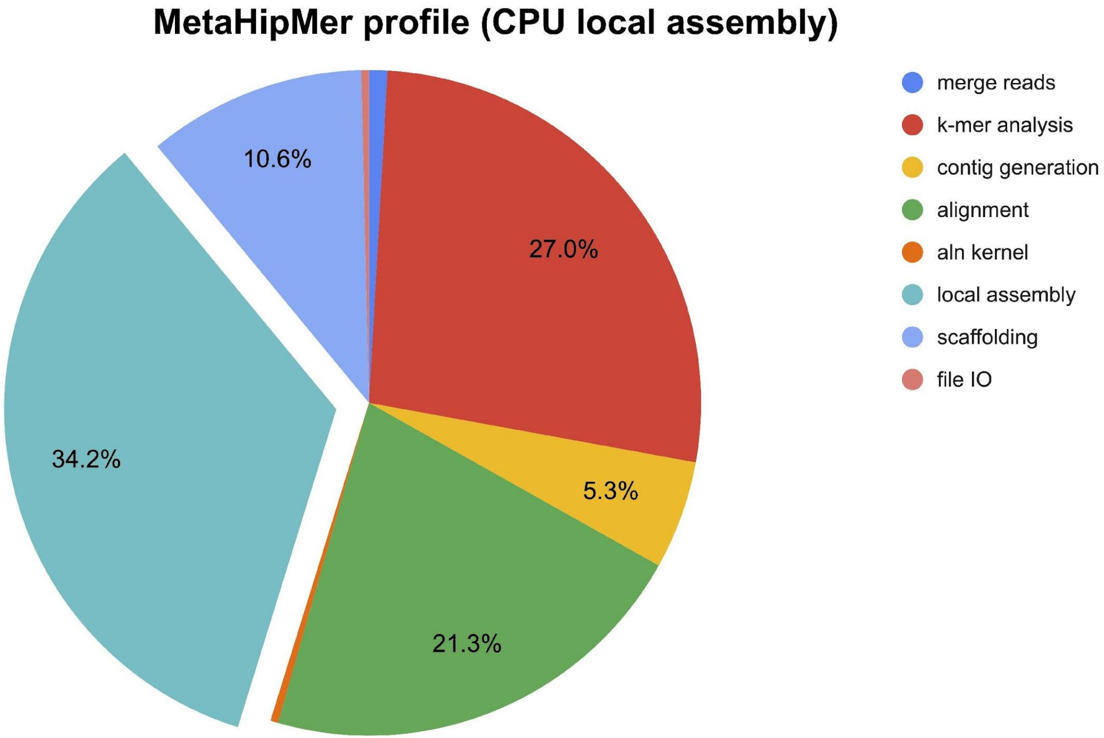

Local Assembly

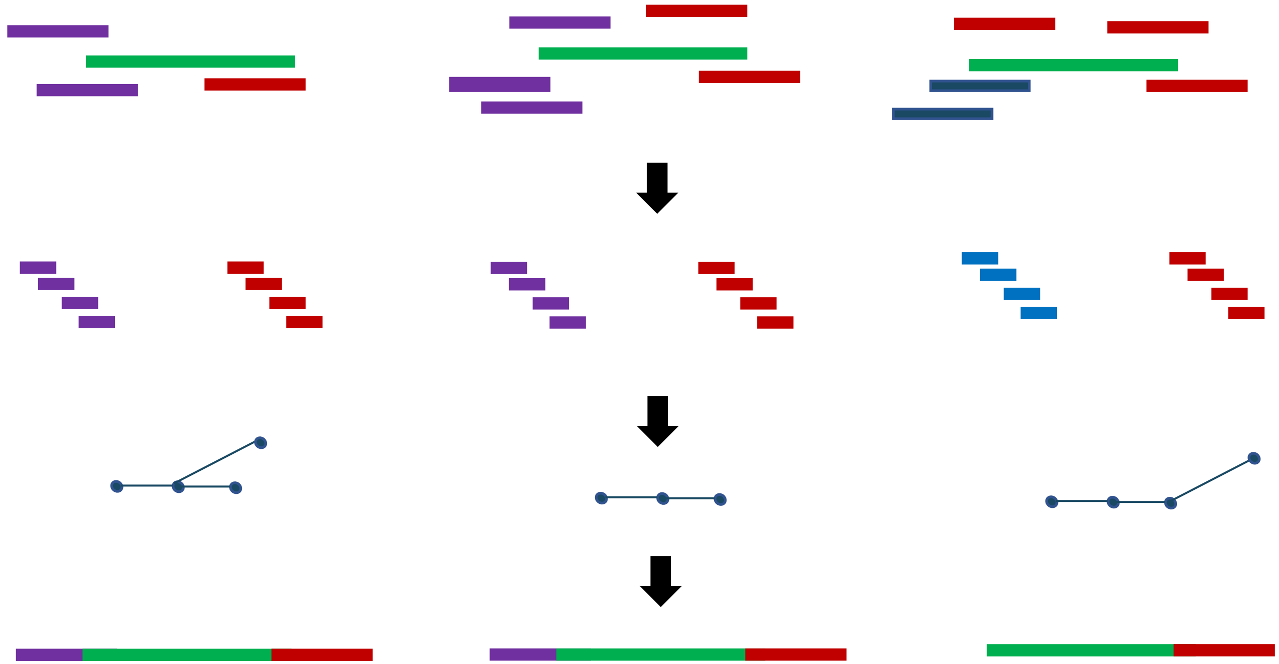

Reads aligning to ends of contigs are obtained from the alignment stage

Reads are broken down into kmers to construct De Bruijn graphs

De Bruijn graphs are traversed to extend contigs

Implementation Challenges on GPU

- Not a typical GPU problem
- No dynamic memory allocation on GPUs and no support for STL containers.
- Length of walks is non-deterministic.
- Static memory allocation can be used.
- Accurate memory usage needs to be known before kernel launch time.

