

Accelerating Genome Analysis

A Primer on an Ongoing Journey

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

6 September 2022

Barcelona Supercomputing Center

SAFARI

ETH zürich

Carnegie Mellon

Overview

- System design for bioinformatics is a critical problem
 - It has large scientific, medical, societal, personal implications
- This talk is about accelerating a key step in bioinformatics:
genome sequence analysis
 - In particular, read mapping
- Many bottlenecks exist in accessing and manipulating huge amounts of genomic data during analysis
- We will cover various recent ideas to accelerate read mapping
 - My personal journey since September 2006

Our Dream (circa 2007)

- An embedded device that can perform comprehensive genome analysis in real time (within a minute)
 - Which of these DNAs does this DNA segment match with?
 - What is the likely genetic disposition of this patient to this drug?
 - What disease/condition might this particular DNA/RNA piece associated with?
 - . . .

A Bright Future for Intelligent Genome Analysis

Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu
[“Accelerating Genome Analysis: A Primer on an Ongoing Journey”](#) IEEE Micro, August 2020.



MinION from ONT

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Sept.-Oct. 2020, pp. 65-75, vol. 40
DOI Bookmark: [10.1109/MM.2020.3013728](https://doi.org/10.1109/MM.2020.3013728)

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41
DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)



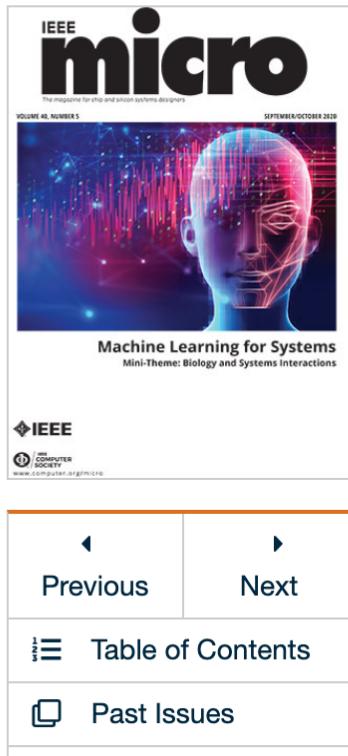
SmidgION from ONT

A Few Overview Readings (I)

Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu

["Accelerating Genome Analysis: A Primer on an Ongoing Journey"](#)

IEEE Micro, August 2020.



[Home](#) / [Magazines](#) / [IEEE Micro](#) / [2020.05](#)

IEEE Micro

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Sept.-Oct. 2020, pp. 65-75, vol. 40

DOI Bookmark: [10.1109/MM.2020.3013728](https://doi.org/10.1109/MM.2020.3013728)

Authors

[Mohammed Alser](#), ETH Zürich

[Zulal Bingol](#), Bilkent University

[Damla Senol Cali](#), Carnegie Mellon University

[Jeremie Kim](#), ETH Zurich and Carnegie Mellon University

[Saugata Ghose](#), University of Illinois at Urbana-Champaign and Carnegie Mellon University

[Can Alkan](#), Bilkent University

[Onur Mutlu](#), ETH Zurich, Carnegie Mellon University, and Bilkent University

A Few Overview Readings (II)

Gagandeep Singh, Mohammed Alser, Damla Senol Cali, Dionysios Diamantopoulos, Juan Gomez-Luna, Henk Corporaal, Onur Mutlu,

"FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications"

IEEE Micro, 2021.

[[Source Code](#)]



Home / Magazines / IEEE Micro / 2021.04

IEEE Micro

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41

DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)

Authors

Gagandeep Singh, ETH Zürich, Zürich, Switzerland

Mohammed Alser, ETH Zürich, Zürich, Switzerland

Damla Senol Cali, Carnegie Mellon University, Pittsburgh, PA, USA

Dionysios Diamantopoulos, Zürich Lab, IBM Research Europe, Rüschlikon, Switzerland

Juan Gomez-Luna, ETH Zürich, Zürich, Switzerland

Henk Corporaal, Eindhoven University of Technology, Eindhoven, The Netherlands

Onur Mutlu, ETH Zürich, Zürich, Switzerland

A Few Overview Readings (III)

Mohammed Alser, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu

"From Molecules to Genomic Variations: Intelligent Algorithms and Architectures for Intelligent Genome Analysis"

Computational and Structural Biotechnology Journal, 2022

[[Source code](#)]



ELSEVIER

01010... 0010101010010
00101010010010101011
10101010010101010111
01010100101010101010
11010100101010101010
10101010010101010111
00101010010101010111
01010101001010101010

COMPUTATIONAL
AND STRUCTURAL
BIOTECHNOLOGY
JOURNAL

journal homepage: www.elsevier.com/locate/csbj



Review

From molecules to genomic variations: Accelerating genome analysis via intelligent algorithms and architectures



Mohammed Alser *, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu *

ETH Zurich, Gloriastrasse 35, 8092 Zürich, Switzerland

SAFARI

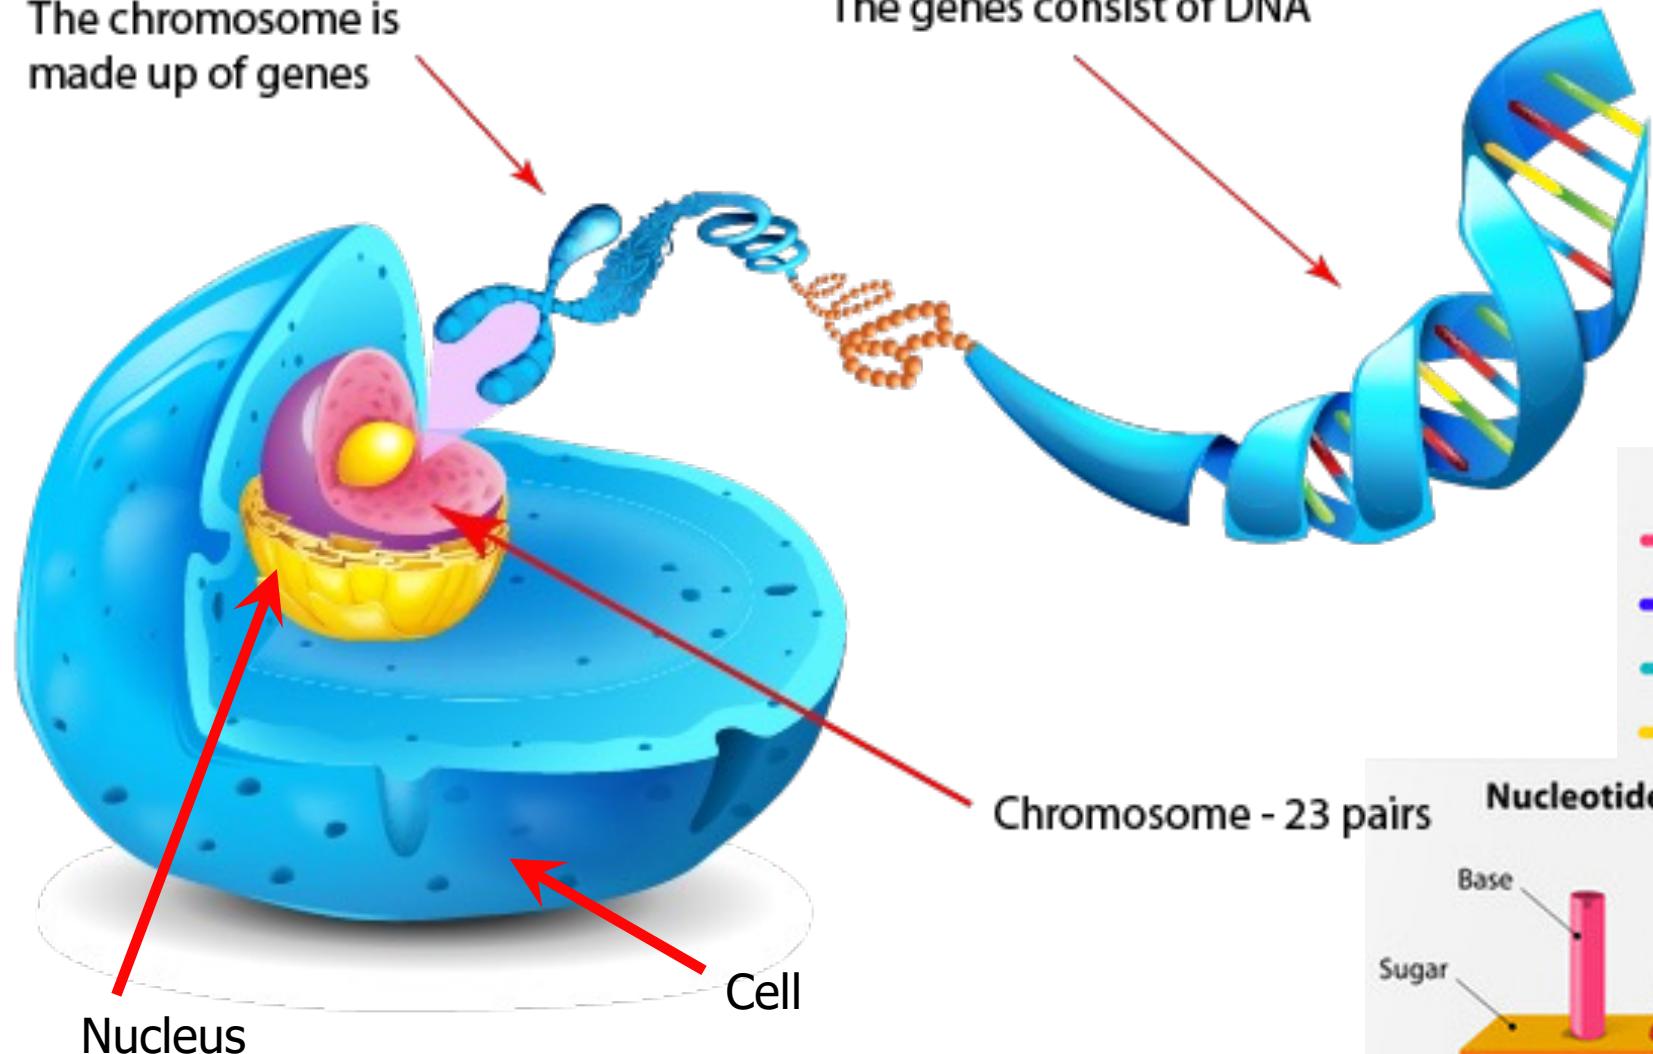
<https://arxiv.org/pdf/2205.07957.pdf>

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

What Is a Genome Made Of?

The chromosome is made up of genes



The genes consist of DNA

Bases

Adenine

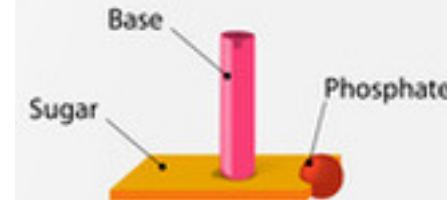
Cytosine

Guanine

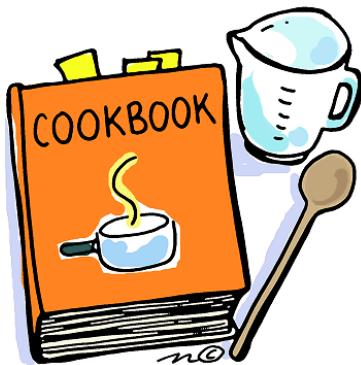
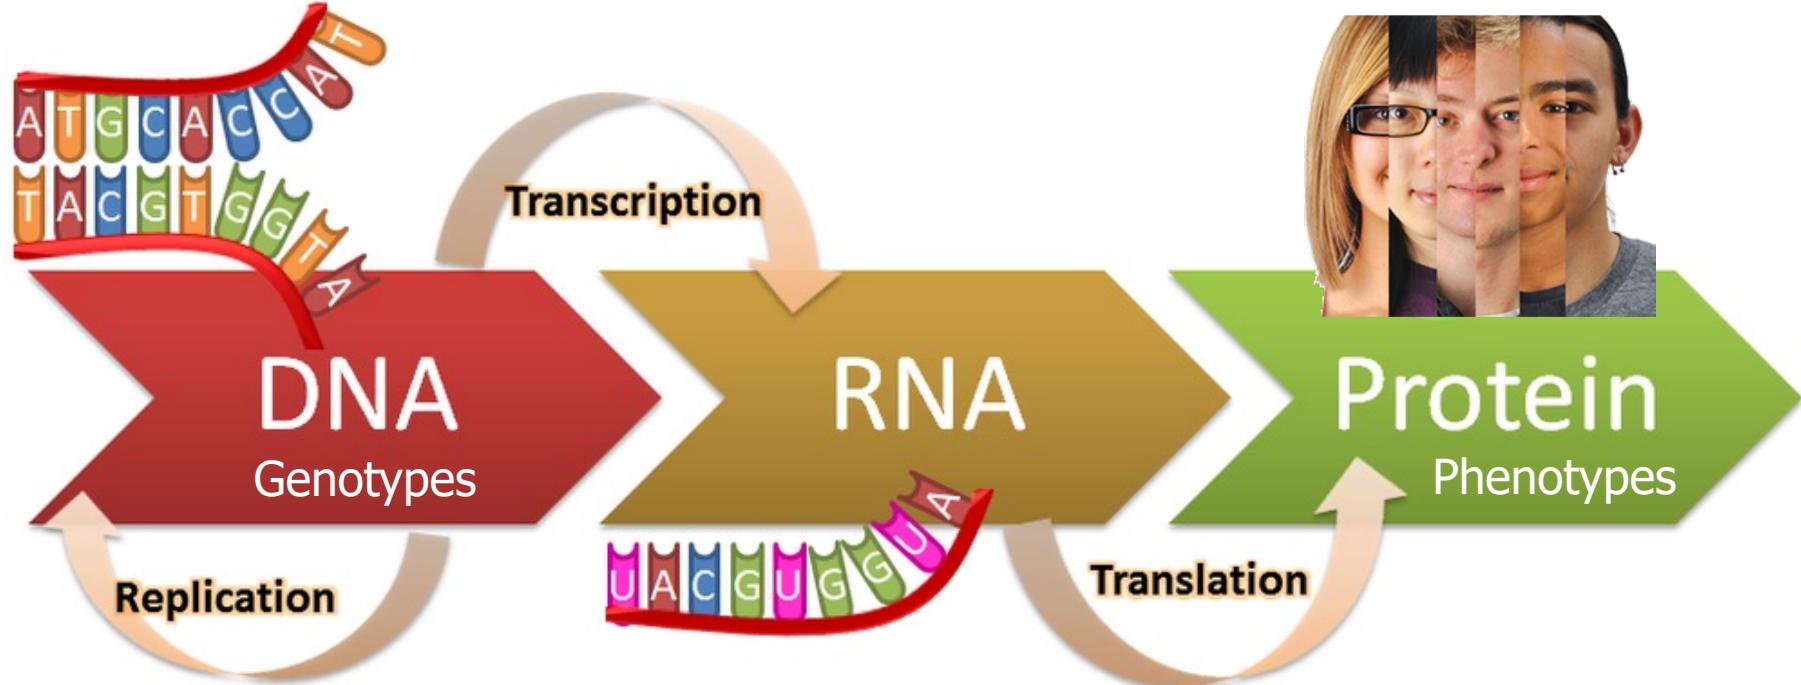
Thymine

Chromosome - 23 pairs

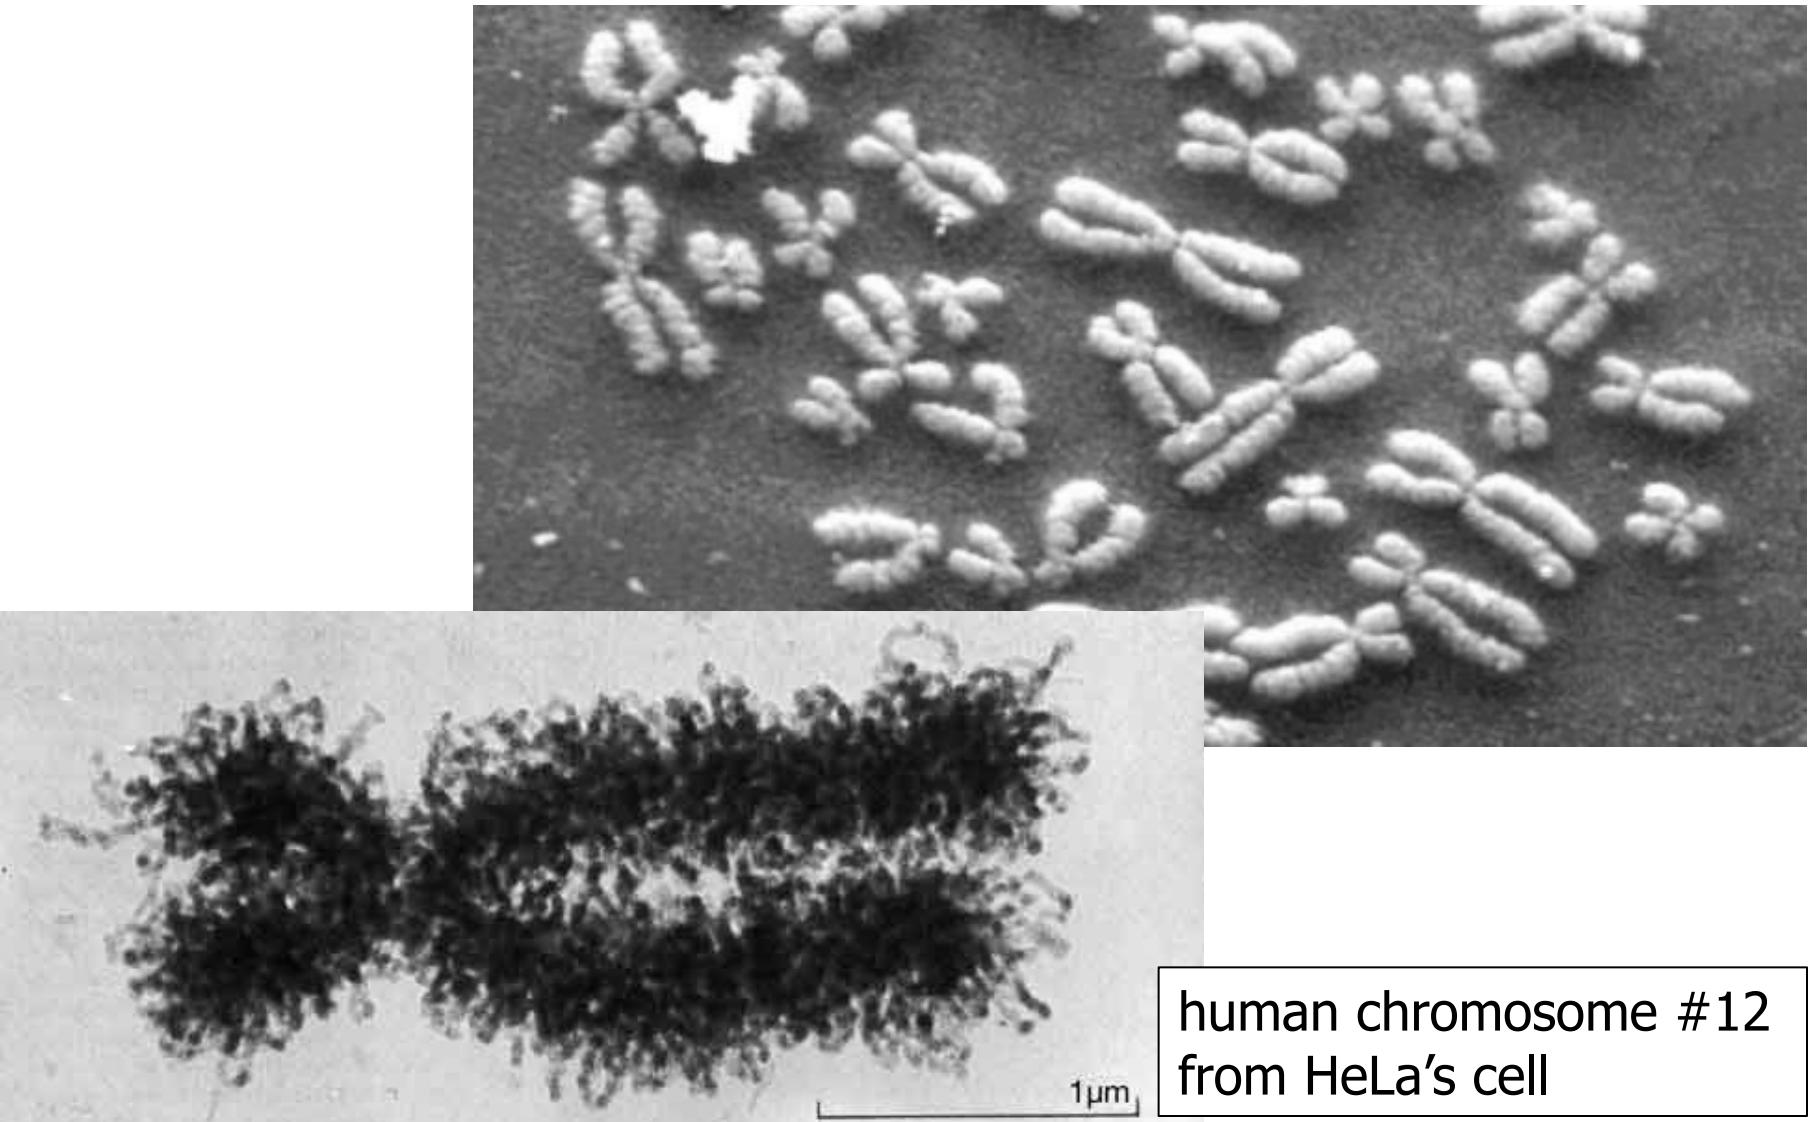
Nucleotide



The Central Dogma of Molecular Biology



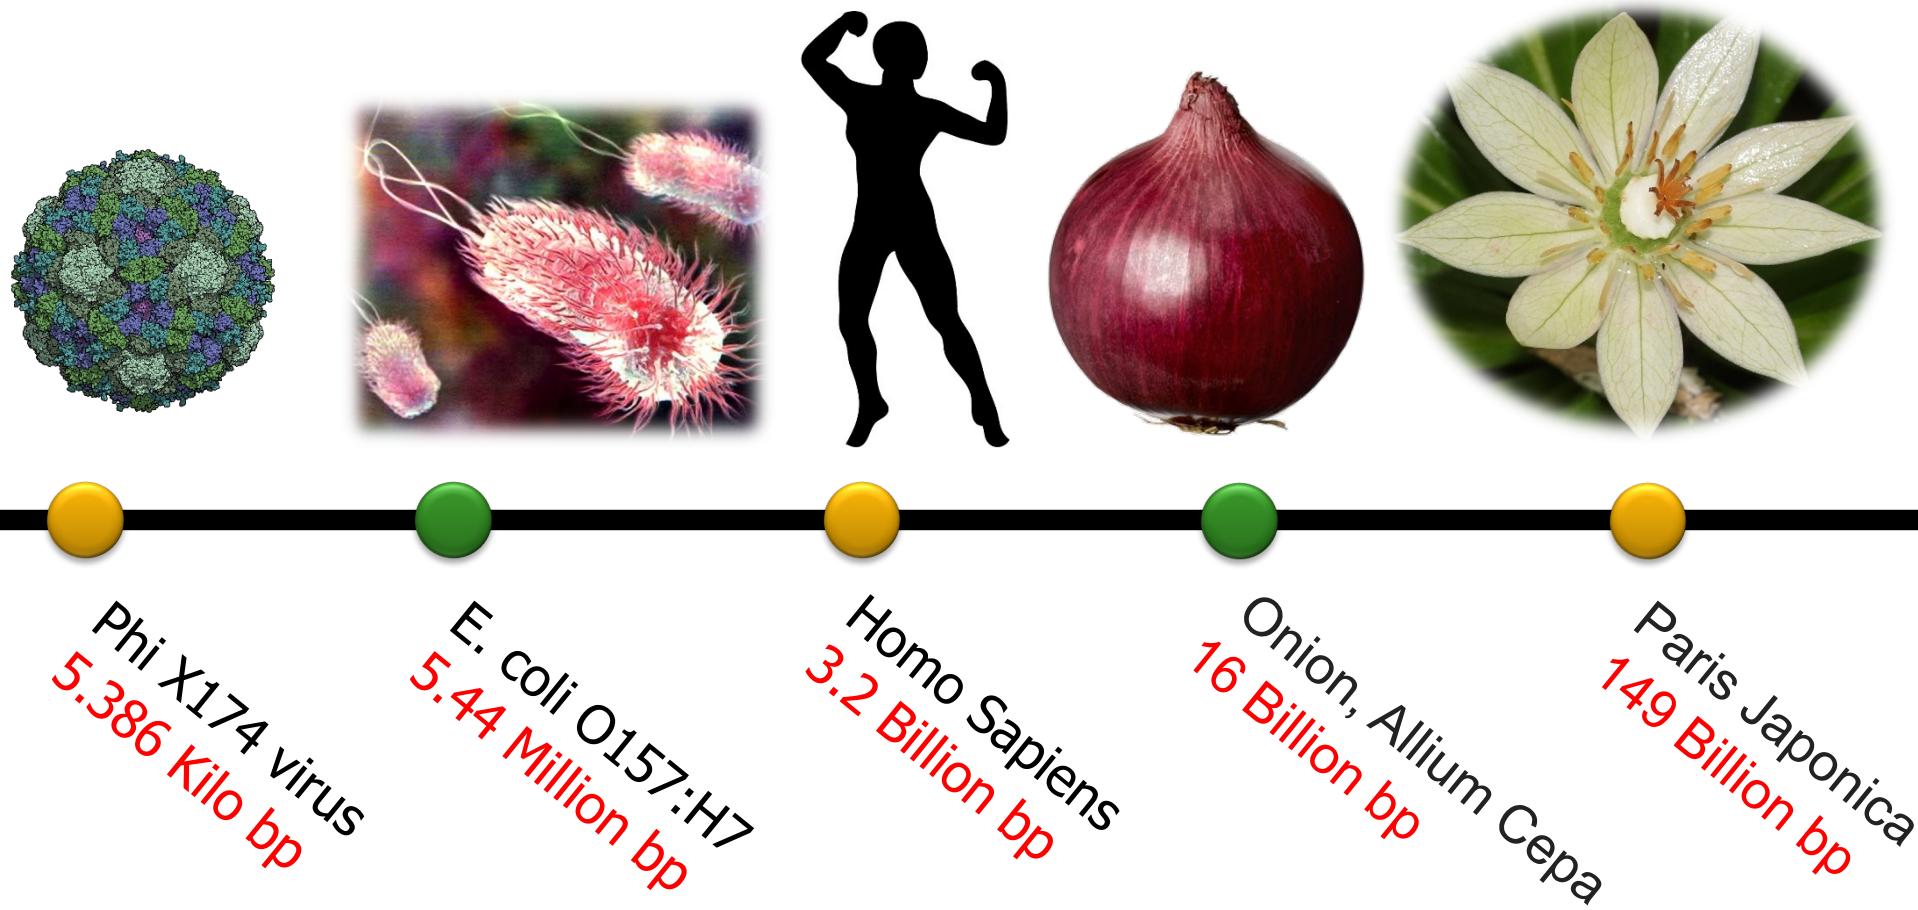
DNA Under Electron Microscope



human chromosome #12
from HeLa's cell

CCTCCTCAGTGCCACCCAGCCCCTGGCAGCTCCAAACA
GGCTCTTATTAAAACACCCTGTTCCCTGCCCTGGAGTG
AGGTGTCAAGGACCTAAACTAAAAAAAAAGAAAAA
AGAAAAAGAAAAAGAATTAAAATTAAAGTAATTCTTGAA
AAAAAACTAATTCTAAGCTTCTCATGTCAAGGACCTAATG
TGCTAAACAGCACTTT**TTGACCATTAT**TTTGGATCTGAAA
GAAATCAAGAATAATGAAGGACTTGATACATTGGAAGA
GGAGAGTCAAGGACCTACAGAAAAAAAAGAAAA
AAGAAAAGAAAAAGA**A**TTTAAAATTAAAGTAATTCTTGAA
AAAAAACTAATTCTAAGCTTCTT**C**ATGTCAAGGACCTAAT
GTCTGTGTTGCAGGTCTTCTGCATTCCCTGTCAAAAAGA
AAAAGAATTAAAATTAAAGTAATTCTTGAAAAAAACTA
ATTCTAAGCTTCTCATGTCAAGGACCTAATGTCAGGCC
GGCTCTTATTAAAACACCCTGTTCCCTGCCCTGGAGTG

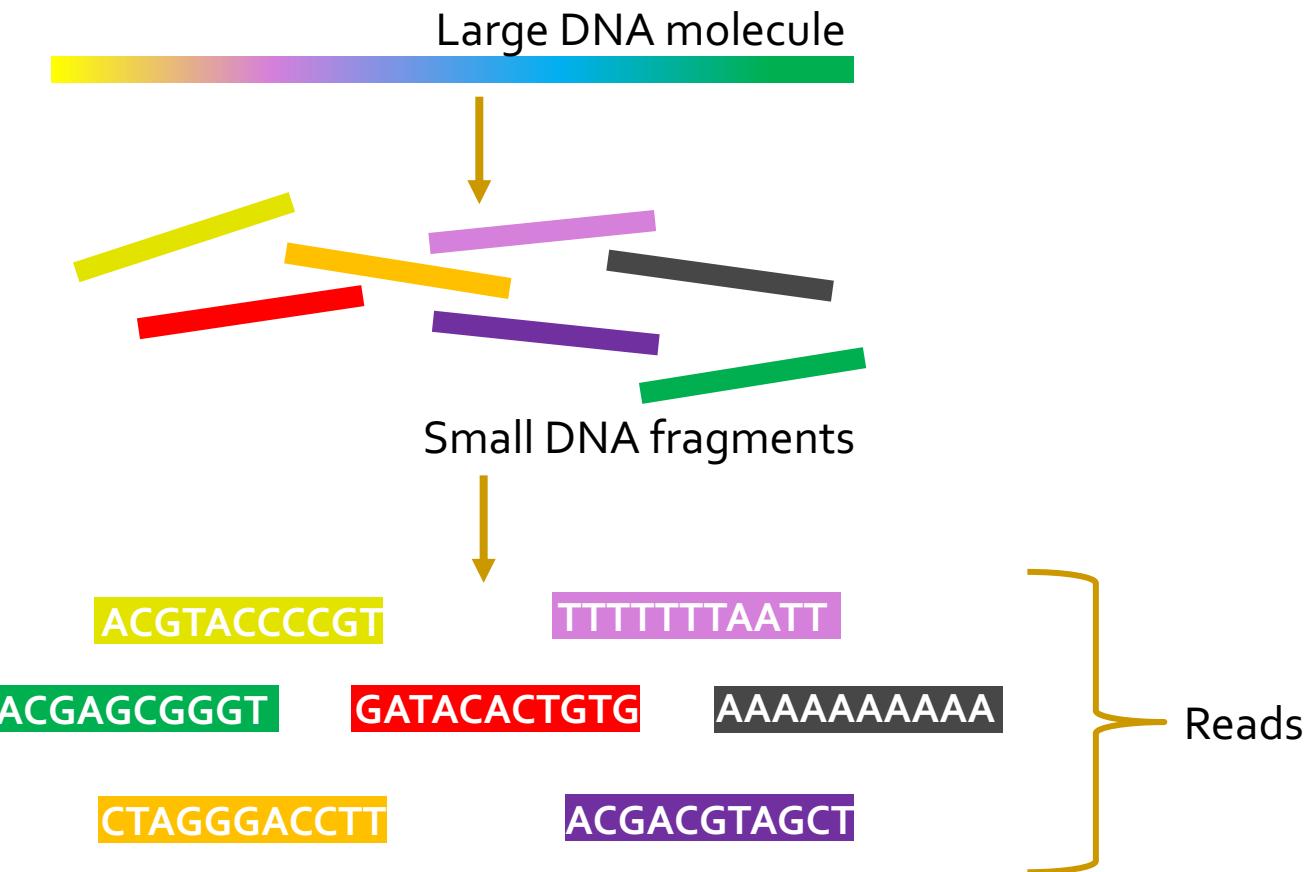
How Large is a Genome?



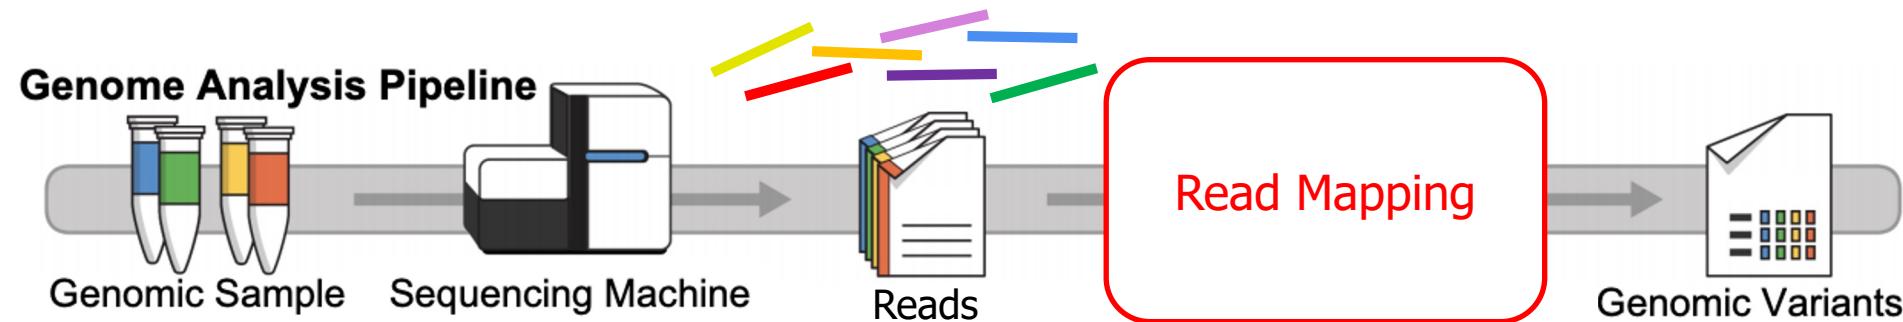
DNA Sequencing

- Goal:
 - Find the complete sequence of A, C, G, T's in an organism's DNA
- Challenge:
 - There is no machine that takes long DNA as an input, and gives the complete sequence as output
 - All sequencing machines chop DNA into pieces and identify relatively small pieces (but not how they fit together)

Genome Sequencing



Genome Sequencing and Analysis



Current sequencing machines provide
small randomized fragments
of the original DNA sequence

Alser+, "[Technology dictates algorithms: Recent developments in read alignment](#)", Genome Biology, 2021

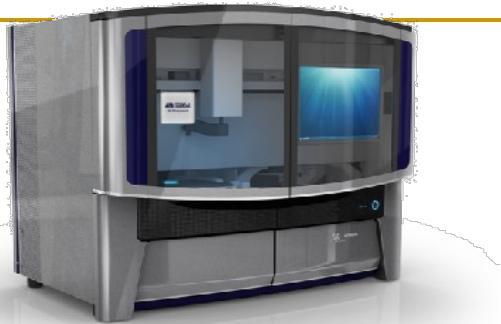
Untangling Yarn Balls & DNA Sequencing



Genome Sequencers



Roche/454



AB SOLiD



Illumina MiSeq



Complete Genomics



Illumina HiSeq 2000



Pacific Biosciences RS



SAFARI

Ion Torrent PGM



Ion Torrent Proton



Oxford Nanopore MinION



Illumina
NovaSeq
6000



Oxford Nanopore GridION

... and more! All produce data with different properties.

High-Throughput Sequencers



Illumina MiSeq



Illumina NovaSeq 6000



Pacific
Biosciences
Sequel II



Pacific Biosciences RS II

Oxford
Nanopore
PromethION



Oxford Nanopore MinION



Oxford
Nanopore
SmidgION

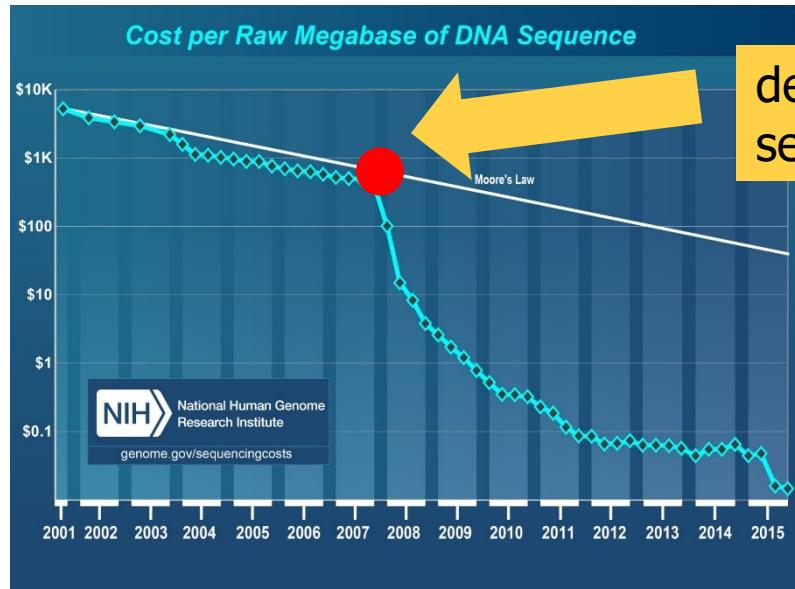
... and more! All produce data with different properties.

The Genomic Era

- 1990-2003: The Human Genome Project (HGP) provides a complete and accurate sequence of all **DNA base pairs** that make up the human genome and finds 20,000 to 25,000 human genes.

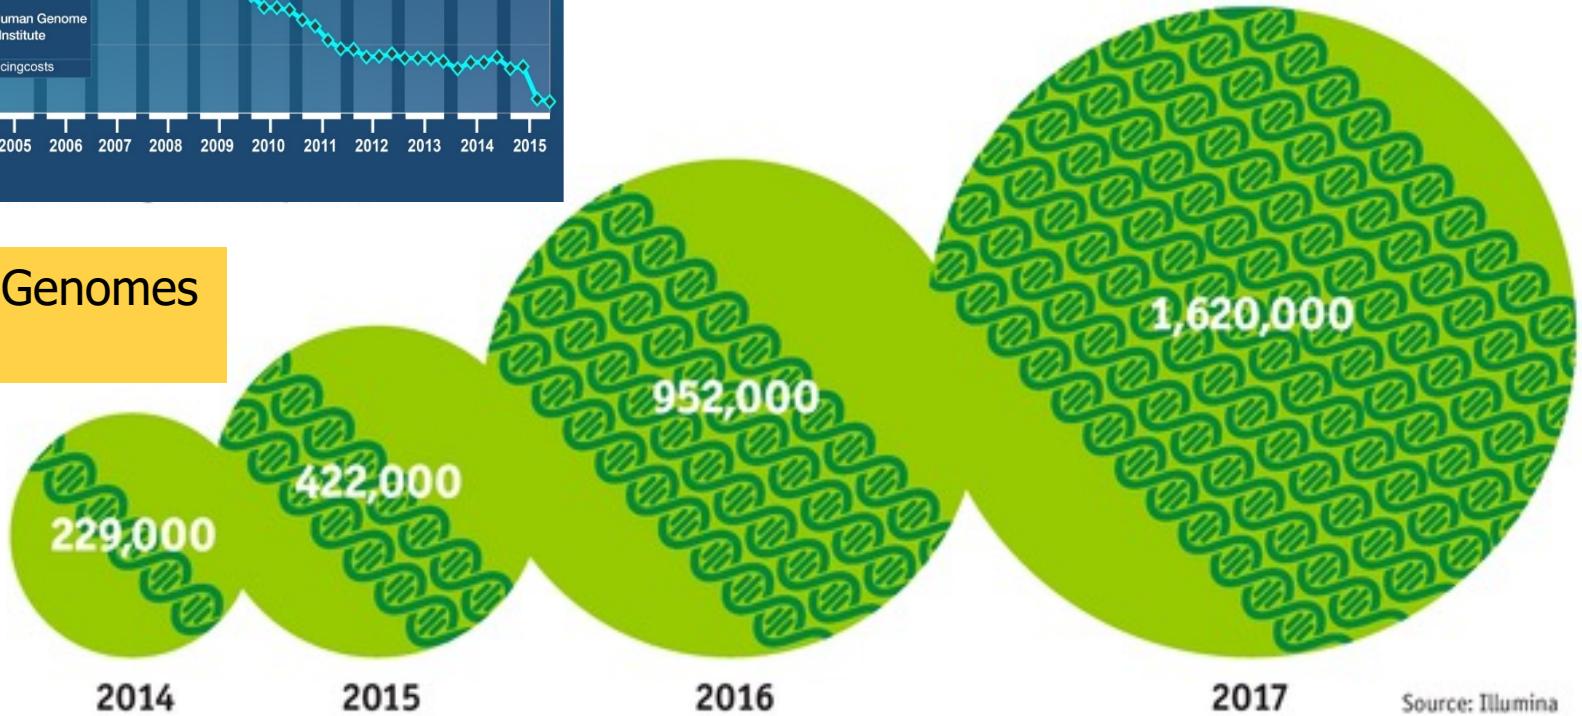


The Genomic Era (continued)



development of high-throughput sequencing (HTS) technologies

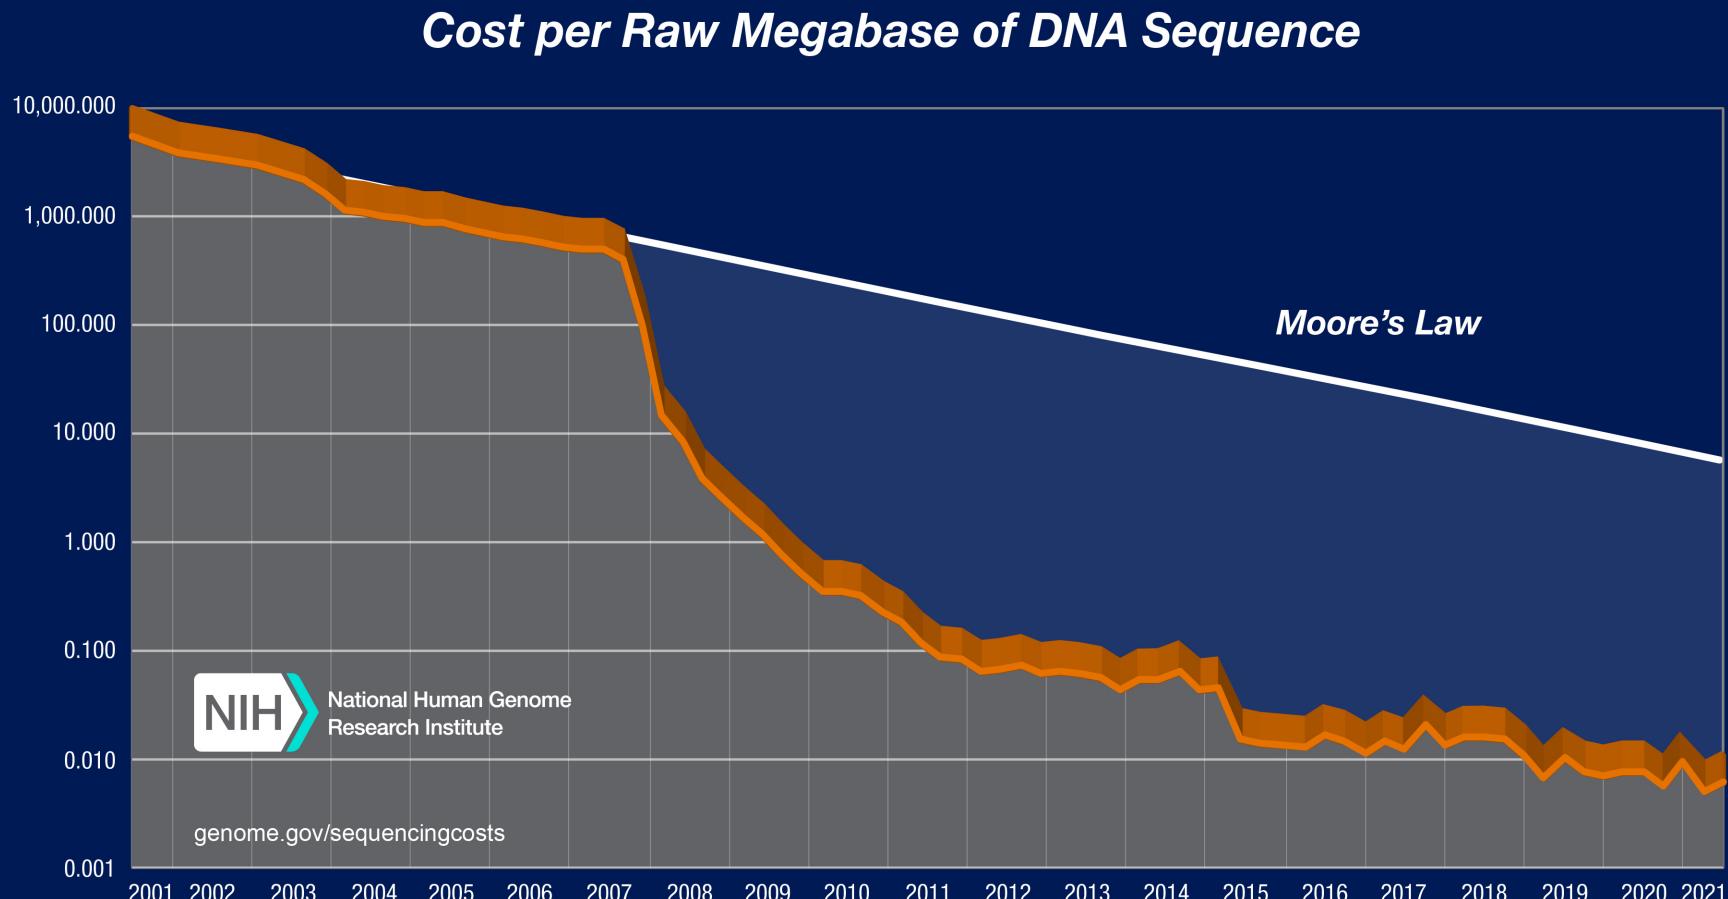
Number of Genomes Sequenced



The Economist

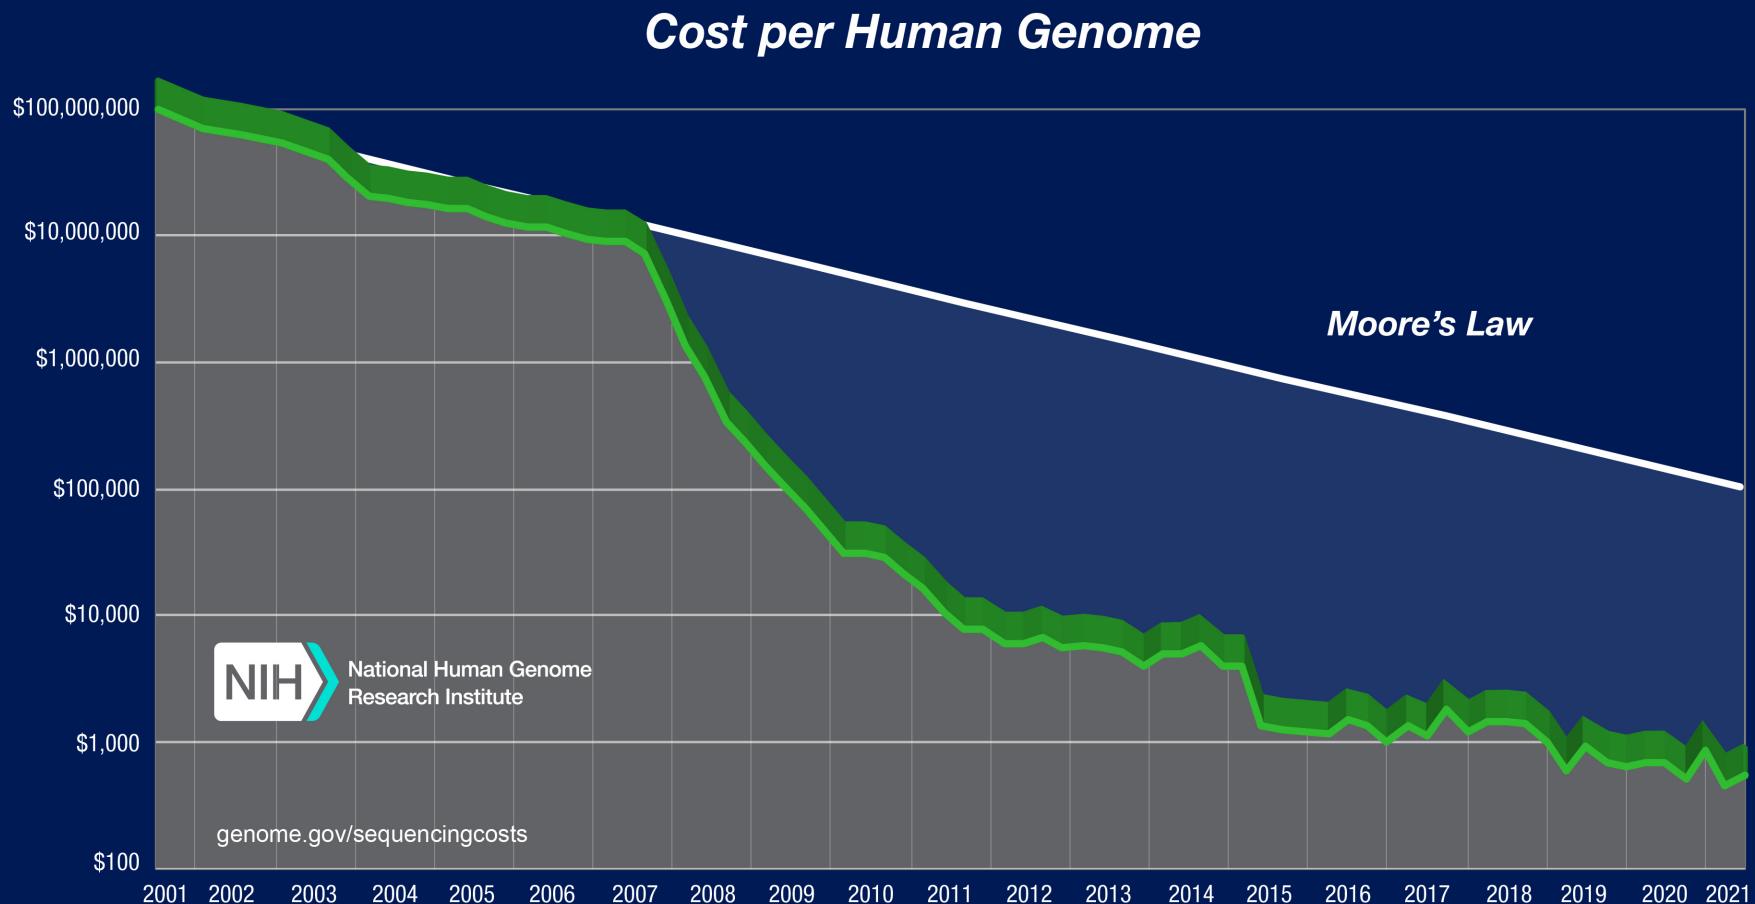
Source: Illumina

Genome Sequencing Cost Is Reducing



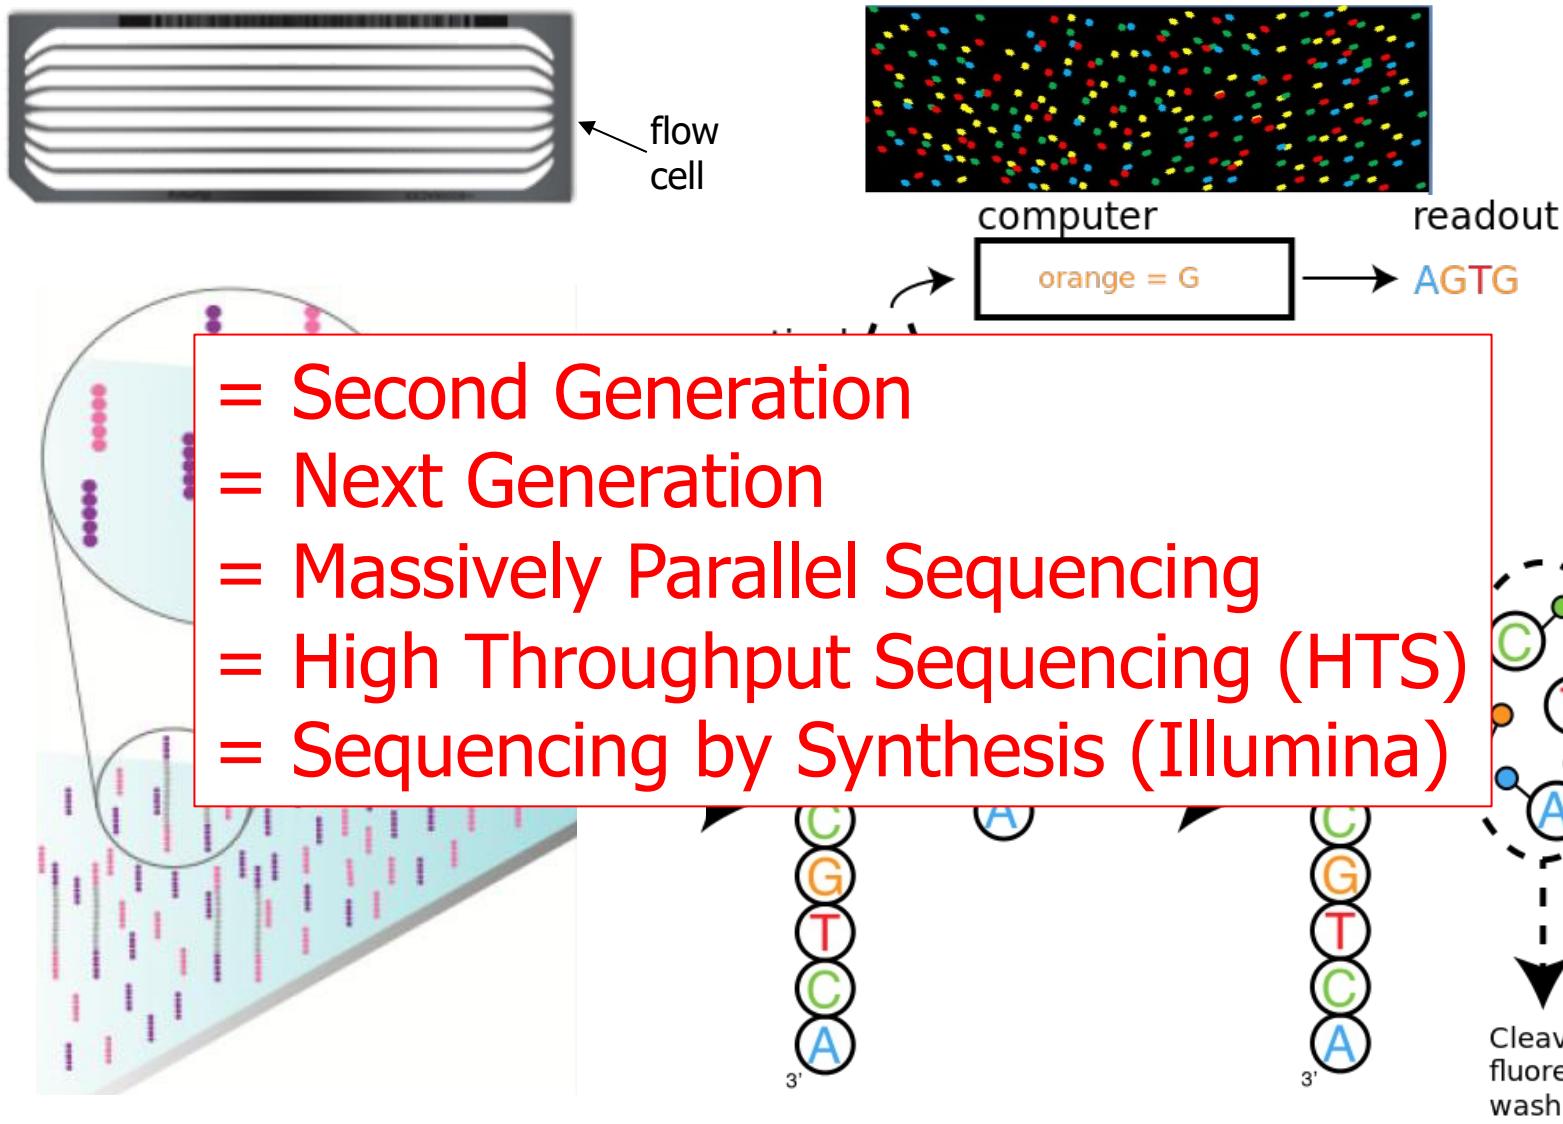
*From NIH (<https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data>)