GPU Local Assembly

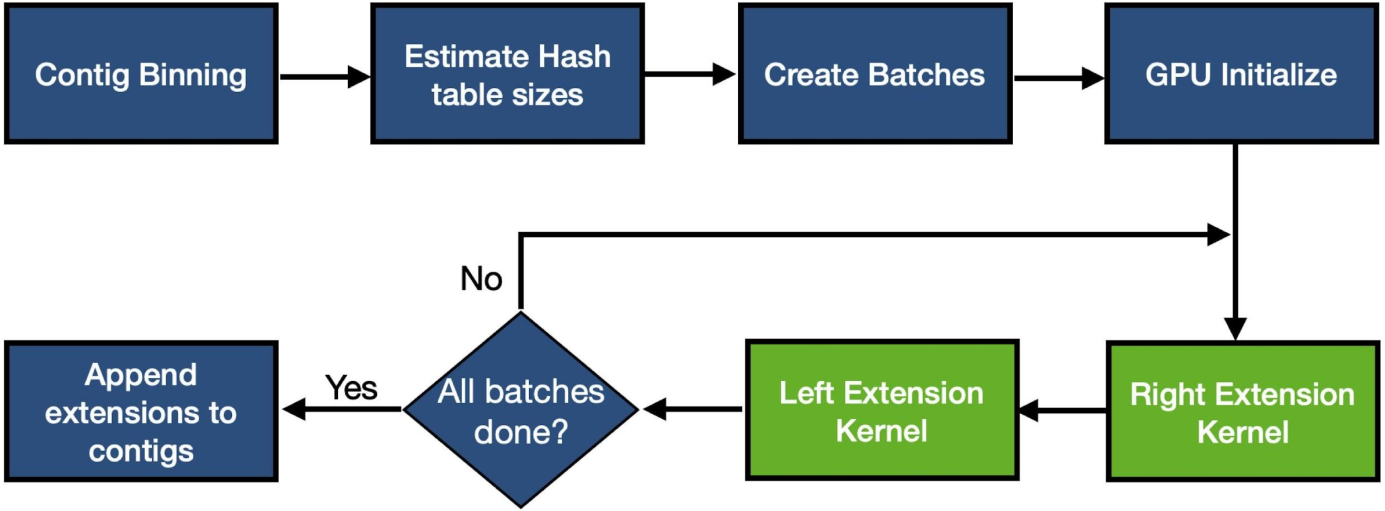

Local Assembly Kernel

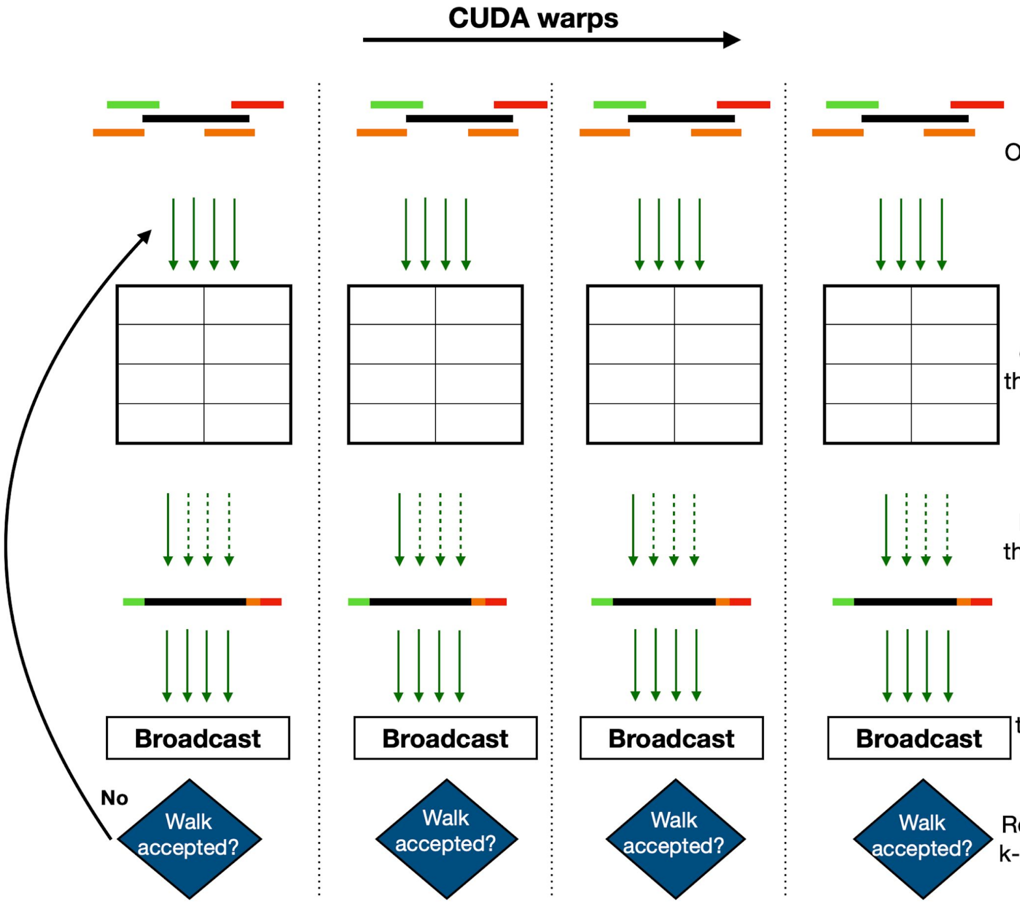

One extension per warp

Hash table construction, all threads participate

DNA walks, one thread participates

Broadcast the state of walk, all threads participate

Repeat with different k-mer size if walk not accepted.

Local Assembly Kernel Integration

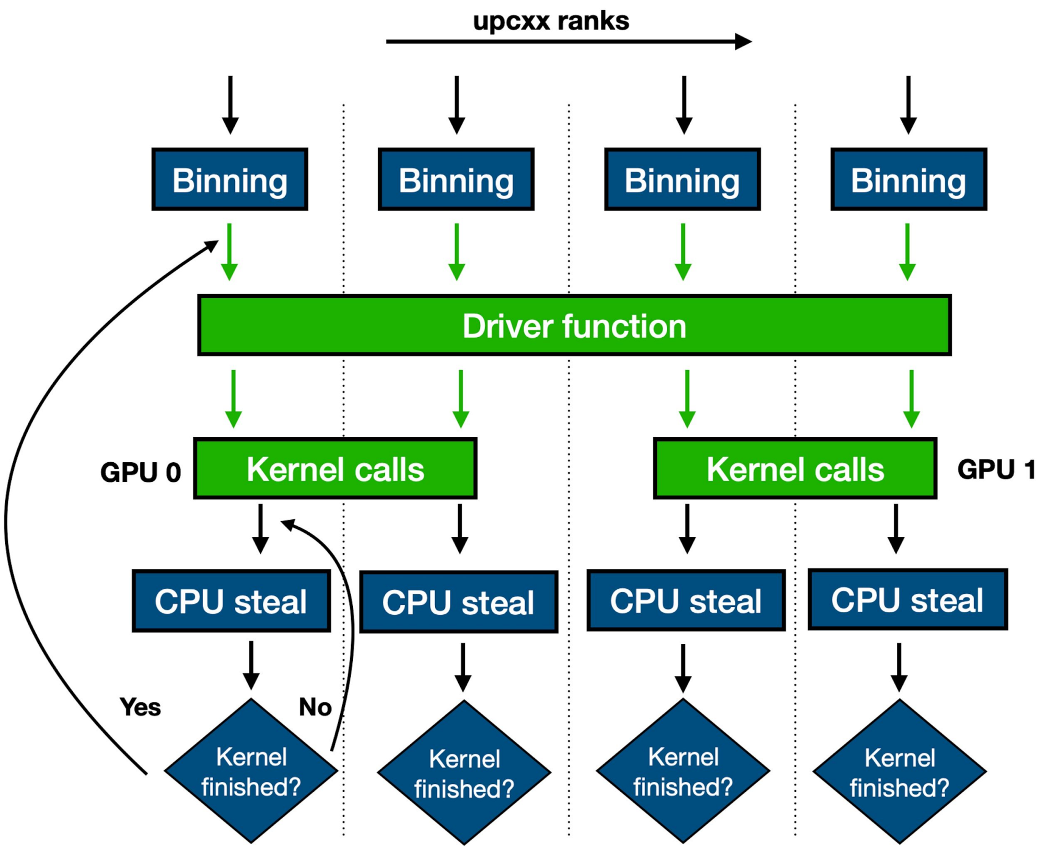

upcxx ranks

Binning

Binning

Binning

Binning

Driver function

GPU 0

Kernel calls

Kernel calls

GPU 1

CPU steal

CPU steal

CPU steal

CPU steal

Yes

No

Kernel finished?

Kernel finished?

Kernel finished?

Kernel finished?