Genome Sequencing Cost Is Reducing

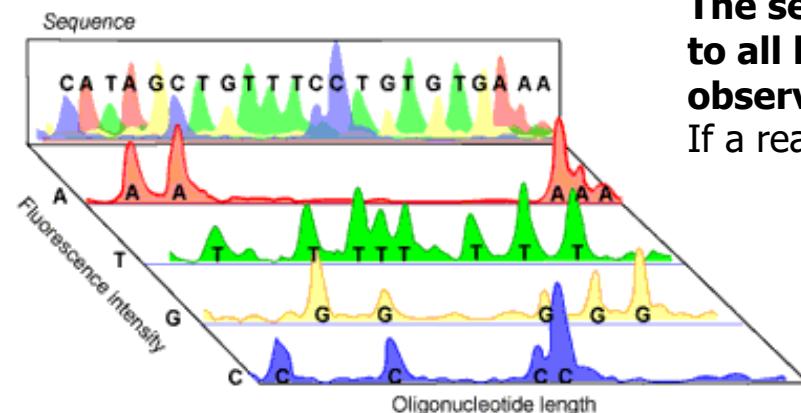


*From NIH (<https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data>)

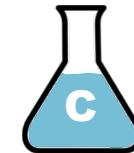
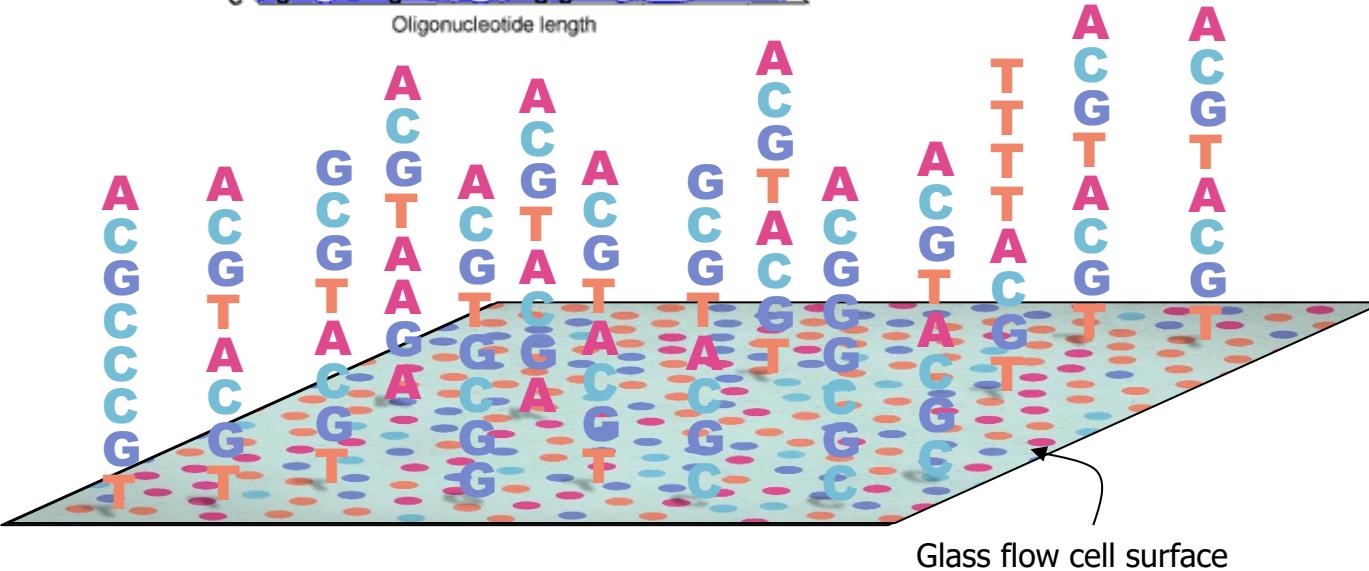
High-Throughput Sequencing (HTS)



High-Throughput Sequencing (HTS)



The sequencer adds the molecule "T" to all bases near the flow cell surface and observes the chemical reaction via a CMOS sensor. If a reaction happens then the base is "A"



As a workaround, HTS technologies sequence random short DNA fragments (75-300 basepairs long) of copies of the original molecule.

High-Throughput Sequencing

- Massively parallel sequencing technology
 - Illumina, Roche 454, Ion Torrent, SOLID...
- Small DNA fragments are first amplified and then sequenced in parallel, leading to
 - High throughput
 - High speed
 - Low cost
 - Short reads
- Sequencing is done by either reading optical signals as each base is added, or by detecting hydrogen ions instead of light, leading to:
 - Low error rates (relatively)
 - Reads lack information about their order and which part of genome they are originated from

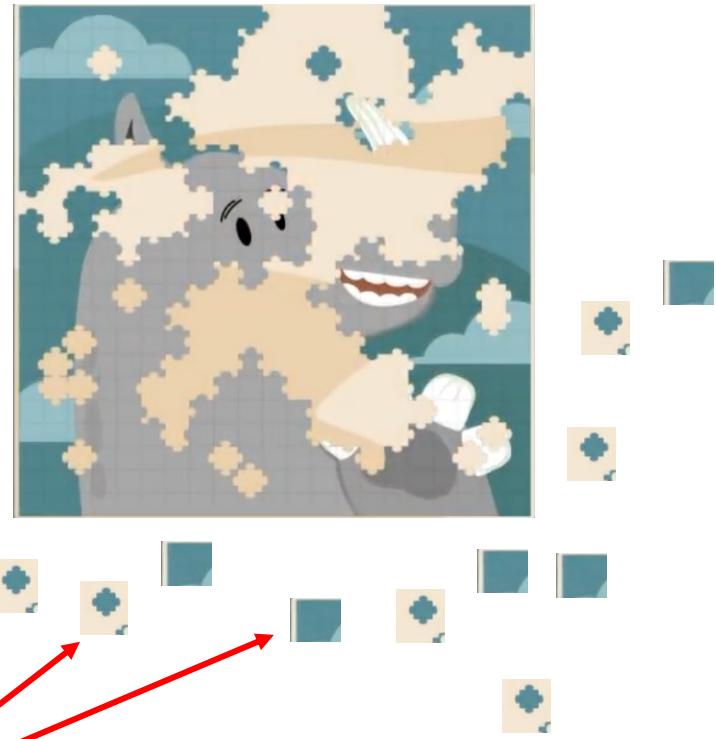
Solving the Puzzle

.FASTA file



Reference genome

.FASTQ file



Sequenced Reads

<https://www.pacb.com/smrt-science/smrt-sequencing/hifi-reads-for-highly-accurate-long-read-sequencing/>

Newer Genome Sequencing Technologies

Nanopore sequencing technology and tools for genome assembly: computational analysis of the current state, bottlenecks and future directions

Damla Senol Cali ✉, Jeremie S Kim, Saugata Ghose, Can Alkan, Onur Mutlu

Briefings in Bioinformatics, bby017, <https://doi.org/10.1093/bib/bby017>

Published: 02 April 2018 Article history ▾

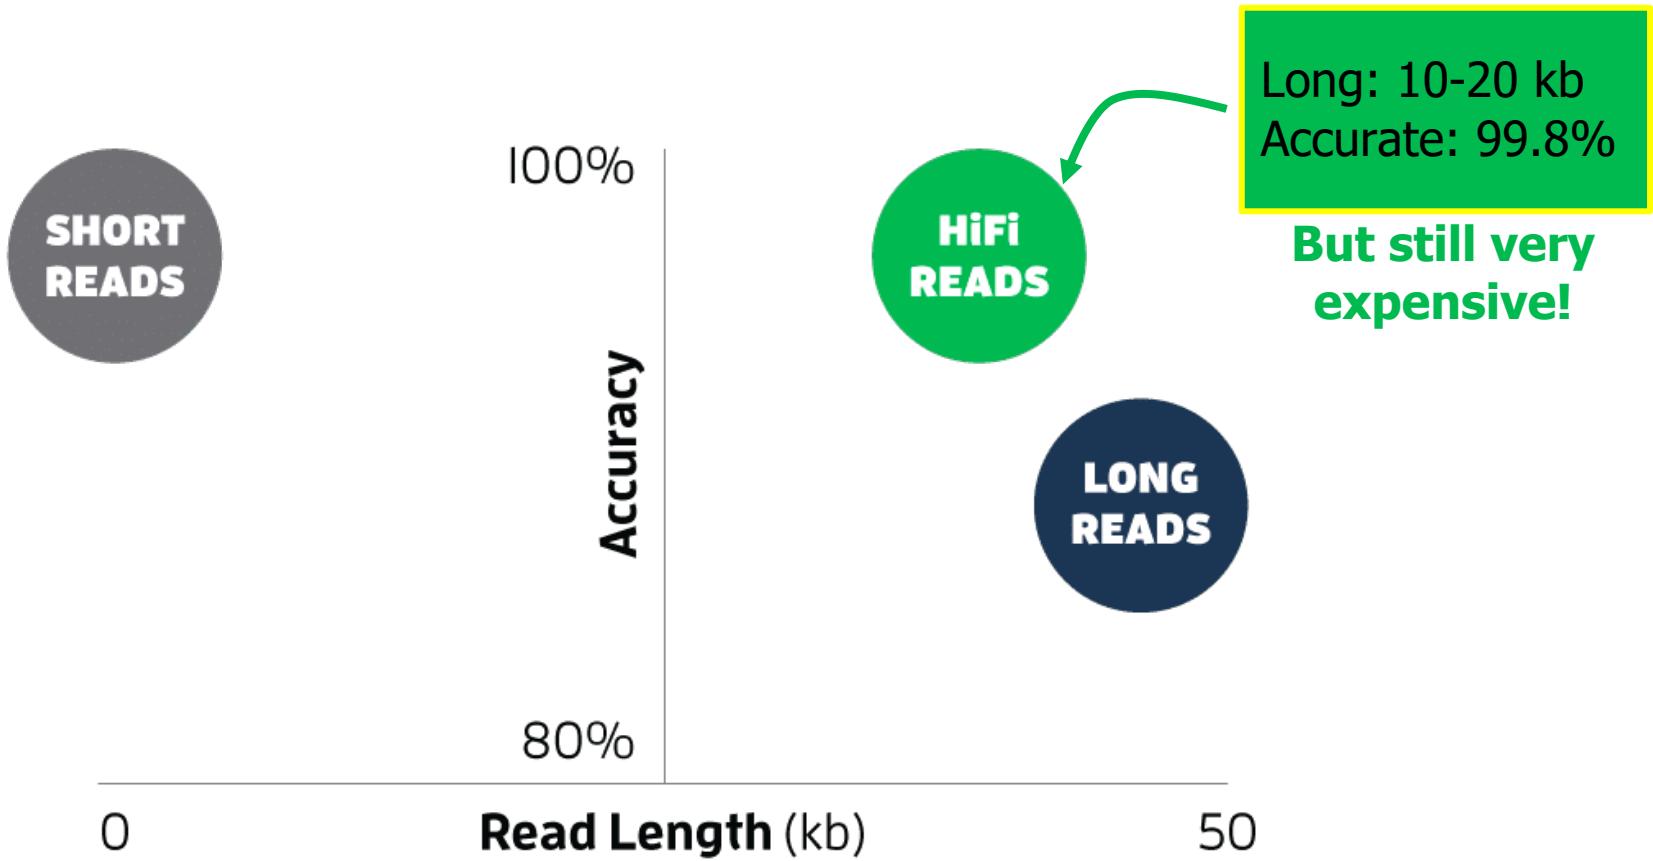


Oxford Nanopore MinION

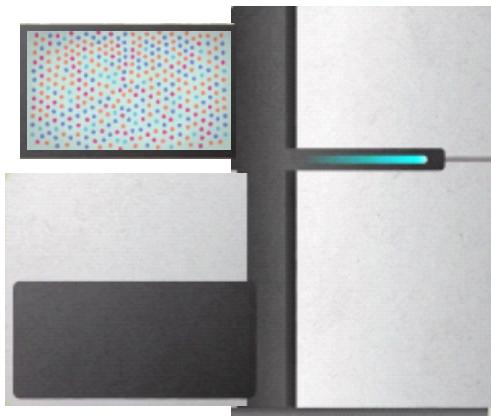
Senol Cali+, “[Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions](#),” *Briefings in Bioinformatics*, 2018.

[Open arxiv.org version] [Slides (pptx) (pdf)] [Talk Video at AACBB 2019]

Types of Genomic Reads



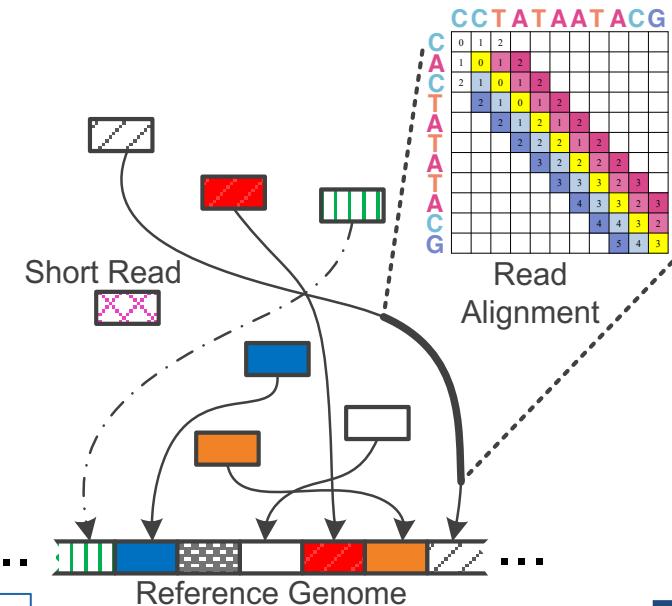
Wenger+, "[Accurate circular consensus long-read sequencing improves variant detection and assembly of a human genome](#)", *Nature Biotechnology*, 2019



Billions of Short Reads

```

ATATATAACGTACGTACGT
TTTAGTACGTACGTACGT
ATACGTACTAGTACGTACGT
ACGCCCCTACGTA
ACGTACTAGTACGT
TTAGTACGTACGTACGT
TACGTACTAAAGTACGT
TACGTACTAGTACGT
TTTAAAAACGTA
CGTACTAGTACGT
GGGAGTACGTACGT
    
```



1 Sequencing

Genome Analysis

2 Read Mapping

reference: TTTATCGCTTCATGACGCAG

read1:	ATCGC ATCC
read2:	TATCGC ATC
read3:	CATCC ATGA
read4:	CGCTTCCAT
read5:	CCATGACGC
read6:	TTCCATGAC



3 Variant Calling

4 Scientific Discovery

Read Mapping Techniques in 111 Pages

In-depth analysis of 107 read mappers (1988-2020)

Mohammed Alser, Jeremy Rotman, Dhrithi Deshpande, Kodi Taraszka, Huwenbo Shi, Pelin Icer Baykal, Harry Taegyun Yang, Victor Xue, Sergey Knyazev, Benjamin D. Singer, Brunilda Balliu, David Koslicki, Pavel Skums, Alex Zelikovsky, Can Alkan, Onur Mutlu, Serghei Mangul

"Technology dictates algorithms: Recent developments in read alignment"

Genome Biology, 2021

[[Source code](#)]

Alser *et al.* *Genome Biology* (2021) 22:249
<https://doi.org/10.1186/s13059-021-02443-7>

Genome Biology

REVIEW

Open Access



Technology dictates algorithms: recent developments in read alignment

Mohammed Alser^{1,2,3†}, Jeremy Rotman^{4†}, Dhrithi Deshpande⁵, Kodi Taraszka⁴, Huwenbo Shi^{6,7}, Pelin Icer Baykal⁸, Harry Taegyun Yang^{4,9}, Victor Xue⁴, Sergey Knyazev⁸, Benjamin D. Singer^{10,11,12}, Brunilda Balliu¹³, David Koslicki^{14,15,16}, Pavel Skums⁸, Alex Zelikovsky^{8,17}, Can Alkan^{2,18}, Onur Mutlu^{1,2,3†} and Serghei Mangul^{5*†}

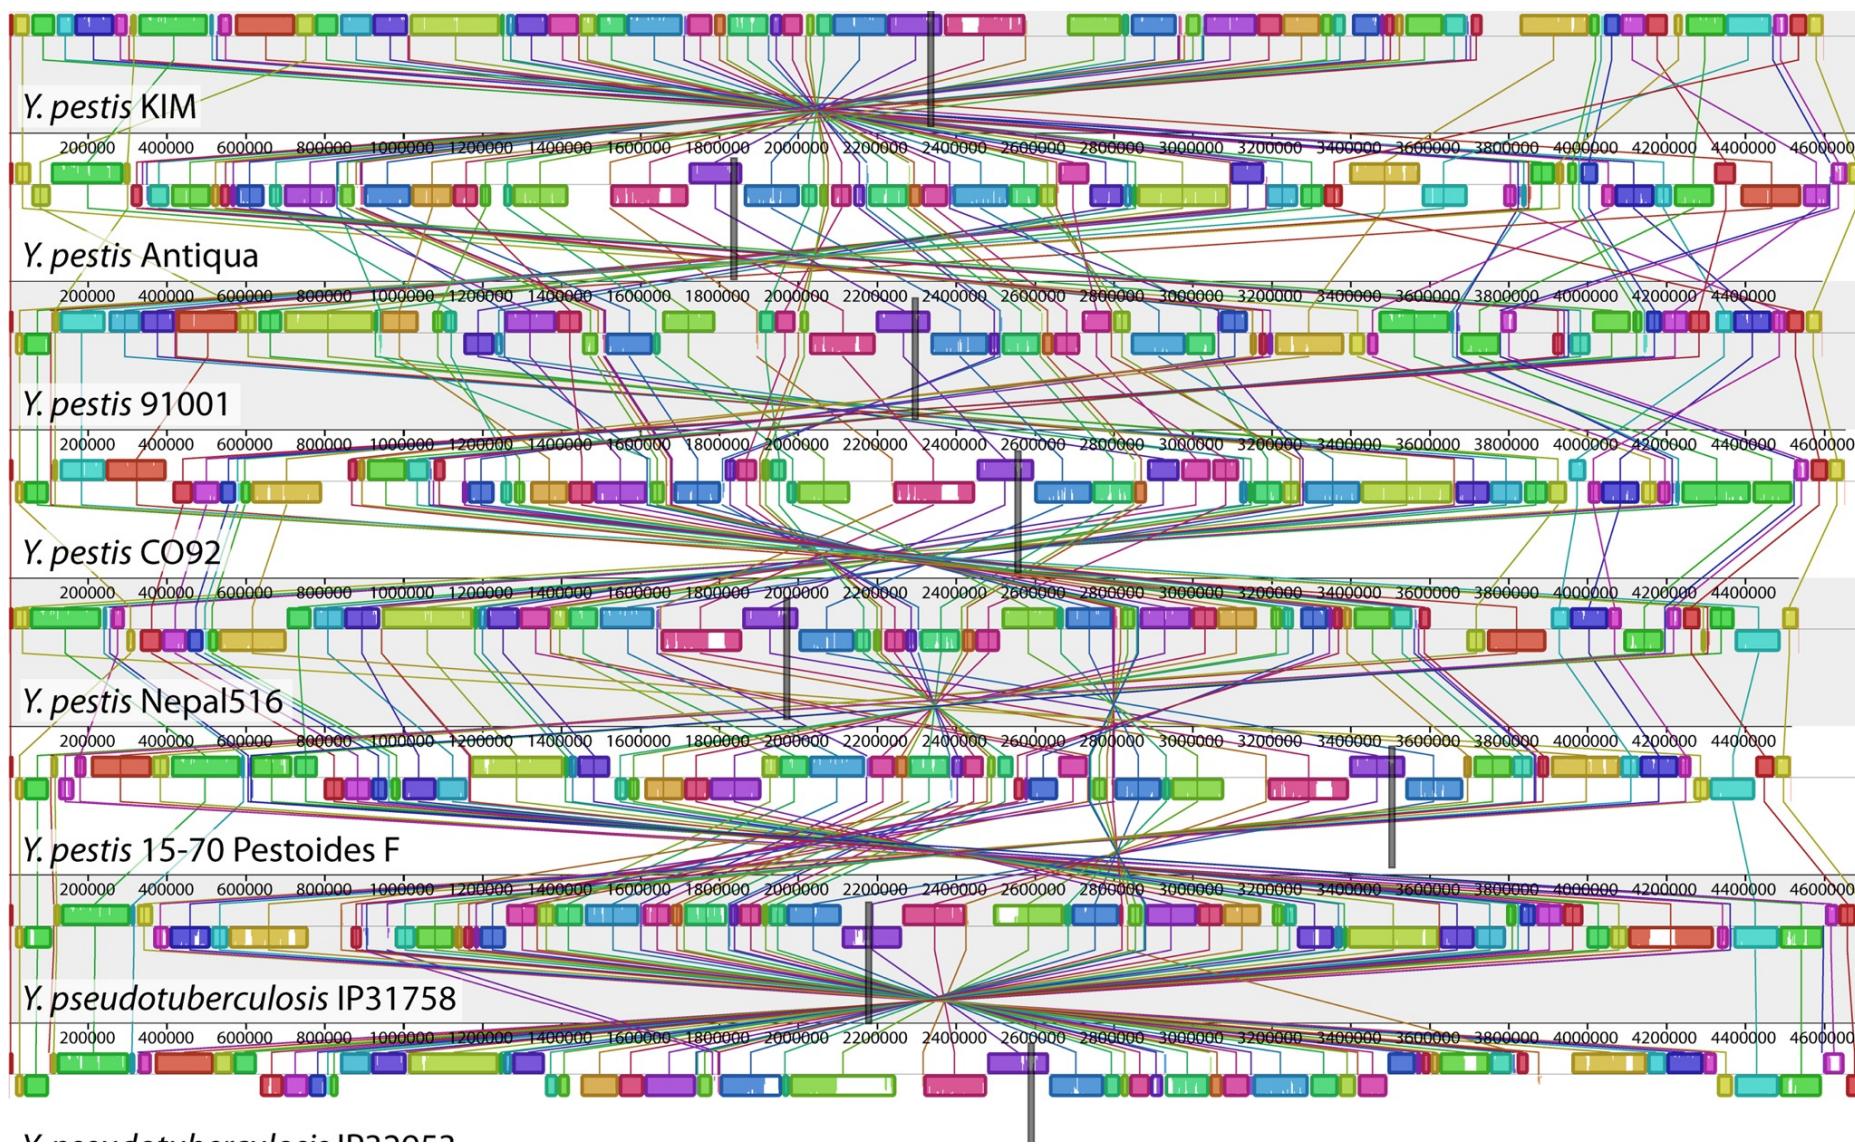
Why Do We Care?

Multiple sequence alignment

PHD Htm				-----MMMMMMMMMMMMMMMM-----	
16082665	<i>T. acid</i>	10	----MASDRKSEG FQSGAGLIRYFEEEEEIKGPAIDPKLVVYMGIAVAIIVEIAKIFWPP	-----	(55)
13541150	<i>T. volc</i>	10	----MASDKKSEG FQSGAGLIRYFEEEEEIKGPAIDPKLVVYIGIAVAIMVELAKIFWPP	-----	(55)
RFAC01077	<i>F. acid</i>	13	-MTSMAKDNQNEN FQSGAGLIRYFNEEEIKGPAIDPKLIYIGIAMGVIVELAKVFWPV	-----	(58)
15791336	<i>H. NRC1</i>	10	----MSSGQNSGGLMSSAGLVRYFDSED SNALQIDPRSVVAVGAAFFGLVVLLAQFFA	-----	(53)
RAG22196	<i>A. fulg</i>	14	MAKAPKGAKTPPLMS SAGIMRYFEE-EKTQIKVSPKTILAAGIVTGVLIILNAYGLWP	-----	(68)
RPO01000	<i>P. aby</i> s	9	----MAKEKTTLPPTGAGLMRFFDE- DTRAIKIPKGAVALTILILIIFEIIIHFVVGPRIFG	-----	(56)
RPH01741	<i>P. hori</i>	9	----MAKEKTTLPPTGAGLMRFFDE- DTRAIKIPKGAIALVLILILIIFEILLHFVVGPRIFG	-----	(56)
AE000914	<i>M. ther</i>	10	----MAKKDKKTLPPSGAGLVRYFEE- ETKGPKLTPEQVVVMSIILAVFCLVLRFSG	-----	(52)
RMJ09857	<i>M. jann</i>	9	----MSKREESTGLATSAGLIRYMD- TSKIRVKPEHVIGVTVAFVIIEAILTYGRFL	-----	(53)
15920503	<i>S. toko</i>	13	-MPSSKKKKSTVPLASMAGLIRYYEE- ENEKIKISPKLIIISIIMVAGVIVASILIPPP	-----	(58)
AE006662	<i>S. solf</i>	11	-MPSSKKKKSTVPMMSAGLIRYYEE- ENEKVKISPKIVIGASLALTIIIVIVITKLF	-----	(55)
RPK02491	<i>P. aero</i>	12	--MARRRKYSGLNPFVAAGLIKFSEE GELEKIKLTPRAAVVISLAIIGLLIAINLLLPL	-----	(58)
RAP00437	<i>A. pern</i>	13	-MSVRRRRERRATPVTAAGLLSFYEE- YEKGKIKISPTEIVVGAAILVSAVVAAPHIFLPAPV	-----	(59)
5803165	<i>H. sapi</i>	49	-----SAGTGGMWRFYTE- DSPGLKVGPPVPLVMSLLFIASVFMLHIWGKYTRS	-----	(96)
13324684	<i>M. musc</i>	49	-----SAGTGGMWRFYTE- DSPGLKVGPPVPLVMSLLFIAAVFMLHIWGKYTRS	-----	(96)
6002114	<i>D. mela</i>	53	-----GAGTGGMWRFYTD- DSPGIKVGPPVPLVMSLLFIASVFMLHIWGKYNR	-----	(100)
14574310	<i>C. eleg</i>	32	-----GGNNNGGLWRFYTE- DSTGLKIGPPVPLVMSLVIASVFVLEIWGKFTR	-----	(81)
10697176	<i>Y. lipo</i>	41	-----GGSSSTM I KLYTD- E SQGLKVDPPVVMVLSLGFIIFSVVALFILAKVSTK	-----	(91)
6320857	<i>S. cere</i>	40	-----GGSSSS I I KLYTD-EANGFRVDSL LVVLFLSVGFI F FSVIALE F LLTKFTHI	-----	(88)
6320932	<i>S. cere</i>	33	-----TNSNN S II KIYSD-EATGLRVDPL LVVLFLAVGFIFSVVVALE V ISKVAGK	-----	(82)

Example Question: If I give you a bunch of sequences, tell me where they are the same and where they are different.

Genome Sequence Alignment: Example



The Genetic Similarity Between Species



Human ~ Human

99.9%



Human ~ Chimpanzee
96%



Human ~ Cat
90%



Human ~ Cow
80%



Human ~ Banana
50-60%

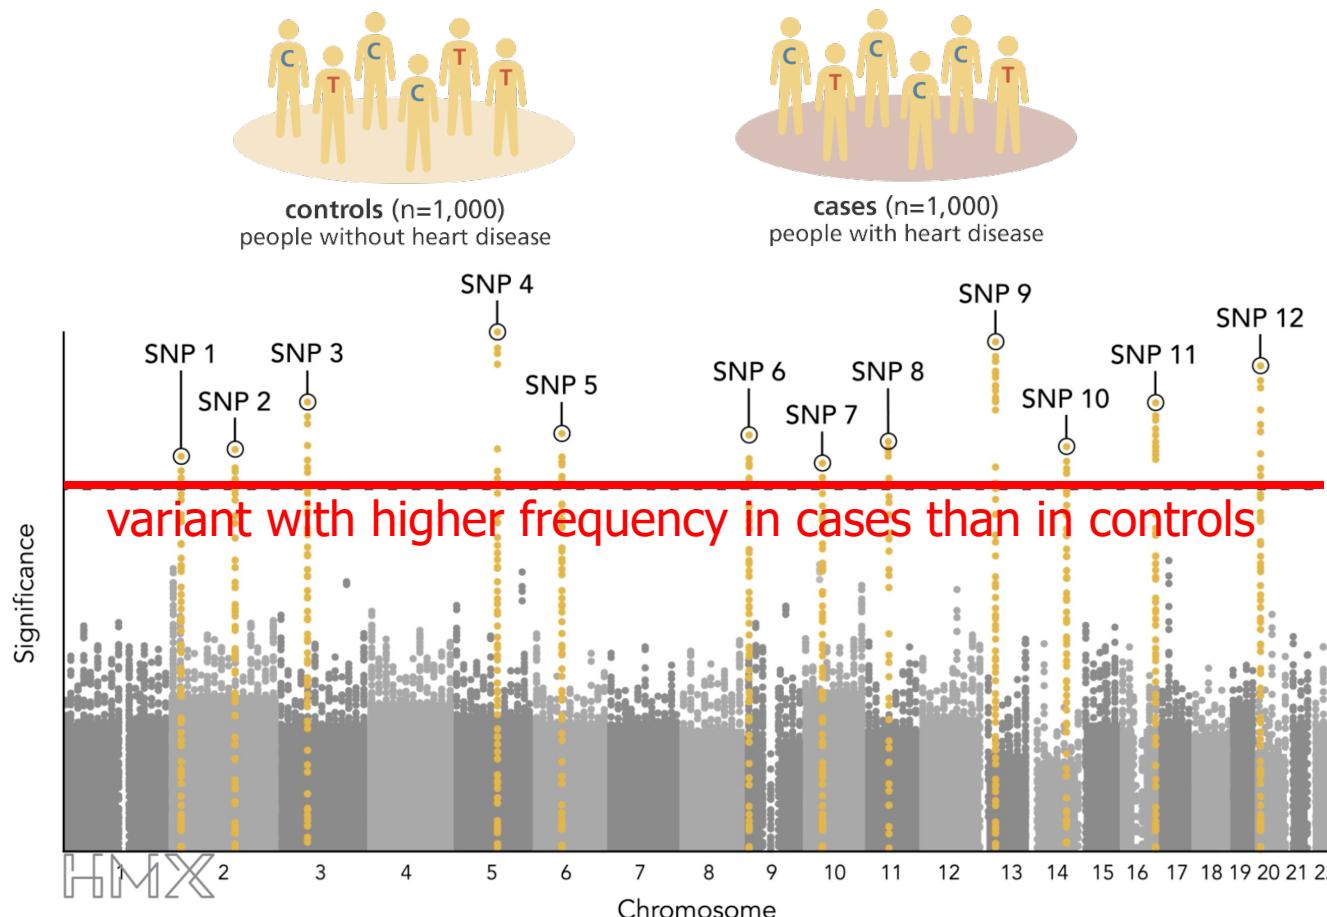
Finding Variations Associated with Traits

	SNP1	SNP2	Blood Pressure
Individual #1	...ACATG CC GACATTTCATA GG CC...		180
Individual #2	...ACATG CC GACATTTCATA A GCC...		175
Individual #3	...ACATG CC GACATTTCATA GG CC...		170
Individual #4	...ACATG CC GACATTTCATA A GCC...		165
Individual #5	...ACATG CC GACATTTCATA GG CC...		160
Individual #6	...ACATG CC GACATTTCATA GG CC...		145
Individual #7	...ACATG CC GACATTTCATA A GCC...		140
Individual #8	...ACATG CC GACATTTCATA A GCC...		130
Individual #9	...ACATG TC GACATTTCATA GG CC...		120
Individual #10	...ACATG TC GACATTTCATA A GCC...		120
Individual #11	...ACATG TC GACATTTCATA GG CC...		115
Individual #12	...ACATG TC GACATTTCATA A GCC...		110
Individual #13	...ACATG TC GACATTTCATA GG CC...		110
Individual #14	...ACATG TC GACATTTCATA A GCC...		110
Individual #15	...ACATG TC GACATTTCATA GG CC...		105
Individual #16	...ACATG TC GACATTTCATA A GCC...		100

SNP: single nucleotide polymorphism

Genome-Wide Association Studies (GWAS)

- Enables detection of genetic variants associated with phenotypes using two groups of people.



SNPs and Personalized Medicine

openSNP Search ≡

SNP rs12979860

Basic Information

Name	rs12979860
Chromosome	19
Position	39248147
Weight of evidence	926

Allele Frequency



A
T
G
C
-
0

Links to SNPedia

Title	Summary
rs12979860 T/T	~20-25% of such hepatitis c patients respond to treatment
rs12979860 C/C	~80% of such hepatitis c patients respond to treatment
rs12979860 C/T	~20-40% of such hepatitis c patients respond to treatment

Much Larger Structural Variations



AUTISM

Weiss, *N Eng J Med* 2008
Deletion of 593 kb



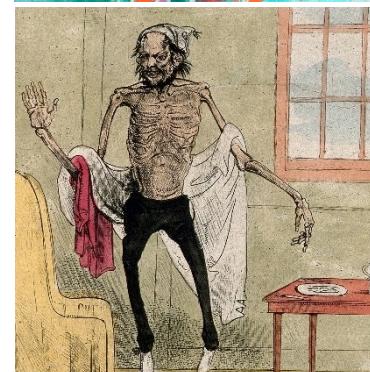
SCHIZOPHRENIA

McCarthy, *Nat Genet* 2009
Duplication of 593 kb



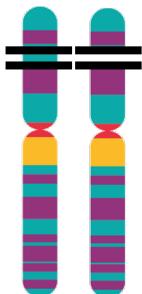
OBESITY

Walters, *Nature* 2010
Deletion of 593 kb

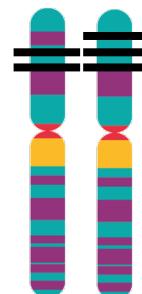


UNDERWEIGHT

Jacquemont, *Nature* 2011
Duplication of 593 kb



Deletion in the short arm
of chromosome 16 (16p11.2)



Duplication in the short arm
of chromosome 16 (16p11.2)

Personalized Medicine for Critically Ill Infants

- rWGS can be performed in 2-day (**costly**) or 5-day time to interpretation.
- Diagnostic rWGS for infants
 - Avoids **morbidity**
 - Reduces **hospital stay length** by 6%-69%
 - Reduces **inpatient cost** by \$800,000-\$2,000,000.

Article | Open Access | Published: 04 April 2018

Rapid whole-genome sequencing decreases infant morbidity and cost of hospitalization

Lauge Farnaes, Amber Hildreth, Nathaly M. Sweeney, Michelle M. Clark, S Chowdhury, Shareef Nahas, Julie A. Cakici, Wendy Benson, Robert H. K Richard Kronick, Matthew N. Bainbridge, Jennifer Friedman, Jeffrey J. Go Ding, Narayanan Veeraraghavan, David Dimmock & Stephen F. Kingsmore

npj Genomic Medicine 3, Article number: 10 (2018) | [Cite this article](#)

Article | Open Access | Published: 05 May 2020

Clinical utility of 24-h rapid trio-exome sequencing for critically ill infants

Huijun Wang, Yanyan Qian, Yulan Lu, Qian Qin, Guoping Lu, Guoqiang Cheng, Ping Zhang, Lin Yang, Bingbing Wu & Wenhao Zhou

npj Genomic Medicine 5, Article number: 20 (2020) | [Cite this article](#)

Recommended Reading

nature reviews genetics

Explore our content ▾

Journal information ▾

nature > nature reviews genetics > review articles > article

Review Article | Published: 15 November 2019

Structural variation in the sequencing era

Steve S. Ho, Alexander E. Urban & Ryan E. Mills 

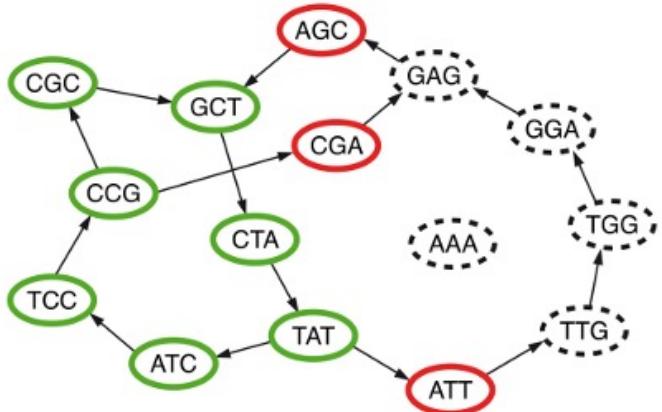
Nature Reviews Genetics 21, 171–189(2020) | Cite this article

15k Accesses | 16 Citations | 309 Altmetric | Metrics

Ho+, "Structural variation in the sequencing era", Nature Reviews Genetics, 2020

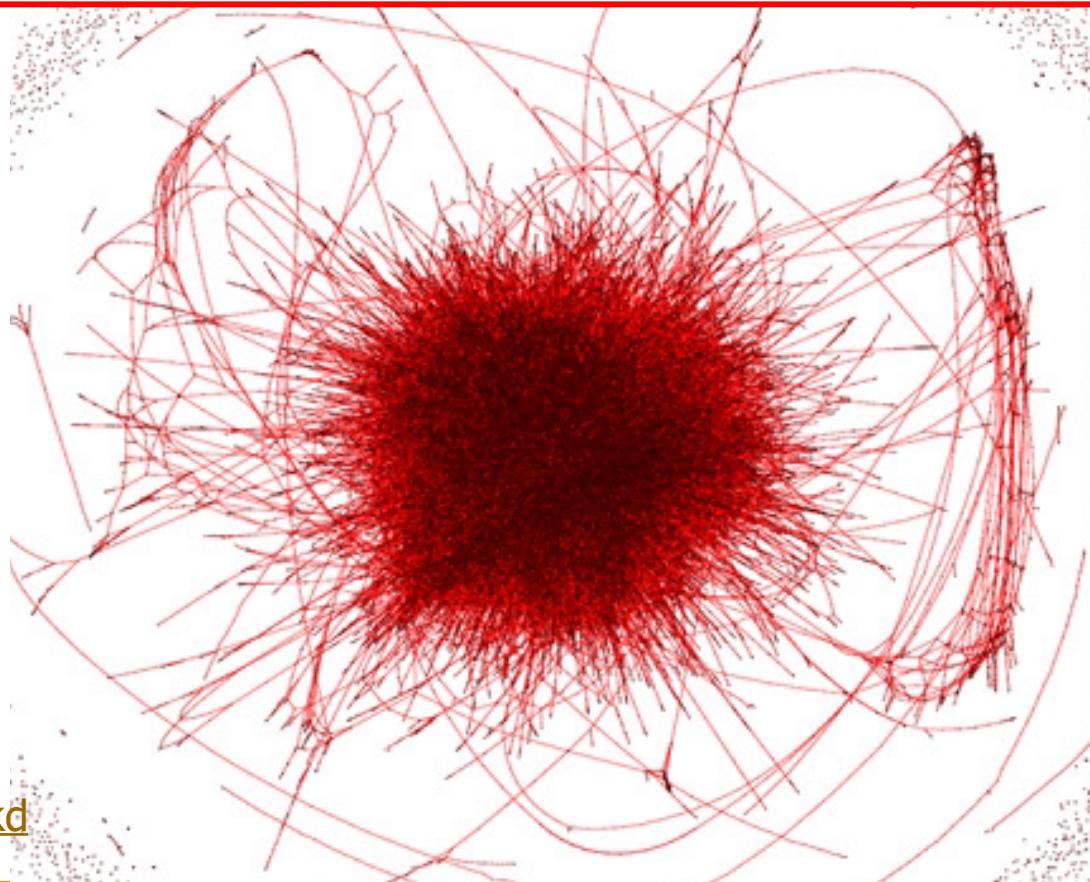
Metagenomics, genome assembly, de novo sequencing

Question 2: Given a bunch of short sequences,
Can you identify the approximate species cluster
for genetically unknown organisms (bacteria)?



uncleaned de Bruijn graph

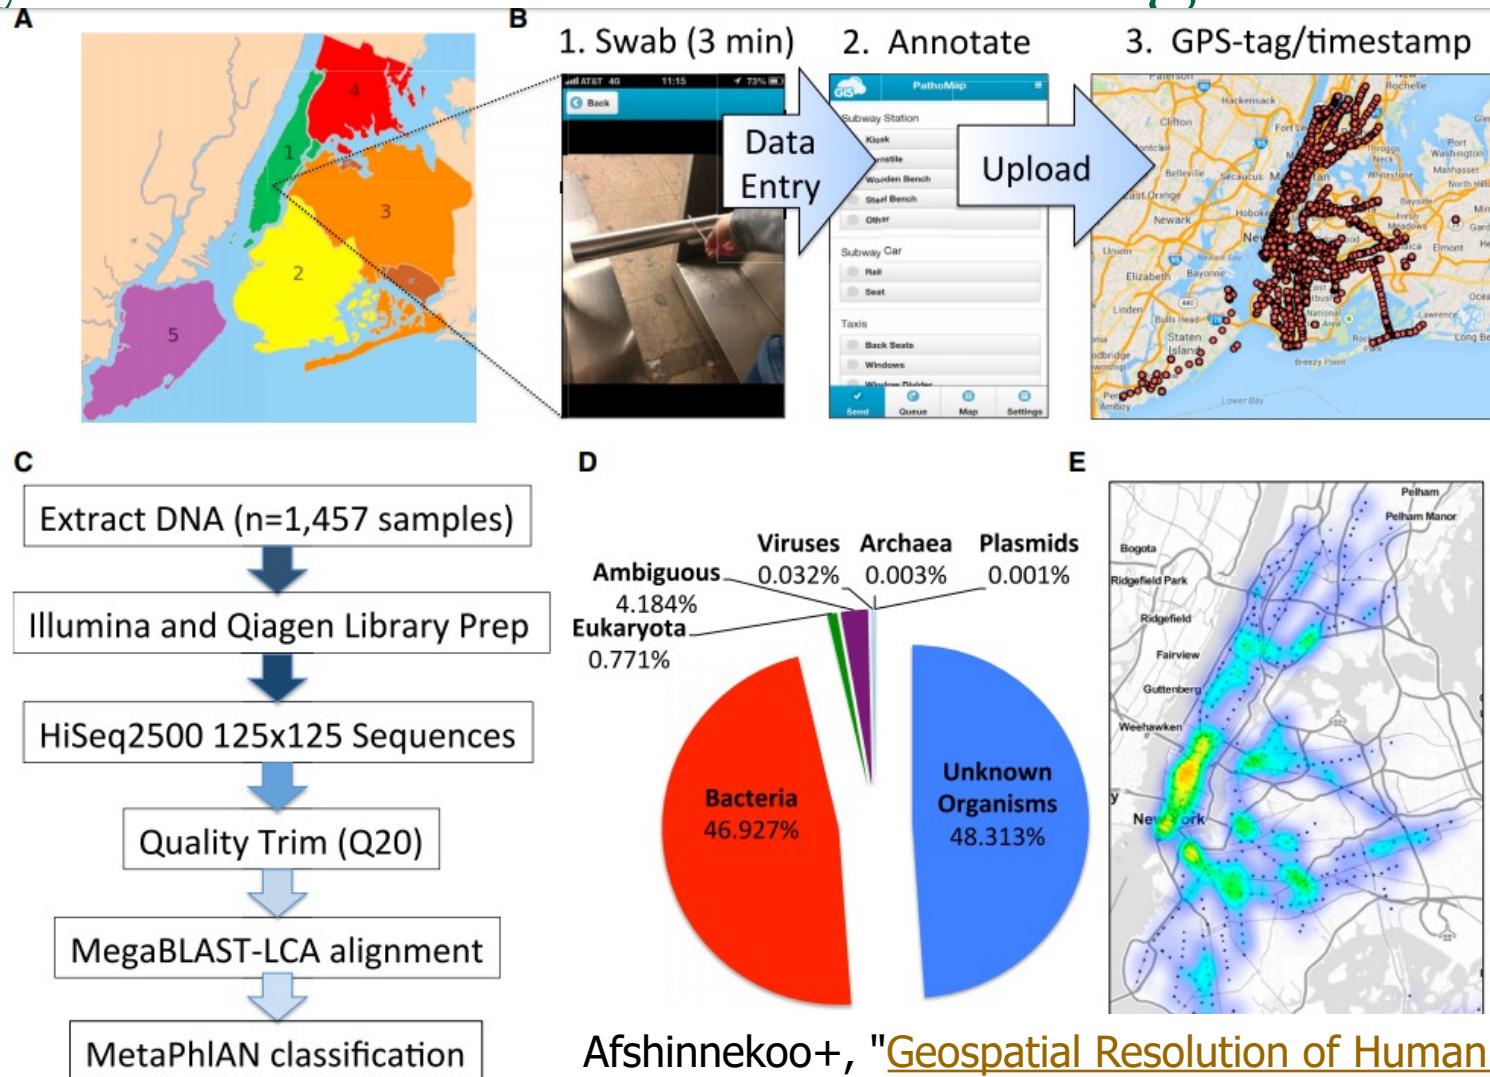
<http://math.oregonstate.edu/~koslickd>



Population-Scale Microbiome Profiling



City-Scale Microbiome Profiling



Afshinnekoo+, "[Geospatial Resolution of Human and Bacterial Diversity with City-Scale Metagenomics](#)", Cell Systems, 2015

Figure 1. The Metagenome of New York City

- (A) The five boroughs of NYC include (1) Manhattan (green)
- (B) The collection from the 466 subway stations of NYC across the 24 subway lines involved three main steps: (1) collection with Copan Elution swabs, (2) data entry into the database, and (3) uploading of the data. An image is shown of the current collection database, taken from <http://pathomap.giscloud.com>.
- (C) Workflow for sample DNA extraction, library preparation, sequencing, quality trimming of the FASTQ files, and alignment with MegaBLAST and MetaPhlAn to discern taxa present.