GPU Accelerated Local Assembly

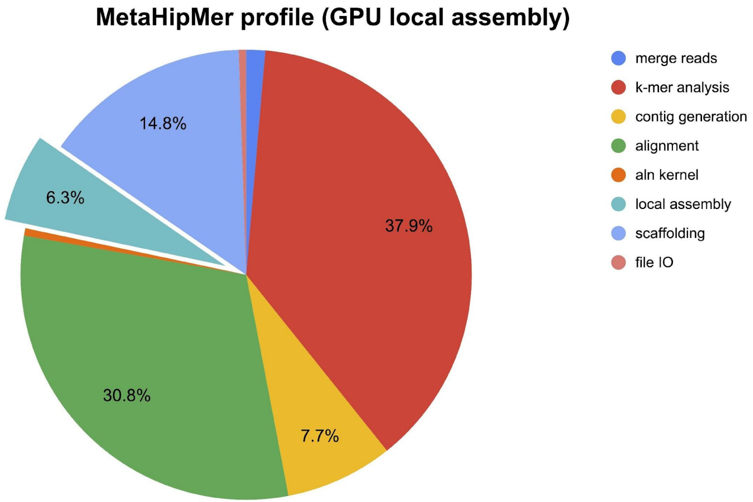

With GPU local assembly

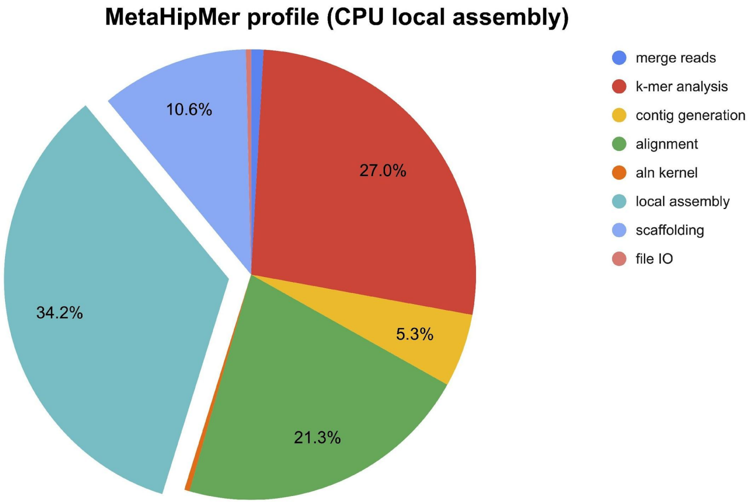

With CPU local assembly

Awan, Muaaz Gul, et al. "Accelerating large scale de novo metagenome assembly using GPUs." Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. 2021.