Global-Scale Microbiome Profiling

Cell

Log in Register Sign up

PDF [9 MB] Figures Save

ARTICLE | ONLINE NOW

A global metagenomic map of urban microbiomes and antimicrobial resistance

David Danko⁶⁸ • Daniela Bezdán⁶⁸ • Evan E. Afshin • ... Sibo Zhu • Christopher E. Mason ✉ 69 ✉ 69

The International MetaSUB Consortium • Show all authors • Show footnotes

Open Access • Published: May 26, 2021 • DOI: <https://doi.org/10.1016/j.cell.2021.05.002>

A

City samples

- 40°42'46"N 74°00'21"W
read1 ACGATCACCA
read2 CGATCGACA
read3 TCGAAGCGG
- 51°30'26"N 00°07'39"W
read1 CGATCGAAAT
read2 GGTGCTGAT
read3 CAAGCATGC
- 35°41'00"N 139°41'00"E
read1 CCGATTATTC
read2 CGTTGTAC
read3 AACAGAATA
- 22°54'30"S 43°11'47"W
read1 CCCGTATAAA
read2 TTCCCTTATA
read3 CGATCTATG

Extracted k-mers

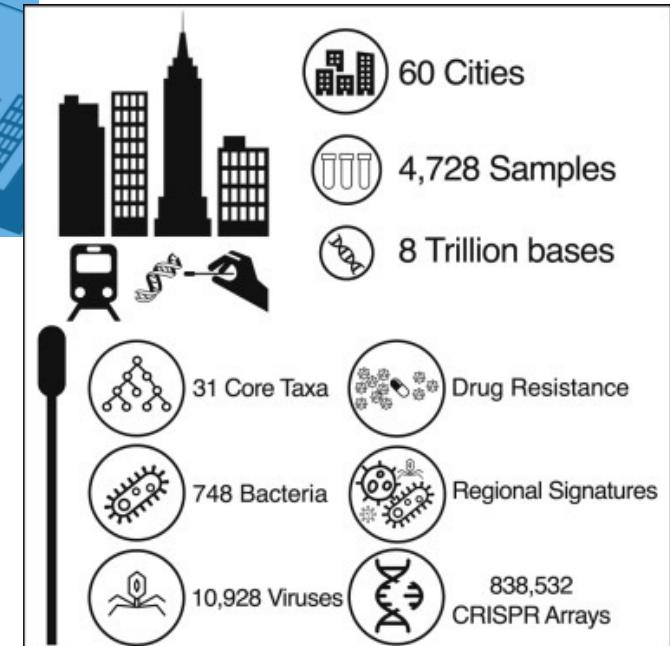
Graph-based multi-sample sequence index

Sequence query

World Map

Query **City** **Score** **Metadata**

Query	City	Score	Metadata
CGATCG...	●	100	line 1, bench, 2018/6/21-10:15am
CGATCG...	●	96	line 2, kiosk, 2018/6/21-2:27pm
CGATCG...	●	95	line 2, kiosk, 2018/6/21-11:49am
TTATAC...	●	93	line h, bench, 2018/6/21-5:18pm



Danko+, "A global metagenomic map of urban microbiomes and antimicrobial resistance", Cell, 2021

A Tsunami of Sequencing Data

A Tera-scale increase in sequencing production in the past 25 years		
Genes & Operons	1990	Kilo = 1,000
Bacterial genomes	1995	Mega = 1,000,000
Human genome	2000	Giga = 1,000,000,000
Human microbiome	2005	Tera = 1,000,000,000,000
50K Microbiomes	2015	Peta = 1,000,000,000,000,000
what is expected for the next 15 years ? (a Giga?)		
200K Microbiomes	2020	Exa = 1,000,000,000,000,000,000
1M Microbiomes	2025	Zetta = 1,000,000,000,000,000,000,000
Earth Microbiome	2030	Yotta = 1,000,000,000,000,000,000,000,000

Source:
[@kyrpides](#)

Another Question: Example from 2020-...

200 Oxford Nanopore sequencers have left UK for China, to support rapid, near-sample coronavirus sequencing for outbreak surveillance

Fri 31st January 2020

Following extensive support of, and collaboration with, public health professionals in China, Oxford Nanopore has shipped an additional 200 MinION sequencers and related consumables to China. These will be used to support the ongoing surveillance of the current coronavirus outbreak, adding to a large number of the devices already installed in the country.



Each MinION sequencer is approximately the size of a stapler, and can provide rapid sequence information about the coronavirus.



700Kg of Oxford Nanopore sequencers and consumables are on their way for use by Chinese scientists in understanding the current coronavirus outbreak.

Example: Scalable SARS-CoV-2 Testing



THE PREPRINT SERVER FOR HEALTH SCIENCES



Cold
Spring
Harbor
Laboratory

BMJ Yale

HOME | ABOL

Search

Comments (1)

Swab-Seq: A high-throughput platform for massively scaled up SARS-CoV-2 testing

Joshua S. Bloom, Eric M. Jones, Molly Gasperini, Nathan B. Lubock, Laila Sathe, Chetan Munugala, A. Sina Booeshaghi, Oliver F. Brandenberg, Longhua Guo, James Boocock, Scott W. Simpkins, Isabella Lin, Nathan LaPierre, Duke Hong, Yi Zhang, Gabriel Oland, Bianca Judy Choe, Sukantha Chandrasekaran, Evann E. Hilt, Manish J. Butte, Robert Damoiseaux, Aaron R. Cooper, Yi Yin, Lior Pachter, Omai B. Garner, Jonathan Flint, Eleazar Eskin, Chongyuan Luo, Sriram Kosuri, Leonid Kruglyak, Valerie A. Arboleda

doi: <https://doi.org/10.1101/2020.08.04.20167874>

Bloom+, "[Swab-Seq: A high-throughput platform for massively scaled up SARS-CoV-2 testing](#)", *medRxiv*, 2020

Example: Rapid Surveillance of Ebola Outbreak

Figure 1: Deployment of the portable genome surveillance system in Guinea.

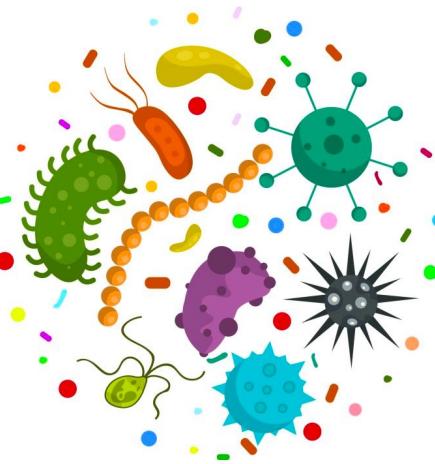


Quick+, "Real-time, portable genome sequencing for Ebola surveillance", *Nature*, 2016

We Need Faster & Scalable Genome Analysis



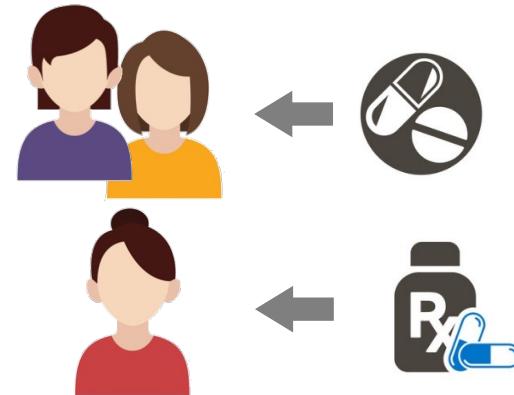
Understanding **genetic variations, species, evolution, ...**



Predicting the presence and relative abundance of **microbes** in a sample

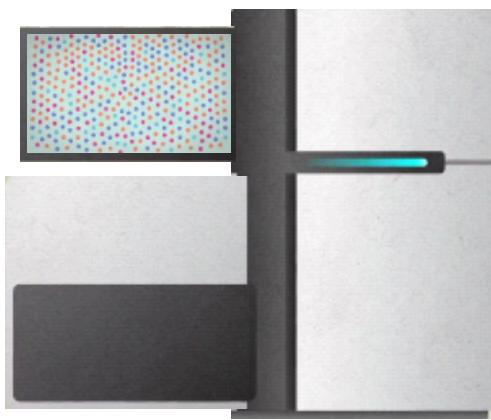


Rapid surveillance of **disease outbreaks**



Developing **personalized medicine**

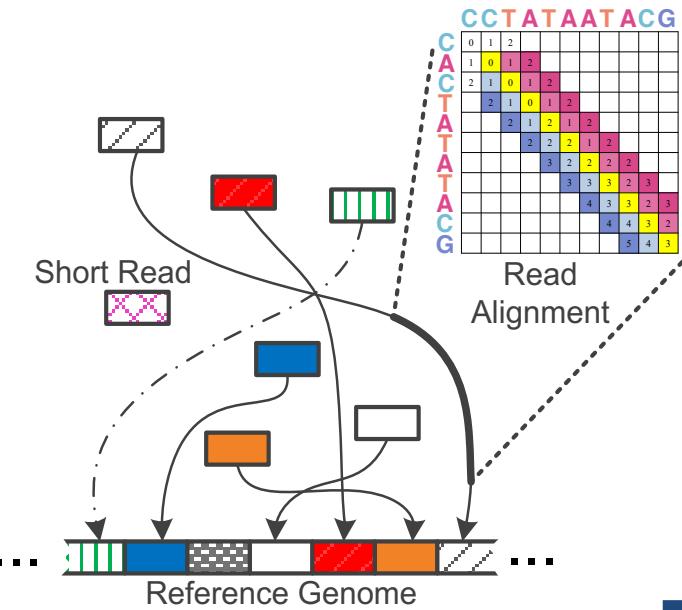
One Problem



Billions of Short Reads

```

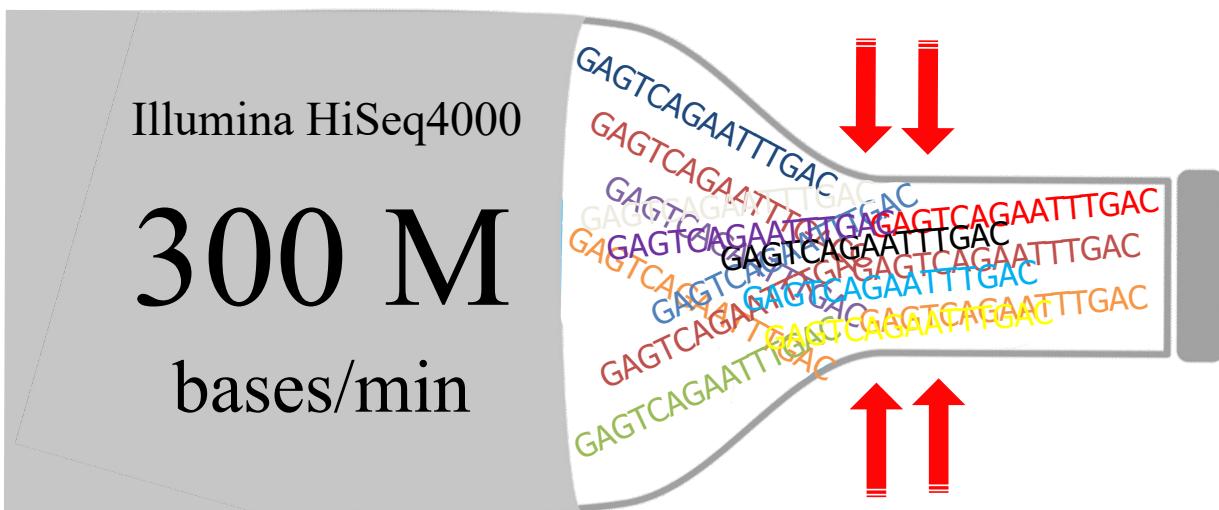
ATATATAACGTACGTACGT
TTTAGTACGTACGTACGT
ATACGTACTAGTACGTACGT
ACGCCCCTACGTA
ACGTACTAGTACGT
TTAGTACGTACGTACGT
TACGTACTAAAGTACGT
TACGTACTAGTACGT
TTTAAAAACGTA
CGTACTAGTACGT
GGGAGTACGTACGT
    
```



1 Sequencing

2 Read Mapping

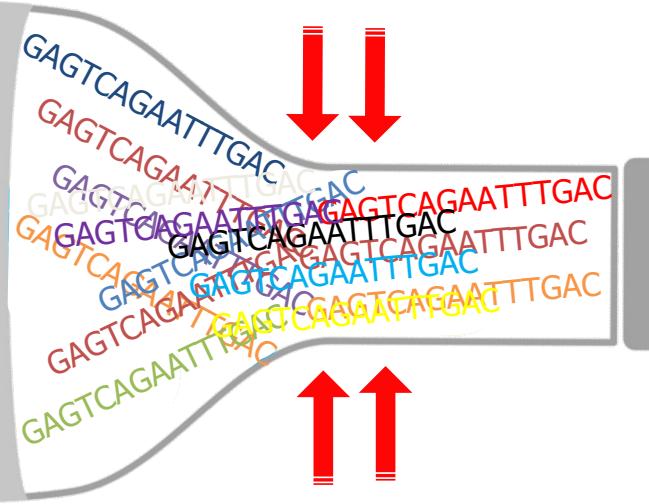
We Are Bottlenecked in Read Mapping



The Read Mapping Bottleneck

300 Million
bases/minute

Read Sequencing **

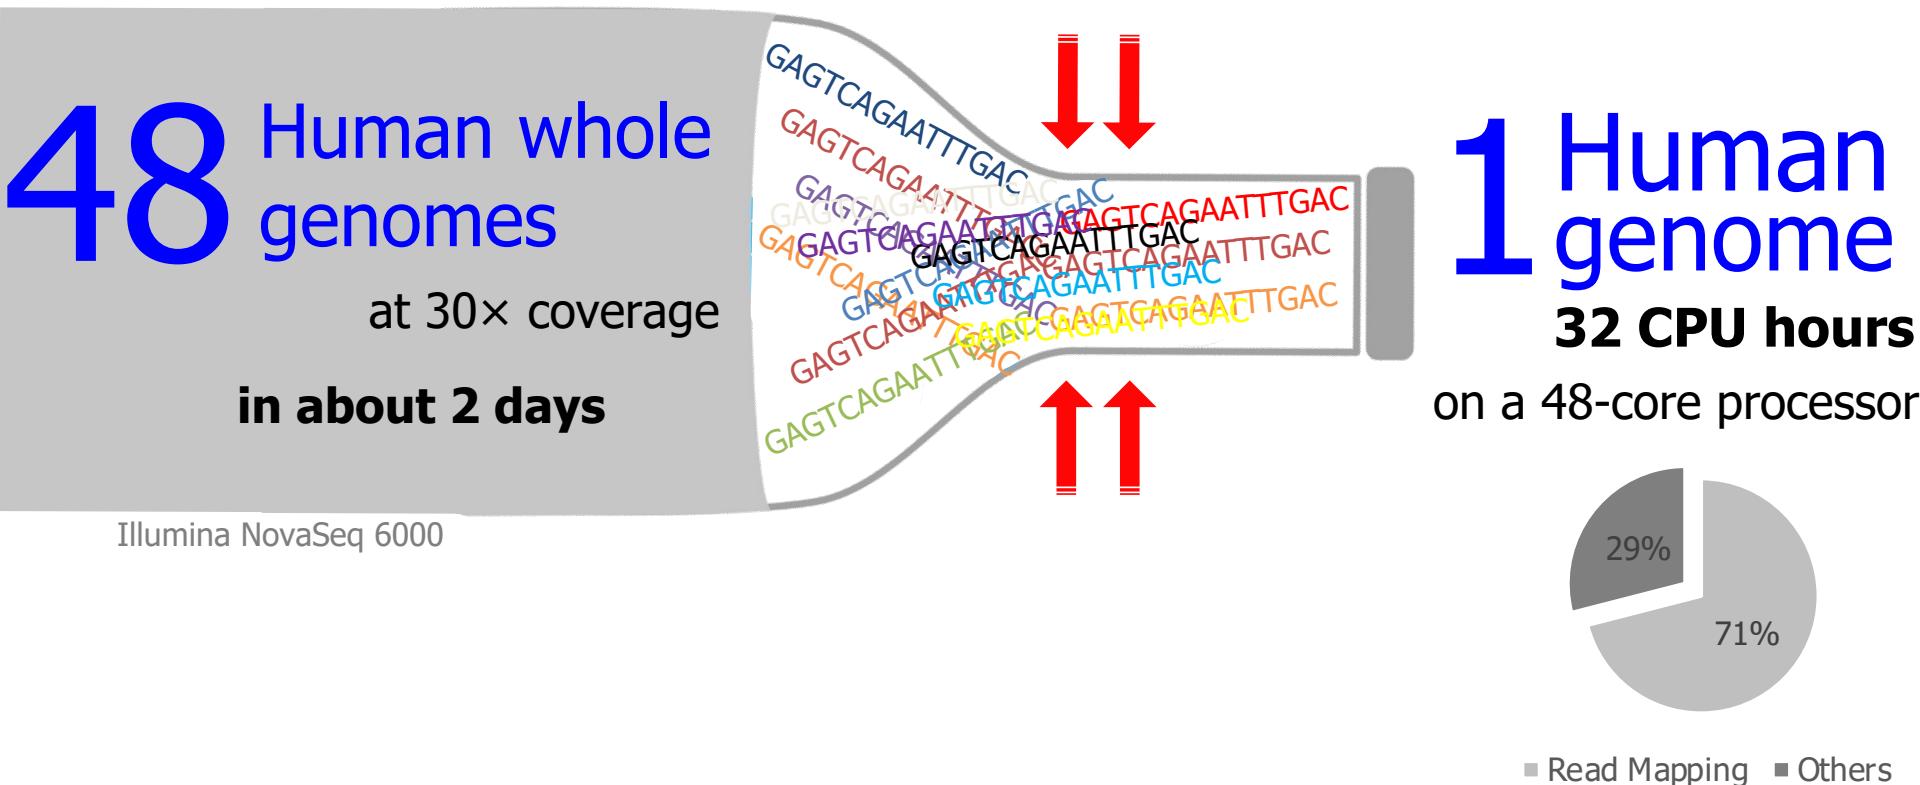


150x slower

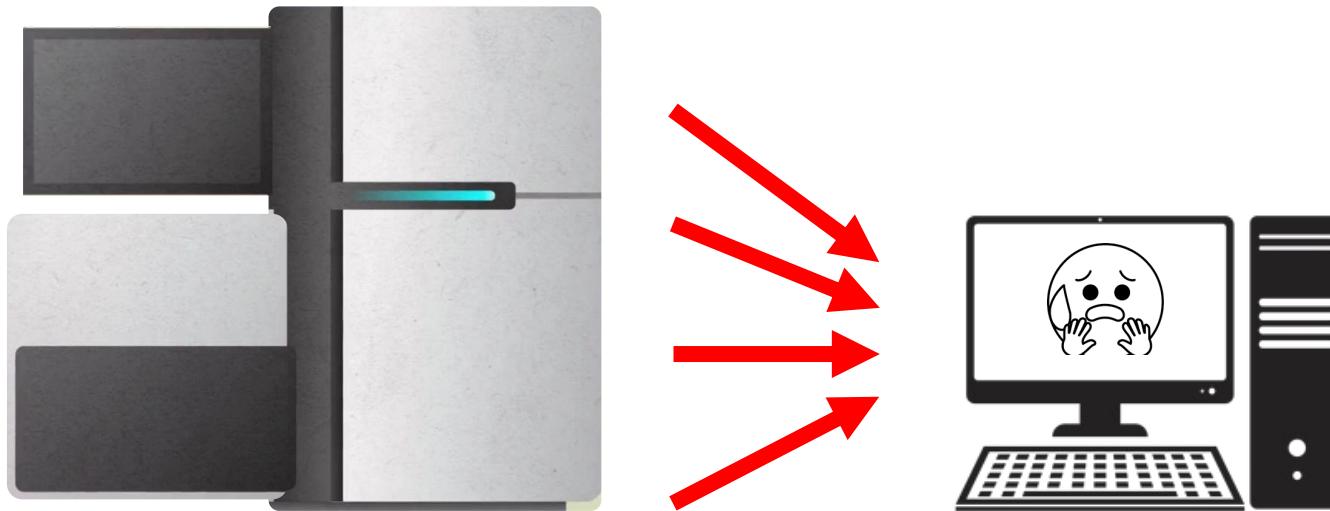
* BWA-MEM

** HiSeqX10, MinION

The Read Mapping Bottleneck



Problem with (Genome) Analysis Today



Special-Purpose Machine
for Data Generation

General-Purpose Machine
for Data Analysis

FAST

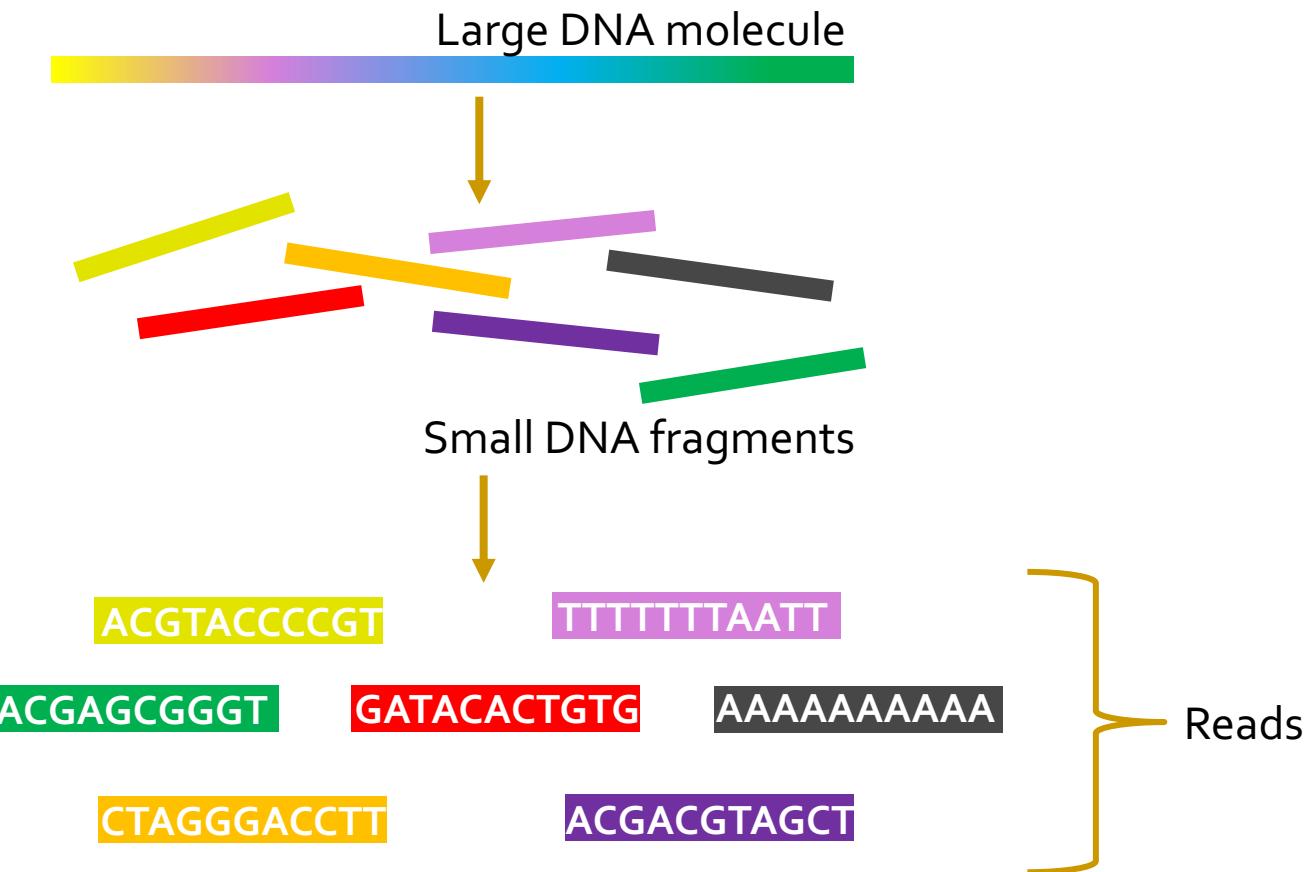
SLOW

Slow and inefficient processing capability

One Problem

**Need to construct
the entire genome
from many sequenced reads**

Genome Sequencing

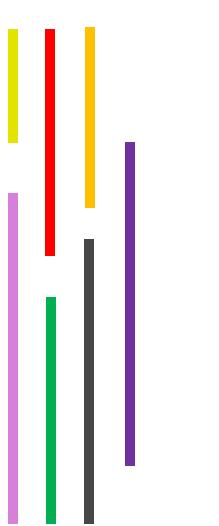


Genome Sequence Analysis

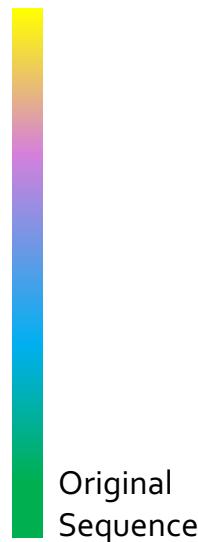
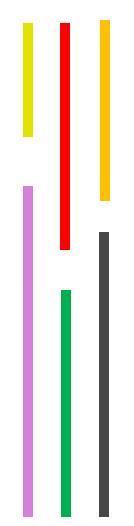
ACGTACCCCGT
ACGAGCGGGT GATACACTGTG AAAAAAAA
CTAGGGACCTT ACGACGTAGCT

Reads

Read Mapping, method of aligning the reads against a **known reference genome** to **detect matches and variations**

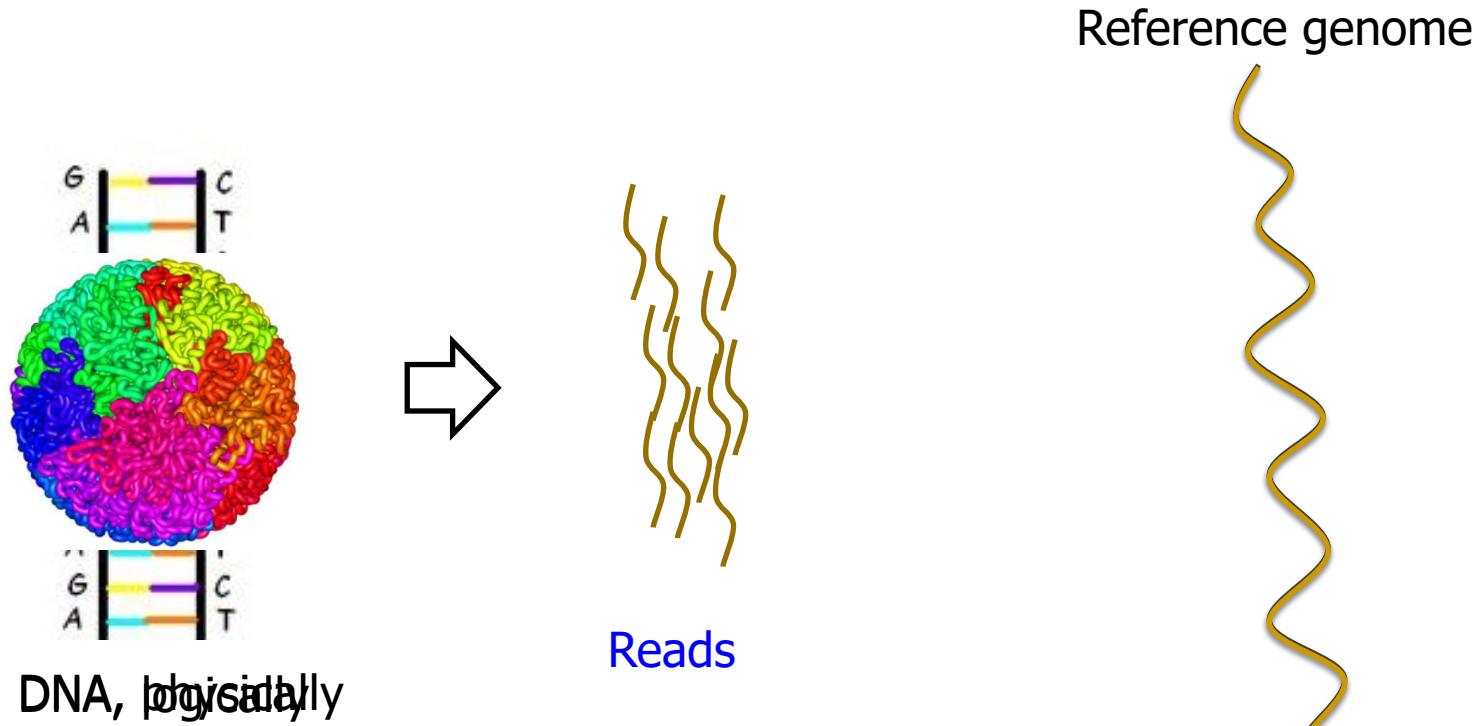


De novo Assembly, method of merging the reads in order to **construct** the original sequence (reference genome)



Read Mapping

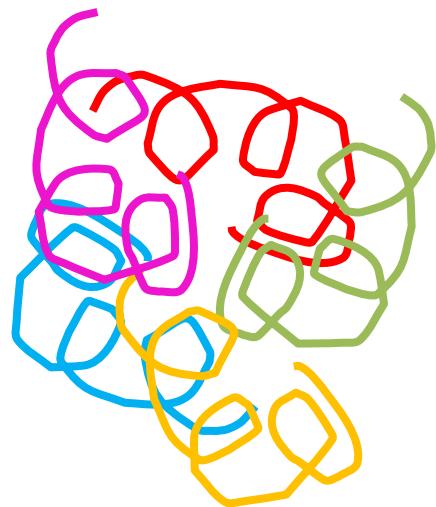
- Map many short DNA fragments (**reads**) to a known reference genome with some differences allowed



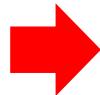
Mapping short reads to reference genome is challenging (billions of 50-300 base pair reads)

Read Mapping for Metagenomic Analysis

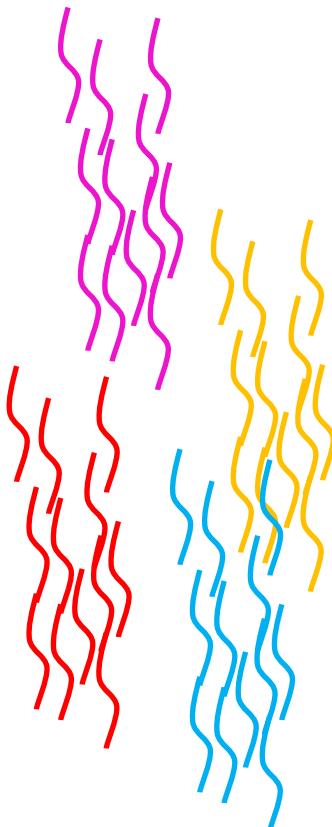
Reads from different **unknown** donors at sequencing time are mapped to **many known reference** genomes



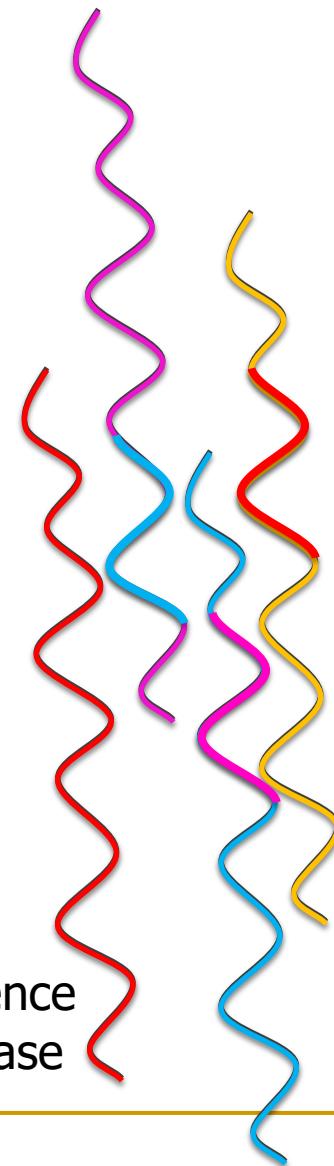
Genetic material recovered
directly from environmental
samples



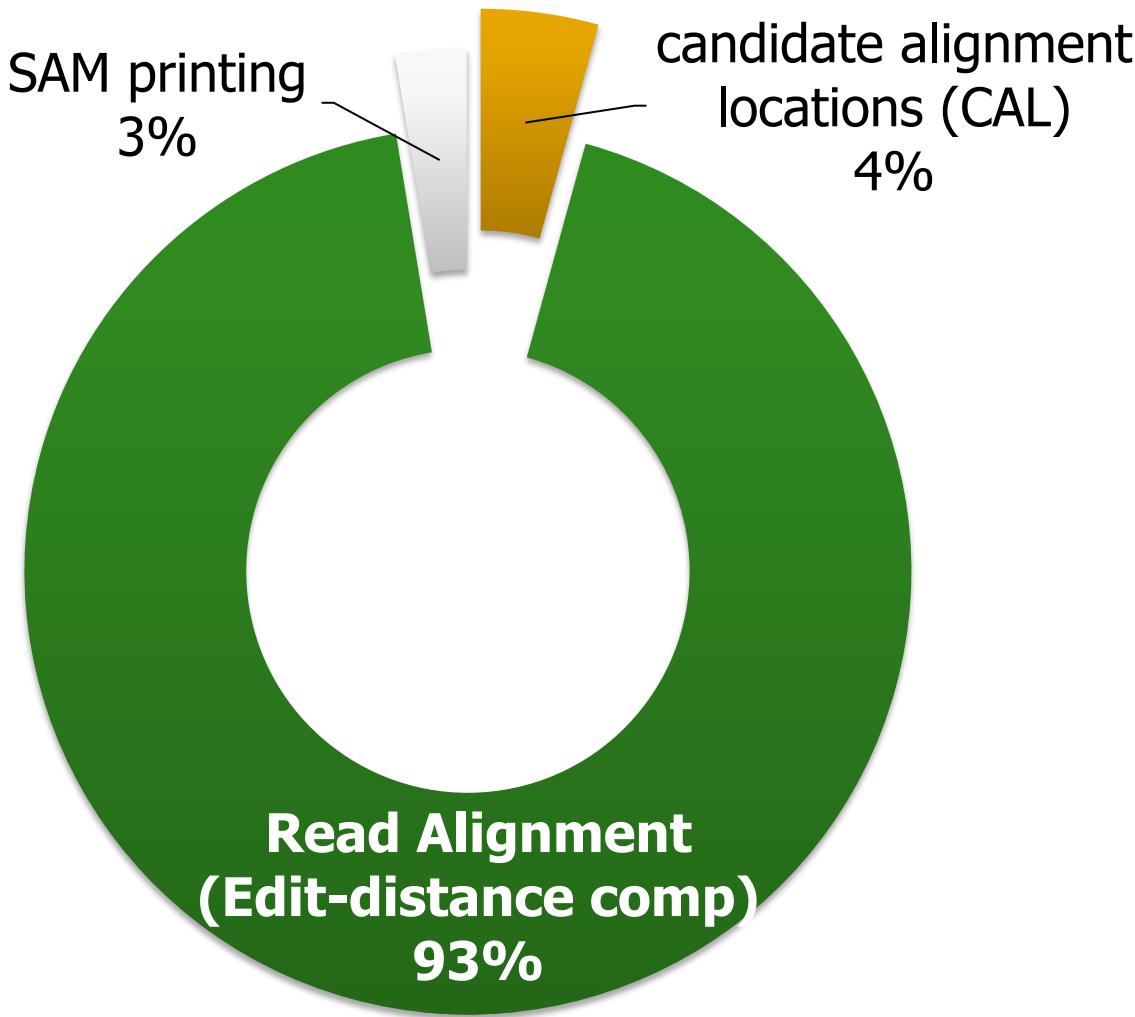
Reads in
“text format”



Reference
Database



Read Mapping Execution Time (Old Times)



Matching Each Read to Reference Genome

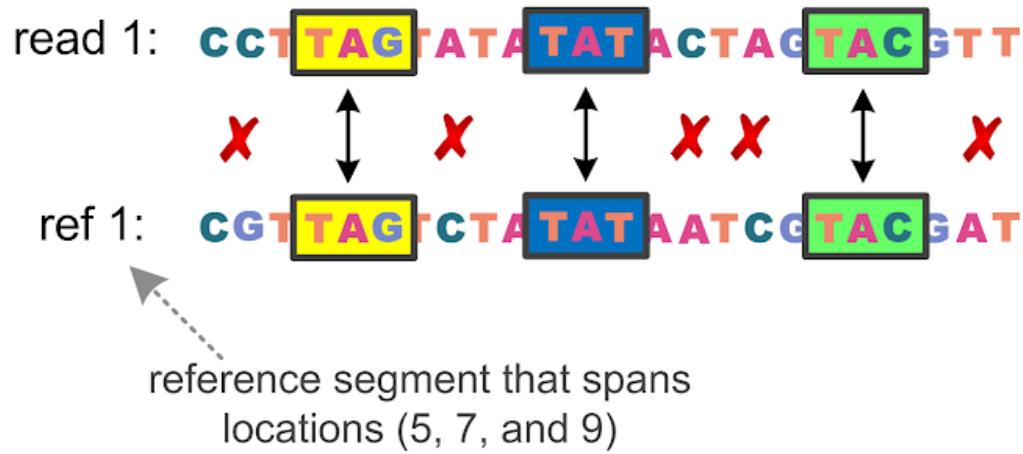
Reference Genome .FASTA file:

```
>NG_008679.1:5001-38170 Homo sapiens paired box 6 (PAX6)
ACCC[T]TCATTGACATTAAACTCTGGGGCAGG[G]GAACGC GGCTGTCA GATCT
GCCACTTCCCCTGCCGAGCGGCGGTGAGAAGTGTGGGAACC CGCGCTGCCAGGCTCACCTGCCTCCCCGC
CCTCCGCTCCCAGGTAACCGCC[CCCCGGCCCGGCTCGGGGCCCGCGGGCCTCTCCGCTG
CCAGCGACTGCTGTCCCCAAATCAAAGCCCGCCCAAGTGGCCCCGGGGCTTGATTTGCTTTAAAAG
GAGGCATACAAAGATGGAAGCGAGTTACTGAGGGAGGGATAGGAAGGGGGTGGAGGAGGGACTTGTCTT
TGCCGAGGTG[C]CAAAAGTAGCA[CTCCTA]TCCAGTCCGGCCCT
GAGCTGGGAGTAGGGGGCGGGAGTCTGCTGCTGTCTGCTAAAGCCACTCGCGACCGCGAAAAATGCA
GGAGGTGGGGACGCACTTGCATCCAGACCTCCTCTGCATCGCAGTTC[CGCTTGGGAAAG
TCCGTACCCGCGCCT[AAGACACCC]TGCCGCGGTGAGGTGCAGCAGAAGTTCCC
GCGGTTGCAAAGTGCAGATGGCTGGACCGCAACAAAGTCTAGAGATGGGTTCGTTCTCAGAAAGACGC
```

Sequenced Reads .FASTQ file:

```
@HWI-EAS209_0006_FC706VJ:5:58:5894:21141#ATCACG/1
T[AATAAAATCT]TTAGATN[NNNNNNNNNTAG
+HWI-EAS209_0006_FC706VJ:5:58:5894:21141#ATCACG/1
efcfffffcfeffffcfffffffdf`feed]`]_Ba_`__[YBBBBBBBBBBRTT
```

Base-by-Base Comparison



Read Alignment/Verification

- **Edit distance** is defined as the minimum number of edits (i.e. insertions, deletions, or substitutions) needed to make the read exactly match the reference segment.

NETHERLANDS x SWITZERLAND

N	E	-	T	H	E	R	L	A	N	D	S
S	W	I	T	Z	E	R	L	A	N	D	-

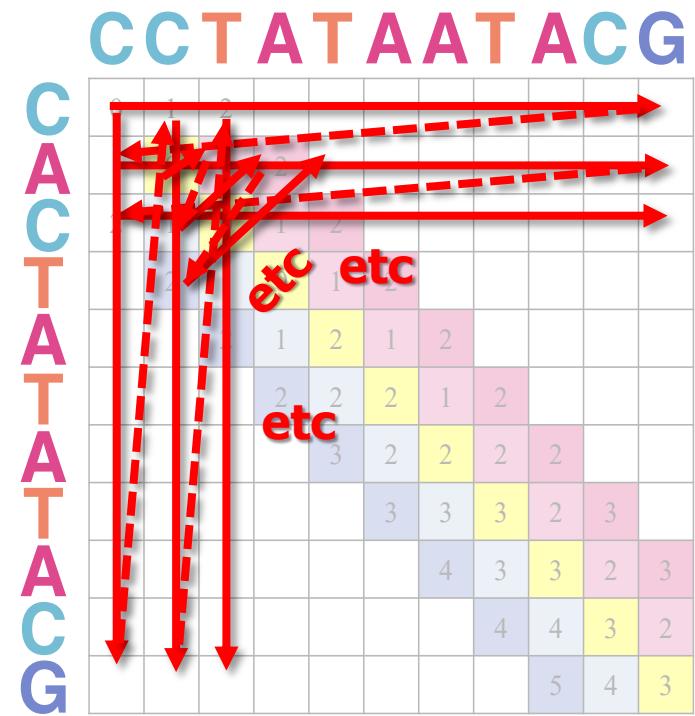


Challenges in Read Mapping

- Need to find many mappings of each read
 - A short read may map to many locations, especially with High-Throughput DNA Sequencing technologies
 - How can we find all mappings efficiently?
- Need to tolerate small variances/errors in each read
 - Each individual is different: Subject's DNA may slightly differ from the reference (Mismatches, insertions, deletions)
 - How can we efficiently map each read with up to e errors present?
- Need to map each read very fast (i.e., performance is important)
 - Human DNA is 3.2 billion base pairs long → Millions to billions of reads (State-of-the-art mappers take weeks to map a human's DNA)
 - How can we design a much higher performance read mapper?

Why Is Read Alignment Slow?

- **Quadratic-time** dynamic-programming algorithm(s)
- **Data dependencies** limit the computation parallelism
- **Entire matrix** computed even though strings may be dissimilar



Read Alignment

Example: Dynamic Programming Table

immediate left,
upper left,
upper entries of its own

		N	E	T	H	E	R	L	A	N	D	S
	S	2	3	4	5	6	7	8	9	10	11	
S	1											
W	2											
I	3											
T	4											
Z	5											
E	6											
R	7											
L	8											
A	9											
N	10											
D	11											



Example: Dynamic Programming Table

NETHERLANDS x SWITZERLAND

	N	E	T	H	E	R	L	A	N	D	S	
	0	1	2	3	4	5	6	7	8	9	10	11
S	1	1	2	3	4	5	6	7	8	9	10	10
W	2	2	2	3	4	5	6	7	8	9	10	11
I	3	3	3	3	4	5	6	7	8	9	10	11
T	4	4	4	3	4	5	6	7	8	9	10	11
Z	5	5	5	4	4	5	6	7	8	9	10	11
E	6	6	5	5	5	4	5	6	7	8	9	10
R	7	7	6	6	6	5	4	5	6	7	8	9
L	8	8	7	7	7	6	5	4	5	6	7	8
A	9	9	8	8	8	7	6	5	4	5	6	7
N	10	9	9	9	9	8	7	6	5	4	5	6
D	11	10	10	10	10	9	8	7	6	5	4	5

- Matrix-filling is $O(mn)$ time and space.
- Backtrace is $O(m + n)$ time.

Example: Dynamic Programming

- **Quadratic-time** dynamic-programming algorithm

WHY?!

Enumerate all possible prefixes

NETHERLANDS x SWITZERLAND

- C NETHERLANDS x S
- C NETHERLANDS x SW
- C NETHERLANDS x SWI
- C NETERLANDS x SWIT
- C NETHERLANDS x SWITZ
- C NETHERLANDS x SWITZE
- C NETHERLANDS x SWITZER
- C NETHERLANDS x SWITZERL
- E NETHERLANDS x SWITZERLA
- E NETHERLANDS x SWITZERLAN
- C NETHERLANDS x SWITZERLAND

	N	E	T	H	E	R	L	A	N	D	S
0	0	1	2	3	4	5	6	7	8	9	10
S	1	1	1	1	1	1	1	1	1	1	10
W	2	2	2	2	2	2	2	2	2	2	10
I	3	3	3	3	3	3	3	3	3	3	10
T	4	4	4	4	4	4	4	4	4	4	10
Z	5	5	5	5	5	5	5	5	5	5	10
E	6	6	6	6	6	6	6	6	6	6	10
R	7	7	7	7	7	7	7	7	7	7	9
L	8	8	8	8	8	8	8	8	8	8	8
A	9	9	8	8	8	8	7	6	5	4	7
N	10	9	9	9	9	9	8	7	6	5	6
D	11	10	10	10	10	10	9	8	7	6	5

Computational Cost is Mathematically Proven

arXiv.org > cs > arXiv:1412.0348

Search...

Help | Advanced

Computer Science > Computational Complexity

[Submitted on 1 Dec 2014 (v1), last revised 15 Aug 2017 (this version, v4)]

Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false)

Arturs Backurs, Piotr Indyk

The edit distance (a.k.a. the Levenshtein distance) between two strings is defined as the minimum number of insertions, deletions or substitutions of symbols needed to transform one string into another. The problem of computing the edit distance between two strings is a classical computational task, with a well-known algorithm based on dynamic programming. Unfortunately, all known algorithms for this problem run in nearly quadratic time.

In this paper we provide evidence that the near-quadratic running time bounds known for the problem of computing edit distance might be tight. Specifically, we show that, if the edit distance can be computed in time $O(n^{2-\delta})$ for some constant $\delta > 0$, then the satisfiability of conjunctive normal form formulas with N variables and M clauses can be solved in time $M^{O(1)}2^{(1-\epsilon)N}$ for a constant $\epsilon > 0$. The latter result would violate the Strong Exponential Time Hypothesis, which postulates that such algorithms do not exist.

Read Mapping Techniques in 111 Pages

In-depth analysis of 107 read mappers (1988-2020)

Mohammed Alser, Jeremy Rotman, Dhrithi Deshpande, Kodi Taraszka, Huwenbo Shi, Pelin Icer Baykal, Harry Taegyun Yang, Victor Xue, Sergey Knyazev, Benjamin D. Singer, Brunilda Balliu, David Koslicki, Pavel Skums, Alex Zelikovsky, Can Alkan, Onur Mutlu, Serghei Mangul

"[Technology dictates algorithms: Recent developments in read alignment](#)"

Genome Biology, 2021

[[Source code](#)]

Alser *et al.* *Genome Biology* (2021) 22:249
<https://doi.org/10.1186/s13059-021-02443-7>

Genome Biology

REVIEW

Open Access



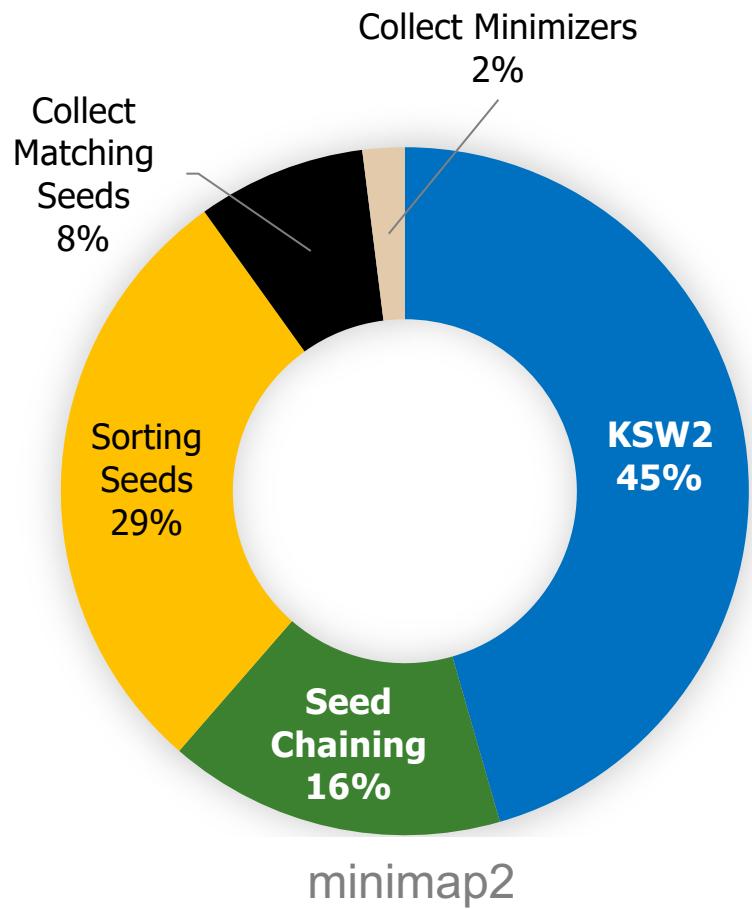
Technology dictates algorithms: recent developments in read alignment

Mohammed Alser^{1,2,3†}, Jeremy Rotman^{4†}, Dhrithi Deshpande⁵, Kodi Taraszka⁴, Huwenbo Shi^{6,7}, Pelin Icer Baykal⁸, Harry Taegyun Yang^{4,9}, Victor Xue⁴, Sergey Knyazev⁸, Benjamin D. Singer^{10,11,12}, Brunilda Balliu¹³, David Koslicki^{14,15,16}, Pavel Skums⁸, Alex Zelikovsky^{8,17}, Can Alkan^{2,18}, Onur Mutlu^{1,2,3†} and Serghei Mangul^{5*†}

Read Mapping Execution Time (Modern)

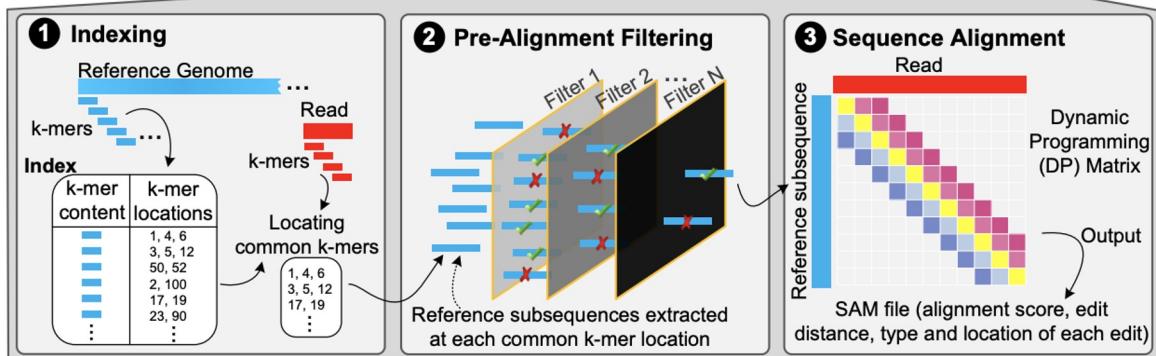
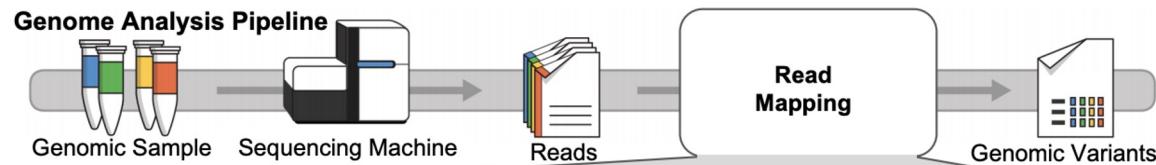
>60%

of the read mapper's execution time is spent in sequence alignment



ONT FASTQ size: 103MB (151 reads), Mean length: 356,403 bp, std: 173,168 bp, longest length: 817,917 bp

Accelerating Read Mapping



Accelerating Indexing

Reducing the number of seeds

Reducing data movement during indexing

Accelerating Pre-Alignment Filtering

q-gram filtering

Pigeonhole principle

Base counting

Sparse DP

Accelerating Alignment

Accurate alignment accelerators

Heuristic-based alignment accelerators

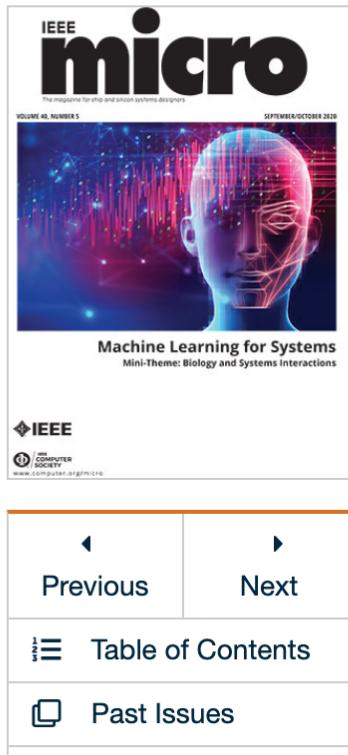
Alser+, "[Accelerating Genome Analysis: A Primer on an Ongoing Journey](#)", IEEE Micro, 2020.

Detailed Analysis of Tackling the Bottleneck

Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu

“Accelerating Genome Analysis: A Primer on an Ongoing Journey”

IEEE Micro, August 2020.



[Home](#) / [Magazines](#) / [IEEE Micro](#) / [2020.05](#)

IEEE Micro

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Sept.-Oct. 2020, pp. 65-75, vol. 40

DOI Bookmark: [10.1109/MM.2020.3013728](https://doi.org/10.1109/MM.2020.3013728)

Authors

[Mohammed Alser](#), ETH Zürich

[Zulal Bingol](#), Bilkent University

[Damla Senol Cali](#), Carnegie Mellon University

[Jeremie Kim](#), ETH Zurich and Carnegie Mellon University

[Saugata Ghose](#), University of Illinois at Urbana-Champaign and Carnegie Mellon University

[Can Alkan](#), Bilkent University

[Onur Mutlu](#), ETH Zurich, Carnegie Mellon University, and Bilkent University

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

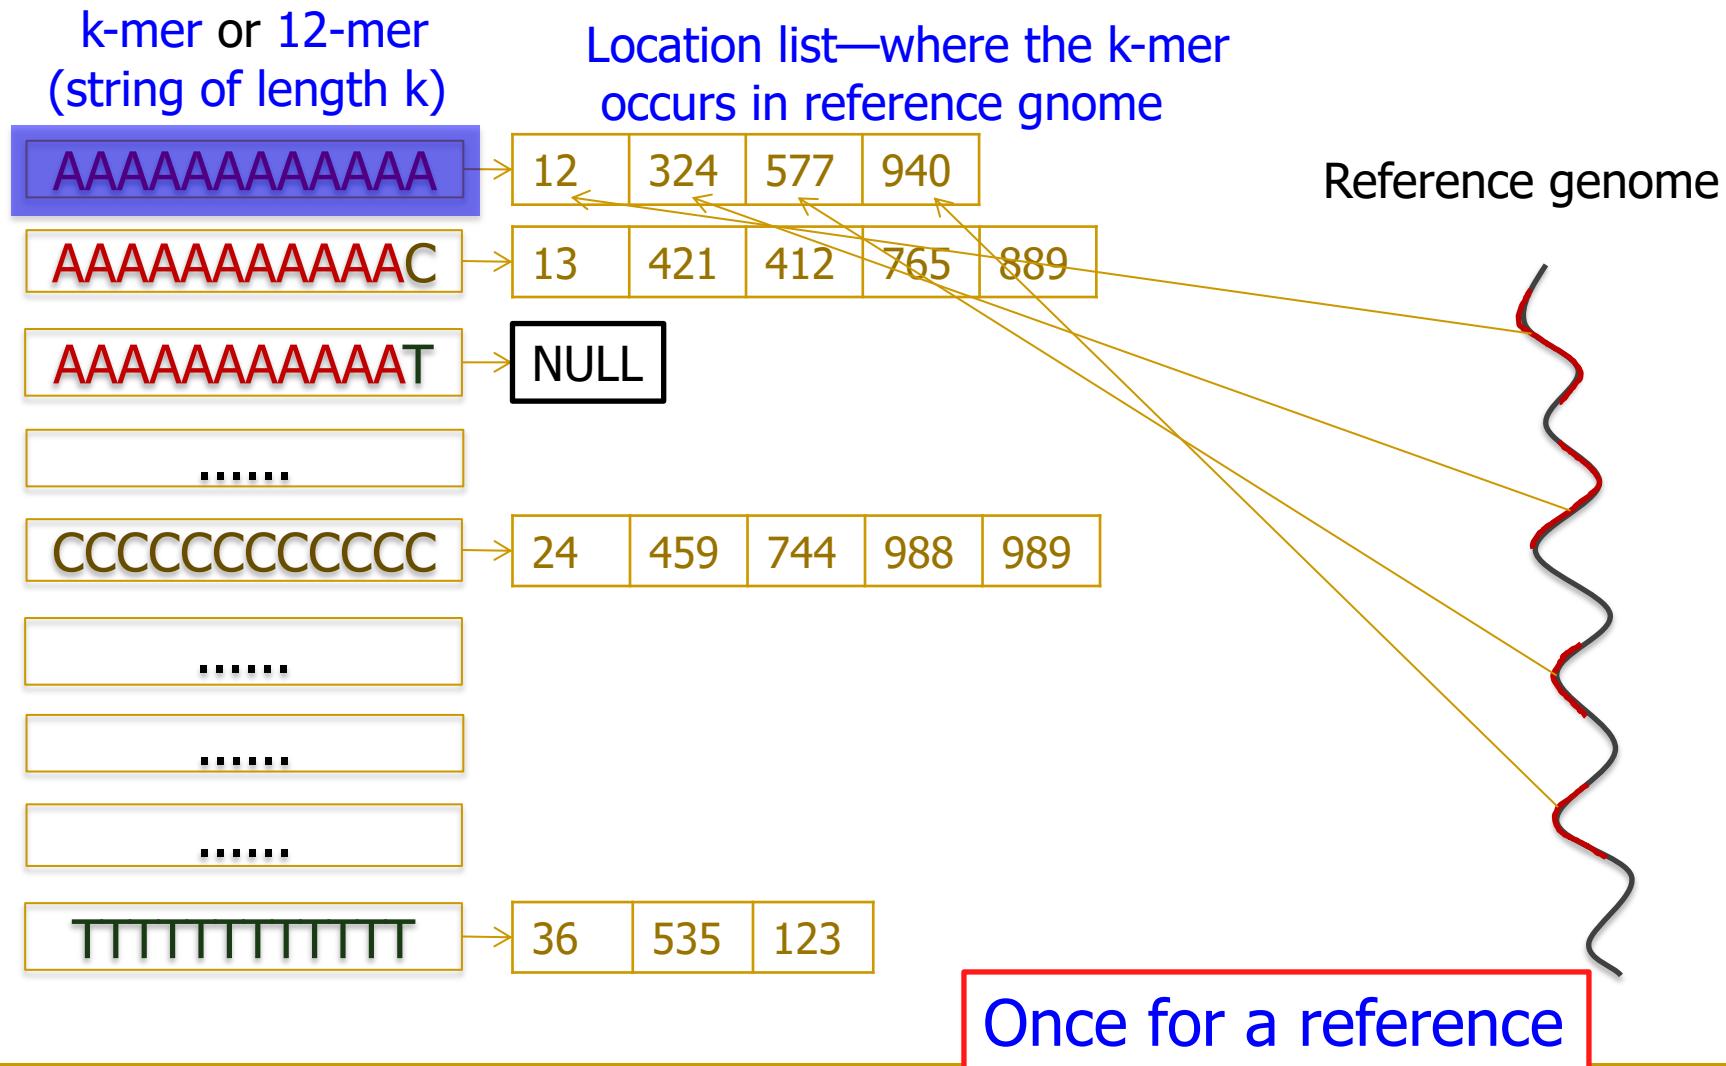
Read Mapping Algorithms: Two Styles

- Hash based seed-and-extend (hash table, suffix array, suffix tree)
 - Index the “k-mers” in the genome into a hash table (pre-processing)
 - When searching a read, find the location of a k-mer in the read; then extend through alignment
 - More sensitive (can find all mapping locations), but slow
 - Requires large memory; this can be reduced with cost to run time
- Burrows-Wheeler Transform & Ferragina-Manzini Index based aligners
 - BWT is a compression method used to compress the genome index
 - Perfect matches can be found very quickly, memory lookup costs increase for imperfect matches
 - Reduced sensitivity

Hash Table Based Read Mappers

- Key Idea
 - Preprocess the reference into a *Hash Table*
 - Use *Hash Table* to map reads

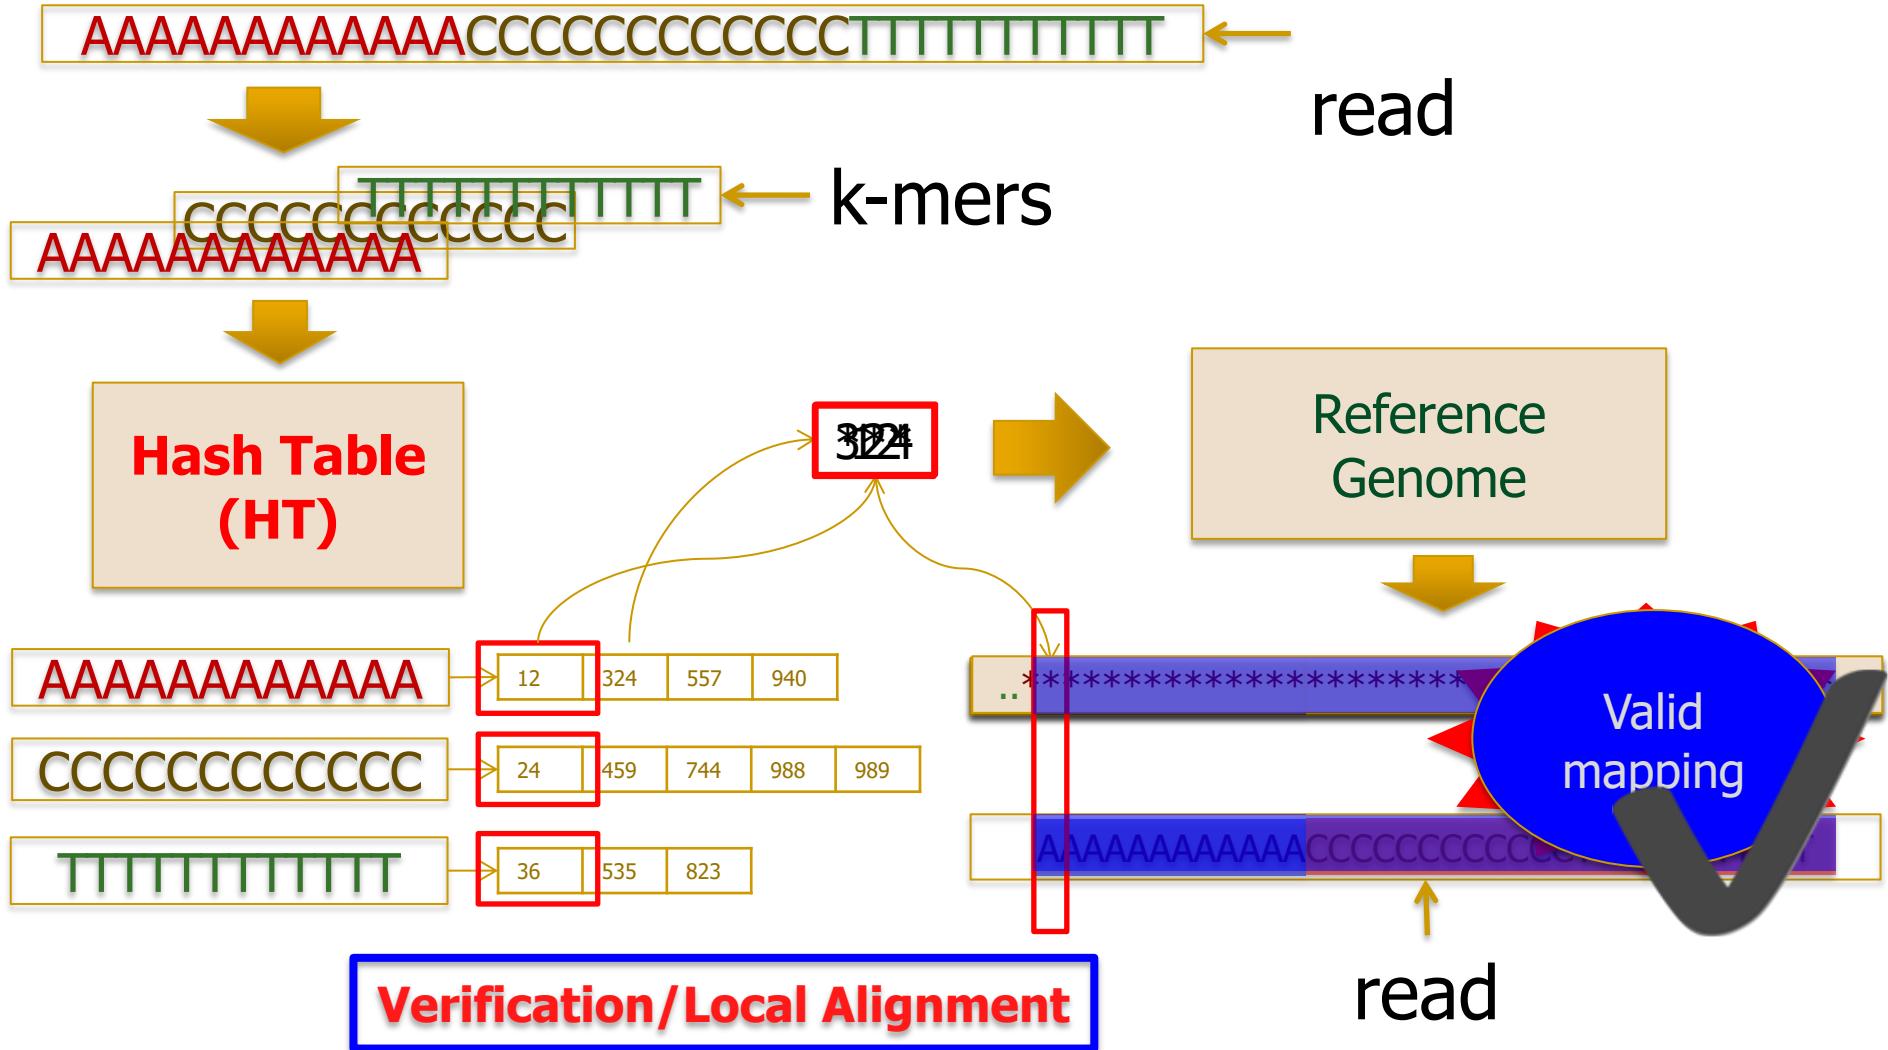
Hash Table-Based Mappers [Alkan+ Nature Gen'09]



Hash Table Based Read Mappers

- Key Idea
 - Preprocess the reference into a *Hash Table*
 - Use *Hash Table* to map reads

Hash Table-Based Mappers [Alkan+ Nature Gen'09]



Our First Step: Comprehensive Mapping

- + Guaranteed to find *all* mappings → sensitive
- + Can tolerate up to e errors



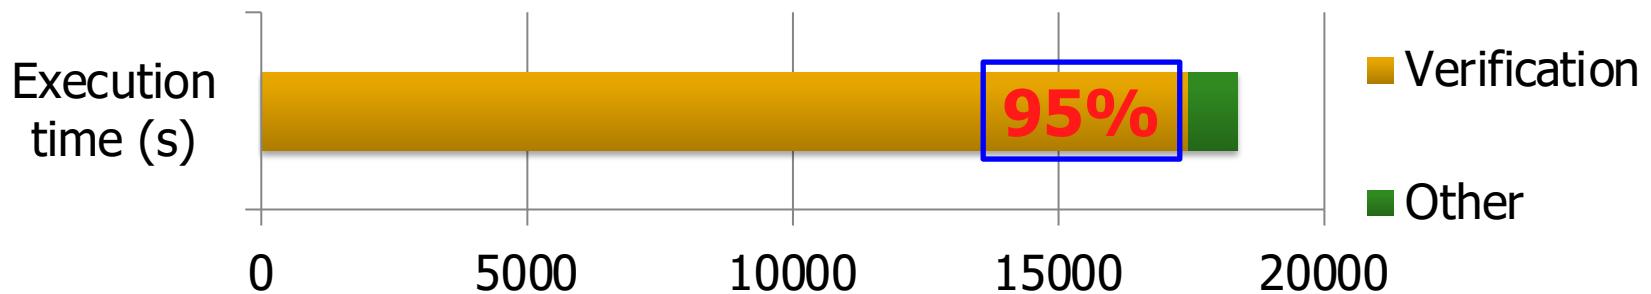
<http://mrfast.sourceforge.net/>

Personalized copy number and segmental duplication maps using next-generation sequencing

Can Alkan^{1,2}, Jeffrey M Kidd¹, Tomas Marques-Bonet^{1,3}, Gozde Aksay¹, Francesca Antonacci¹, Fereydoun Hormozdiari⁴, Jacob O Kitzman¹, Carl Baker¹, Maika Malig¹, Onur Mutlu⁵, S Cenk Sahinalp⁴, Richard A Gibbs⁶ & Evan E Eichler^{1,2}

Problem and Goal

- Poor performance of existing read mappers: Very slow
 - Verification/alignment takes too long to execute
 - Verification requires a memory access for reference genome + many base-pair-wise comparisons between the reference and the read (edit distance computation)



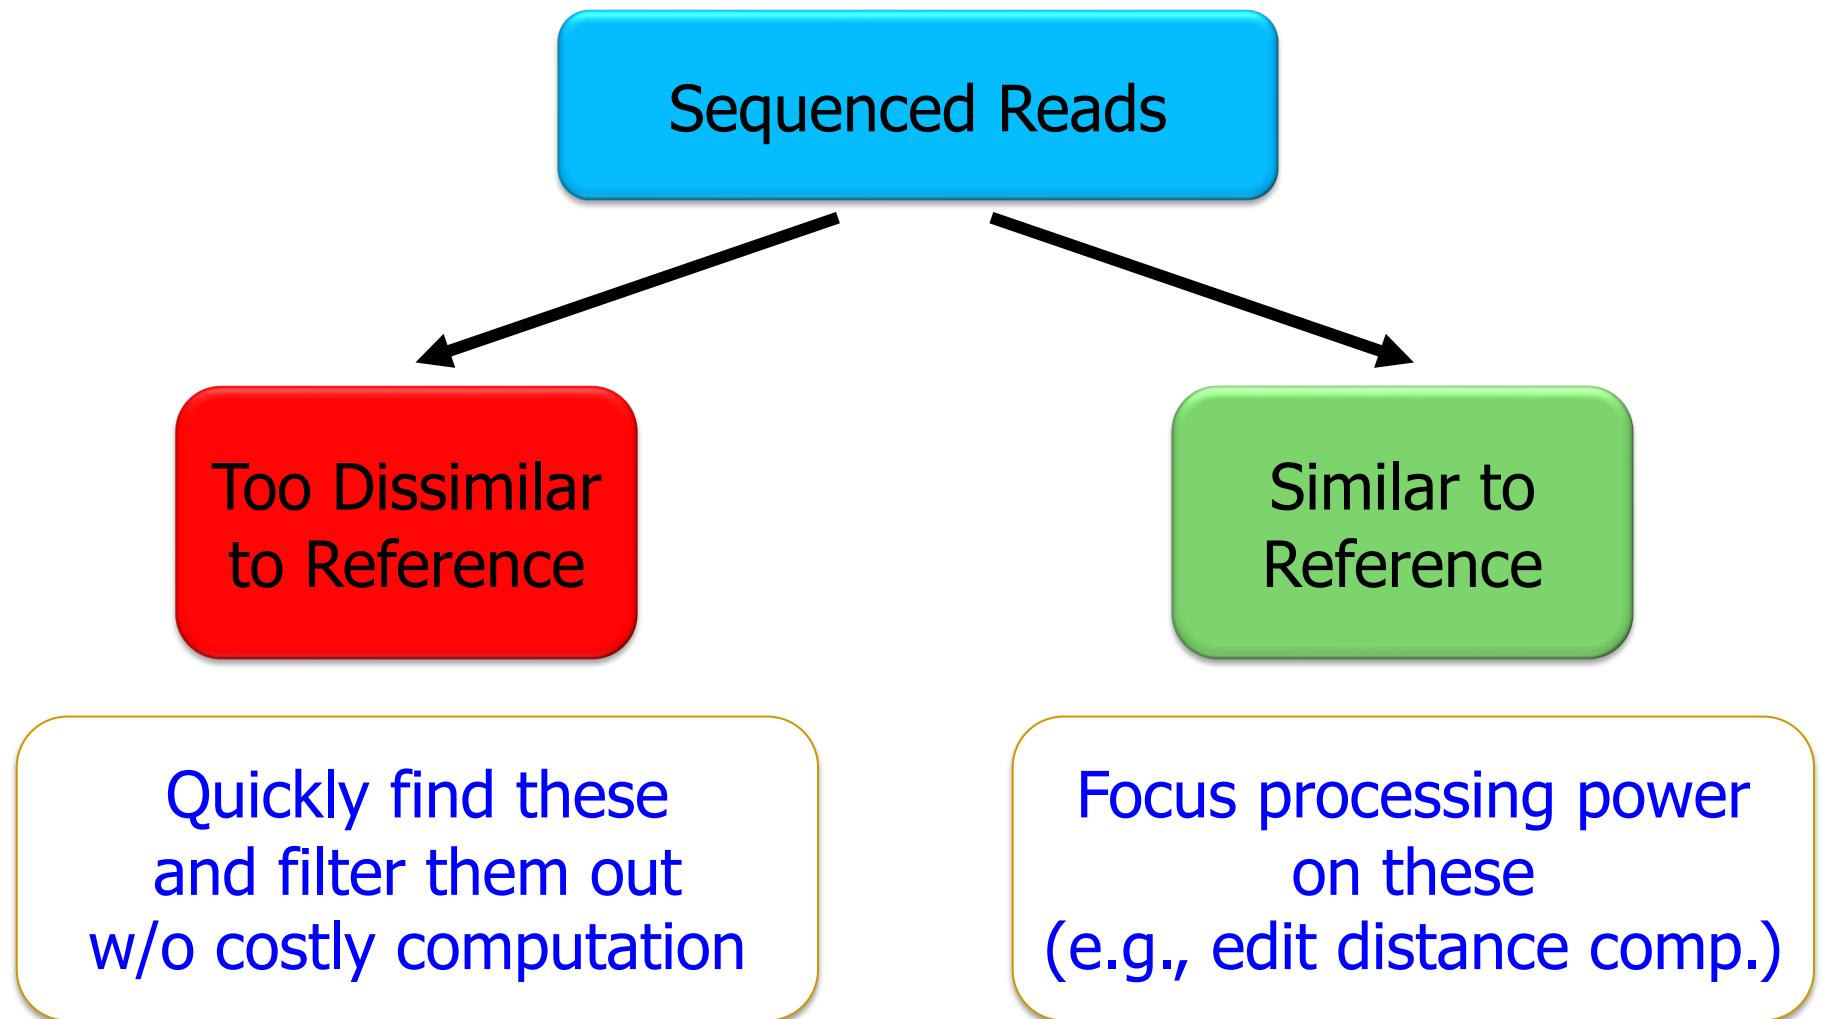
- Goal: Speed up the mapper by reducing the cost of verification

Overarching Key Idea

Filter fast before you align

**Minimize costly
edit distance computations**
("approximate string comparisons")

Overarching Key Idea



Accelerating Genome Analysis: Overview

- Mohammed Alser, Zulal Bingol, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, and Onur Mutlu,

"Accelerating Genome Analysis: A Primer on an Ongoing Journey"

IEEE Micro (IEEE MICRO), Vol. 40, No. 5, pages 65-75, September/October 2020.

[[Slides \(pptx\)\(pdf\)](#)]

[[Talk Video \(1 hour 2 minutes\)](#)]

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Mohammed Alser
ETH Zürich

Zülal Bingöl
Bilkent University

Damla Senol Cali
Carnegie Mellon University

Jeremie Kim
ETH Zurich and Carnegie Mellon University

Saugata Ghose
University of Illinois at Urbana–Champaign and
Carnegie Mellon University

Can Alkan
Bilkent University

Onur Mutlu
ETH Zurich, Carnegie Mellon University, and
Bilkent University

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

Our First Filter: Pure Software Approach

- Download the source code and try for yourself
 - [Download link to FastHASH](#)

Xin *et al.* BMC Genomics 2013, **14**(Suppl 1):S13
http://www.biomedcentral.com/1471-2164/14/S1/S13



PROCEEDINGS

Open Access

Accelerating read mapping with FastHASH

Hongyi Xin¹, Donghyuk Lee¹, Farhad Hormozdiari², Samihan Yedkar¹, Onur Mutlu^{1*}, Can Alkan^{3*}

From The Eleventh Asia Pacific Bioinformatics Conference (APBC 2013)
Vancouver, Canada. 21-24 January 2013

Reducing the Cost of Verification

- Most verification (edit distance computation) calculations are unnecessary
 - 1 out of 1000 potential locations passes the verification process
- We can get rid of unnecessary verification calculations by
 - *Detecting and rejecting *early* invalid mappings (filtering)*
 - *Reducing the *number* of potential mappings to examine*

Key Observations [Xin+, BMC Genomics 2013]

- Observation 1
 - Adjacent k-mers in the read should also be adjacent in the reference genome
 - Read mapper can quickly reject mappings that do **not** satisfy this property

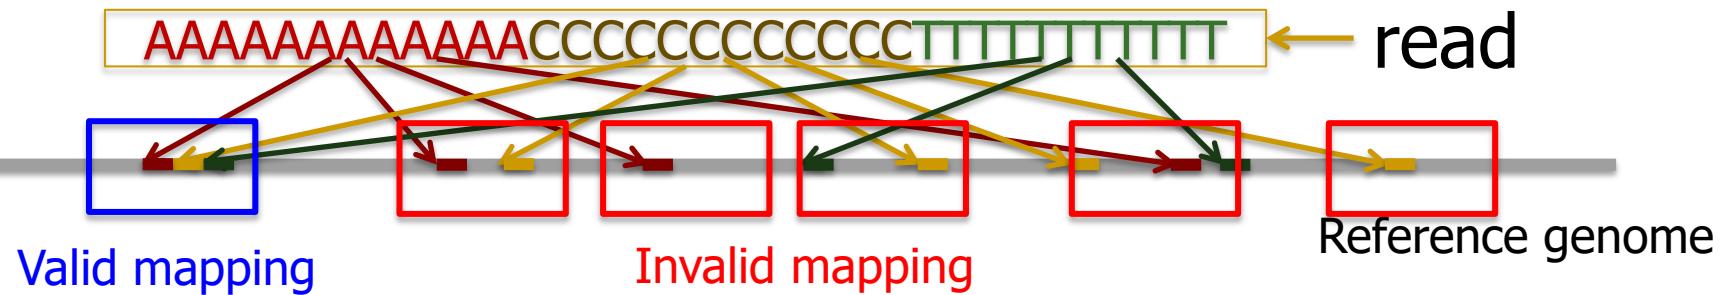
- Observation 2
 - Some k-mers are **cheaper** to verify than others because they have shorter location lists (they occur less frequently in the reference genome)
 - Mapper needs to examine only $e+1$ k-mers' locations to tolerate e errors
 - Read mapper can choose the cheapest $e+1$ k-mers and verify their locations

FastHASH Mechanisms [Xin+, BMC Genomics 2013]

- **Adjacency Filtering (AF):** Rejects obviously invalid mapping locations at early stage to avoid unnecessary verifications
- **Cheap K-mer Selection (CKS):** Reduces the absolute number of potential mapping locations to verify

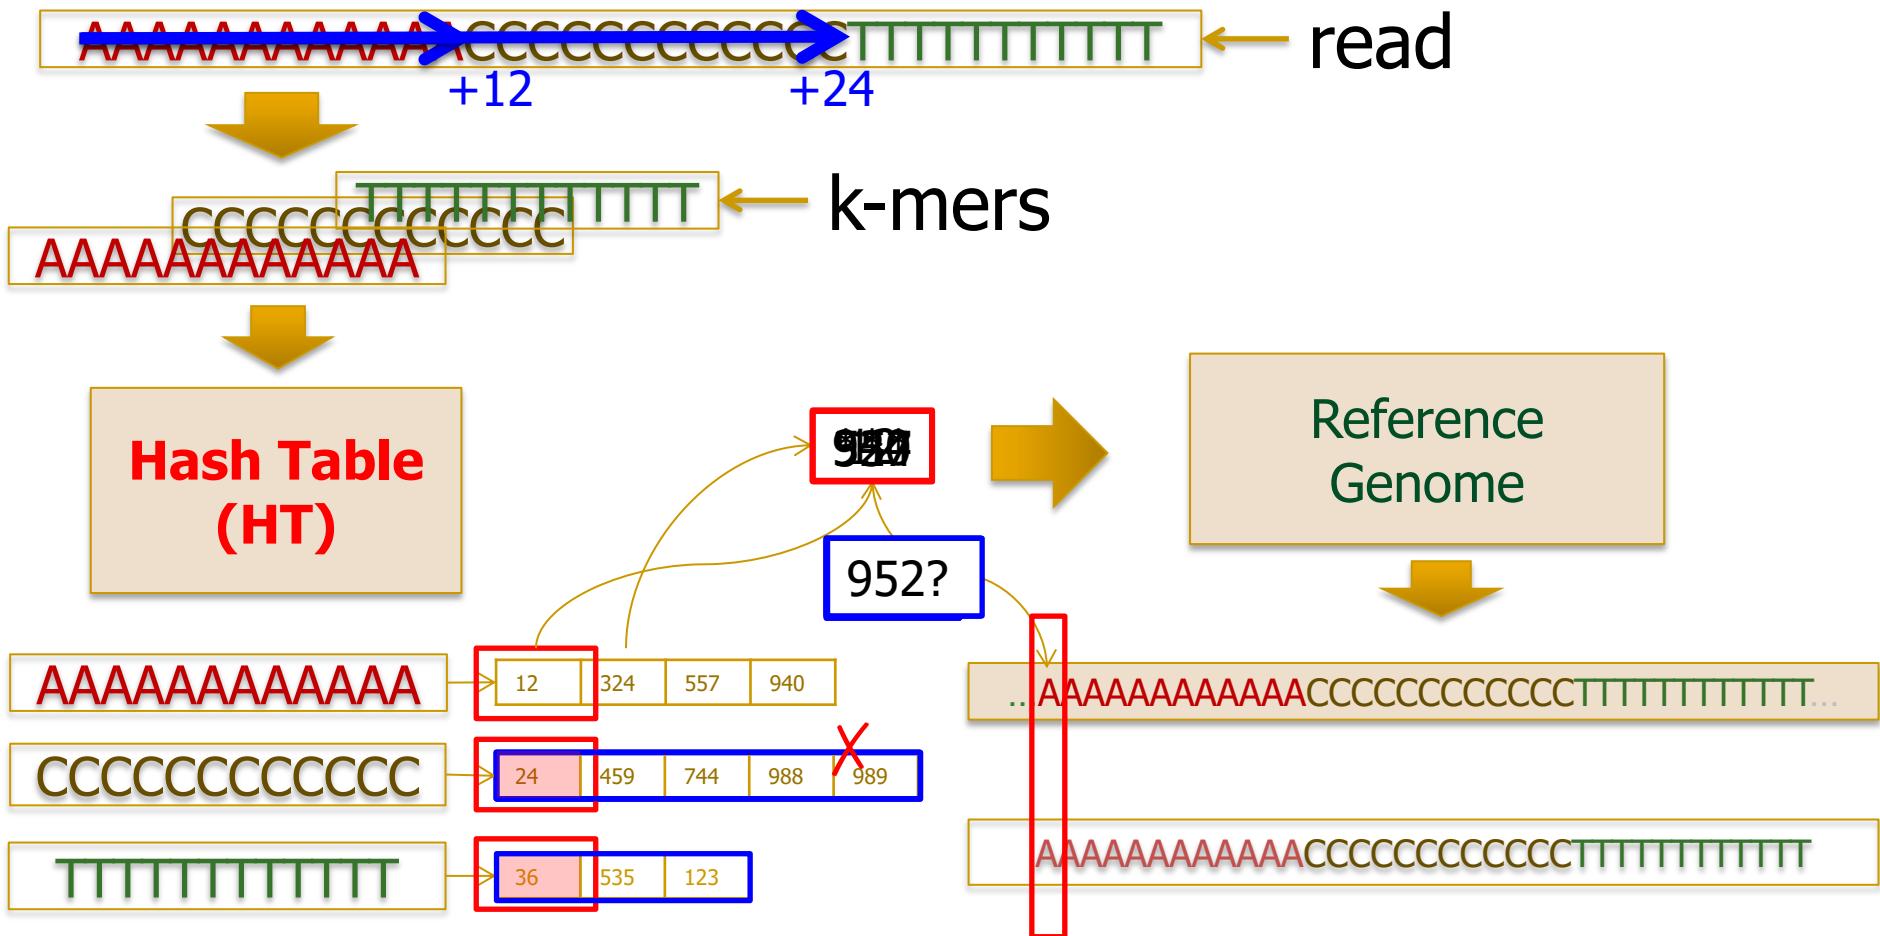
Adjacency Filtering (AF)

- **Goal:** detect and filter out invalid mappings at early stage
- **Key Insight:** For a valid mapping, adjacent k-mers in the read are also adjacent in the reference genome



- **Key Idea:** search for adjacent locations in the k-mers' location lists (in the index)
 - If more than e k-mers fail \rightarrow there must be more than e errors \rightarrow invalid mapping

Adjacency Filtering (AF)



FastHASH Mechanisms [Xin+, BMC Genomics 2013]

- **Adjacency Filtering (AF):** Rejects obviously invalid mapping locations at early stage to avoid unnecessary verifications
- **Cheap K-mer Selection (CKS):** Reduces the absolute number of potential mapping locations to verify

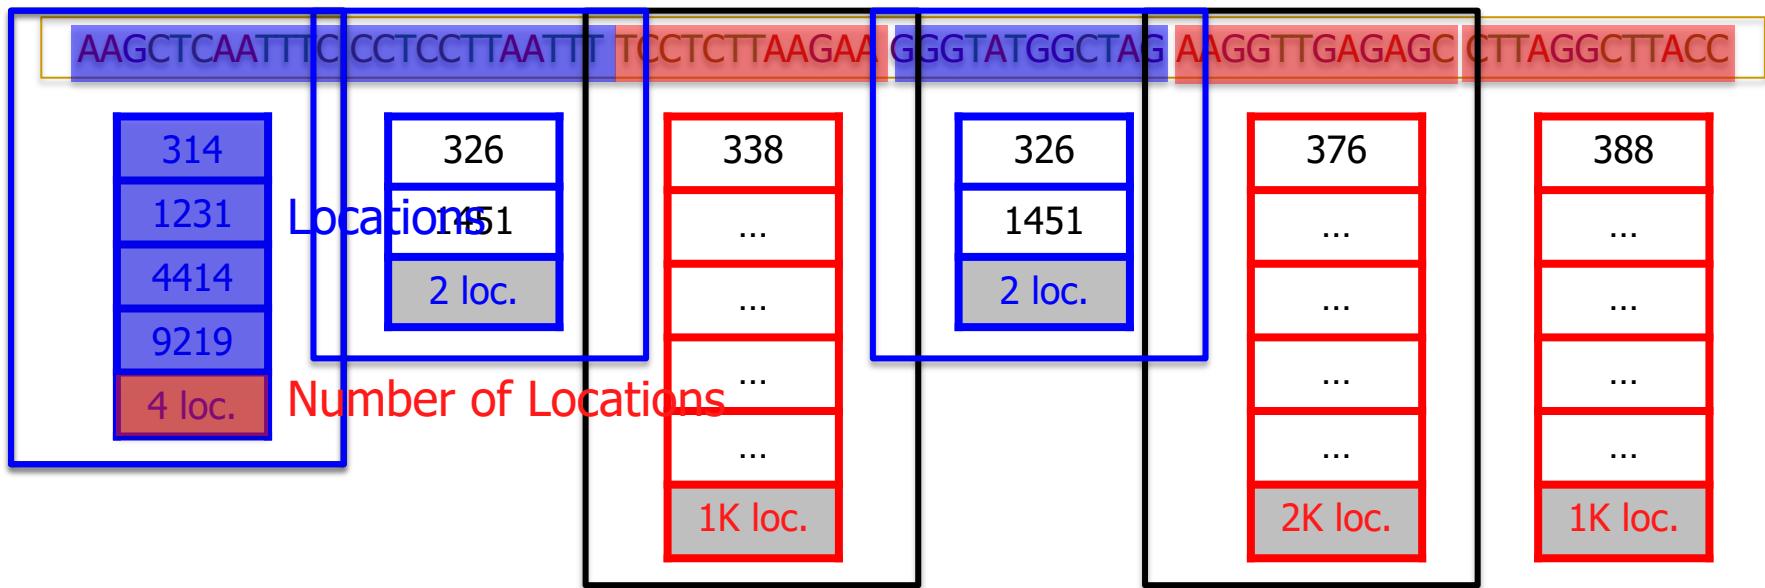
Cheap K-mer Selection (CKS)

- **Goal:** Reduce the number of potential mappings to examine
- **Key insight:**
 - K-mers have different **cost** to examine: Some k-mers are *cheaper* as they have fewer locations than others (occur less frequently in reference genome)
- **Key idea:**
 - Sort the k-mers based on their number of locations
 - **Select the k-mers with the fewest number locations to verify**

Cheap K-mer Selection

- $e=2$ (examine 3 k-mers)

read



Expensive 3kmers

Previous work needs
to verify:

3004 locations

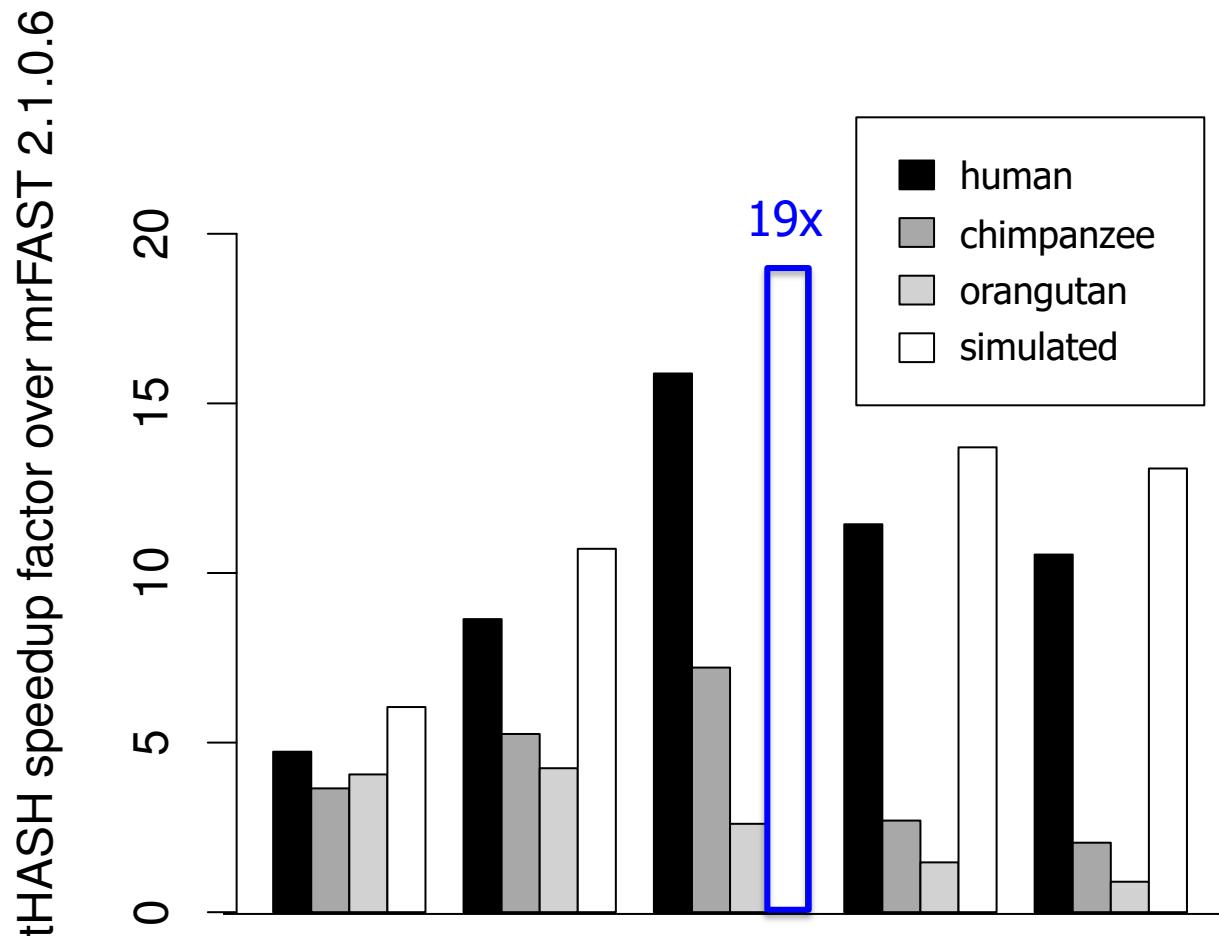
FastHASH verifies only:

8 locations

Methodology

- Implemented **FastHASH** on top of state-of-the-art mapper: **mrFAST**
 - New version **mrFAST-2.5.0.0** over mrFAST-2.1.0.6
- Tested with real read sets generated from Illumina platform
 - 1M reads of a human (160 base pairs)
 - 500K reads of a chimpanzee (101 base pairs)
 - 500K reads of a orangutan (70 base pairs)
- Tested with simulated reads generated from reference genome
 - 1M simulated reads of human (180 base pairs)
- Evaluation system
 - Intel Core i7 Sandy Bridge machine
 - 16 GB of main memory

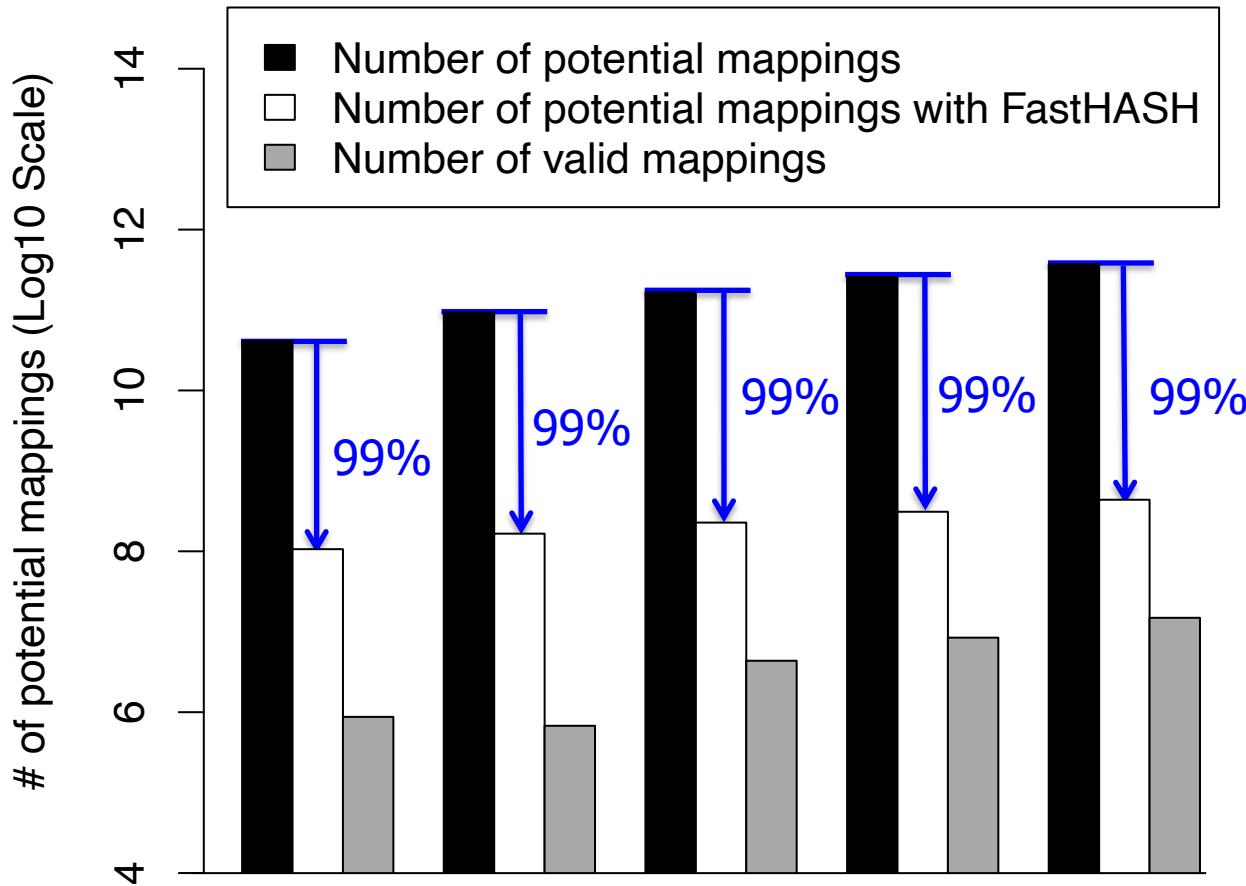
FastHASH Speedup: Entire Read Mapper



With FastHASH, new mrFAST obtains up to 19x speedup over previous version, without losing valid mappings

Analysis

■ Reduction of potential mappings with FastHASH



FastHASH filters out over 99% of the potential mappings without sacrificing any valid mappings

FastHASH Summary & Conclusion

- Problem: Existing read mappers perform poorly, especially in the presence of errors
 - Observation: Most of the verification (edit distance) calculations are unnecessary → filter them out
 - Key Idea: Exploit the structure of the genome to
 - Reject invalid mappings early (Adjacency Filtering)
 - Reduce the number of possible mappings to examine (Cheap K-mer Selection)
 - Key Result: FastHASH obtains up to 19x speedup over the state-of-the-art mapper without losing valid mappings
-

More on FastHASH

- Download source code and try for yourself
 - [Download link to FastHASH](#)

Xin *et al.* BMC Genomics 2013, **14**(Suppl 1):S13
<http://www.biomedcentral.com/1471-2164/14/S1/S13>



PROCEEDINGS

Open Access

Accelerating read mapping with FastHASH

Hongyi Xin¹, Donghyuk Lee¹, Farhad Hormozdiari², Samihan Yedkar¹, Onur Mutlu^{1*}, Can Alkan^{3*}

From The Eleventh Asia Pacific Bioinformatics Conference (APBC 2013)
Vancouver, Canada. 21-24 January 2013

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
 - Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
 - Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
 - Future Opportunities: New Technologies & Applications
-

Shifted Hamming Distance: SIMD Acceleration

<https://github.com/CMU-SAFARI/Shifted-Hamming-Distance>

Bioinformatics, 31(10), 2015, 1553–1560

doi: 10.1093/bioinformatics/btu856

Advance Access Publication Date: 10 January 2015

Original Paper

OXFORD

Sequence analysis

Shifted Hamming distance: a fast and accurate SIMD-friendly filter to accelerate alignment verification in read mapping

Hongyi Xin^{1,*}, John Greth², John Emmons², Gennady Pekhimenko¹,
Carl Kingsford³, Can Alkan^{4,*} and Onur Mutlu^{2,*}

Xin+, [**"Shifted Hamming Distance: A Fast and Accurate SIMD-friendly Filter to Accelerate Alignment Verification in Read Mapping"**](#), Bioinformatics 2015.