GPU Accelerated Local Assembly

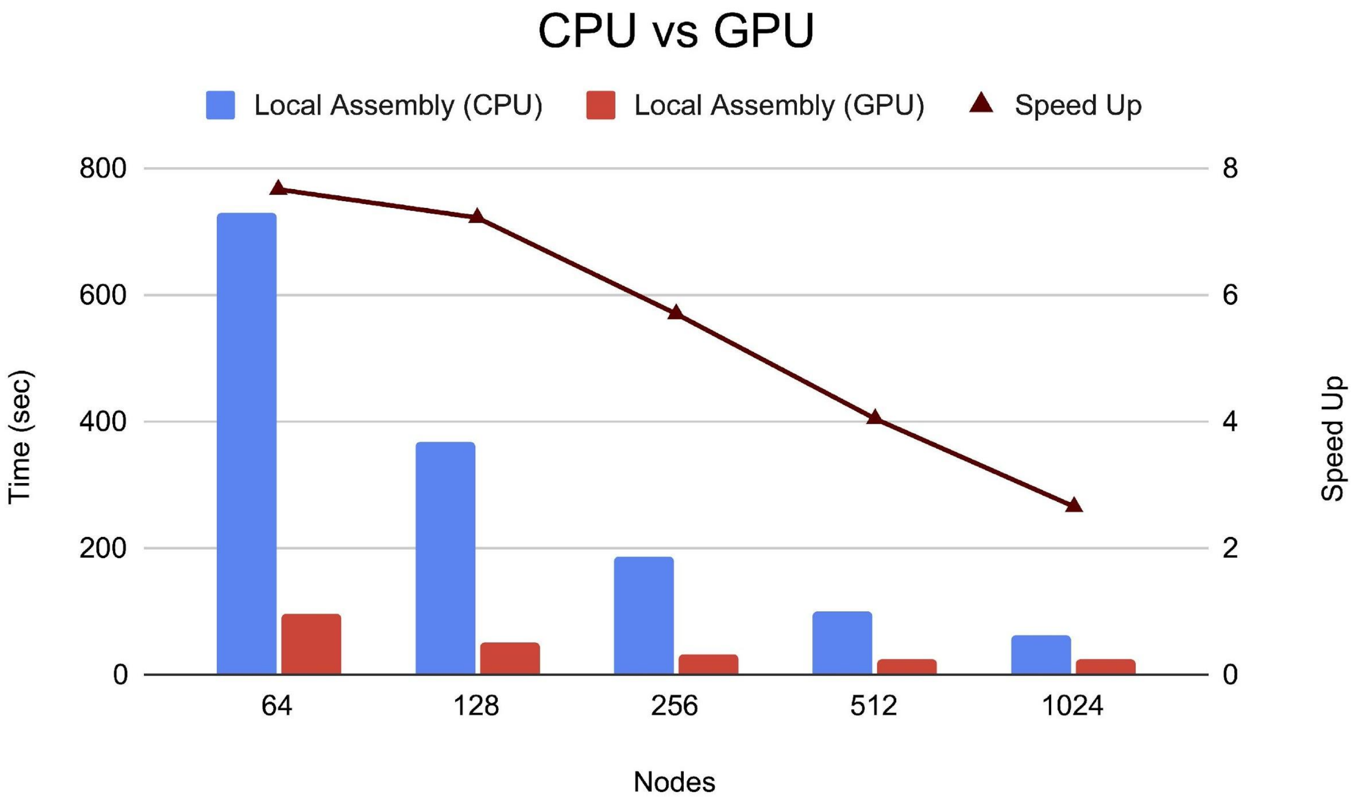

Awan, Muaaz Gul, et al. "Accelerating large scale de novo metagenome assembly using GPUs." Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. 2021.