Shifted Hamming Distance

■ Key observation:

- If two strings differ by E edits, then every bp match can be aligned in at most $2E$ shifts (of one of the strings).
 - Insight: Shifting a string by one “corrects” for one “error”

■ Key idea:

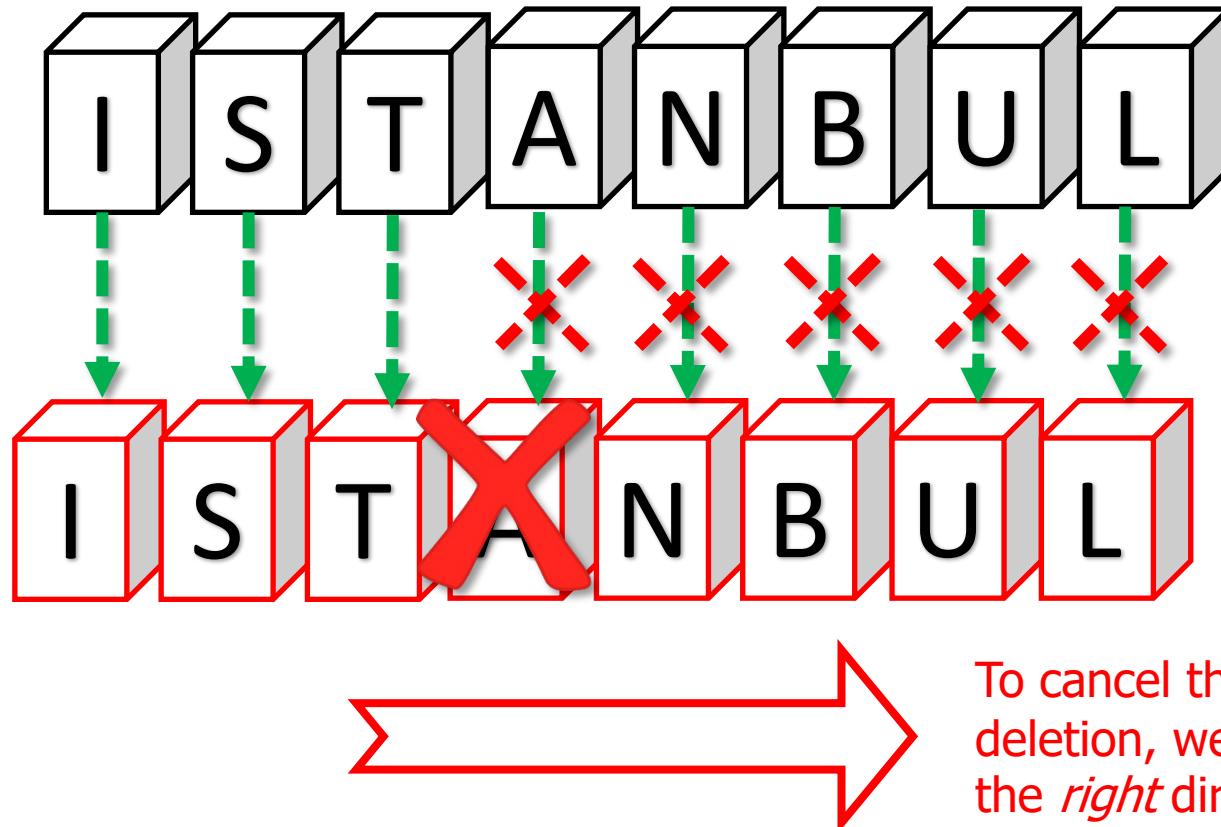
- Compute “Shifted Hamming Distance”: **AND of 2E Hamming Distances of two strings**, to filter out invalid mappings
 - Uses bit-parallel operations that nicely map to SIMD instructions

■ Key result:

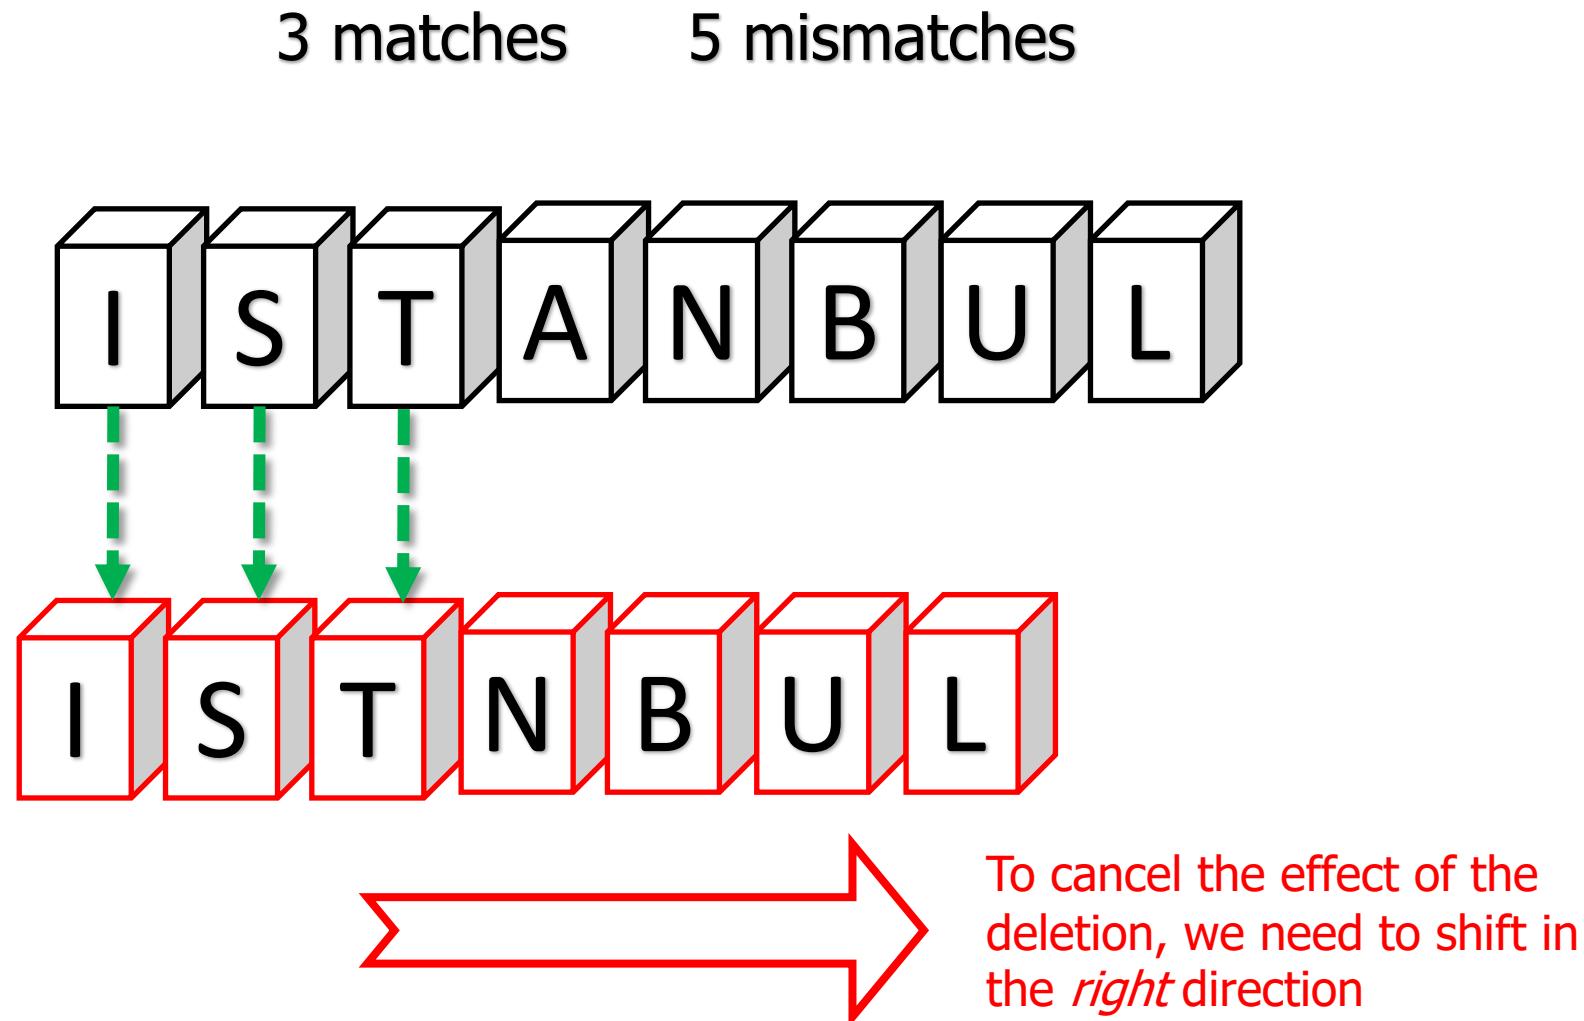
- SHD is 3x faster than SeqAn (the best implementation of Gene Myers’ bit-vector algorithm), with only a 7% false positive rate
 - The **fastest CPU-based filtering (pre-alignment) mechanism**
-

Hamming Distance ($\Sigma \oplus$)

3 matches 5 mismatches
Edit = 1 Deletion



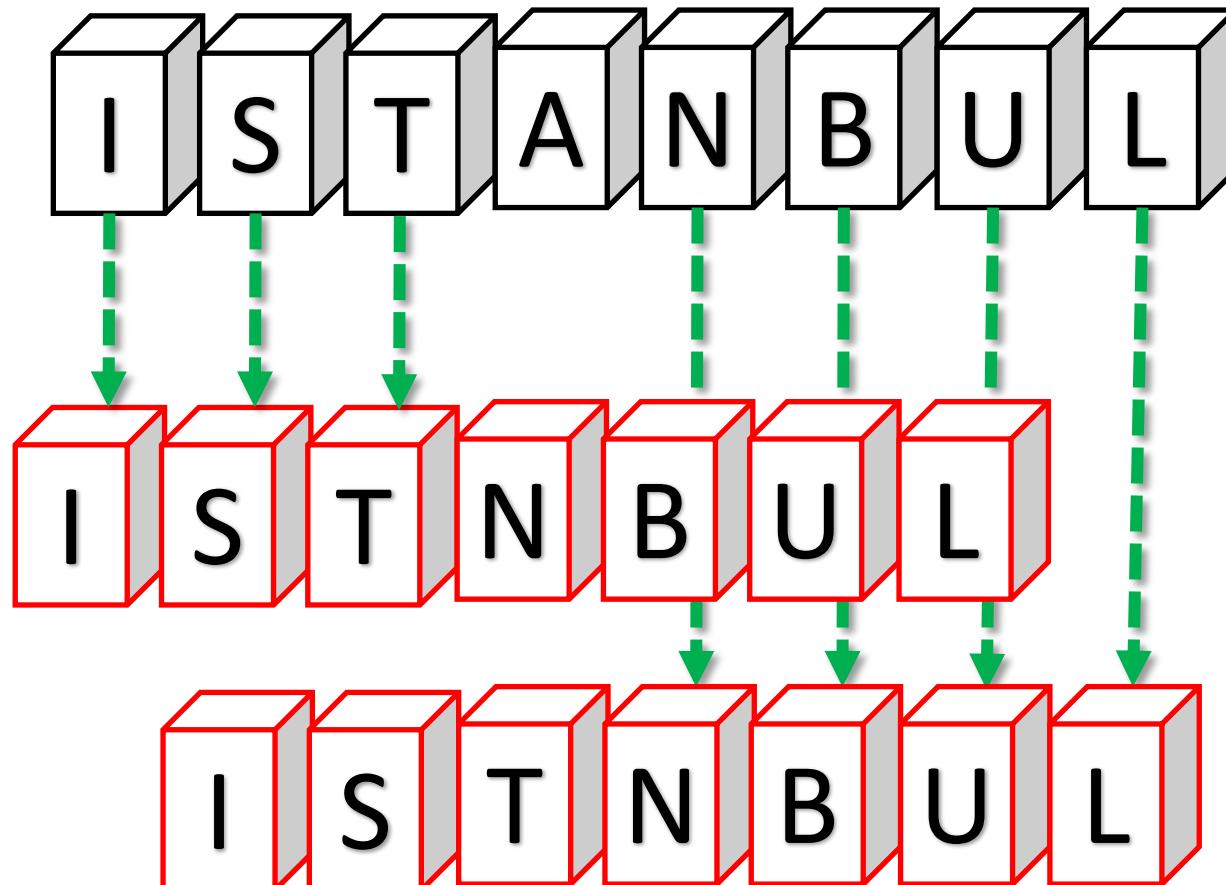
Insight: Shifting a String Helps Similarity Search



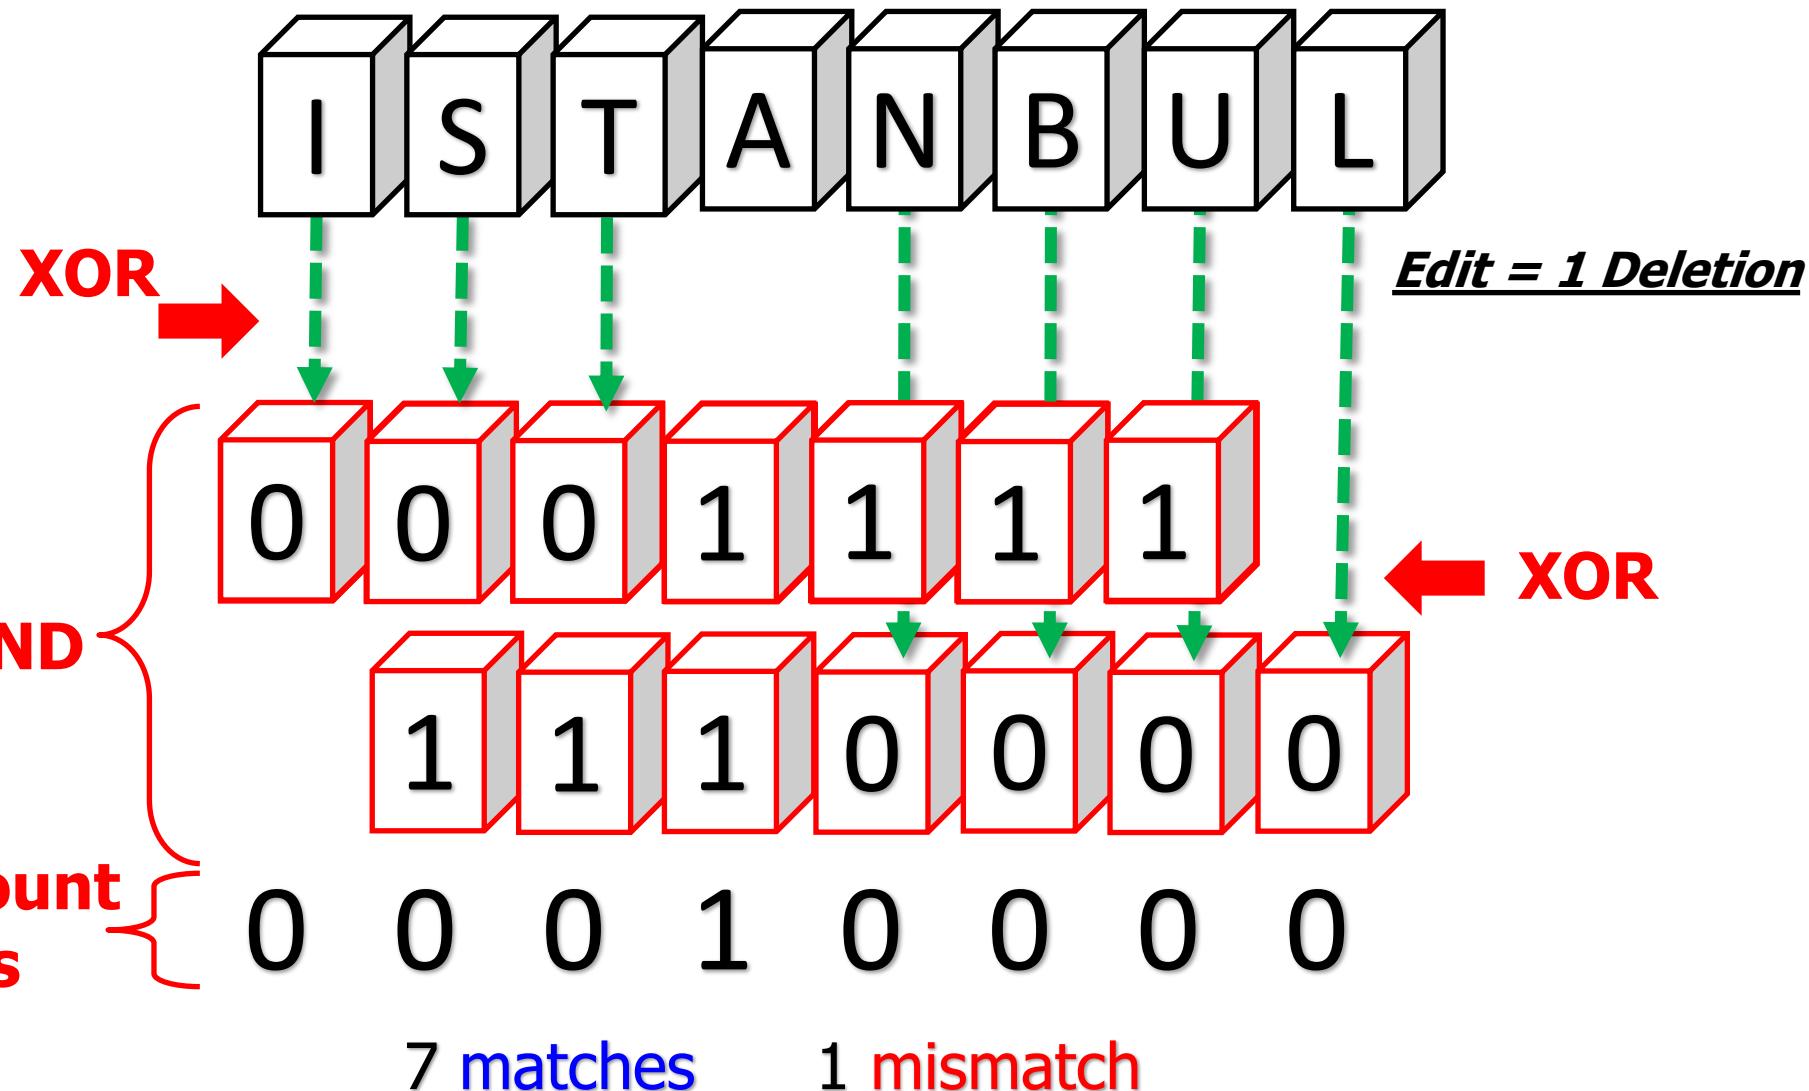
Insight: Shifting a String Helps Similarity Search

7 matches

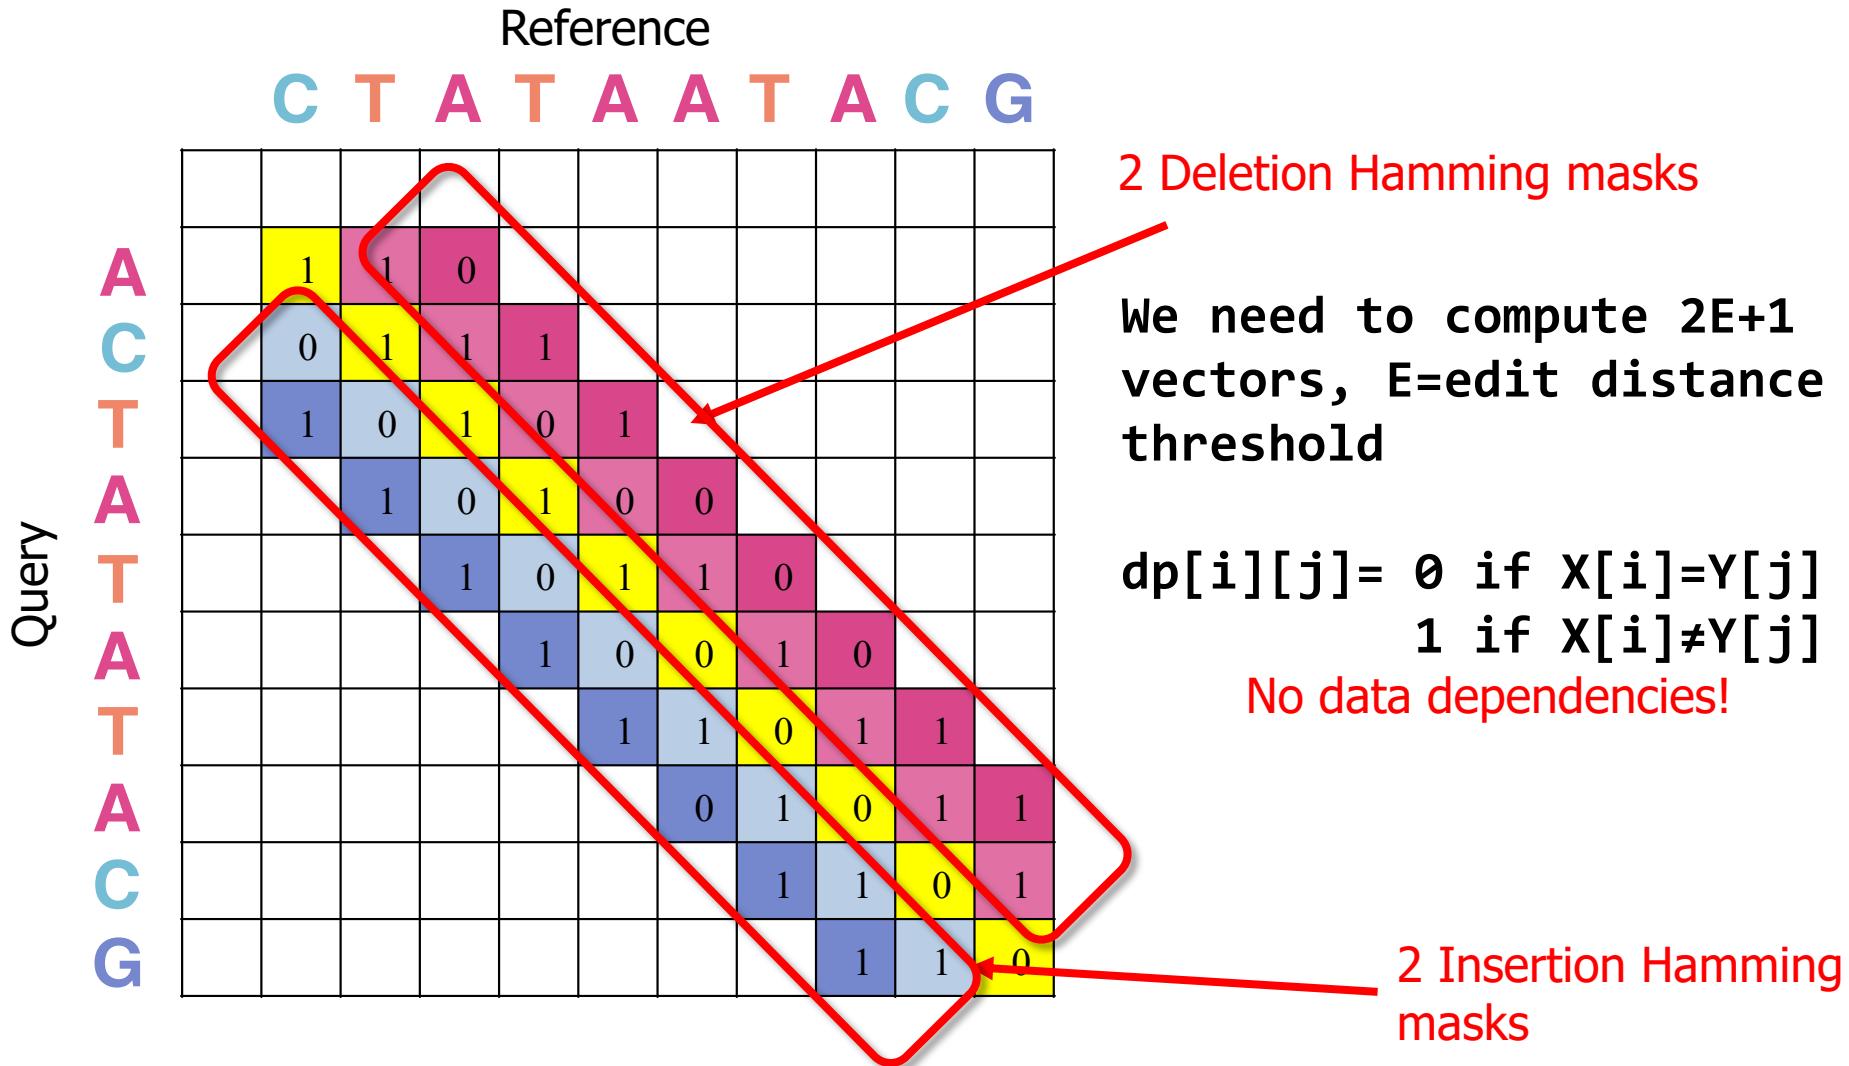
1 mismatch



Shifted Hamming Distance



Highly Parallel Matrix Computation



Key Idea of SHD Filtering



Amend random zeros:
101 → 111 & 1001 → 1111

AND all masks,
ACCEPT iff number of '1' \leq Threshold

Query :GAGAGAGATATTAGTGGCAGCACTACAAACACAAAAGAGGACCAACTACGTGTCTAAAAGGGGGAACATTGTTGGGCCGGA
Reference :GAGAGAGATAGTTAGTGGCAGCCACTACAAACACAAAAGAGGACCAACTACGTGTCTAAAAGGGGGAGACATTGTTGGGCCGG

--- Masks after amendment ---

Needleman-Wunsch Alignment : 

Alignment vs. Pre-alignment (Filtering)

Needleman-Wunsch

	C	T	A	T	A	A	T	A	C	G
A	0	-1	-2							
C	-1	1	-1	-2						
T	-2	-2	-2	-1	-2					
A	-2	-3	-2	-1	-2					
T	-3	-3	-2	-1	-2					
A	-4	-3	-2	-1	-2					
T	-4	-3	-2	-2	-2					
A	-4	-3	-2	-2	-3	-2				
T	-4	-3	-2	-2	-3	-3				
A	-4	-3	-2	-2	-3	-3	-2			
C										
G										

$|dp[i][j-1]$ -1 // Inser.

Neighborhood Map

	C	T	A	T	A	A	T	A	C	G
A										
C										
T										
A										
T										
A										
T										
A										
C										
G										

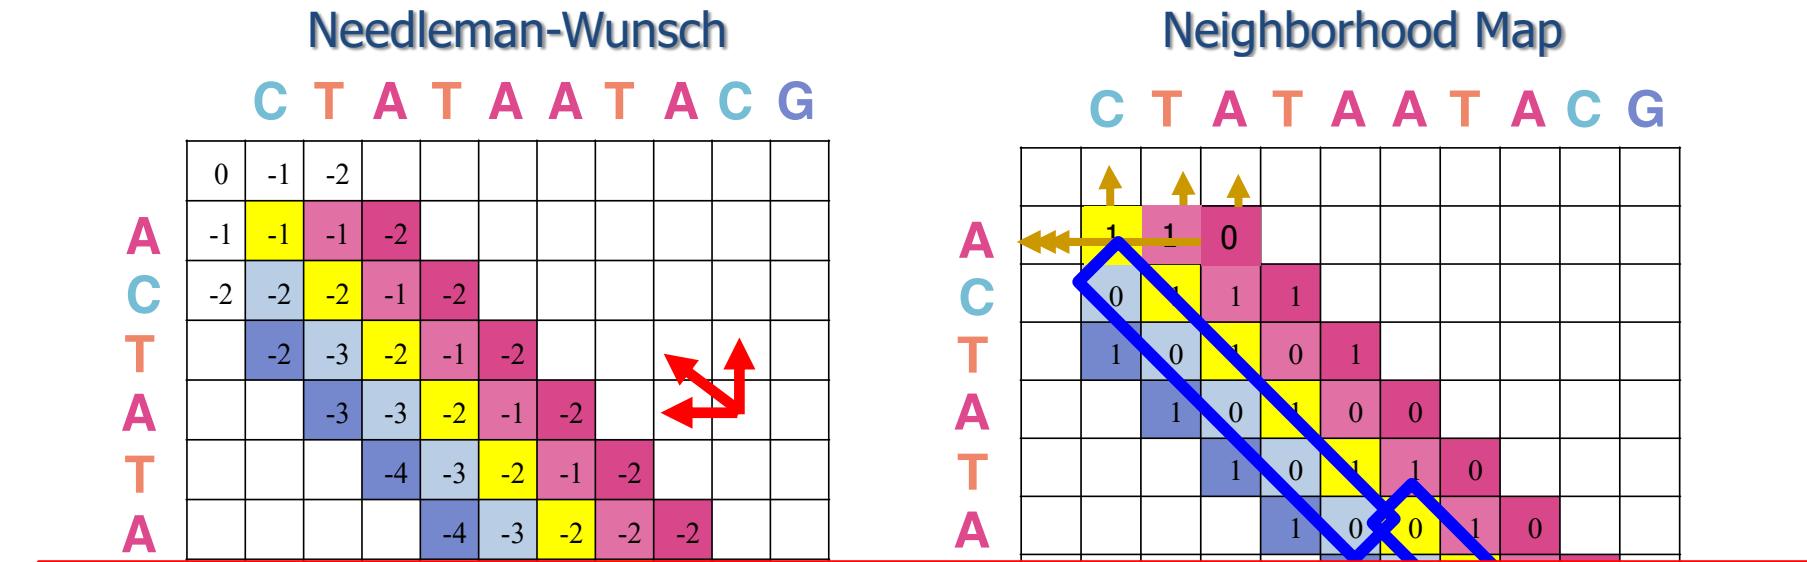
$dp[i][j]=0$ if $x[i]=y[j]$

Our goal is to track the diagonally consecutive matches
in the neighborhood map

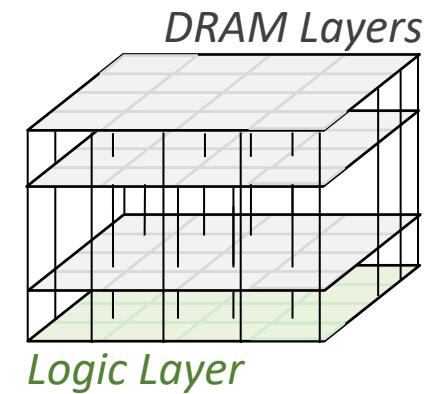
Each cell depends on one or more pre-computed cells!

NO data dependencies!

Alignment Matrix vs. Neighborhood Map



Independent vectors can be processed in parallel using hardware technologies



New Bottleneck: Filtering (Pre-Alignment)

Sequencing generates many reads, each of which potentially mapping to many locations



Filtering (Pre-alignment) eliminates the need to verify/align read to invalid mapping locations



Alignment/verification (costly edit distance computation) is performed **only** on reads that pass the filter

- New bottleneck in read mapping becomes the “filtering (pre-alignment)” step

More on Shifted Hamming Distance

<https://github.com/CMU-SAFARI/Shifted-Hamming-Distance>

Bioinformatics, 31(10), 2015, 1553–1560

doi: 10.1093/bioinformatics/btu856

Advance Access Publication Date: 10 January 2015

Original Paper

OXFORD

Sequence analysis

Shifted Hamming distance: a fast and accurate SIMD-friendly filter to accelerate alignment verification in read mapping

Hongyi Xin^{1,*}, John Greth², John Emmons², Gennady Pekhimenko¹,
Carl Kingsford³, Can Alkan^{4,*} and Onur Mutlu^{2,*}

Xin+, [**"Shifted Hamming Distance: A Fast and Accurate SIMD-friendly Filter to Accelerate Alignment Verification in Read Mapping"**](#), Bioinformatics 2015.

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

Location Filtering (Pre-alignment)

- **Alignment** is **expensive**
 - We need to align millions to billions of reads
- Modern read mappers reduce the time spent on alignment for increased performance. Can be done in two ways:
 1. Optimize the algorithm for alignment
 2. Reduce the number of alignments necessary by **filtering** out mismatches quickly
- Both methods are used by mappers today, but **filtering** has replaced alignment as the bottleneck [Xin+, BMC Genomics 2013]

Location Filtering (Pre-alignment)

- **Alignment** is expensive
 - We need to align millions to billions of reads

- Mappers are responsible for filtering out mismatches quickly

Our goal is to accelerate read mapping by improving the filtering step

out mismatches quickly

- Both methods are used by mappers today, but filtering has replaced alignment as the bottleneck [Xin+, BMC Genomics 2013]

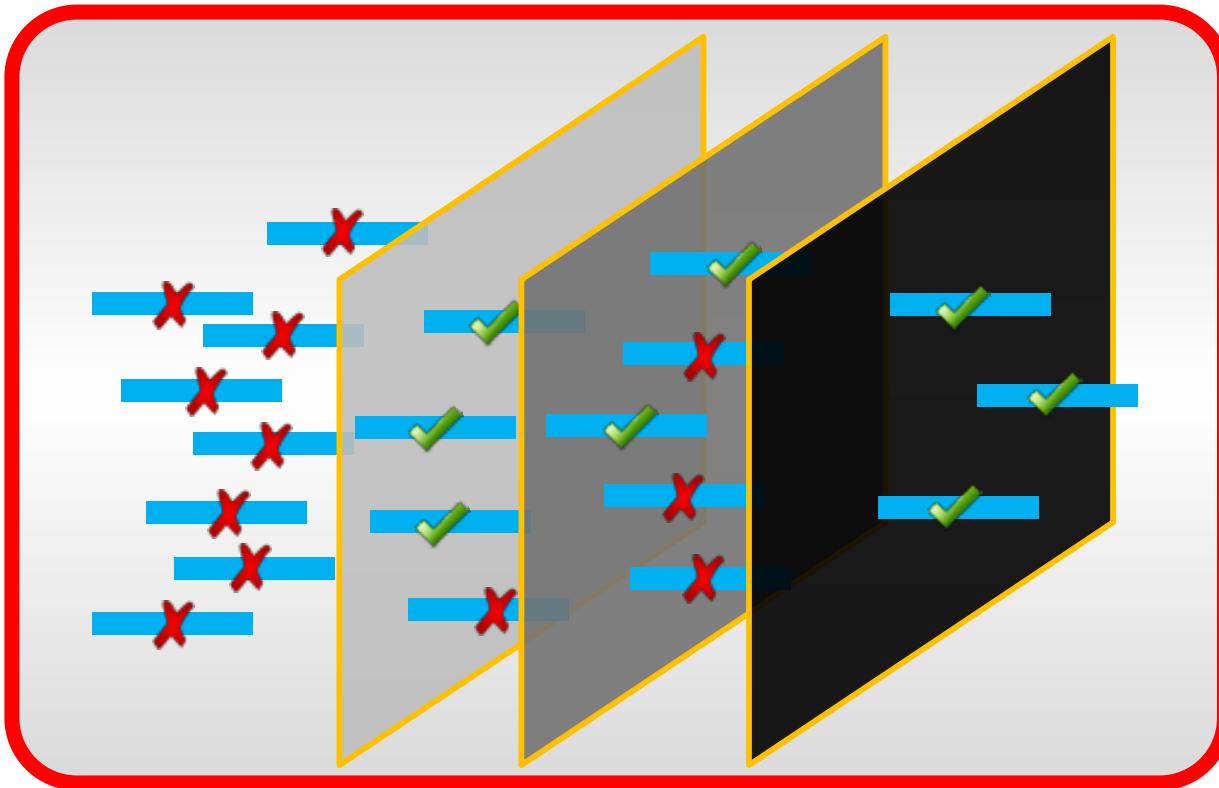
Ideal Location Filtering Algorithm

Step 2

Query
the
Index

Step 3

Read
Alignment



1. Filters out most of the incorrect mappings
2. Preserves all correct mappings
3. Does this quickly

Location Filtering Example

Read Sequence (100 bp)



AMatchg...



Mismatch.

False
Accept

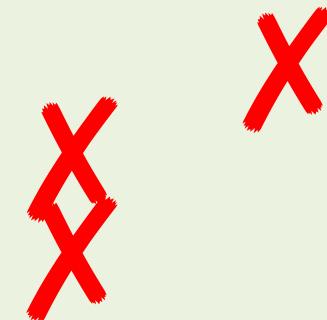
Hash Table

37 140
894 1203
1564



Reference Genome

Filter



Alignment vs. Pre-alignment (Filtering)

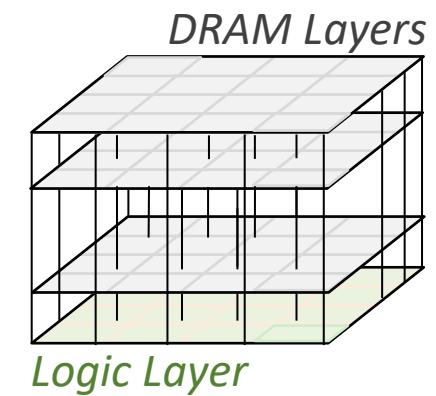
Needleman-Wunsch

		C	T	A	T	A	A	T	A	C	G	
		0	1	2								
A C T A T A	1	0	1	2								
	2	1	0	1	2							
	2	1	0	1	2							
	2	1	0	1	2							
	2	1	2	1	2							
			2	2	2	1	2					
				3	2	2	2	2	2			

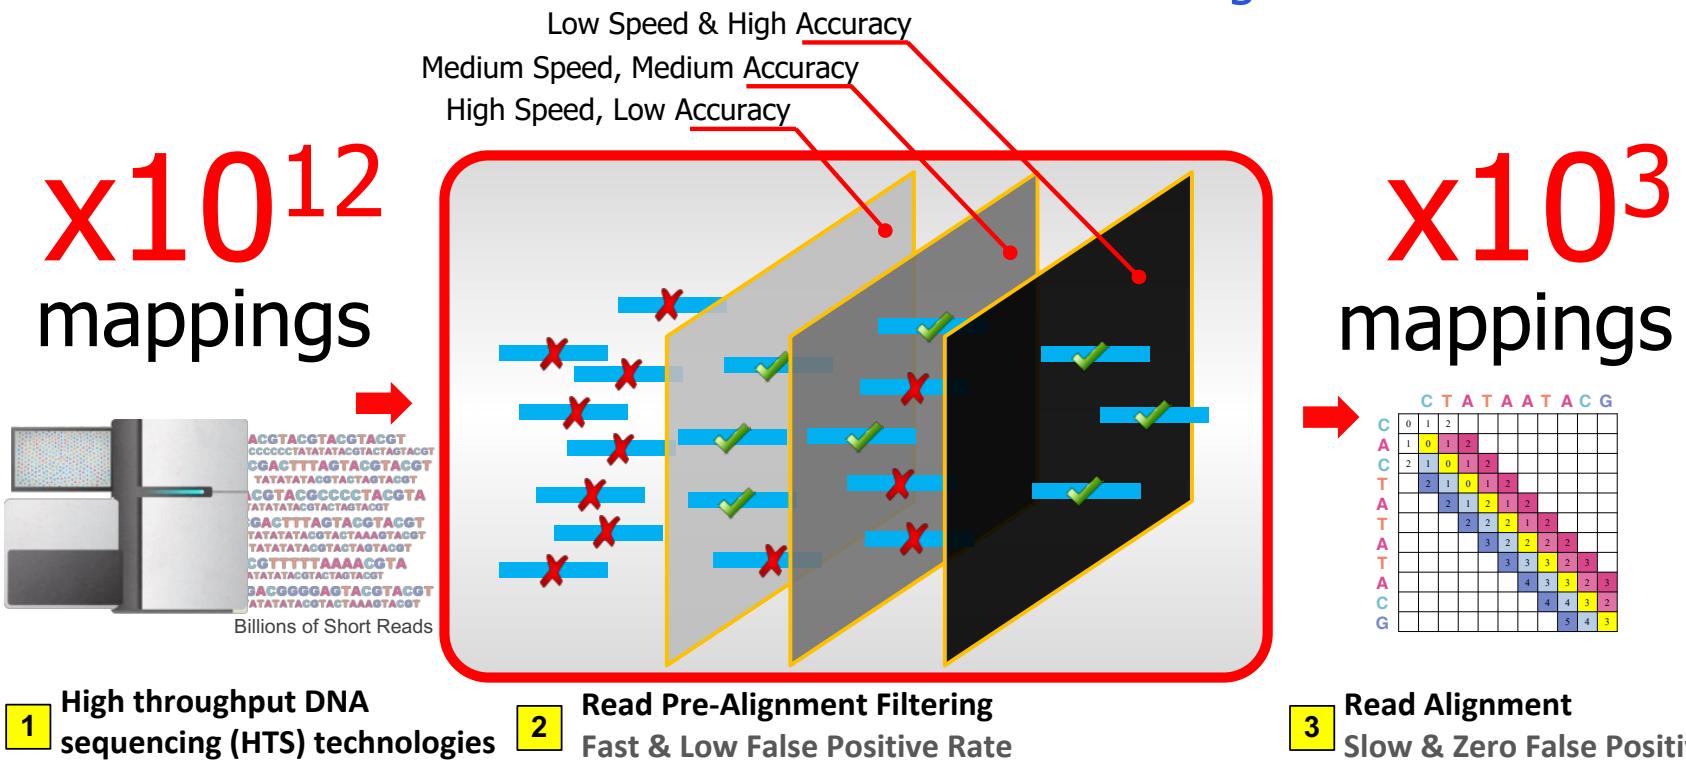
SHD

		C	T	A	T	A	T	A	C	G		
		0	1	2	0	1	0	1	0	1		
A C T A T A	1	0	1	1	0	1	0	1	0	1		
	2	0	1	1	0	1	0	1	0	1		
	1	0	1	0	1	0	1	0	1	0		
	1	0	1	0	1	0	1	0	1	0		
	1	0	1	0	1	0	1	0	1	0		
	1	0	1	0	1	0	1	0	1	0		
	1	0	1	0	1	0	1	0	1	0		

Independent vectors can be processed in parallel using hardware technologies



GateKeeper: FPGA-Based Alignment Filtering



GateKeeper: FPGA-Based Alignment Filtering

- Mohammed Alser, Hasan Hassan, Hongyi Xin, Oguz Ergin, Onur Mutlu, and Can Alkan

"GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping"

Bioinformatics, [published online, May 31], 2017.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

GateKeeper: a new hardware architecture for accelerating pre-alignment in DNA short read mapping

Mohammed Alser , Hasan Hassan, Hongyi Xin, Oğuz Ergin, Onur Mutlu , Can Alkan 

Bioinformatics, Volume 33, Issue 21, 1 November 2017, Pages 3355–3363,

<https://doi.org/10.1093/bioinformatics/btx342>

Published: 31 May 2017 **Article history ▾**

GateKeeper Walkthrough

Generate $2E+1$ masks

Amend random zeros:
 $101 \rightarrow 111$ & $1001 \rightarrow 1111$

AND all masks,
ACCEPT iff number of '1' ≤ Threshold

Query :GAGAGAGATATTAGTGTGCAACTACAAACACAAAAGAGGACCAACTACGTGTCTAAAAGGGGGAACATTGTTGGGCCGGA
Reference :GAGAGAGATAGTTAGTGTGCACTACAAACACAAAAGAGGACCAACTACGTGTCTAAAAGGGGAGACATTGTTGGGCCGG

g Mask : 00000000001000000000001111111011110001110110101101111111110001000001111011010010101

2-Deletion Mask : 000000001011011100111111111111110111100011101101011011111111100010001001101101101001010

1-Insertion Mask :11111111110111110111110111011000100100111111111110010110011000101011101111101111101

2-Insertion Mask :00000010011111001111111100100011010101001101011111111111011100111111000111101100

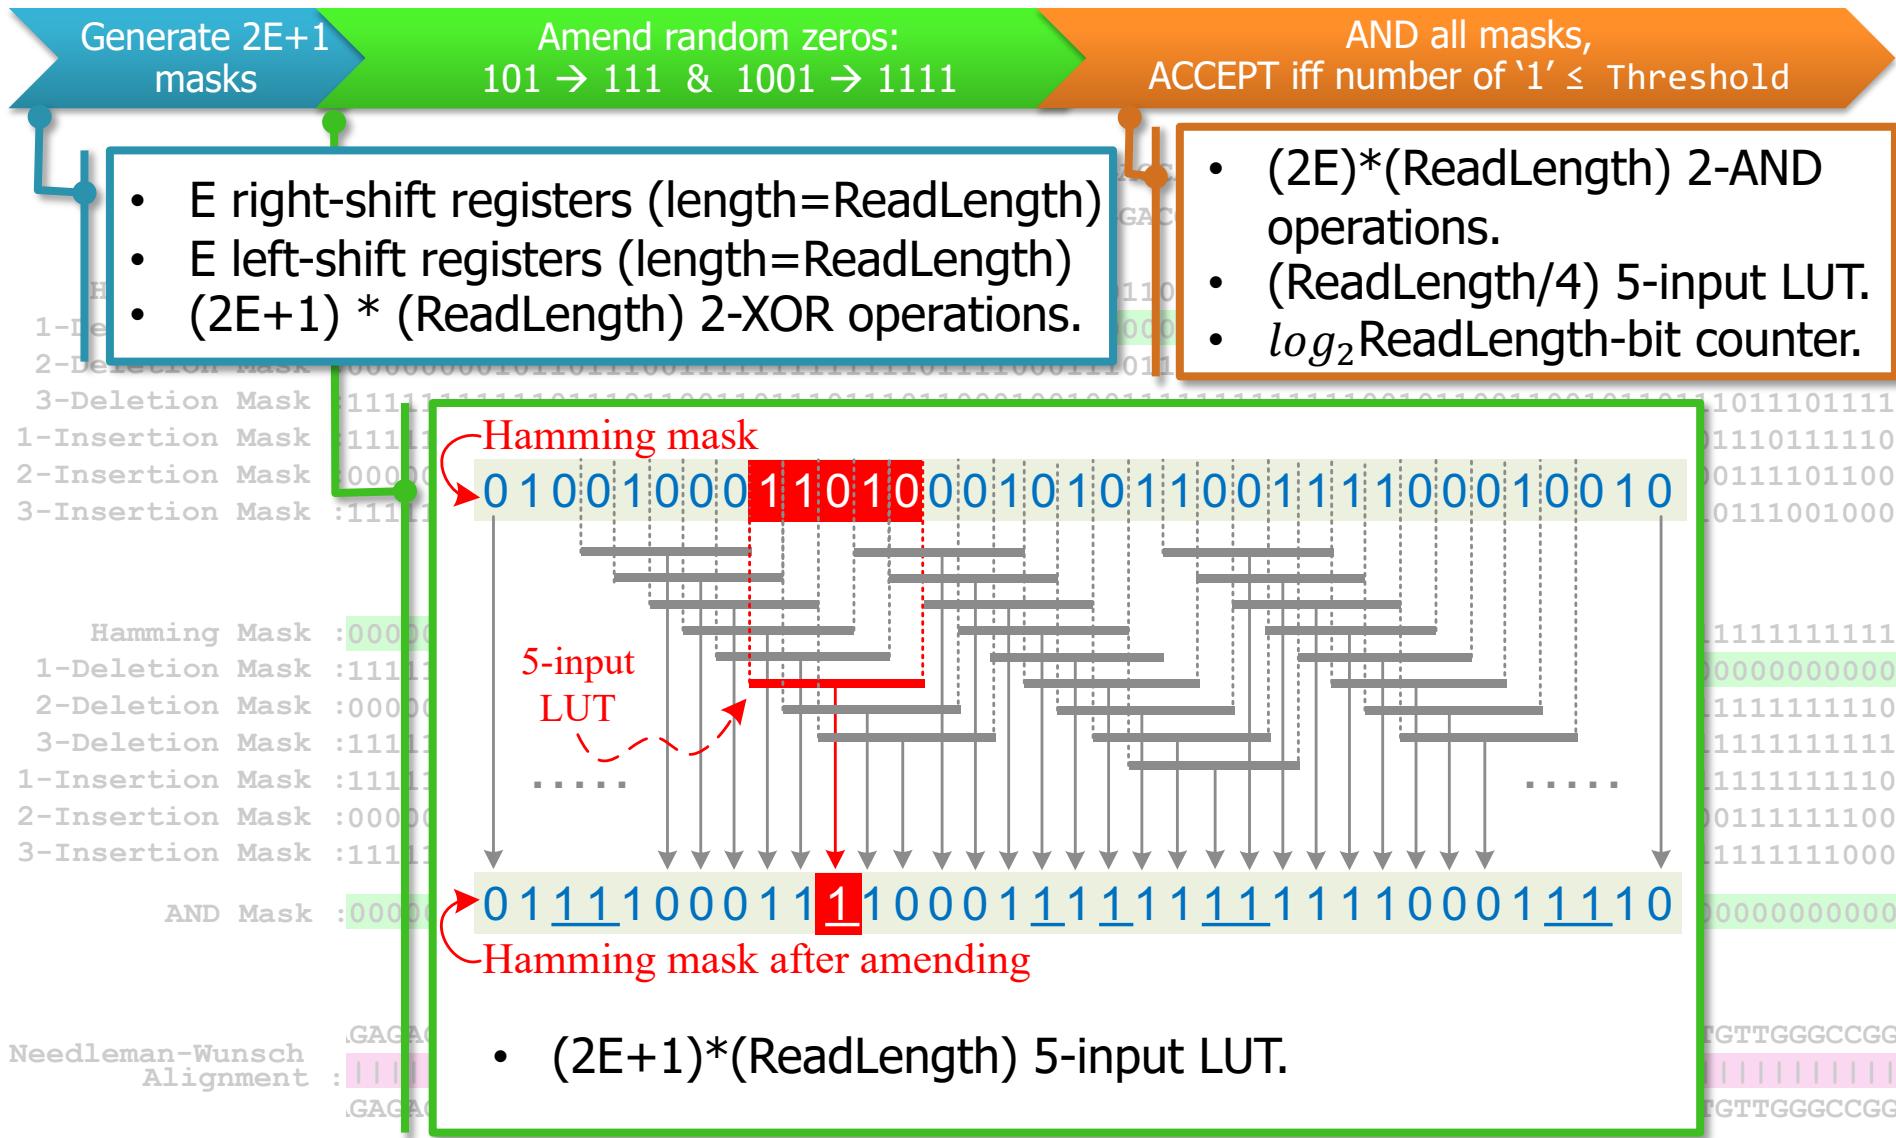
3-Insertion Mask : 111111110111011001100011111110101101111110011001011101111111011101111010111001000

--- Masks after amendment ---

Needleman-Wunsch Alignment : ||| : ||| : ||| : : |||

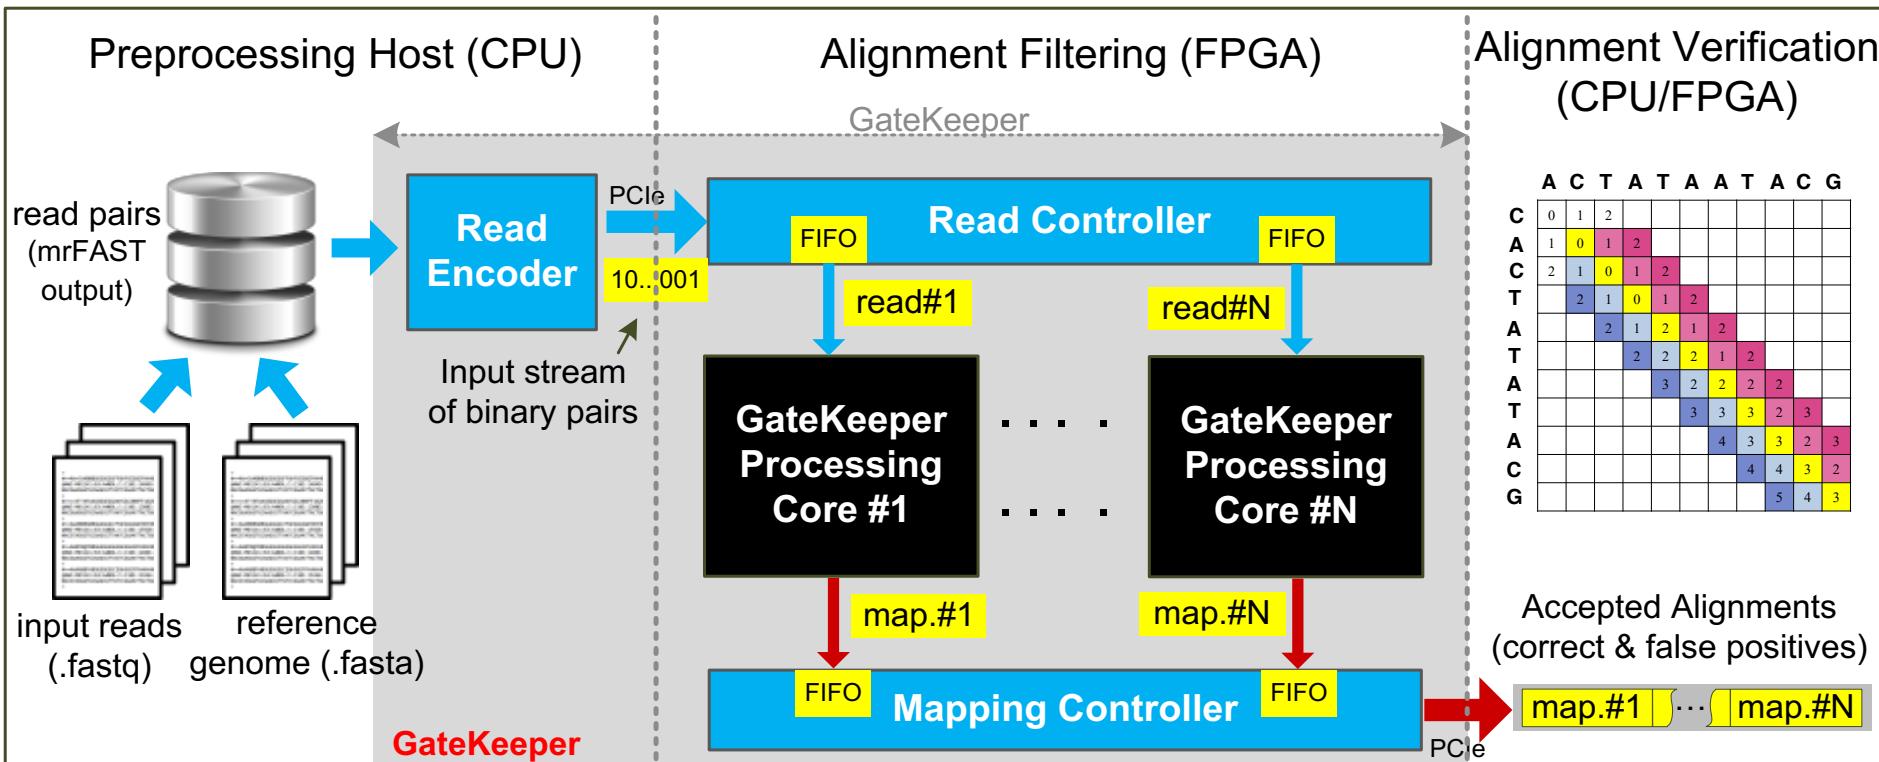
.GAGAGAGATAGTTAGTGTGCAGCCACTACAACACAAAAGAGGGACCAACTTACGTGTCTAAAGGGGAGACATTGTTGGGCCGG

GateKeeper Walkthrough (cont'd)

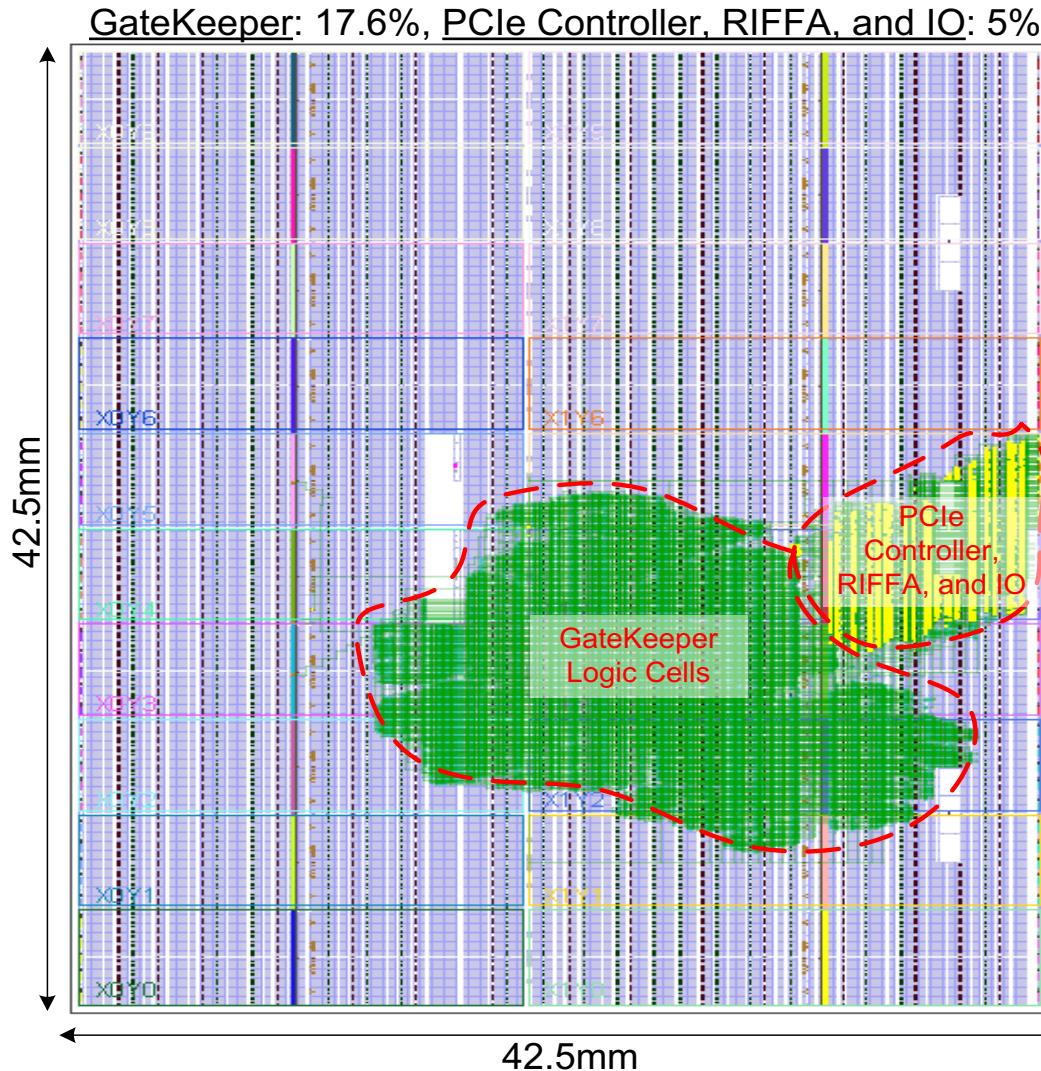


GateKeeper Accelerator Architecture

- **Maximum data throughput** = ~ 13.3 billion bases/sec
- Can examine **8 (300 bp) or 16 (100 bp) mappings concurrently** at 250 MHz
- **Occupies 50%** (100 bp) to **91%** (300 bp) of the FPGA slice LUTs and registers



FPGA Chip Layout



Read length:

300 bp

Error threshold:

E=15

GateKeeper vs. SHD

GateKeeper

- FPGA (Xilinx VC709)
- Multi-core (parallel)
- Examines a single mapping @ 125 MHz
- Limited to PCIe Gen3(4x) transfer rate (128 bits @ 250MHz)
- Amending requires:
 - (2E+1) 5-input LUT.