25 TB Assembly

- Assembled a metagenomic dataset of 25 TB on 1500 Summit nodes in 47 minutes.
- Largest microbiome dataset ever assembled.
- This JGI user data contained freshwater lake samples from a 17 year period.
- Offered new insights into the evolution of microbial communities over time.
- Scientific questions this will help answer:
  - Characterize microbial components of freshwater lake food webs.
  - Help study the effects of climate and land-use change on freshwater carbon cycles.
  - Help model systems for studying diversification and dynamics of well-defined microbial populations

Lake Mendota algal bloom (https://www.nsf.gov/)

Performance on Perlmutter

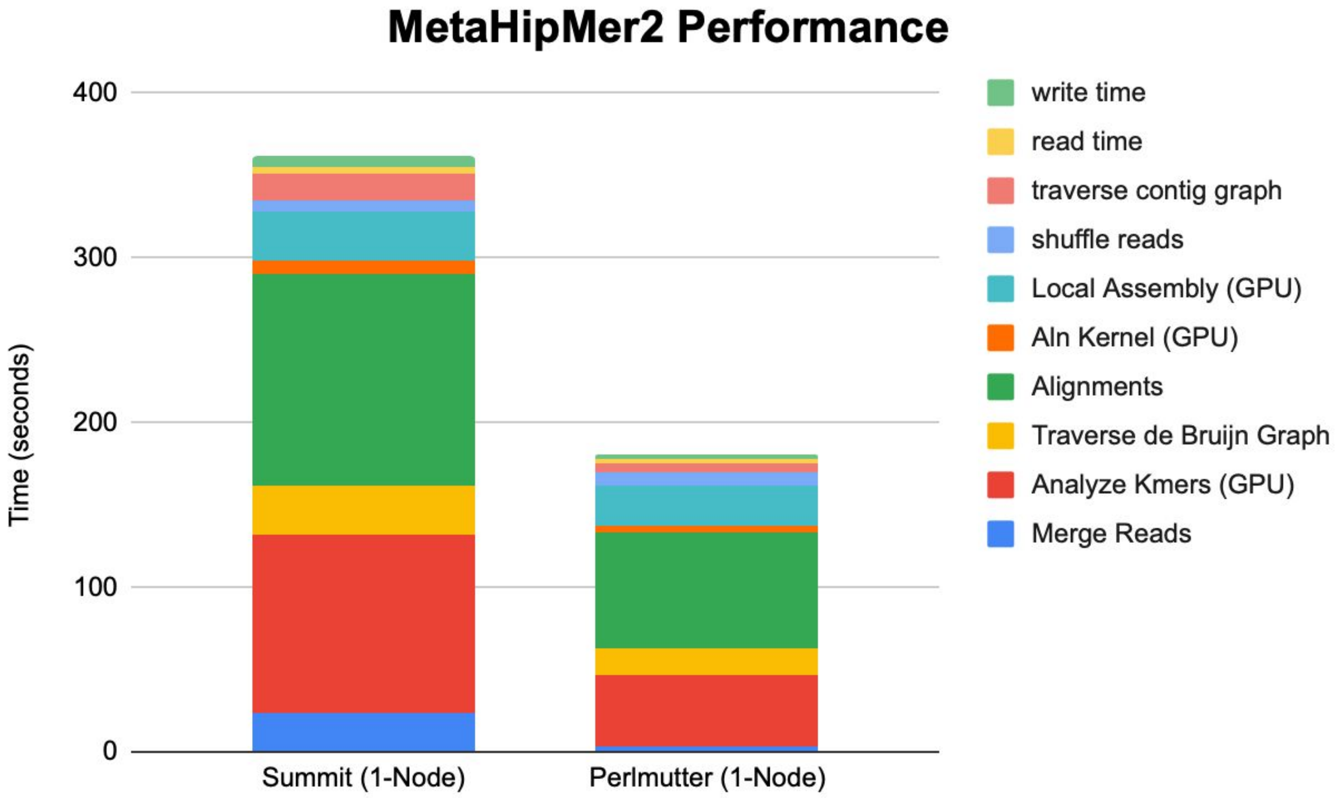

Acknowledgements


Rob Egan


Steven Hofmeyr


Jack Deslippe


Oguz
Selvitopi


Aydin Buluc

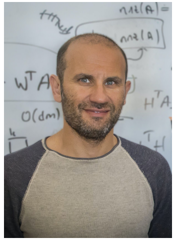
Leonid Oliker


Katherine Yelick

BERKELEY LAB

JGI
JOINT GENOME INSTITUTE
DEPARTMENT OF ENERGY

Berkeley
UNIVERSITY OF CALIFORNIA

NERSC

OAK RIDGE
National Laboratory
LEADERSHIP
COMPUTING
FACILITY

ECP
EXASCALE COMPUTING PROJECT

U.S. DEPARTMENT OF
ENERGY
Office of
Science

Thank you :)