SHD

- Intel SIMD
- Single-core (sequential)
- Examines a single mapping @ ~2MHz
- Limited to a read length of 128 bp (SSE register size)
- Amending requires:
 - 4(2E+1) bitwise OR.
 - 4(2E+1) packed shuffle.
 - 3(2E+1) shift.

GateKeeper: Speed & Accuracy Results

90x-130x faster filter

than SHD (Xin et al., 2015) and the Adjacency Filter (Xin et al., 2013)

4x lower false accept rate

than the Adjacency Filter (Xin et al., 2013)

10x speedup in read mapping

with the addition of GateKeeper to the mrFAST mapper (Alkan et al., 2009)

Freely available online

github.com/BilkentCompGen/GateKeeper

GateKeeper Conclusions

- **FPGA-based** pre-alignment **greatly** speeds up read mapping
 - **10x speedup** of a state-of-the-art mapper (mrFAST)
- FPGA-based pre-alignment can be **integrated** with the **sequencer**
 - It can help to hide the complexity and details of the FPGA
 - Enables **real-time filtering** while sequencing
 - Paves the way to **on-device genome analysis**

More on GateKeeper

- Mohammed Alser, Hasan Hassan, Hongyi Xin, Oguz Ergin, Onur Mutlu, and Can Alkan

"GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping"

Bioinformatics, [published online, May 31], 2017.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

GateKeeper: a new hardware architecture for accelerating pre-alignment in DNA short read mapping

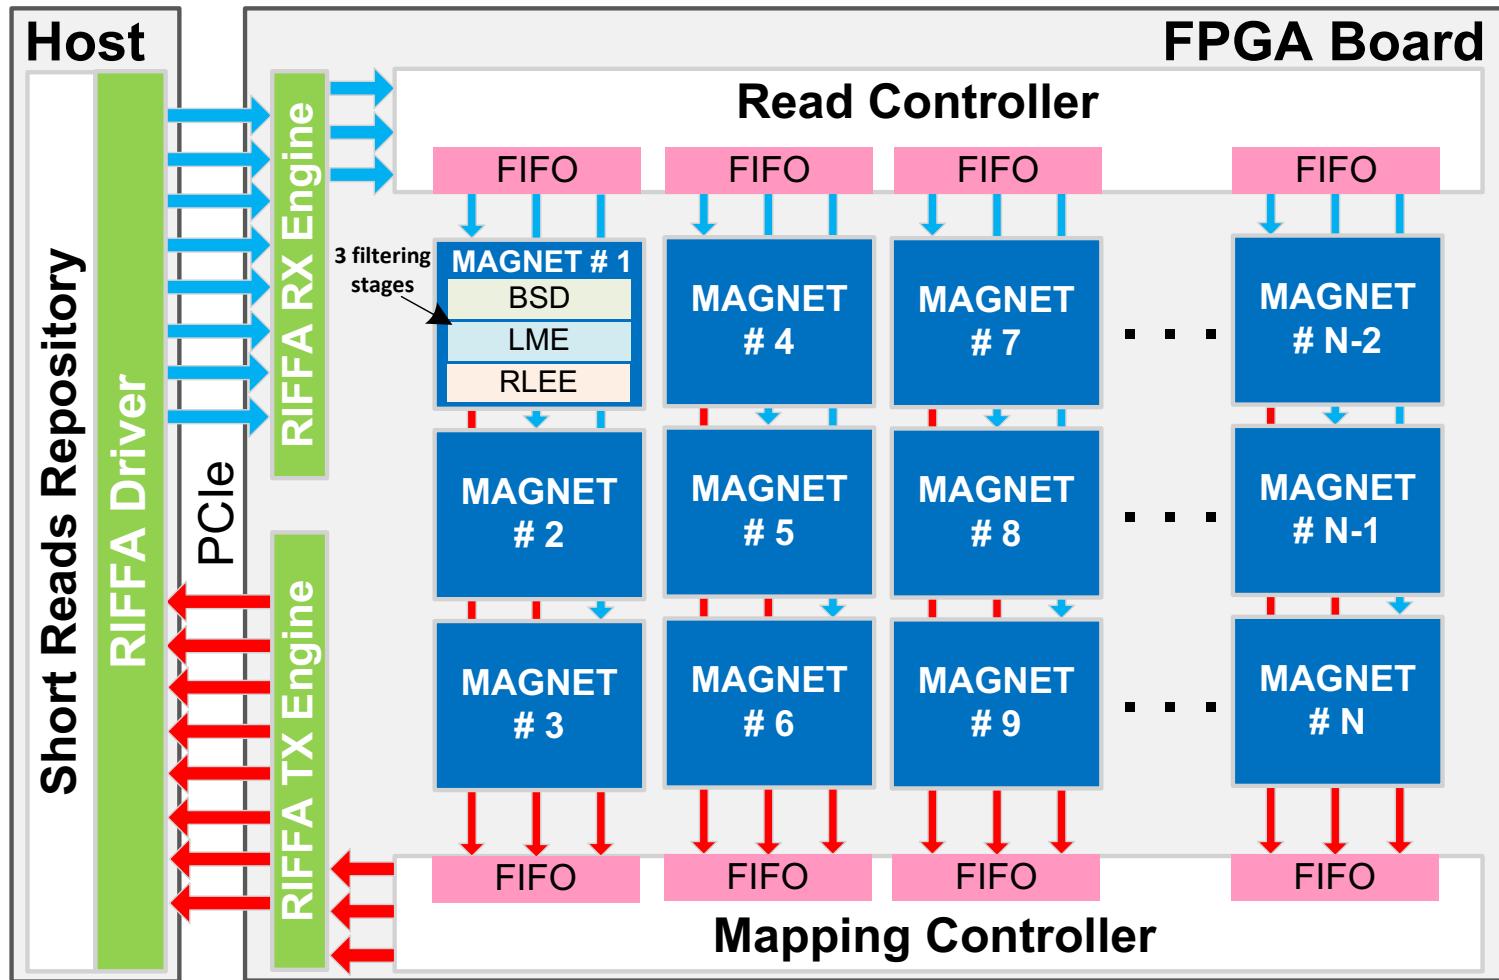
Mohammed Alser , Hasan Hassan, Hongyi Xin, Oğuz Ergin, Onur Mutlu , Can Alkan 

Bioinformatics, Volume 33, Issue 21, 1 November 2017, Pages 3355–3363,

<https://doi.org/10.1093/bioinformatics/btx342>

Published: 31 May 2017 **Article history ▾**

MAGNET Accelerator [Alser+, TIR 2017]

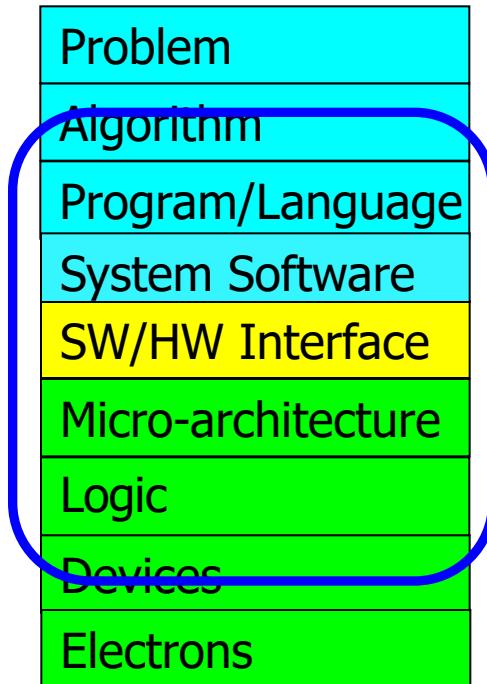


Can We Do Better?

Faster, More Accurate,
More Scalable
Pre-Alignment Filtering

Algorithm-Arch-Device Co-Design is Critical

Computer Architecture
(expanded view)



Shouji (障子) [Alser+, Bioinformatics 2019]

Mohammed Alser, Hasan Hassan, Akash Kumar, Onur Mutlu, and Can Alkan,
"Shouji: A Fast and Efficient Pre-Alignment Filter for Sequence Alignment"
Bioinformatics, [published online, March 28], 2019.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

Bioinformatics, 2019, 1–9
doi: 10.1093/bioinformatics/btz234
Advance Access Publication Date: 28 March 2019
Original Paper



Sequence alignment

Shouji: a fast and efficient pre-alignment filter for sequence alignment

Mohammed Alser^{1,2,3,*}, Hasan Hassan¹, Akash Kumar², Onur Mutlu^{1,3,*} and Can Alkan^{3,*}

¹Computer Science Department, ETH Zürich, Zürich 8092, Switzerland, ²Chair for Processor Design, Center For Advancing Electronics Dresden, Institute of Computer Engineering, Technische Universität Dresden, 01062 Dresden, Germany and ³Computer Engineering Department, Bilkent University, 06800 Ankara, Turkey

*To whom correspondence should be addressed.

Associate Editor: Inanc Birol

Received on September 13, 2018; revised on February 27, 2019; editorial decision on March 7, 2019; accepted on March 27, 2019

Shouji

■ Key observation:

- ❑ Correct alignment always includes long identical subsequences
- ❑ Processing the entire sequence at once is ineffective for hardware design

■ Key idea:

- ❑ Use an **overlapping sliding window** approach to quickly and accurately find all long identical subsequences (consecutive zeros)

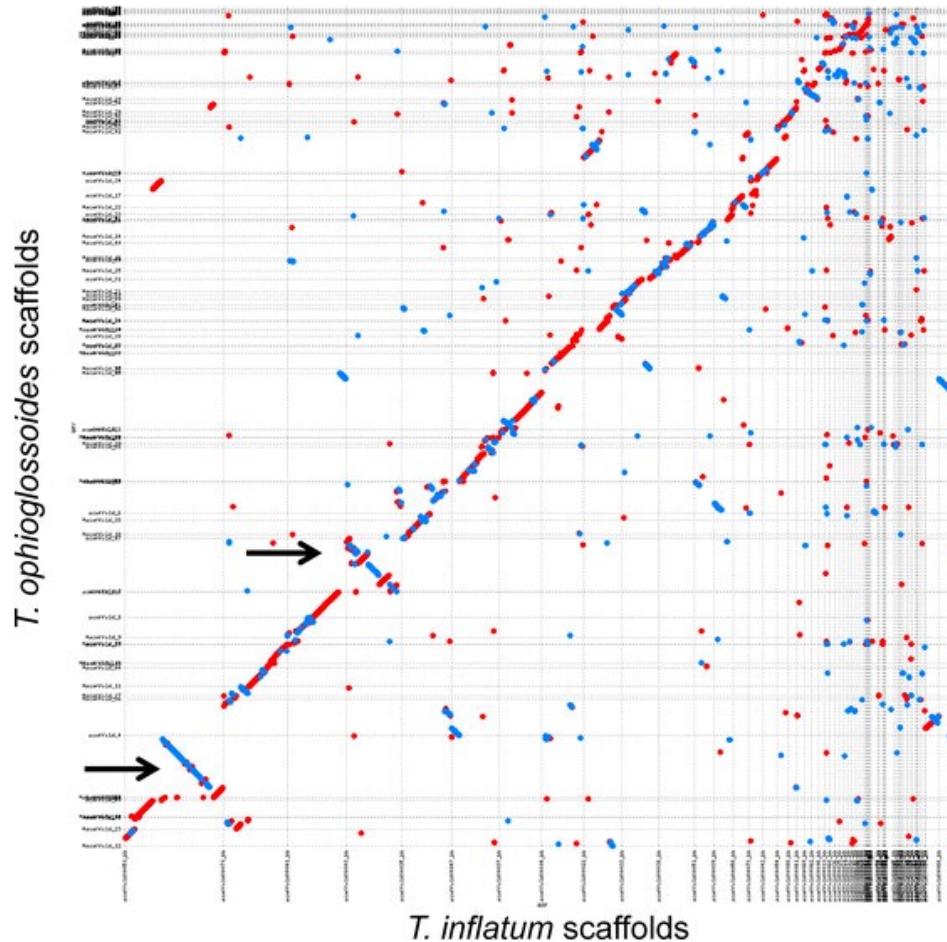
■ Key result:

- ❑ Shouji accelerates the **best-performing CPU read aligner** Edlib (Bioinformatics 2017) by up to 18.8x using 16 filtering units that work in parallel
- ❑ Shouji on FPGA is up to 10,000x faster than on CPU
- ❑ Shouji is 2.4x to 467x more accurate than GateKeeper (Bioinformatics 2017) and SHD (Bioinformatics 2015)

Shouji

■ Key observation:

- Correct alignment always includes long identical subsequences



Dot plot, dot matrix
(Lipman and Pearson, 1985)

Shouji Walkthrough

Build the Neighborhood Map

Find all common subsequences
(diagonal segments of consecutive zeros)
shared between two given sequences

j	1	2	3	4	5	6	7	8	9	10	11	12
i	G	G	T	G	C	A	G	A	G	C	T	C
	0	0	1	0	0	0	0	0	0	0	0	0
1	G	0	0	1	0	0	0	0	0	0	0	0
2	G	0	0	1	0	0	0	0	0	0	0	0
3	T	1	0	0	0	0	0	0	0	0	0	0
4	G	0	0	1	0	0	0	0	0	0	0	0
5	A	1	0	1	1	1	1	1	0	0	0	0
6	G	0	0	1	1	1	1	1	0	1	0	0
7	A	1	0	1	1	1	1	1	1	1	1	1
8	G	0	0	1	1	1	1	1	0	1	0	1
9	T	1	1	1	1	1	1	1	0	0	0	0
10	T	1	1	1	1	1	1	1	0	0	0	1
11	G	1	0	1	1	1	1	1	1	1	1	1
12	T	1	1	0	0	0	0	0	0	1	1	0

Store longest subsequence
in Shouji Bit-vector

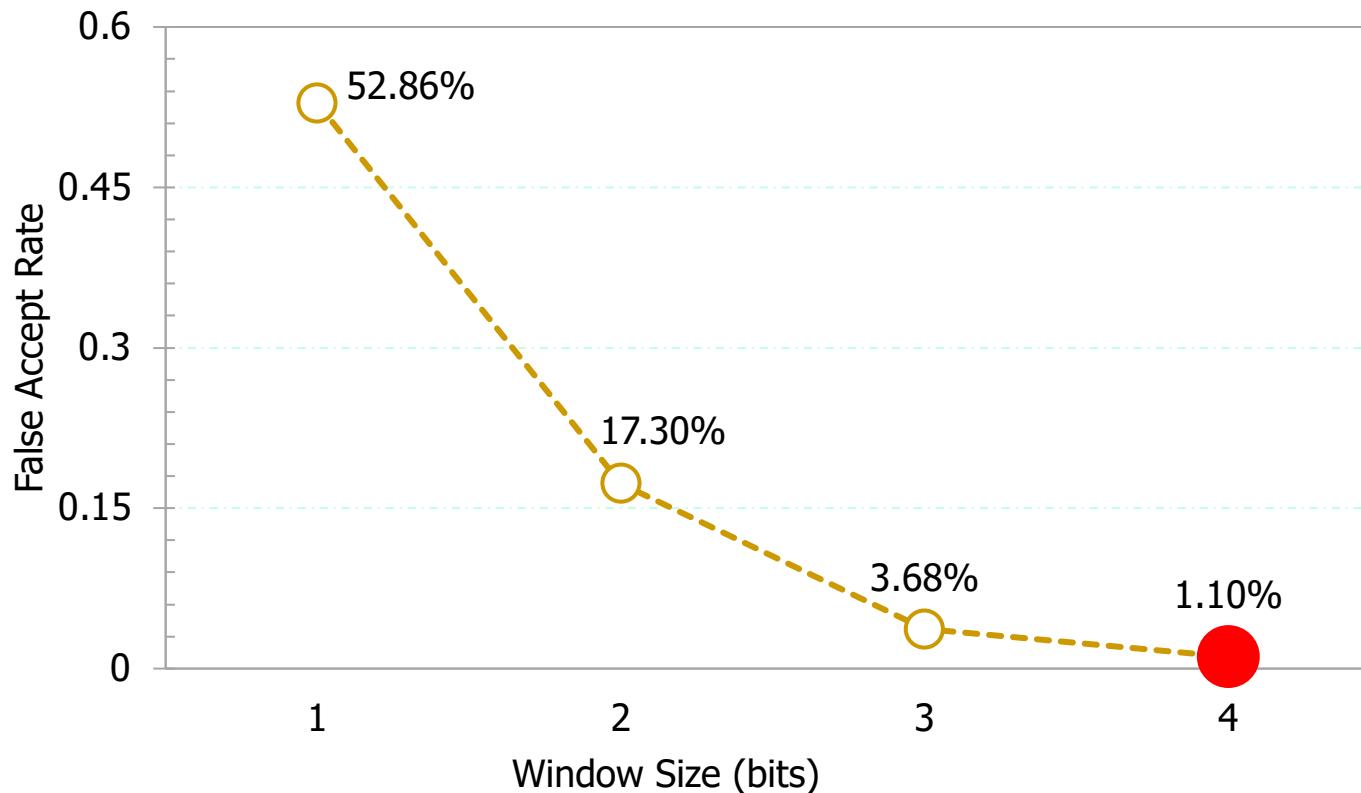
0 0 0 0 1 0 0 0 0 1 0 1

ACCEPT iff number of '1's \leq Threshold

[Shouji: a fast and efficient pre-alignment filter for sequence alignment](#), *Bioinformatics* 2019,
<https://doi.org/10.1093/bioinformatics/btz234>

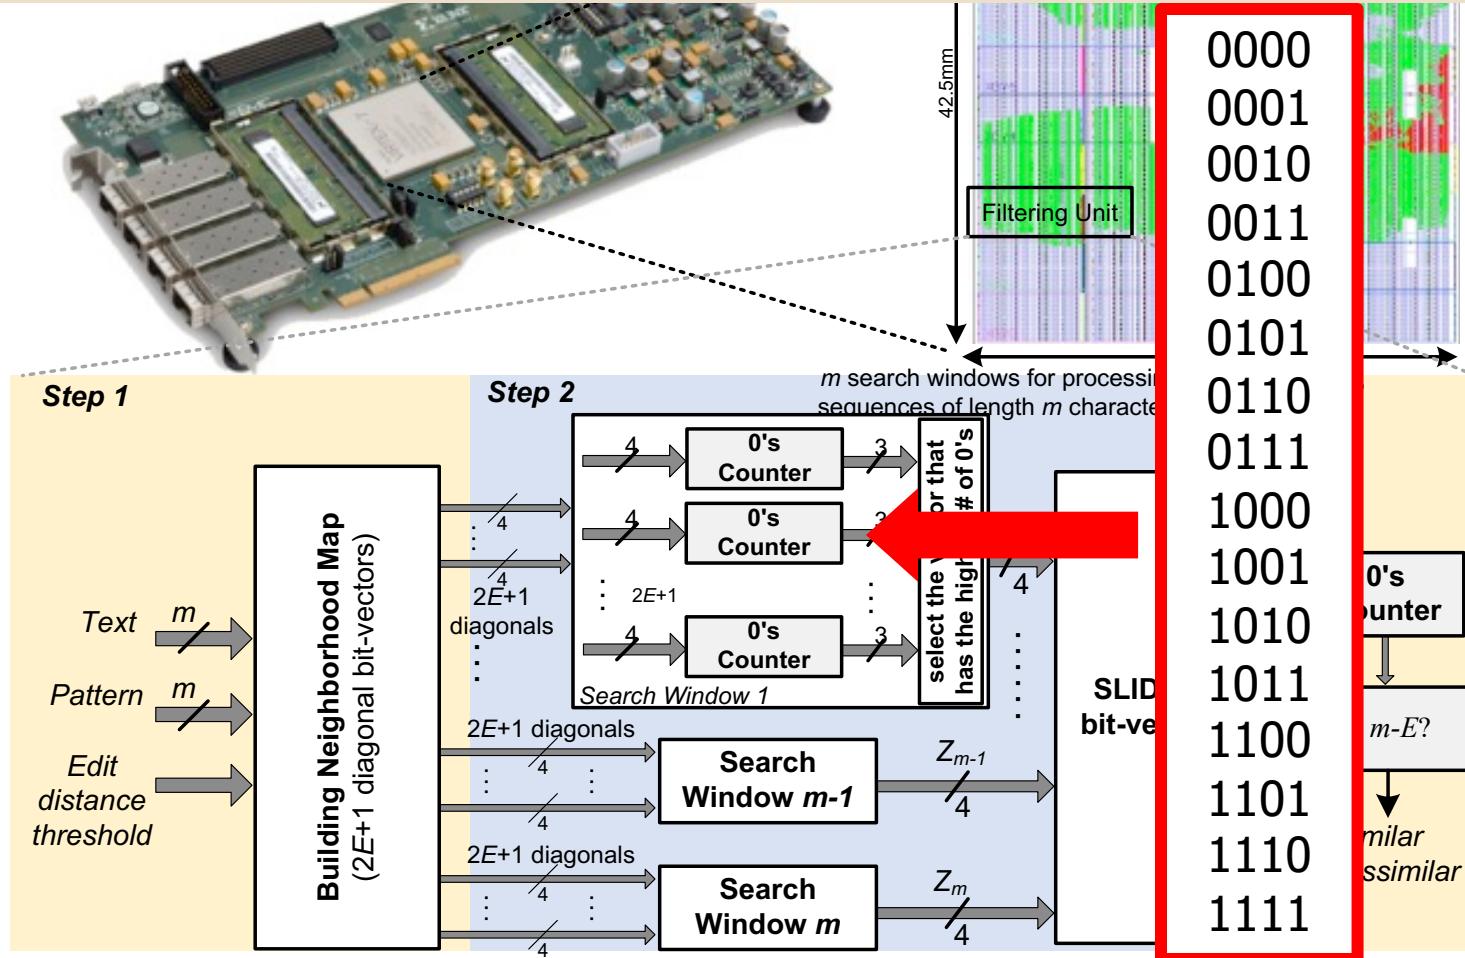
Effect of Sliding Window Size

- Large enough window to accurately capture longer streaks of matches → lower false positives
- Small enough window to perform fast computation



Hardware Implementation

Counting is performed **concurrently** for **all bit-vectors** and **all sliding windows** in a single clock cycle using **multiple 4-input LUTs**



More on Shouji (障子) [Alser+, Bioinformatics 2019]

Mohammed Alser, Hasan Hassan, Akash Kumar, Onur Mutlu, and Can Alkan,

"Shouji: A Fast and Efficient Pre-Alignment Filter for Sequence Alignment"

Bioinformatics, [published online, March 28], 2019.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

Bioinformatics, 2019, 1–9
doi: 10.1093/bioinformatics/btz234
Advance Access Publication Date: 28 March 2019
Original Paper



Sequence alignment

Shouji: a fast and efficient pre-alignment filter for sequence alignment

Mohammed Alser^{1,2,3,*}, Hasan Hassan¹, Akash Kumar², Onur Mutlu^{1,3,*} and Can Alkan^{3,*}

¹Computer Science Department, ETH Zürich, Zürich 8092, Switzerland, ²Chair for Processor Design, Center For Advancing Electronics Dresden, Institute of Computer Engineering, Technische Universität Dresden, 01062 Dresden, Germany and ³Computer Engineering Department, Bilkent University, 06800 Ankara, Turkey

*To whom correspondence should be addressed.

Associate Editor: Inanc Birol

Received on September 13, 2018; revised on February 27, 2019; editorial decision on March 7, 2019; accepted on March 27, 2019

SneakySnake [Alser+, Bioinformatics 2020]

Mohammed Alser, Taha Shahroodi, Juan-Gomez Luna, Can Alkan, and Onur Mutlu,
"SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs"

Bioinformatics, to appear in 2020.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

Bioinformatics
doi.10.1093/bioinformatics/xxxxxx
Advance Access Publication Date: Day Month Year
Manuscript Category



Subject Section

SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs

Mohammed Alser^{1,2,*}, Taha Shahroodi¹, Juan Gómez-Luna^{1,2},
Can Alkan^{4,*}, and Onur Mutlu^{1,2,3,4,*}

¹Department of Computer Science, ETH Zurich, Zurich 8006, Switzerland

²Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8006, Switzerland

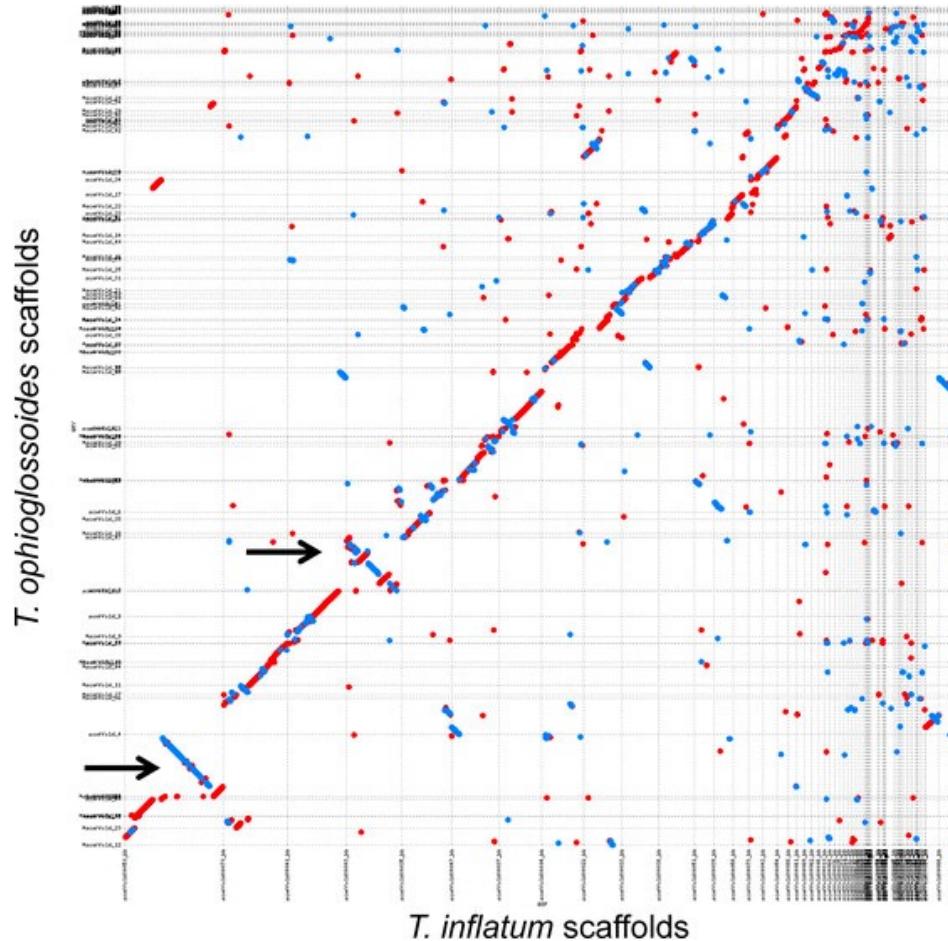
³Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh 15213, PA, USA

⁴Department of Computer Engineering, Bilkent University, Ankara 06800, Turkey

SneakySnake

■ Key observation:

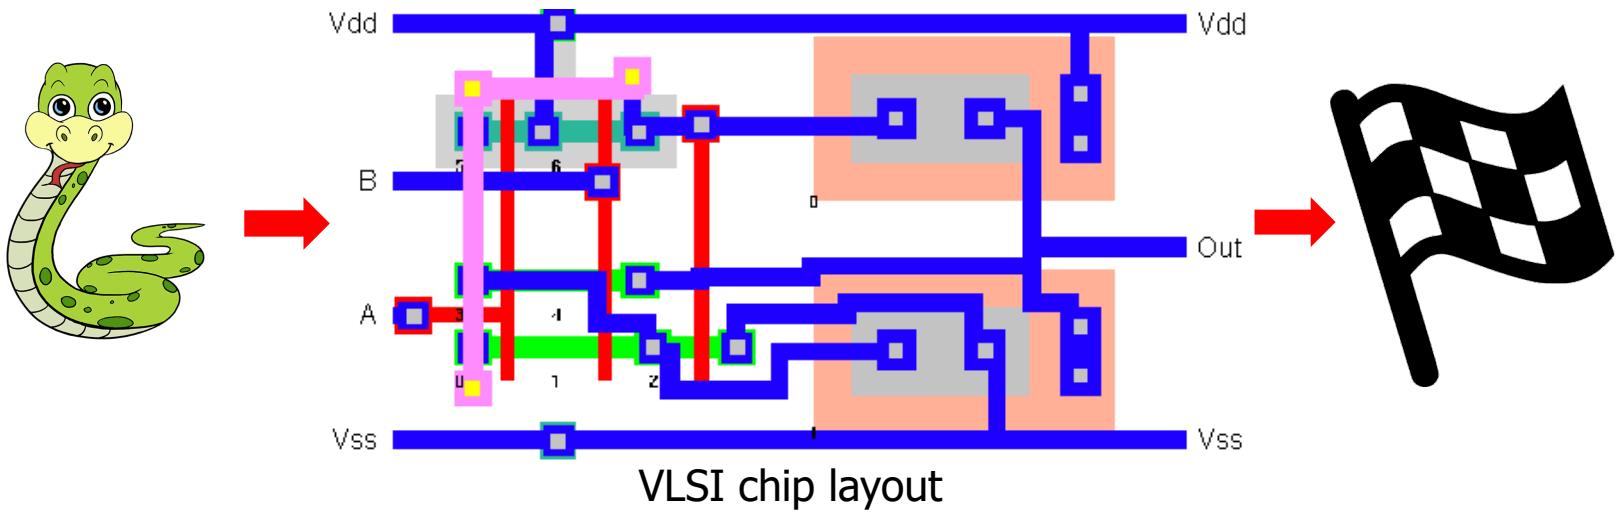
- Correct alignment is a sequence of non-overlapping long matches



Dot plot, dot matrix
(Lipman and Pearson, 1985)

SneakySnake

- **Key observation:**
 - Correct alignment is a sequence of non-overlapping long matches
- **Key idea:**
 - Reduce the approximate string matching problem to the **Single Net Routing problem** in VLSI chip layout



SneakySnake

■ Key observation:

- Correct alignment is a sequence of non-overlapping long matches

■ Key idea:

- Reduce the approximate string matching problem to the Single Net Routing problem in VLSI chip layout

■ Key result:

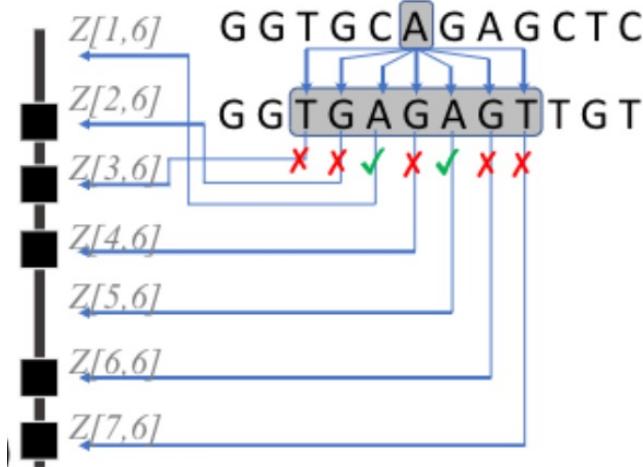
- SneakySnake is up to four orders of magnitude more accurate than Shouji (Bioinformatics'19) and GateKeeper (Bioinformatics'17)
- SneakySnake greatly accelerates state-of-the-art CPU sequence aligners, Edlib (Bioinformatics'17) and Parasail (BMC Bioinformatics'16)
 - by up to 37.7× and 43.9× (>12× on average), on CPUs
 - by up to 413× and 689× (>400× on average) with FPGAs/GPUs

SneakySnake Walkthrough

Building Neighborhood Map

Finding the Optimal Routing Path

Examining the Snake Survival



$$E = 3$$

	<i>column</i>	1	2	3	4	5	6	7	8	9	10	11	12
<i>3rd Upper Diagonal</i>	1	1	1	0	1	1	0	0	0	1	1	1	1
<i>2nd Upper Diagonal</i>	1	1	1	0	1	1	1	1	1	1	0	1	1
<i>1st Upper Diagonal</i>	1	0	1	1	1	0	0	0	0	1	0	1	1
<i>Main Diagonal</i>	0	0	0	0	1	1	1	1	1	1	1	1	1
<i>1st Lower Diagonal</i>	0	1	1	1	1	0	0	1	1	1	0	1	1
<i>2nd Lower Diagonal</i>	1	0	1	0	1	1	1	1	0	1	1	1	1
<i>3rd Lower Diagonal</i>	0	1	1	1	1	1	1	1	1	1	1	1	1

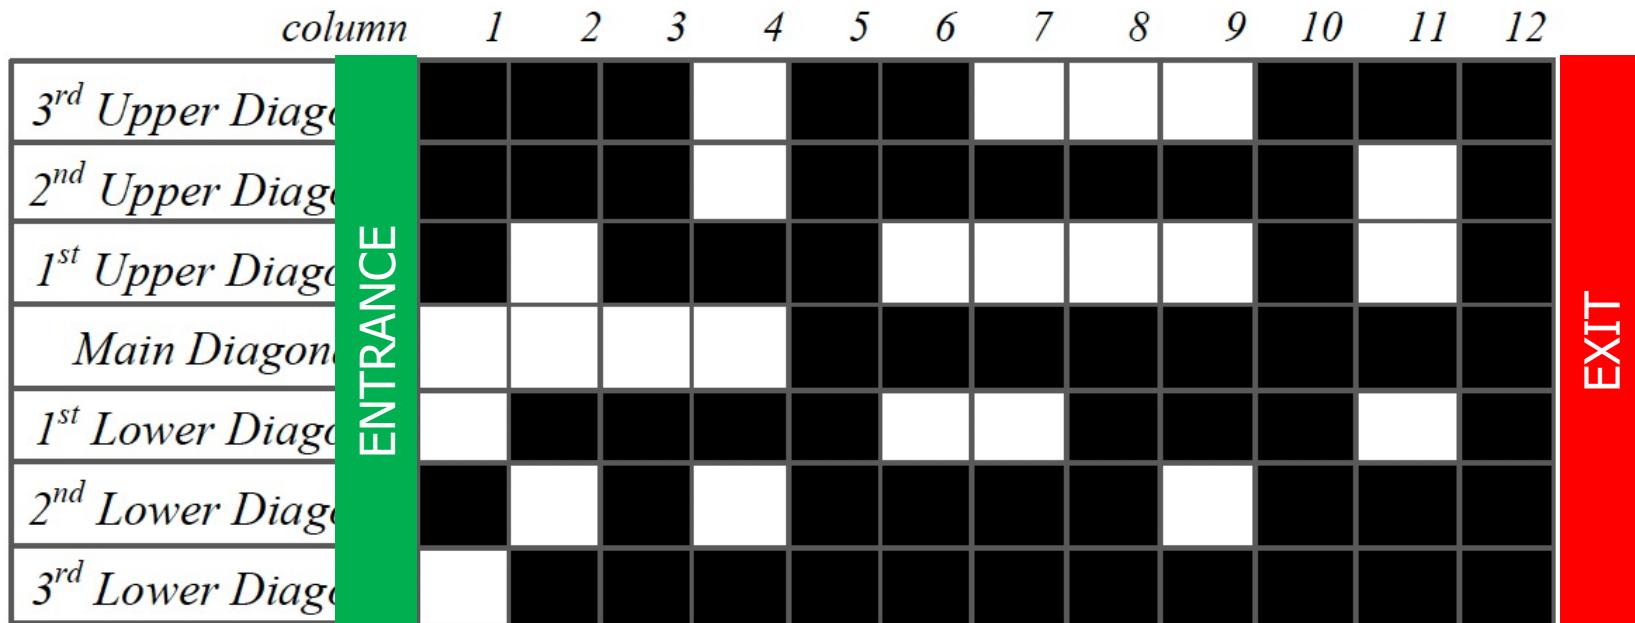
SneakySnake Walkthrough

Building Neighborhood Map

Finding the Optimal Routing Path

Examining the Snake Survival

$$E = 3$$

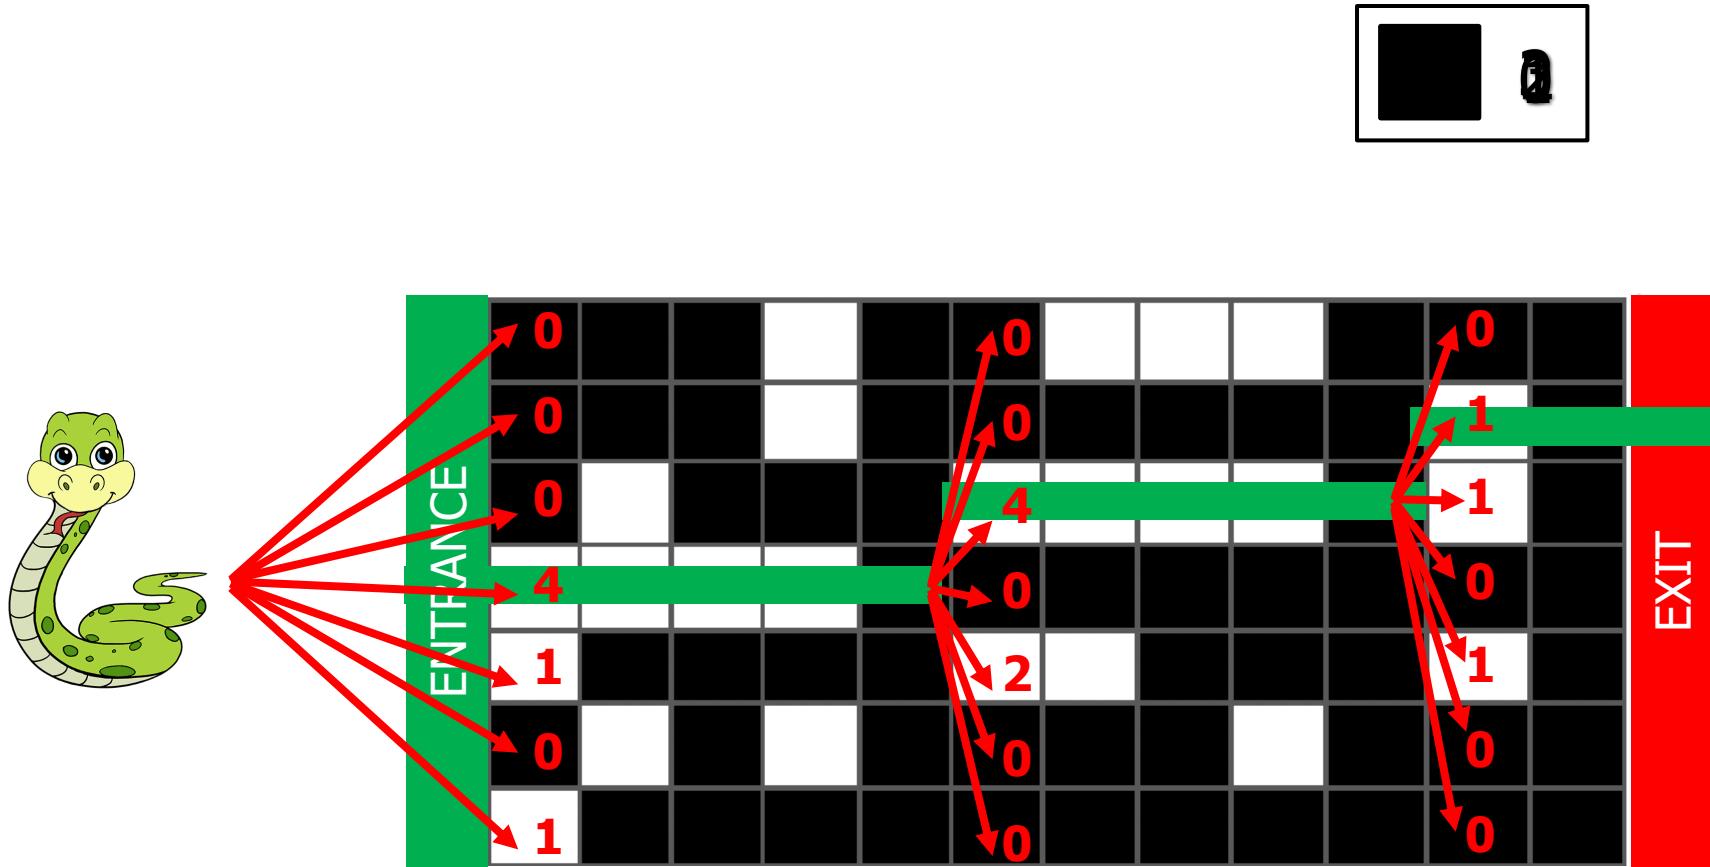


SneakySnake Walkthrough

Building Neighborhood Map

Finding the Optimal Routing Path

Examining the Snake Survival



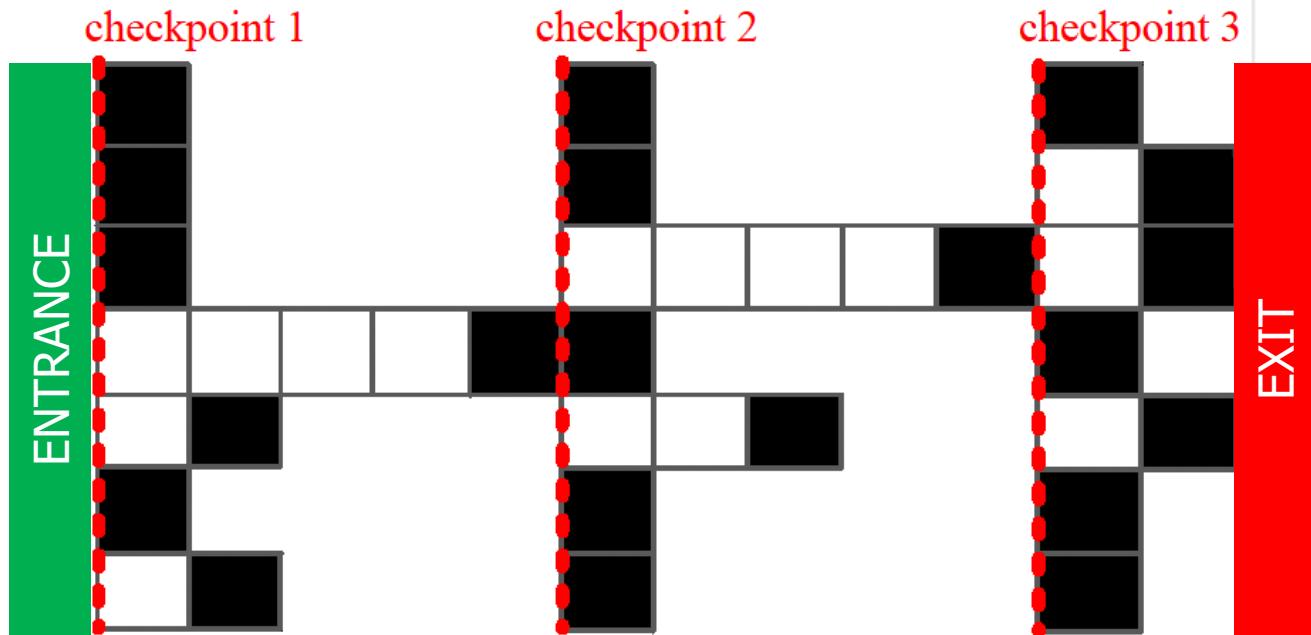
SneakySnake Walkthrough

Building Neighborhood Map

Finding the Routing Travel Path

Examining the Snake Survival

**This is what you actually need to build
and it can be done on-the-fly!**



FPGA Resource Analysis

- FPGA resource usage for a single filtering unit of GateKeeper, Shouji, and Snake-on-Chip for a sequence length of 100 and under different edit distance thresholds (E).

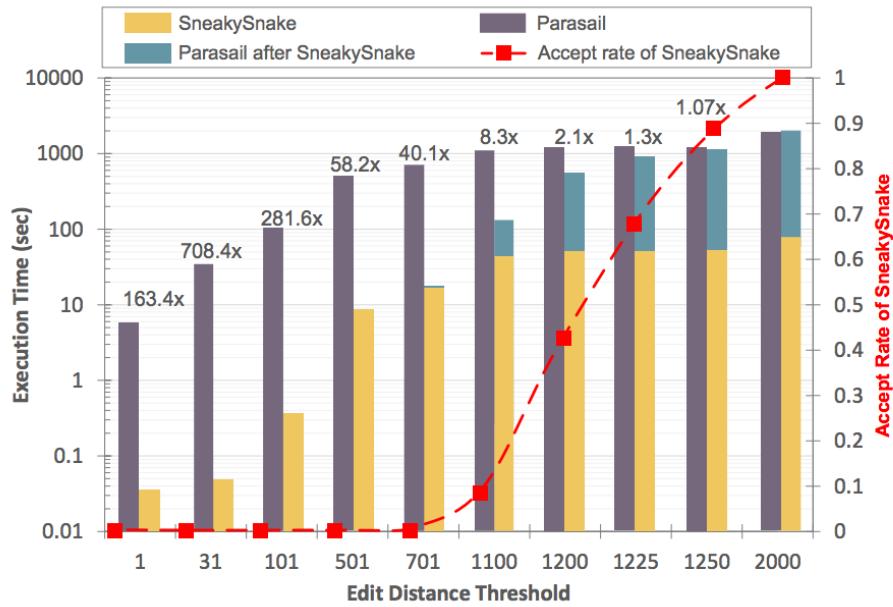
	E (bp)	Slice LUT	Slice Register	No. of Filtering Units
GateKeeper	2	0.39%	0.01%	16
	5	0.71%	0.01%	16
Shouji	2	0.69%	0.08%	16
	5	1.72%	0.16%	16
Snake-on-Chip	2	0.68%	0.16%	16
	5	1.42%	0.34%	16

Key Results of SneakySnake

- ❑ SneakySnake is up to **four orders of magnitude more accurate** than **Shouji** (Bioinformatics'19) and **GateKeeper** (Bioinformatics'17)
- ❑ Short reads:
 - ❑ SneakySnake **accelerates Edlib** (Bioinformatics'17) and **Parasail** (BMC Bioinformatics'16) by
 - up to **37.7× and 43.9×** ($>12\times$ on average), on CPUs
 - up to **413× and 689×** ($>400\times$ on average) using **FPGAs/GPUs**
- ❑ Long reads:
 - ❑ SneakySnake **accelerates Parasail** and **KSW2** by **140.1× and 17.1×** on average, respectively, on CPUs

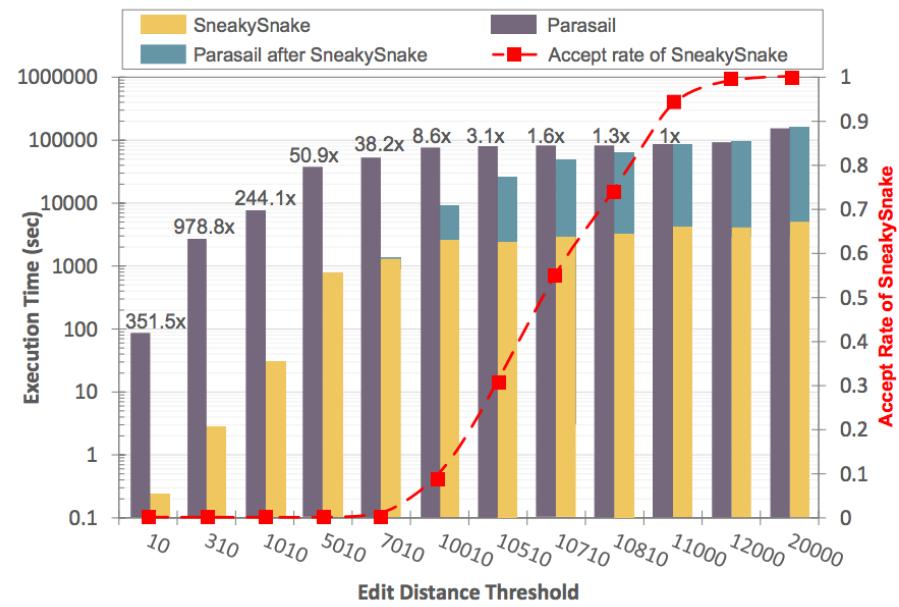
Long Read Mapping (SneakySnake vs Parasail)

10K bp reads



(a)

100K bp reads

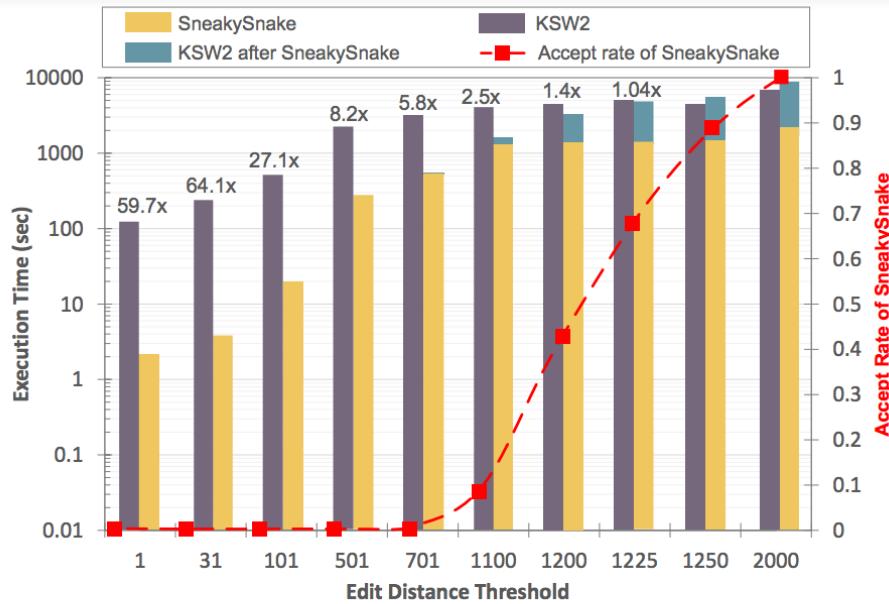


(b)

Fig. 10: The execution time of SneakySnake, Parasail, and SneakySnake integrated with Parasail using long sequences, (a) 10Kbp and (b) 100Kbp, and 40 CPU threads. The left y-axes of (a) and (b) are on a logarithmic scale. For each edit distance threshold value, we provide in the right y-axes of (a) and (b) the rate of accepted pairs (out of 100,000 pairs for 10Kbp and out of 74,687 pairs for 100Kbp) by SneakySnake that are passed to Parasail. We present the end-to-end speedup values obtained by integrating SneakySnake with Parasail.

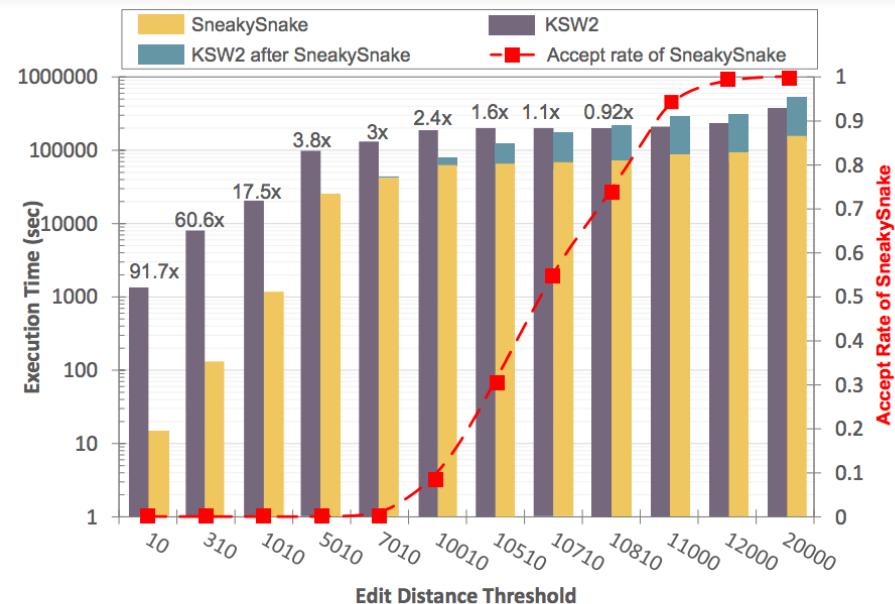
Long Read Mapping (SneakySnake vs KSW2)

10K bp reads



(a)

100K bp reads



(b)

Fig. 11: The execution time of SneakySnake, KSW2, and SneakySnake integrated with KSW2 using long sequences, (a) 10Kbp and (b) 100Kbp, and a single CPU thread. The left y-axes of (a) and (b) are on a logarithmic scale. For each edit distance threshold value, we provide in the right y-axes of (a) and (b) the rate of accepted pairs (out of 100,000 pairs for 10Kbp and out of 74,687 pairs for 100Kbp) by SneakySnake that are passed to KSW2. We present the end-to-end speedup values obtained by integrating SneakySnake with KSW2.

More on SneakySnake [Alser+, Bioinformatics 2020]

Mohammed Alser, Taha Shahroodi, Juan-Gomez Luna, Can Alkan, and Onur Mutlu,

"SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs"

Bioinformatics, to appear in 2020.

[[Source Code](#)]

[[Online link at Bioinformatics Journal](#)]

Bioinformatics
doi.10.1093/bioinformatics/xxxxxx
Advance Access Publication Date: Day Month Year
Manuscript Category



Subject Section

SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs

**Mohammed Alser^{1,2,*}, Taha Shahroodi¹, Juan Gómez-Luna^{1,2},
Can Alkan^{4,*}, and Onur Mutlu^{1,2,3,4,*}**

¹Department of Computer Science, ETH Zurich, Zurich 8006, Switzerland

²Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8006, Switzerland

³Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh 15213, PA, USA

⁴Department of Computer Engineering, Bilkent University, Ankara 06800, Turkey

GenASM Framework [MICRO 2020]

- Damla Senol Cali, Gurpreet S. Kalsi, Zulal Bingol, Can Firtina, Lavanya Subramanian, Jeremie S. Kim, Rachata Ausavarungnirun, Mohammed Alser, Juan Gomez-Luna, Amirali Boroumand, Anant Nori, Allison Scibisz, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,

["GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis"](#)

Proceedings of the 53rd International Symposium on Microarchitecture (MICRO), Virtual, October 2020.

[[Lightning Talk Video](#) (1.5 minutes)]

[[Lightning Talk Slides \(pptx\)](#) ([pdf](#))]

[[Talk Video](#) (18 minutes)]

[[Slides \(pptx\)](#) ([pdf](#))]

GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis

Damla Senol Cali[†][✉] Gurpreet S. Kalsi[✉] Zülal Bingöl[▽] Can Firtina[◊] Lavanya Subramanian[‡] Jeremie S. Kim^{◊†}
Rachata Ausavarungnirun[○] Mohammed Alser[◊] Juan Gomez-Luna[◊] Amirali Boroumand[†] Anant Nori[✉]
Allison Scibisz[†] Sreenivas Subramoney[✉] Can Alkan[▽] Saugata Ghose^{★†} Onur Mutlu^{◊†▽}

[†]*Carnegie Mellon University* [✉]*Processor Architecture Research Lab, Intel Labs* [▽]*Bilkent University* [◊]*ETH Zürich*

[‡]*Facebook* [○]*King Mongkut's University of Technology North Bangkok* [★]*University of Illinois at Urbana-Champaign*

Problem & Our Goal

- Multiple steps of read mapping require *approximate string matching*
 - ASM enables read mapping to account for sequencing errors and genetic variations in the reads
- ASM makes up a significant portion of read mapping (more than 70%)
- One of the major bottlenecks of genome sequence analysis

Our Goal:

Accelerate approximate string matching by designing a fast and flexible framework, which can be used to accelerate *multiple steps* of the genome sequence analysis pipeline

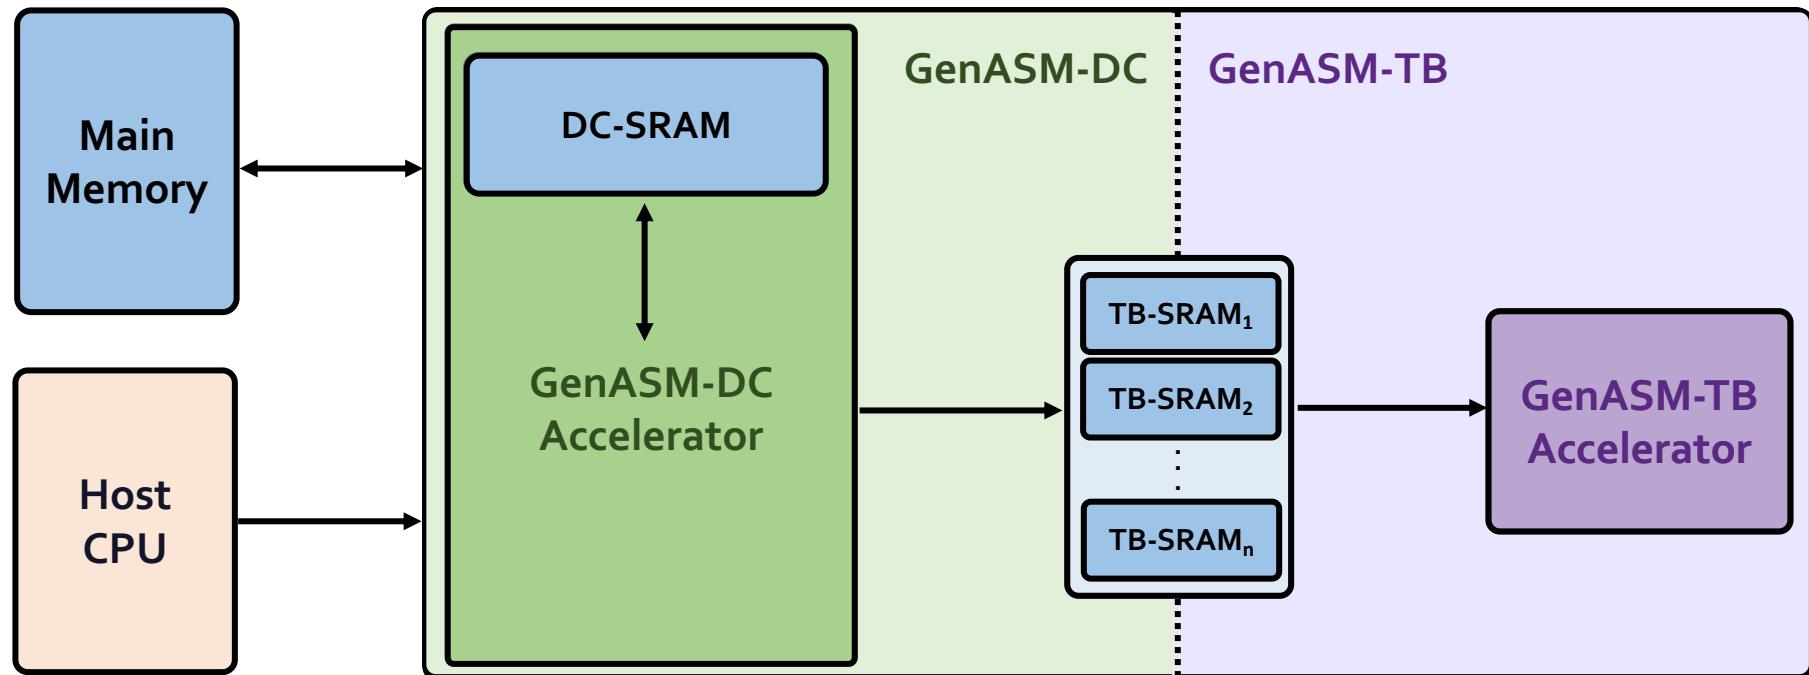
GenASM: ASM Framework for GSA

Our Goal:

Accelerate approximate string matching
by designing a fast and flexible framework,
which can accelerate *multiple steps* of genome sequence analysis

- **GenASM:** First ASM acceleration framework for GSA
 - Based on the *Bitap* algorithm
 - Uses **fast and simple bitwise operations** to perform ASM
 - Modified and extended ASM algorithm
 - **Highly-parallel Bitap** with long read support
 - Bitvector-based **novel algorithm** to perform *traceback*
 - Co-design of our modified **scalable** and memory-efficient algorithms with **low-power** and area-efficient hardware accelerators

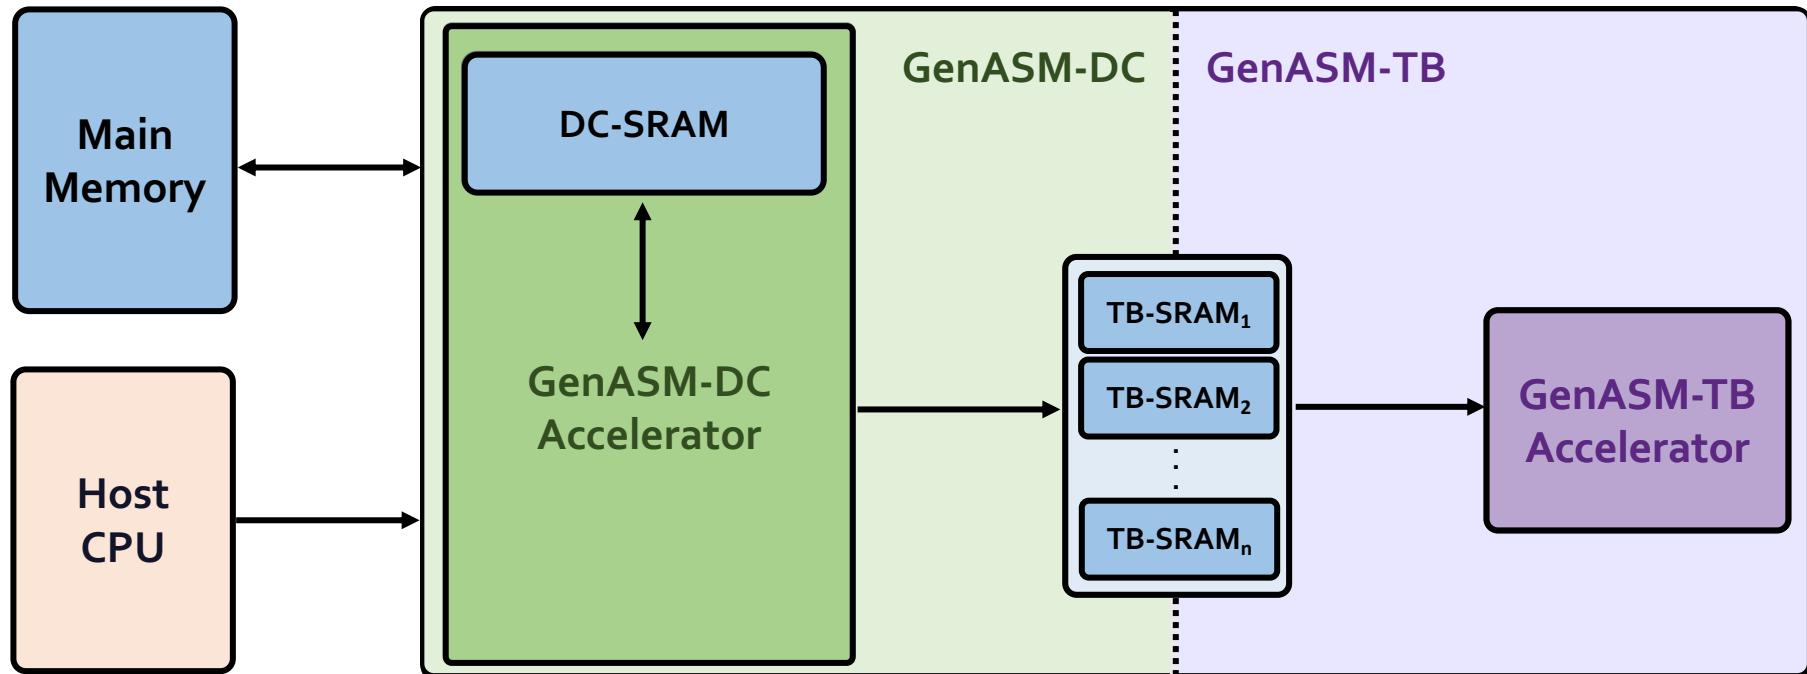
GenASM: Hardware Design



GenASM-DC:
generates bitvectors
and performs edit
Distance Calculation

GenASM-TB:
performs TraceBack
and assembles the
optimal alignment

GenASM: Hardware Design

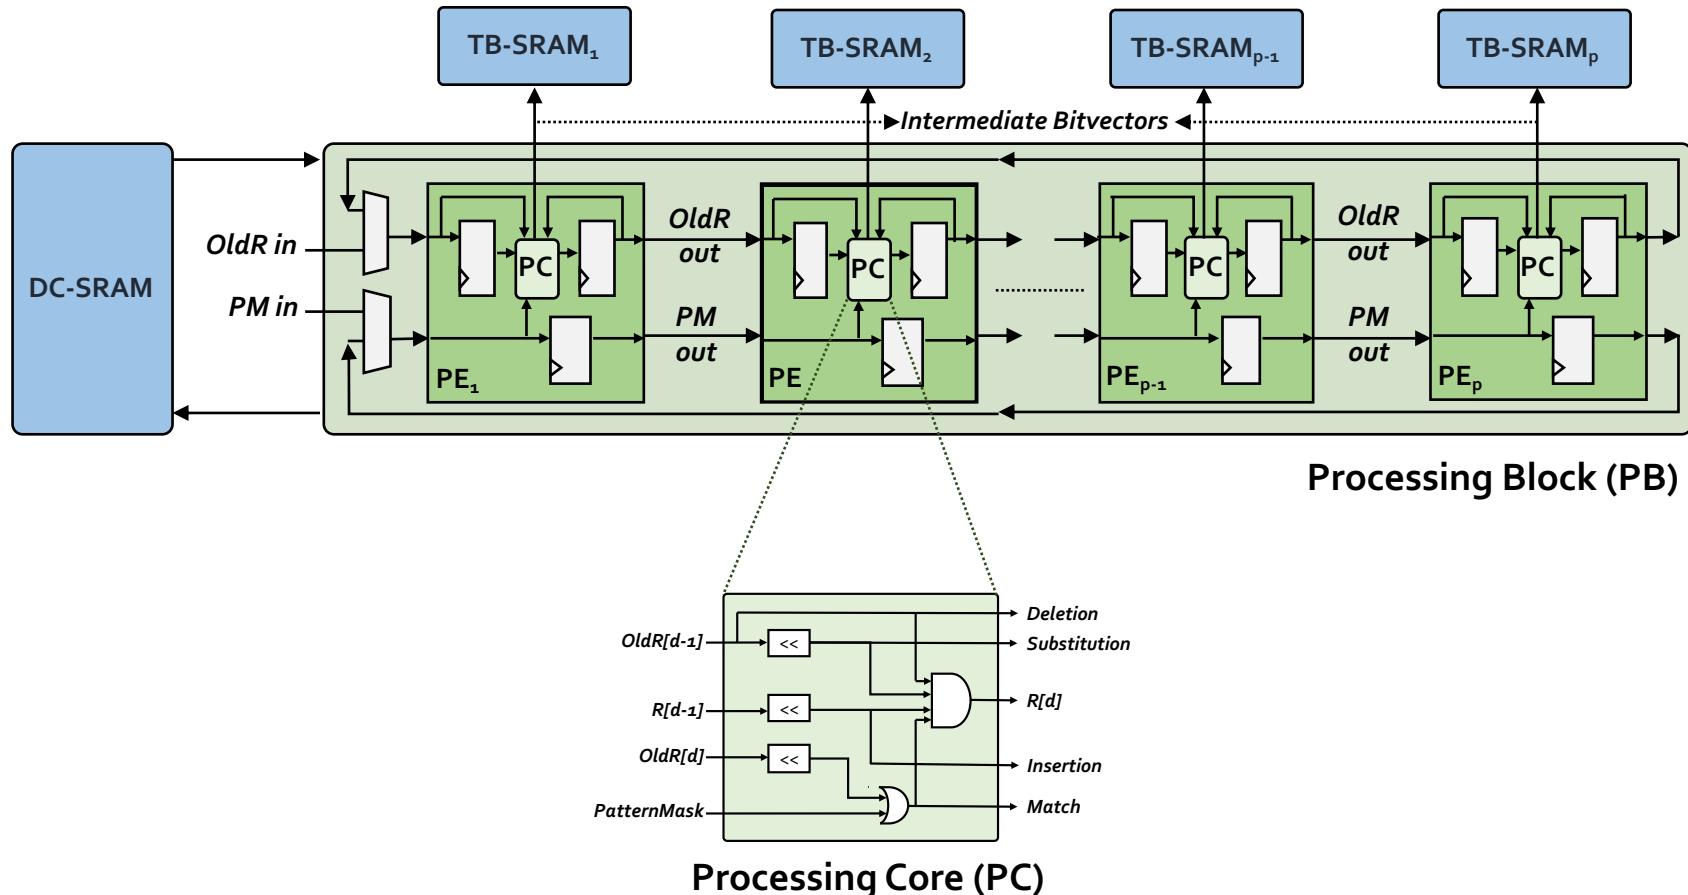


Our *specialized compute units* and *on-chip SRAMs* help us to:

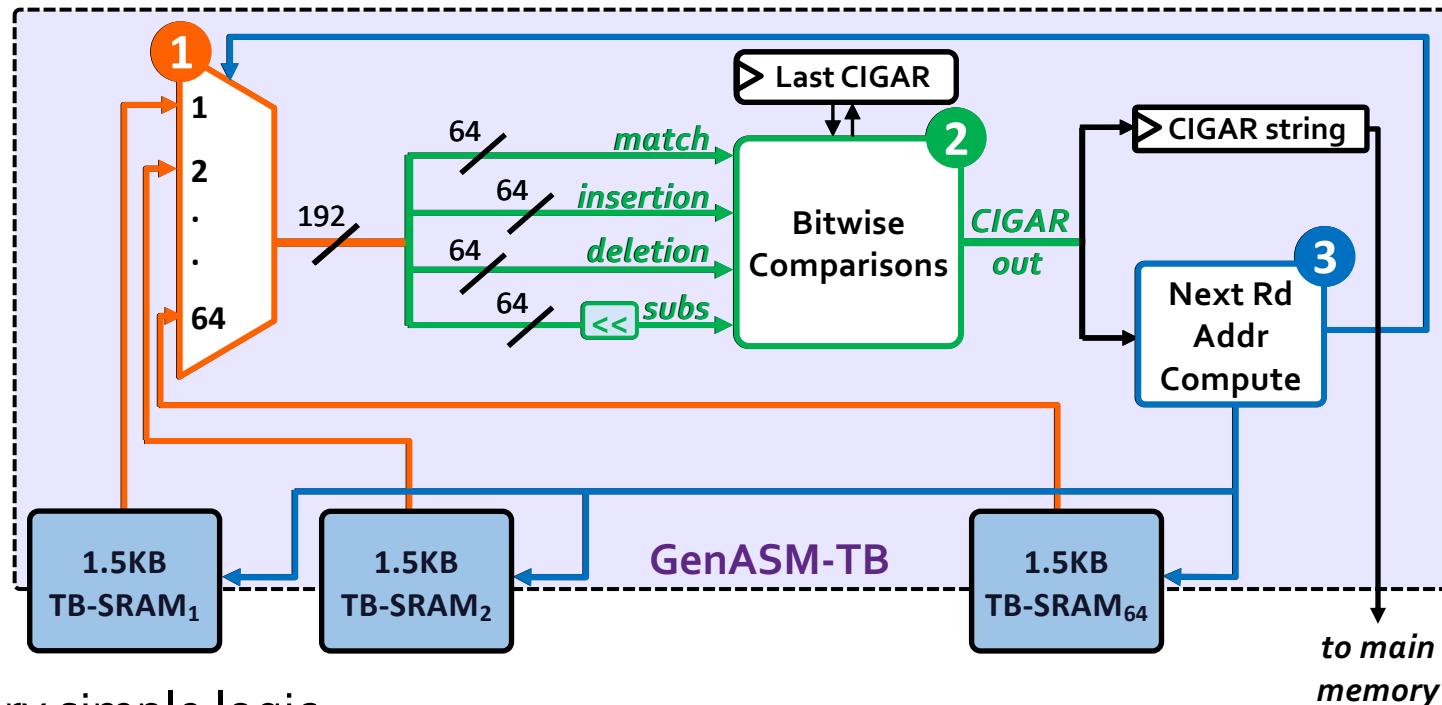
- Match **the rate of computation** with **memory capacity and bandwidth**
 - Achieve **high performance and power efficiency**
 - Scale **linearly in performance** with
the number of parallel compute units that we add to the system

GenASM-DC: Hardware Design

- Linear cyclic systolic array based accelerator
 - Designed to maximize parallelism and minimize memory bandwidth and memory footprint



GenASM-TB: Hardware Design



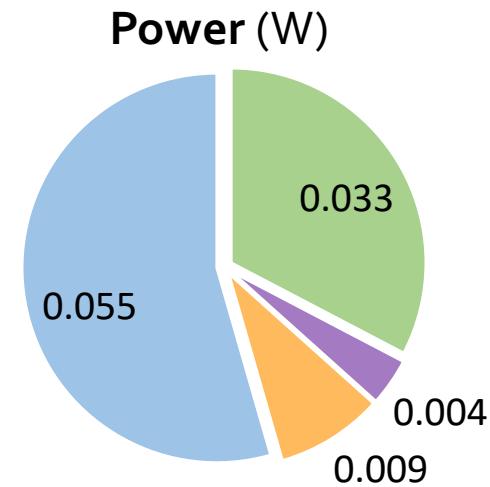
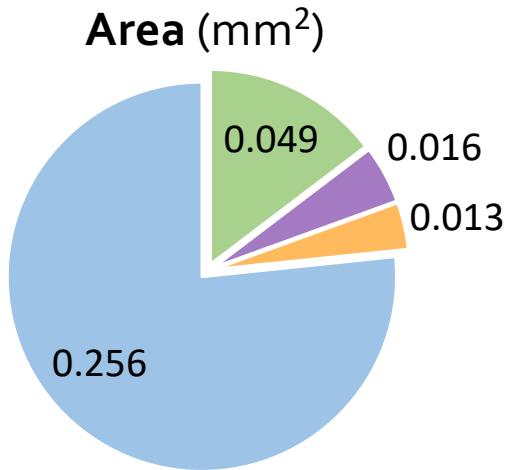
□ Very simple logic:

- ① Reads the bitvectors from one of the TB-SRAMs using the computed address
- ② Performs the required bitwise comparisons to find the traceback output for the current position
- ③ Computes the next TB-SRAM address to read the new set of bitvectors

Key Results – Area and Power

- Based on our **synthesis** of **GenASM-DC** and **GenASM-TB** accelerator datapaths using the Synopsys Design Compiler with a **28nm** LP process:
 - Both GenASM-DC and GenASM-TB operate **@ 1GHz**

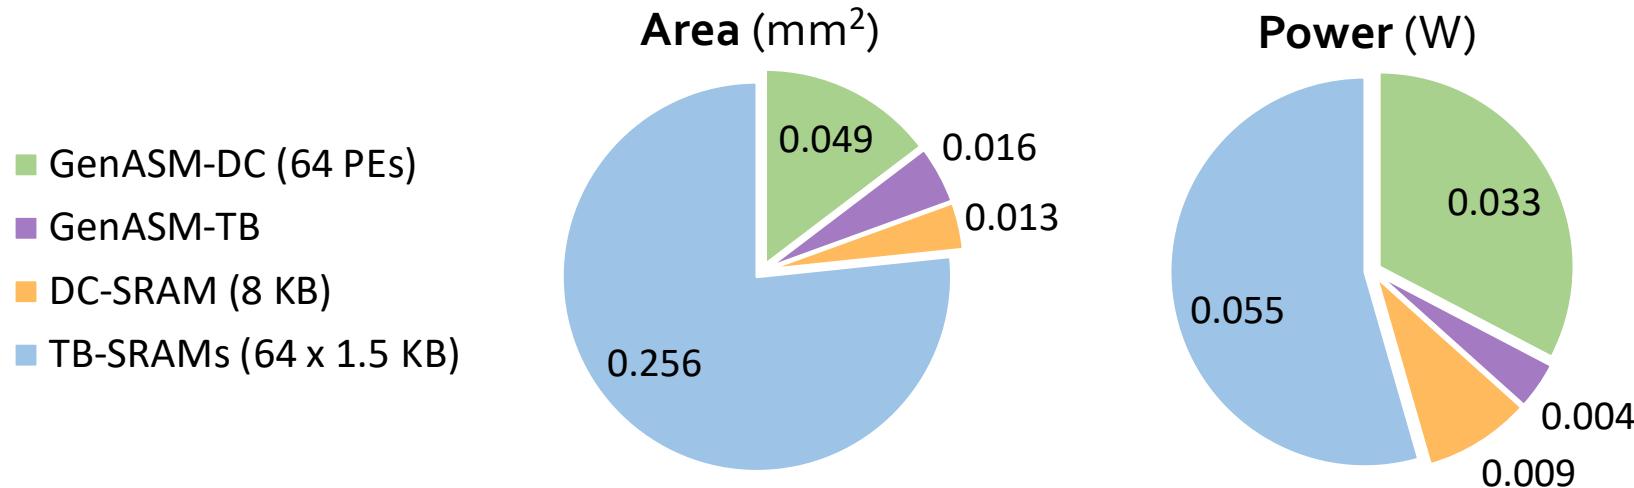
- GenASM-DC (64 PEs)
- GenASM-TB
- DC-SRAM (8 KB)
- TB-SRAMs (64 x 1.5 KB)



Total (1 vault):	0.334 mm ²	0.101 W
Total (32 vaults):	10.69 mm ²	3.23 W
% of a Xeon CPU core:	1%	1%

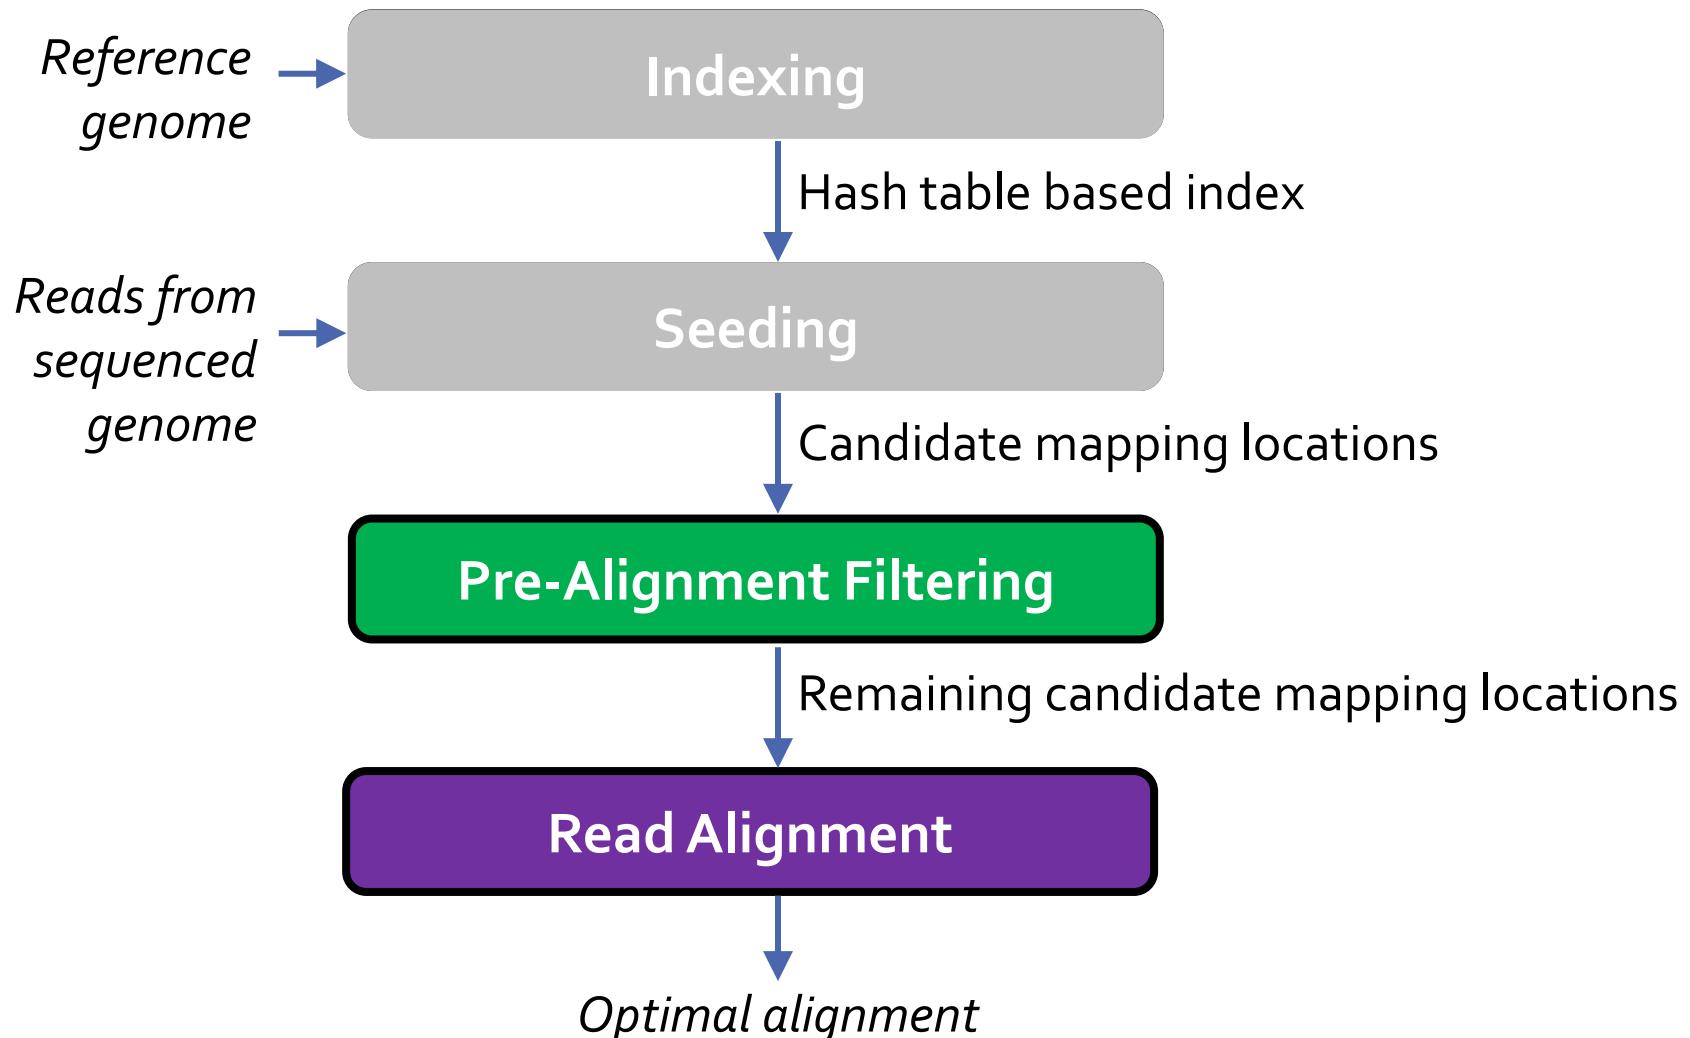
Key Results – Area and Power

- ❑ Based on our **synthesis** of **GenASM-DC** and **GenASM-TB** accelerator datapaths using the Synopsys Design Compiler with a **28nm** LP process:
 - Both GenASM-DC and GenASM-TB operate **@ 1GHz**



GenASM has low area and power overheads

Use Cases of GenASM



Use Cases of GenASM (cont'd.)

(1) Read Alignment Step of Read Mapping

- Find the optimal alignment of how reads map to candidate reference regions

(2) Pre-Alignment Filtering for Short Reads

- Quickly identify and filter out the unlikely candidate reference regions for each read

(3) Edit Distance Calculation

- Measure the similarity or distance between two sequences
- We also discuss other possible use cases of GenASM in our paper:
 - Read-to-read overlap finding, hash-table based indexing, whole genome alignment, generic text search

Key Results

(1) Read Alignment

- ❑ **116×** speedup, **37×** less power than **Minimap2** (state-of-the-art **SW**)
- ❑ **111×** speedup, **33×** less power than **BWA-MEM** (state-of-the-art **SW**)
- ❑ **3.9×** better throughput, **2.7×** less power than **Darwin** (state-of-the-art **HW**)
- ❑ **1.9×** better throughput, **82%** less logic power than **GenAx** (state-of-the-art **HW**)

(2) Pre-Alignment Filtering

- ❑ **3.7×** speedup, **1.7×** less power than **Shouji** (state-of-the-art **HW**)

(3) Edit Distance Calculation

- ❑ **22–12501×** speedup, **548–582×** less power than **Edlib** (state-of-the-art **SW**)
- ❑ **9.3–400×** speedup, **67×** less power than **ASAP** (state-of-the-art **HW**)

More on GenASM Framework [MICRO 2020]

- Damla Senol Cali, Gurpreet S. Kalsi, Zulal Bingol, Can Firtina, Lavanya Subramanian, Jeremie S. Kim, Rachata Ausavarungnirun, Mohammed Alser, Juan Gomez-Luna, Amirali Boroumand, Anant Nori, Allison Scibisz, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,

["GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis"](#)

Proceedings of the 53rd International Symposium on Microarchitecture (MICRO), Virtual, October 2020.

[[Lightning Talk Video](#) (1.5 minutes)]

[[Lightning Talk Slides \(pptx\)](#) ([pdf](#))]

[[Talk Video](#) (18 minutes)]

[[Slides \(pptx\)](#) ([pdf](#))]

GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis

Damla Senol Cali^{†✉} Gurpreet S. Kalsi[✉] Zülal Bingöl[▽] Can Firtina[◊] Lavanya Subramanian[‡] Jeremie S. Kim^{◊†}
Rachata Ausavarungnirun[○] Mohammed Alser[◊] Juan Gomez-Luna[◊] Amirali Boroumand[†] Anant Nori[✉]
Allison Scibisz[†] Sreenivas Subramoney[✉] Can Alkan[▽] Saugata Ghose^{★†} Onur Mutlu^{◊†▽}

[†]*Carnegie Mellon University* [✉]*Processor Architecture Research Lab, Intel Labs* [▽]*Bilkent University* [◊]*ETH Zürich*

[‡]*Facebook* [○]*King Mongkut's University of Technology North Bangkok* [★]*University of Illinois at Urbana-Champaign*

Accelerating Sequence-to-Graph Mapping

- Damla Senol Cali, Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo, Can Firtina, Meryem Banu Cavlak, Jeremie Kim, Nika MansouriGhiasi, Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,

"SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping"

Proceedings of the 49th International Symposium on Computer Architecture (ISCA), New York, June 2022.

[[arXiv version](#)]

SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping

Damla Senol Cali¹ Konstantinos Kanellopoulos² Joël Lindegger² Zülal Bingöl³
Gurpreet S. Kalsi⁴ Ziyi Zuo⁵ Can Firtina² Meryem Banu Cavlak² Jeremie Kim²
Nika Mansouri Ghiasi² Gagandeep Singh² Juan Gómez-Luna² Nour Almadhoun Alserr²
Mohammed Alser² Sreenivas Subramoney⁴ Can Alkan³ Saugata Ghose⁶ Onur Mutlu²

¹Bionano Genomics ²ETH Zürich ³Bilkent University ⁴Intel Labs

⁵Carnegie Mellon University ⁶University of Illinois Urbana-Champaign

SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping

Damla Senol Cali, Ph.D.

damlasenolcali@gmail.com

<https://damlasenolcali.github.io/>

Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo,
Can Firtina, Meryem Banu Cavlak, Jeremie S. Kim, Nika Mansouri Ghiasi,
Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser,
Sreenivas Subramoney, Can Alkan, Saugata Ghose, Onur Mutlu

Carnegie Mellon

ETH zürich



Bilkent University

intel®

UNIVERSITY OF
ILLINOIS
URBANA-CHAMPAIGN

SAFARI

Genome Sequence Analysis

- Mapping the reads to a reference genome (i.e., *read mapping*) is a *critical step* in genome sequence analysis

Linear Reference: ACGTACGT

Read: ACGG

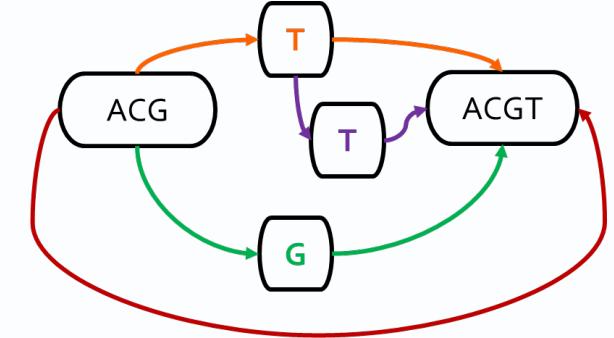
Alternative Sequence: ACGGACGT

Alternative Sequence: ACGTTACGT

Alternative Sequence: ACG–ACGT

Sequence-to-Sequence (S2S) Mapping

Graph-based Reference:



Read: ACGG

Sequence-to-Graph (S2G) Mapping

Sequence-to-graph mapping results in **notable quality improvements**.

However, it is a **more difficult** computational problem,
with **no prior hardware design**.

SeGraM: First Graph Mapping Accelerator

Our Goal:

Specialized, high-performance, scalable, and low-cost algorithm/hardware co-design that alleviates bottlenecks in **multiple steps** of sequence-to-graph mapping

SeGraM: First universal algorithm/hardware co-designed genomic mapping accelerator that can effectively and efficiently support:

- Sequence-to-graph mapping
- Sequence-to-sequence mapping
- Both short and long reads

Use Cases & Key Results

(1) Sequence-to-Graph (S₂G) Mapping

- ❑ **5.9×/106×** speedup, **4.1×/3.0×** less power than **GraphAligner**
for long and short reads, respectively (state-of-the-art **SW**)
- ❑ **3.9×/742×** speedup, **4.4×/3.2×** less power than **vg**
for long and short reads, respectively (state-of-the-art **SW**)

(2) Sequence-to-Graph (S₂G) Alignment

- ❑ **41×–539×** speedup over **PaSGAL** with AVX-512 support (state-of-the-art **SW**)

(3) Sequence-to-Sequence (S₂S) Alignment

- ❑ **1.2×/4.8×** higher throughput than **GenASM** and **GACT** of **Darwin**
for long reads (state-of-the-art **HW**)
- ❑ **1.3×/2.4×** higher throughput than **GenASM** and **SillaX** of **GenAX**
for short reads (state-of-the-art **HW**)

SeGram Talk Video

Recording

Sequence-to-Graph Mapping Pipeline

Linear reference genome 0.1 **Genome Graph Construction**
(construct the graph using a linear reference genome and variations)

Known genetic variations 0.2 **Indexing**
(index the nodes of the graph)

Hash-table-based index (of graph nodes)

Reads from sequenced genome 1 **Seeding**
(query the index & find the seed matches)

Candidate mapping locations (subgraphs)

2 **Filtering/Chaining/Clustering**
(filter out dissimilar query read and subgraph pairs)

Remaining candidate mapping locations (subgraphs)

3 **S2G Alignment**
(perform distance/score calculation & traceback)

Optimal alignment between read & subgraph

Pre-Processing Steps (Offline)

Seed-and-Extend Steps (Online)

Damla Senol Cali SAFARI 14

5:56 / 21:29

SeGram: A Universal HW Accelerator for Genomic Sequence-to-Graph Mapping - Damla Senol Cali (ISCA)

136 views • Premiered 21 hours ago

12

DISLIKE

SHARE

DOWNLOAD

CLIP SAVE

...



Onur Mutlu Lectures
26.9K subscribers

ANALYTICS

EDIT VIDEO

Accelerating Sequence-to-Graph Mapping

- Damla Senol Cali, Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo, Can Firtina, Meryem Banu Cavlak, Jeremie Kim, Nika MansouriGhiasi, Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,

"SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping"

Proceedings of the 49th International Symposium on Computer Architecture (ISCA), New York, June 2022.

[[arXiv version](#)]

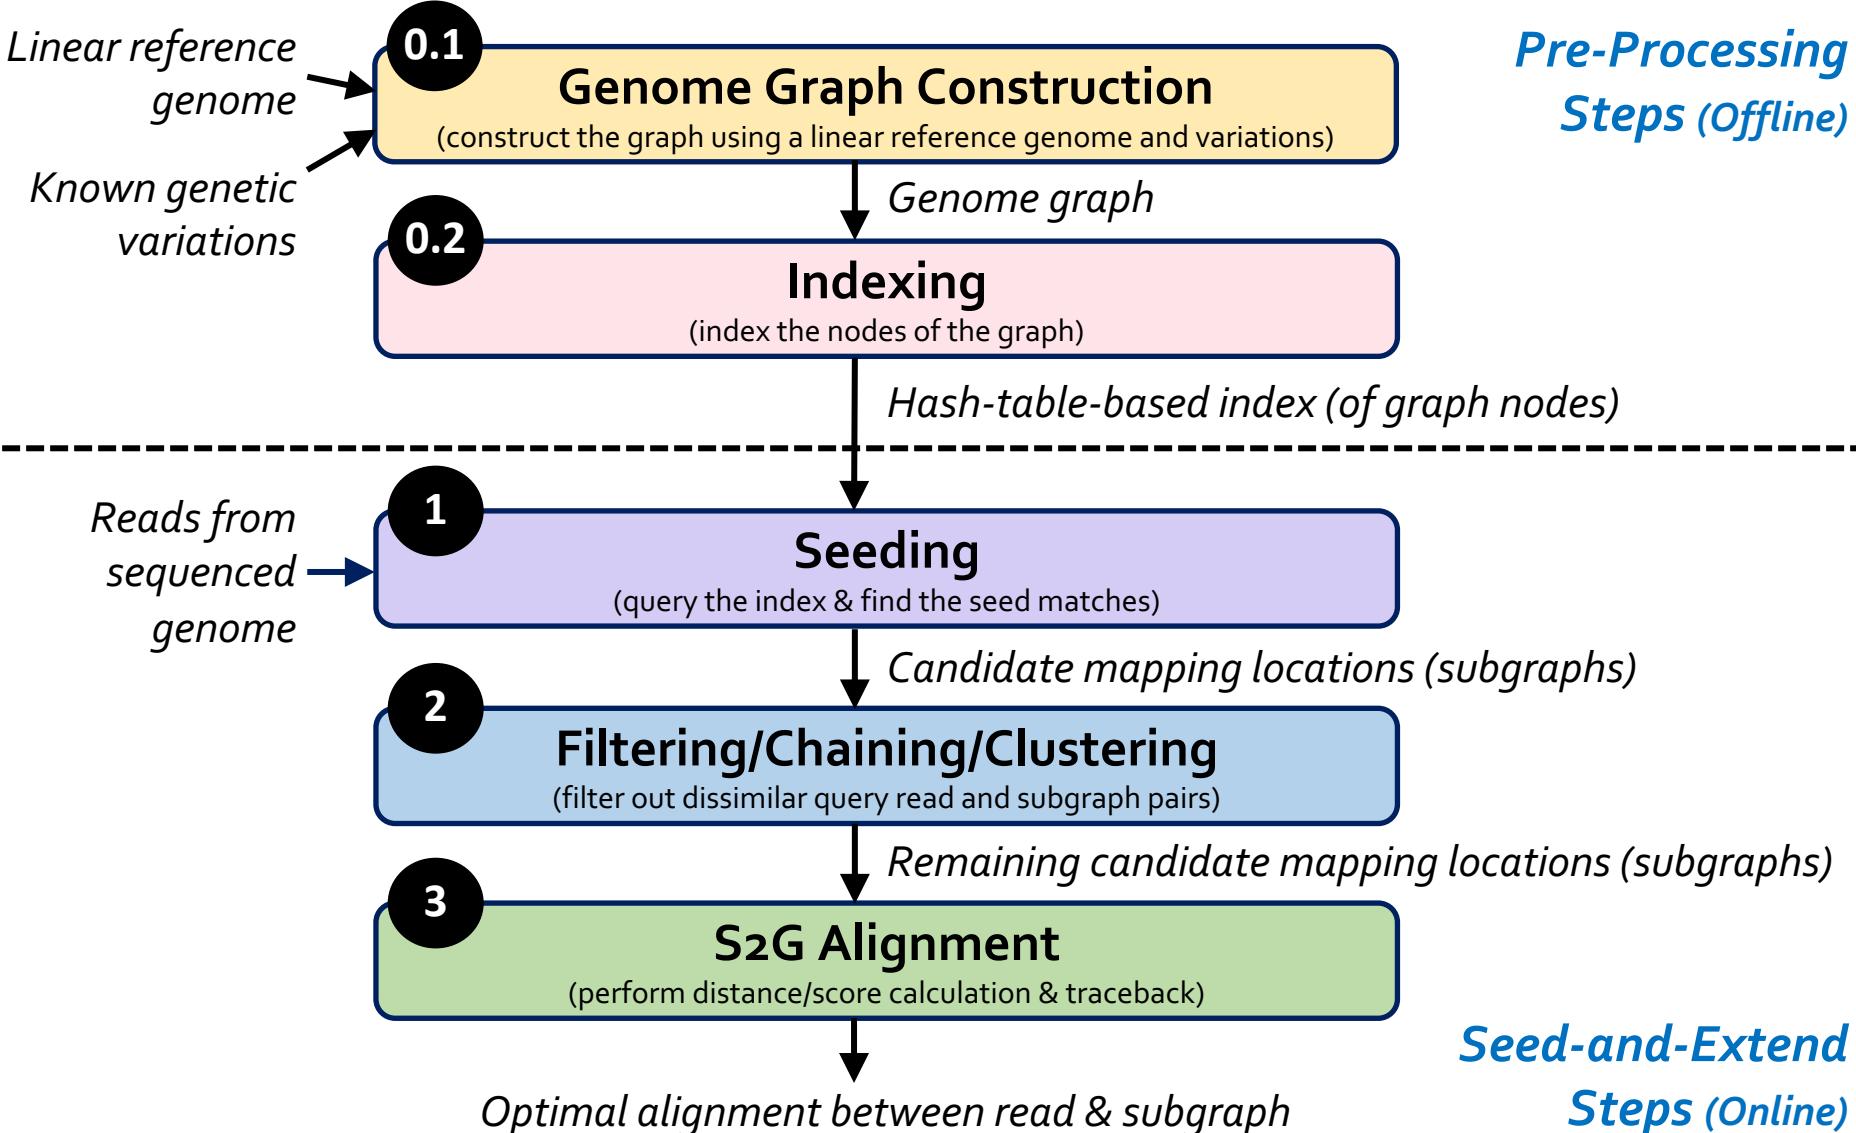
SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping

Damla Senol Cali¹ Konstantinos Kanellopoulos² Joël Lindegger² Zülal Bingöl³
Gurpreet S. Kalsi⁴ Ziyi Zuo⁵ Can Firtina² Meryem Banu Cavlak² Jeremie Kim²
Nika Mansouri Ghiasi² Gagandeep Singh² Juan Gómez-Luna² Nour Almadhoun Alserr²
Mohammed Alser² Sreenivas Subramoney⁴ Can Alkan³ Saugata Ghose⁶ Onur Mutlu²

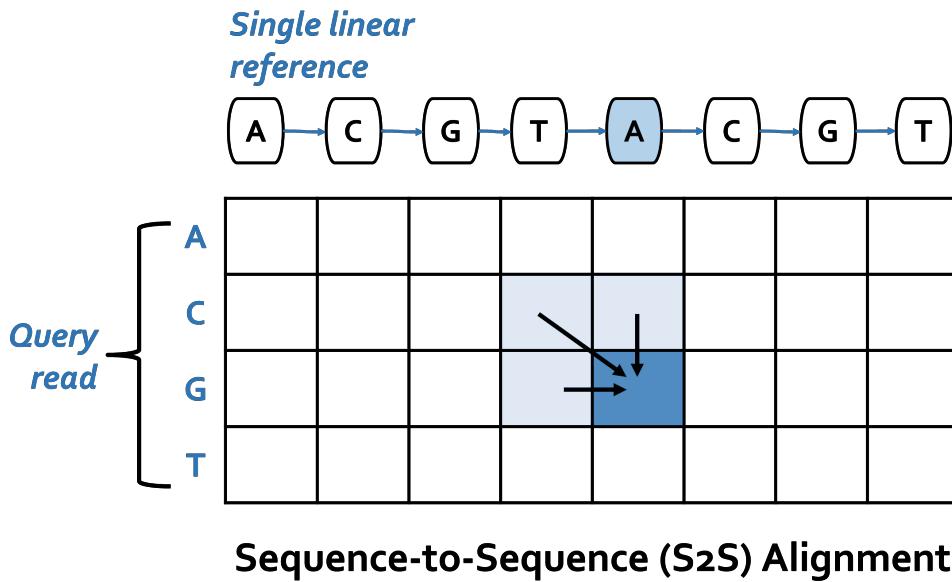
¹Bionano Genomics ²ETH Zürich ³Bilkent University ⁴Intel Labs

⁵Carnegie Mellon University ⁶University of Illinois Urbana-Champaign

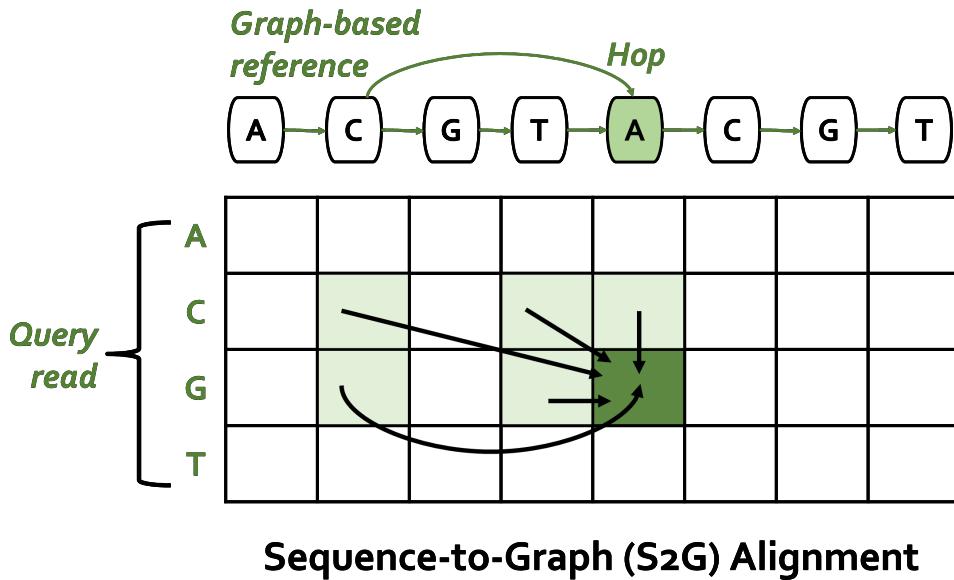
Sequence-to-Graph Mapping Pipeline



S₂S vs. S₂G Alignment



S₂S vs. S₂G Alignment



In contrast to S₂S alignment,
S₂G alignment must incorporate **non-neighboring characters**
as well whenever there is an edge (i.e., **hop**)
from the non-neighboring character to the current character

Analysis of State-of-the-Art Tools

Based on our analysis with **GraphAligner** and **vg**:

SW

Observation 1: Alignment step is the bottleneck

Observation 2: Alignment suffers from high cache miss rates

Observation 3: Seeding suffers from the DRAM latency bottleneck

Observation 4: Baseline tools scale sublinearly

Observation 5: Existing S2S mapping accelerators are unsuitable
for the S2G mapping problem

HW

Observation 6: Existing graph accelerators are unable to handle
S2G alignment

SeGraM: Universal Genomic Mapping Accelerator

- ***First universal genomic mapping accelerator*** that can support *both* sequence-to-graph mapping and sequence-to-sequence mapping, for *both* short and long reads
- ***First algorithm/hardware co-design*** for accelerating sequence-to-graph mapping

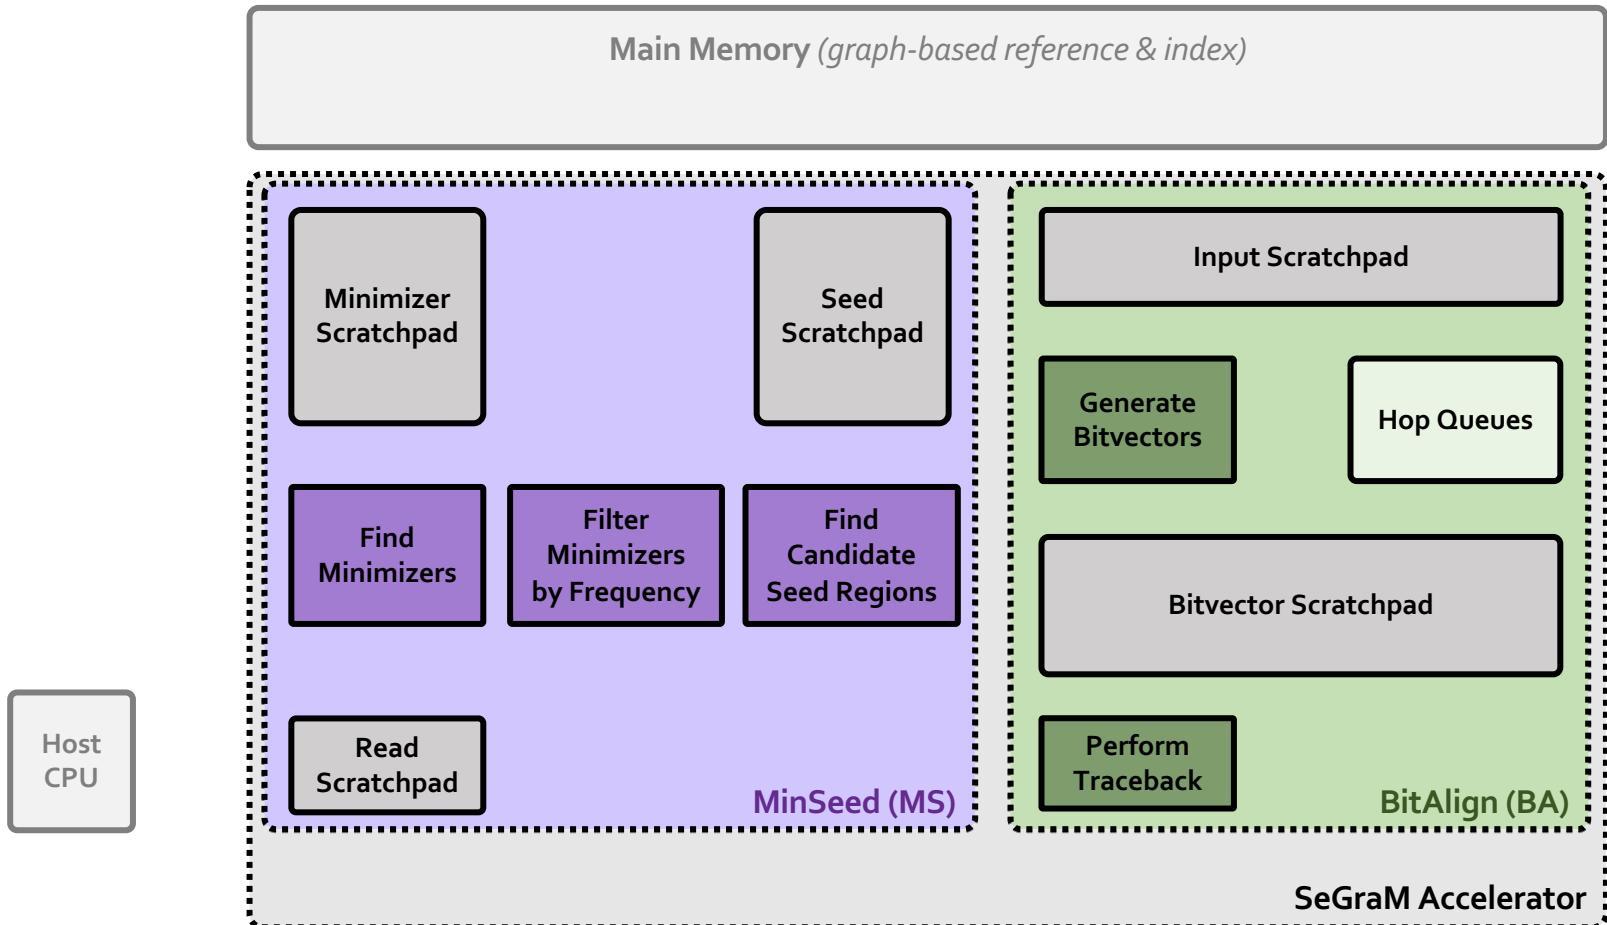
- We base SeGraM upon a **minimizer-based seeding algorithm**
- We propose a **novel bitvector-based alignment algorithm** to perform approximate string matching between a read and a graph-based reference genome

SW

- We co-design both algorithms with **high-performance, scalable, and efficient hardware accelerators**

HW

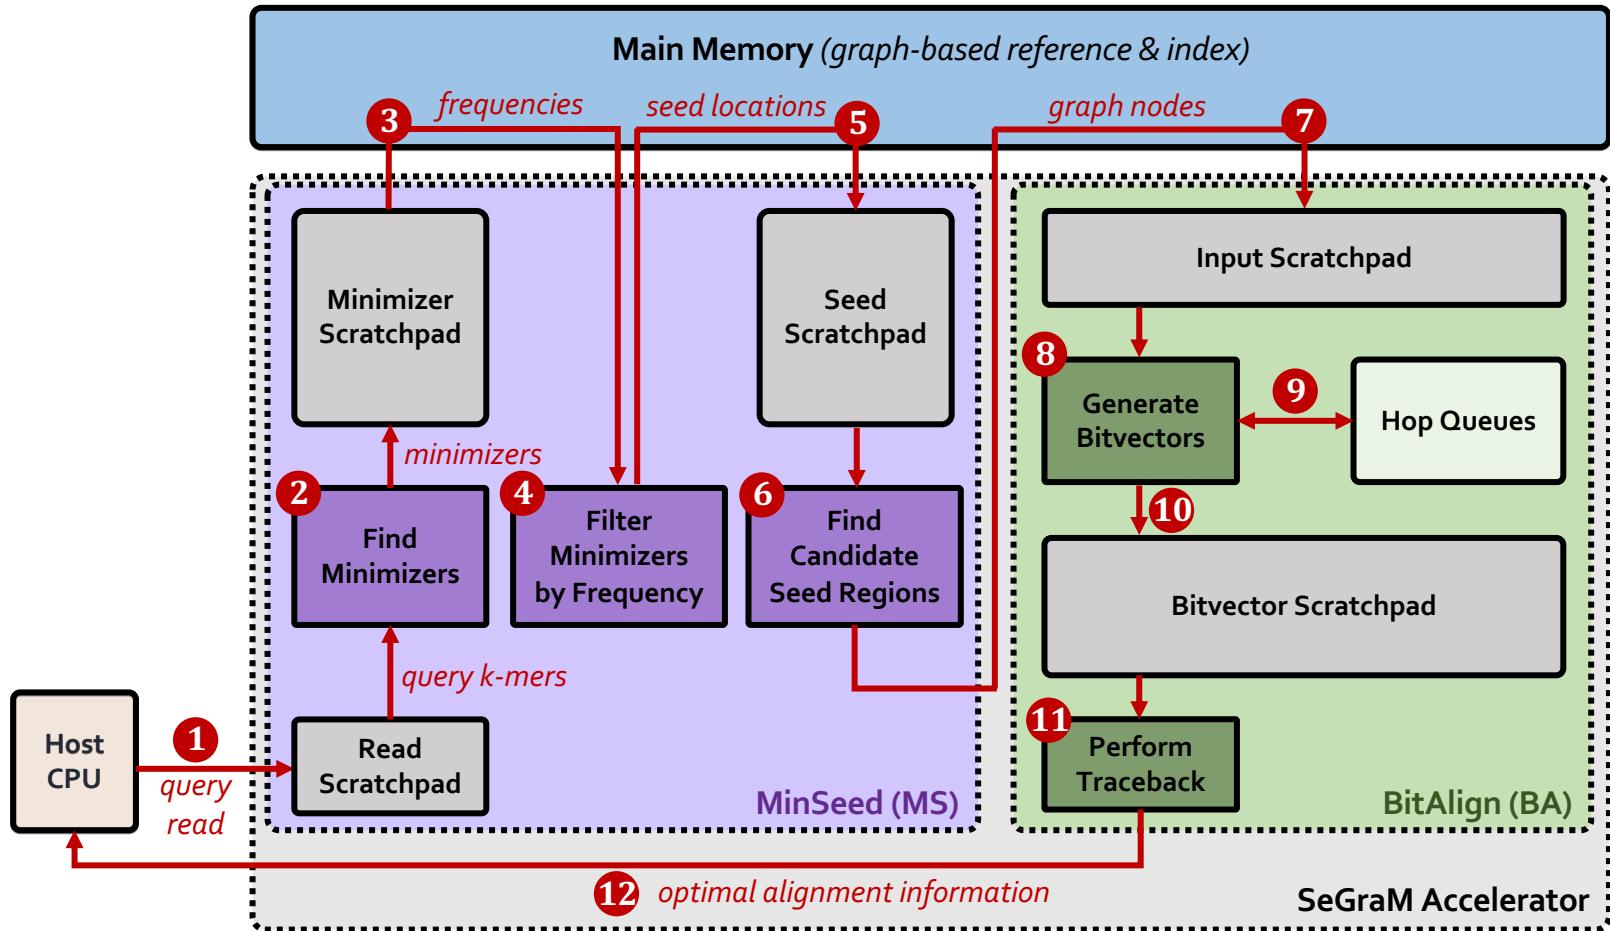
SeGraM Hardware Design



MinSeed: first hardware accelerator for **Minimizer-based Seeding**

BitAlign: first hardware accelerator for (**Bitvector-based**) sequence-to-graph **Alignment**

SeGram Hardware Design

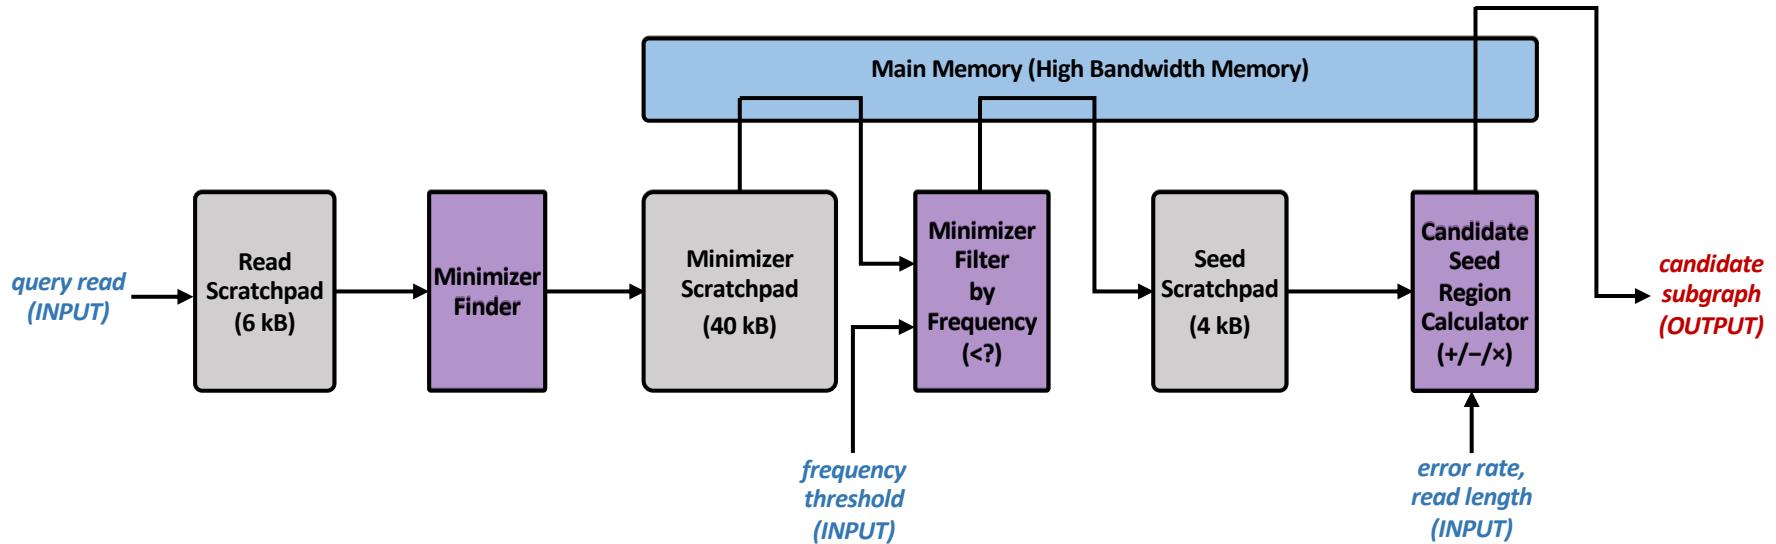


MinSeed: first hardware accelerator for **Minimizer-based Seeding**

BitAlign: first hardware accelerator for (**Bitvector-based**) sequence-to-graph **Alignment**

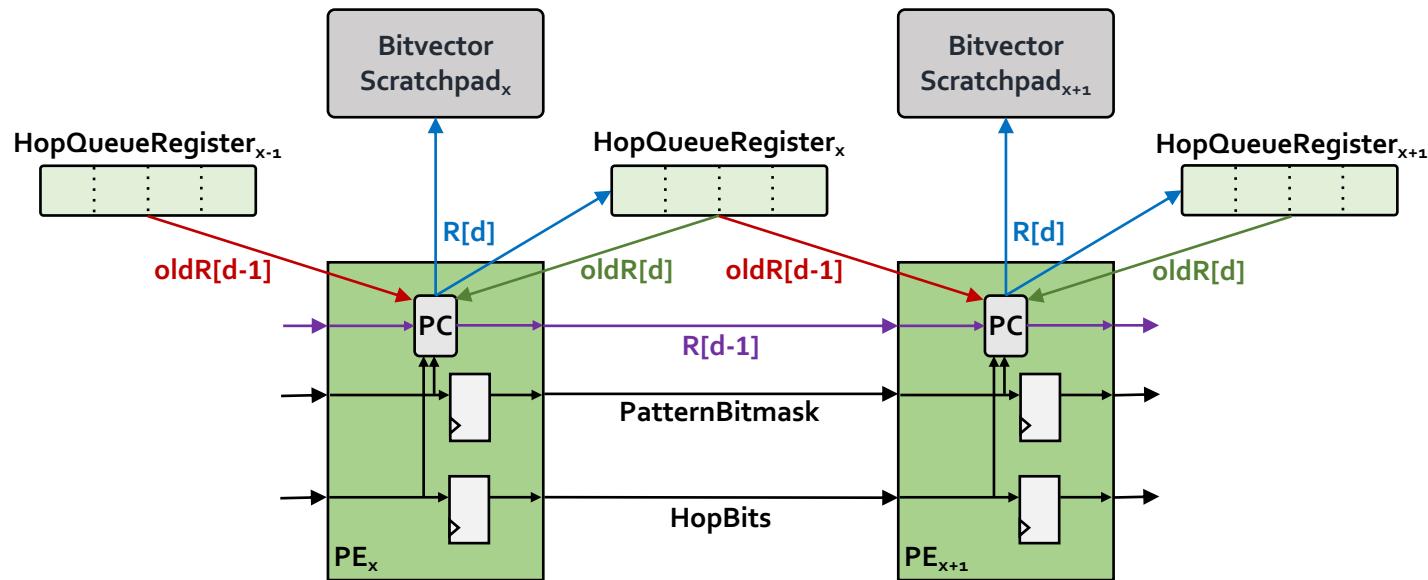
MinSeed HW

- MinSeed = 3 computation modules + 3 scratchpads + memory interface
 - Computation modules: Implemented with simple logic
 - Scratchpads: 50kB in total; employ double buffering technique to hide the latency of MinSeed
 - High-Bandwidth Memory (HBM): Enables low-latency and highly-parallel memory access



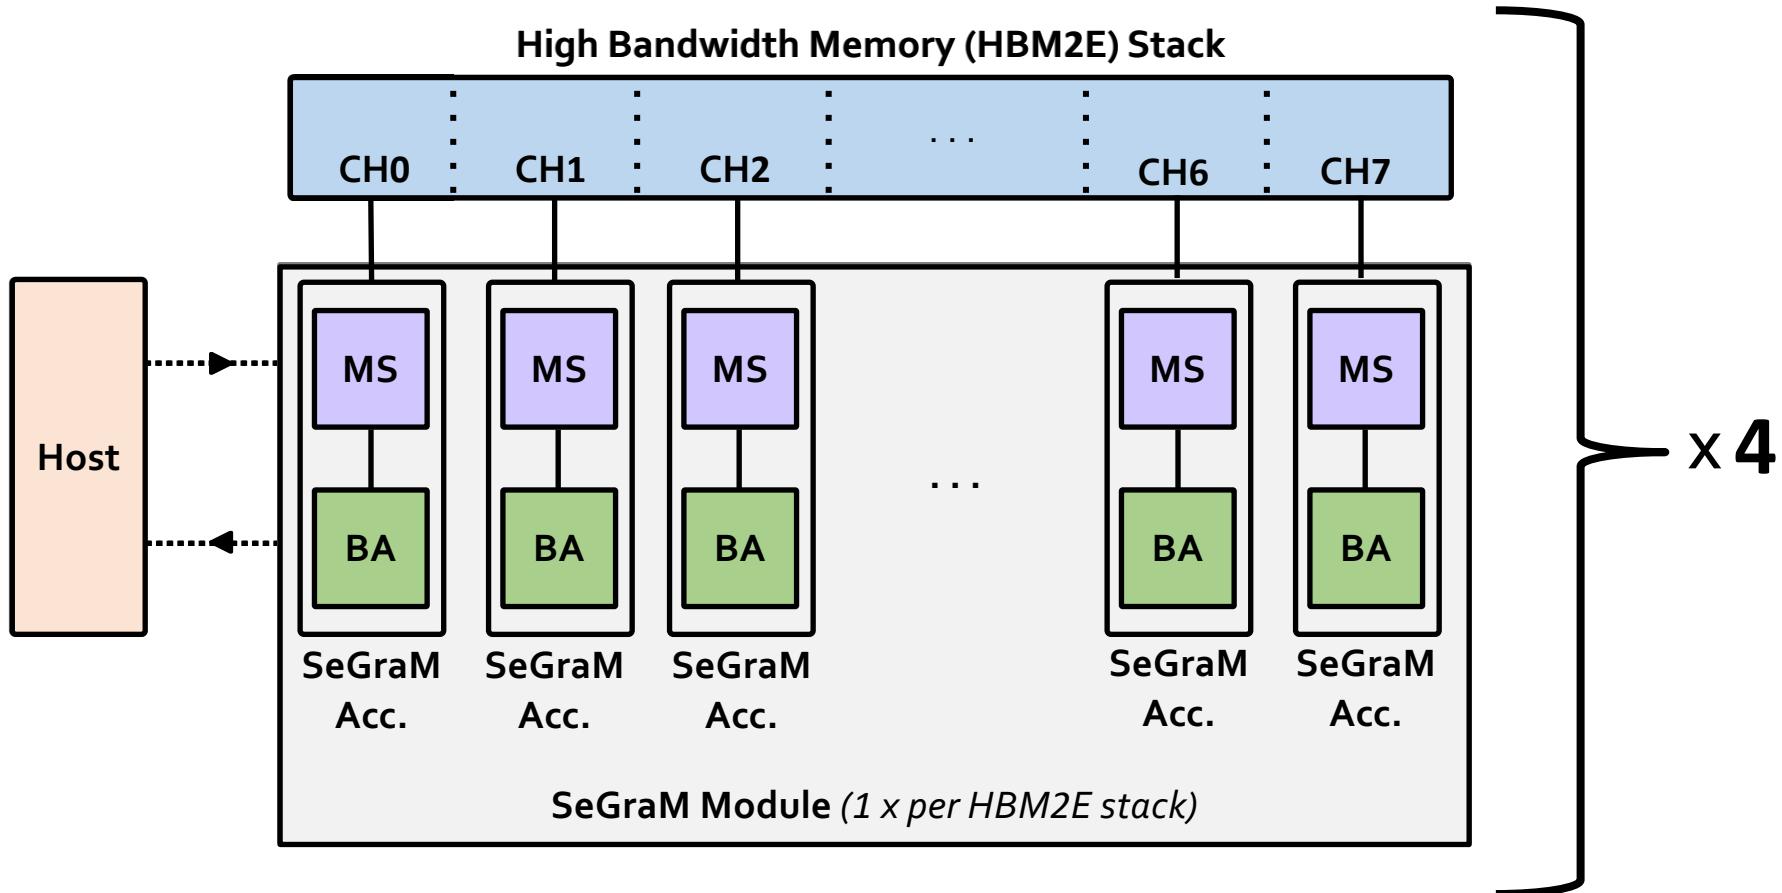
BitAlign HW

- ❑ Linear cyclic systolic array-based accelerator
- ❑ Based on the GenASM hardware design*
- ❑ Incorporates *hop queue registers* to feed the bitvectors of non-neighboring characters/nodes (i.e., *hops*)



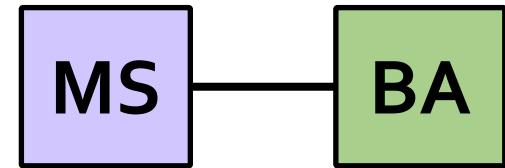
[*] D. Senol Cali *et al.* "[GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis](#)" (MICRO'20)

Overall System Design of SeGram

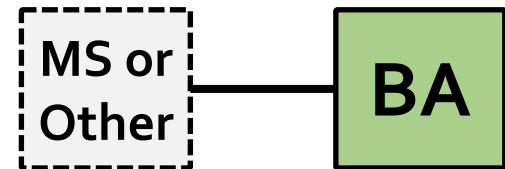


Use Cases of SeGram

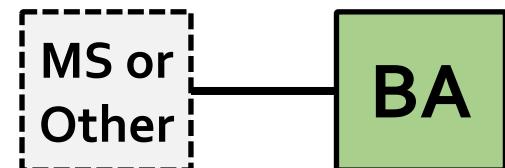
(1) Sequence-to-Graph
Mapping



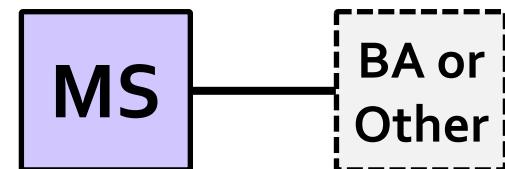
(2) Sequence-to-Graph
Alignment



(3) Sequence-to-Sequence
Alignment



(4) Seeding

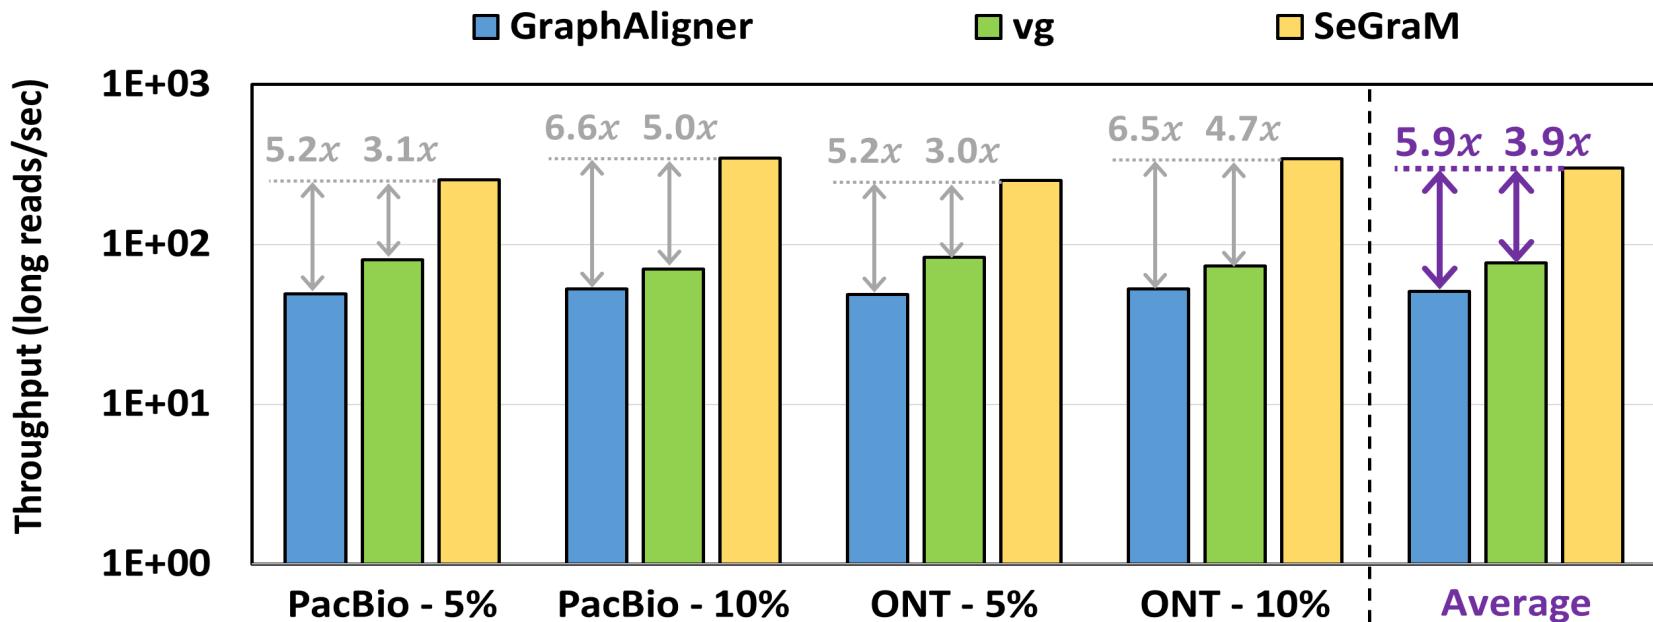


Key Results – Area & Power

- Based on our **synthesis** of **MinSeed** and **BitAlign** accelerator datapaths using the Synopsys Design Compiler with a **28nm** process (**@ 1GHz**):

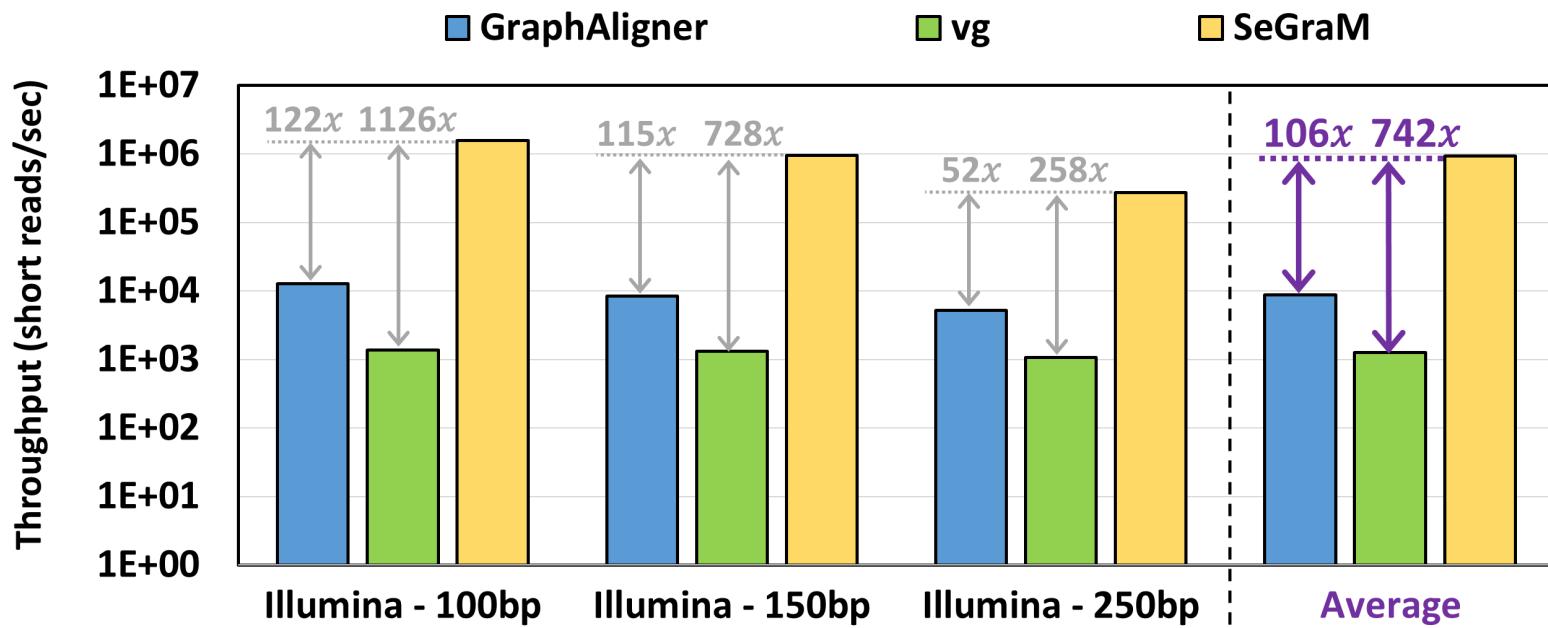
Component	Area (mm ²)	Power (mW)
MinSeed – Logic	0.017	10.8
Read Scratchpad (6 kB)	0.012	7.9
Minimizer Scratchpad (40 kB)	0.055	22.7
Seed Scratchpad (4 kB)	0.008	6.4
BitAlign – Edit Distance Calculation Logic with Hop Queue Registers (64 PEs)	0.393	378.0
BitAlign – Traceback Logic	0.020	2.7
Input Scratchpad (24 kB)	0.033	13.3
Bitvector Scratchpads (128 kB)	0.329	316.2
Total – 1 SeGram Accelerator	0.867	758.0 (0.8 W)
Total – 4 SeGram Modules (32 SeGram Accelerators)	27.744	24.3 W
HBM2E (4 stacks)	--	3.8 W

Key Results – SeGraM with Long Reads



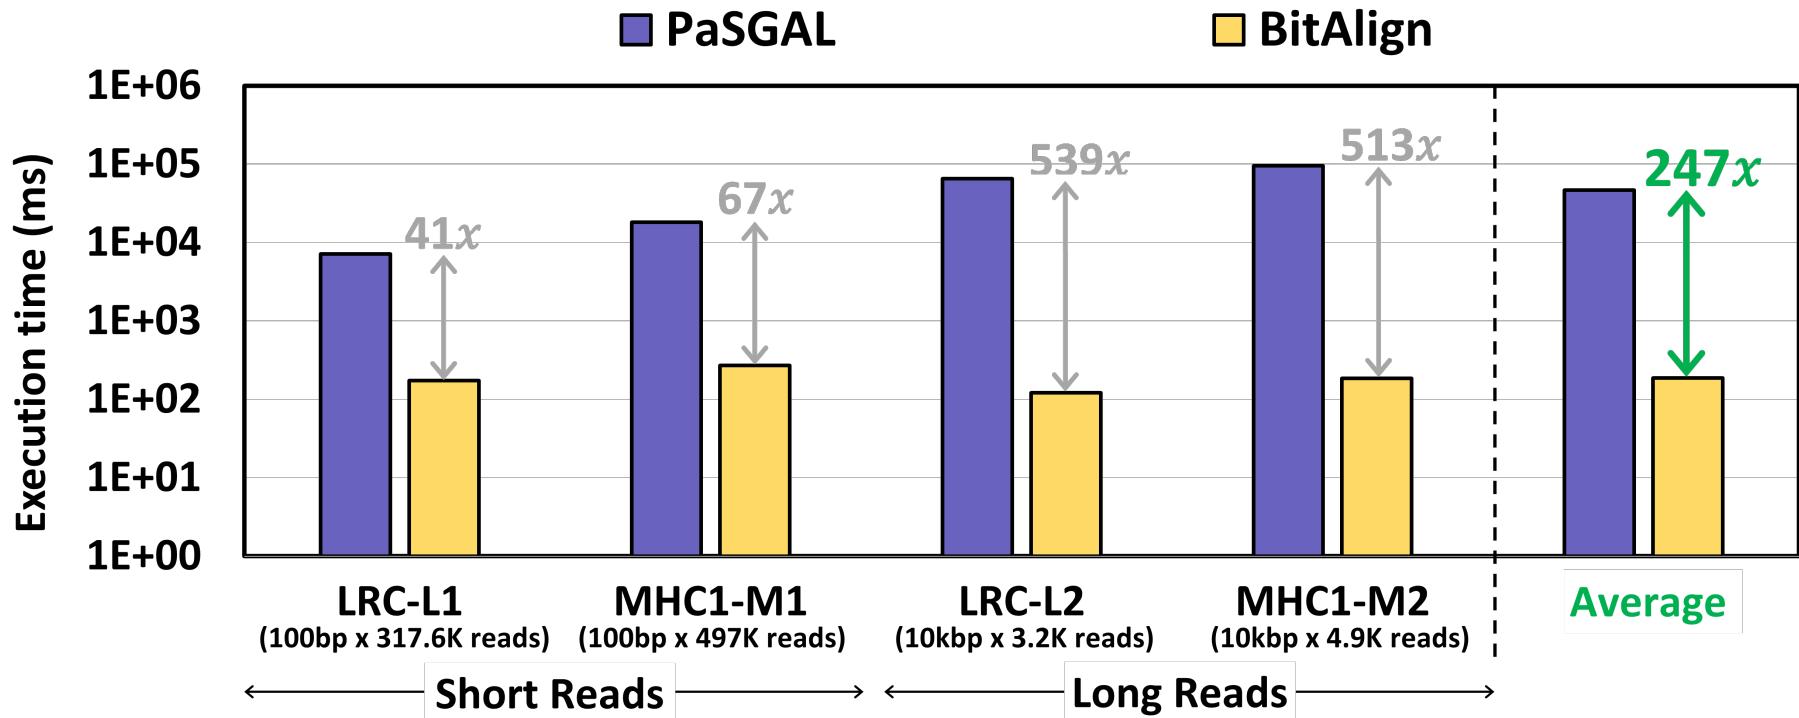
SeGraM provides **5.9x and 3.9x throughput improvement**
over GraphAligner and vg,
while **reducing the power consumption by 4.1x and 4.4x**

Key Results – SeGraM with Short Reads



SeGraM provides **106x and 742x throughput improvement** over GraphAligner and vg,
while **reducing the power consumption by 3.0x and 3.2x**

Key Results – BitAlign (S2G Alignment)



BitAlign provides **41x-539x speedup** over PaSGAL

Conclusion

- **SeGraM:** *First universal algorithm/hardware co-designed genomic mapping accelerator that supports:*
 - Sequence-to-graph (S₂G) & sequence-to-sequence (S₂S) mapping
 - Short & long reads
- **MinSeed:** *First minimizer-based seeding accelerator*
- **BitAlign:** *First (bitvector-based) S₂G alignment accelerator*
- SeGraM **supports multiple use cases:**
 - End-to-end S₂G mapping
 - S₂G alignment
 - S₂S alignment
 - Seeding
- SeGraM **outperforms state-of-the-art software & hardware solutions**

Accelerating Sequence-to-Graph Mapping

- Damla Senol Cali, Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo, Can Firtina, Meryem Banu Cavlak, Jeremie Kim, Nika MansouriGhiasi, Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,

"SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping"

Proceedings of the 49th International Symposium on Computer Architecture (ISCA), New York, June 2022.

[[arXiv version](#)]

SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping

Damla Senol Cali¹ Konstantinos Kanellopoulos² Joël Lindegger² Zülal Bingöl³
Gurpreet S. Kalsi⁴ Ziyi Zuo⁵ Can Firtina² Meryem Banu Cavlak² Jeremie Kim²
Nika Mansouri Ghiasi² Gagandeep Singh² Juan Gómez-Luna² Nour Almadhoun Alserr²
Mohammed Alser² Sreenivas Subramoney⁴ Can Alkan³ Saugata Ghose⁶ Onur Mutlu²

¹Bionano Genomics ²ETH Zürich ³Bilkent University ⁴Intel Labs

⁵Carnegie Mellon University ⁶University of Illinois Urbana-Champaign

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

Read Mapping & Filtering

- Problem: Heavily bottlenecked by Data Movement
- GateKeeper, Shouji, SneakySnake performance limited by DRAM bandwidth [Alser+, Bioinformatics 2017,2019,2020]
- Ditto for SHD [Xin+, Bioinformatics 2015]
- Solution: Processing-in-memory can alleviate the bottleneck
- We need to design mapping & filtering algorithms to fit processing-in-memory

Read Mapping & Filtering in Memory

We need to design
mapping & filtering algorithms
that fit processing-in-memory

Near-Memory Pre-Alignment Filtering

Gagandeep Singh, Mohammed Alser, Damla Senol Cali, Dionysios Diamantopoulos, Juan Gomez-Luna, Henk Corporaal, Onur Mutlu,

"FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications"

IEEE Micro, 2021.

[[Source Code](#)]



[Home](#) / [Magazines](#) / [IEEE Micro](#) / [2021.04](#)

IEEE Micro

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41

DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)

Authors

[Gagandeep Singh](#), ETH Zürich, Zürich, Switzerland

[Mohammed Alser](#), ETH Zürich, Zürich, Switzerland

[Damla Senol Cali](#), Carnegie Mellon University, Pittsburgh, PA, USA

[Dionysios Diamantopoulos](#), Zürich Lab, IBM Research Europe, Rüschlikon, Switzerland

[Juan Gomez-Luna](#), ETH Zürich, Zürich, Switzerland

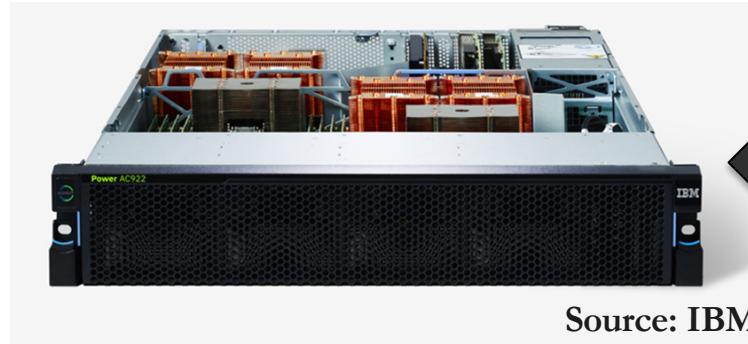
[Henk Corporaal](#), Eindhoven University of Technology, Eindhoven, The Netherlands

[Onur Mutlu](#), ETH Zürich, Zürich, Switzerland

Near-Memory SneakySnake

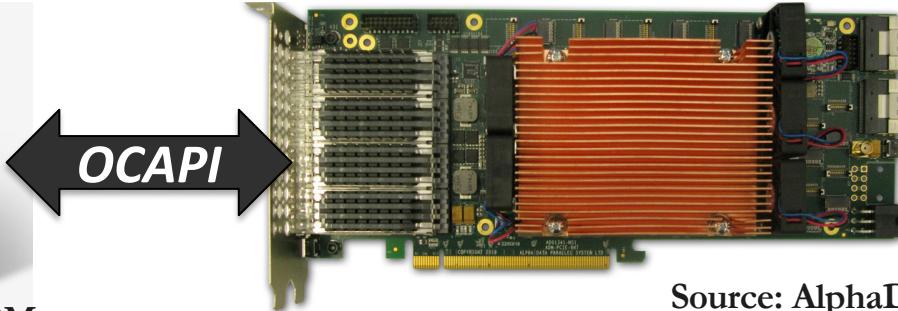
- Problem: Read mapping is heavily bottlenecked by data movement from main memory
- Solution: Perform read mapping near where data resides using specialized logic
- We carefully redesign the accelerator logic of SneakySnake to exploit near-memory computation capability on real FPGA boards that use HBM (high-bandwidth memory)
- **Near-memory SneakySnake** improves **performance** and **energy efficiency** by 27.4 \times and 133 \times , respectively, over a 16-core (64-thread) IBM POWER9 CPU

Near-Memory Acceleration using FPGAs



IBM POWER9 CPU

Source: IBM



HBM-based FPGA board

Source: AlphaData

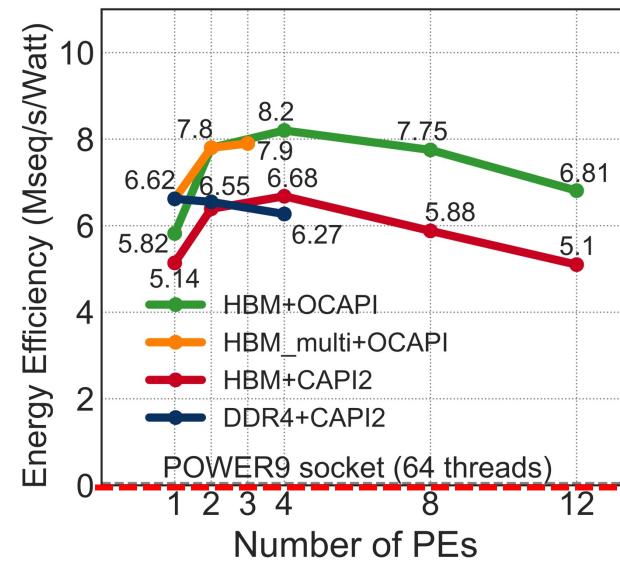
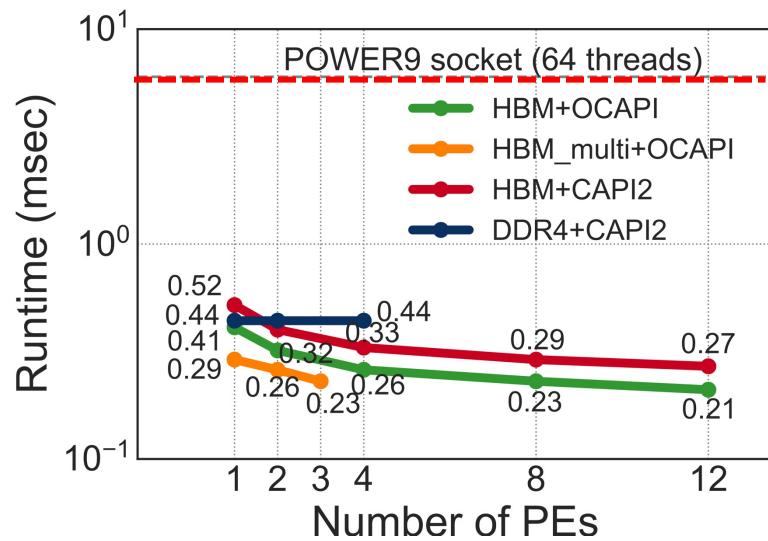
Near-HBM FPGA-based accelerator

Two communication technologies: CAPI2 and OCAPI

Two memory technologies: DDR4 and HBM

Two workloads: Weather Modeling and Genome Analysis

Performance & Energy Greatly Improve



5-27x performance vs. a 16-core (64-thread) IBM POWER9 CPU

12-133x energy efficiency vs. a 16-core (64-thread) IBM POWER9 CPU

HBM alleviates memory bandwidth contention vs. DDR4

More On Near-Memory SneakySnake

Gagandeep Singh, Mohammed Alser, Damla Senol Cali, Dionysios Diamantopoulos, Juan Gomez-Luna, Henk Corporaal, Onur Mutlu,

"FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications"

IEEE Micro, 2021.

[[Source Code](#)]



[Home](#) / [Magazines](#) / [IEEE Micro](#) / [2021.04](#)

IEEE Micro

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41

DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)

Authors

[Gagandeep Singh](#), ETH Zürich, Zürich, Switzerland

[Mohammed Alser](#), ETH Zürich, Zürich, Switzerland

[Damla Senol Cali](#), Carnegie Mellon University, Pittsburgh, PA, USA

[Dionysios Diamantopoulos](#), Zürich Lab, IBM Research Europe, Rüschlikon, Switzerland

[Juan Gomez-Luna](#), ETH Zürich, Zürich, Switzerland

[Henk Corporaal](#), Eindhoven University of Technology, Eindhoven, The Netherlands

[Onur Mutlu](#), ETH Zürich, Zürich, Switzerland

Location Filtering in 3D-Stacked PIM

- Jeremie S. Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, and Onur Mutlu,
"GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies"

BMC Genomics, 2018.

Proceedings of the 16th Asia Pacific Bioinformatics Conference (APBC), Yokohama, Japan, January 2018.

[[Slides \(pptx\)](#) ([pdf](#))]

[[Source Code](#)]

[[arxiv.org Version \(pdf\)](#)]

[[Talk Video at AACBB 2019](#)]

Research | [Open Access](#) | Published: 09 May 2018

GRIM-Filter: Fast seed location filtering in DNA read mapping using processing-in-memory technologies

[Jeremie S. Kim](#) , [Damla Senol Cali](#), [Hongyi Xin](#), [Donghyuk Lee](#), [Saugata Ghose](#), [Mohammed Alser](#), [Hasan Hassan](#), [Oguz Ergin](#), [Can Alkan](#)  & [Onur Mutlu](#) 

[BMC Genomics](#) **19**, Article number: 89 (2018) | [Cite this article](#)

4340 Accesses | 39 Citations | 9 Altmetric | [Metrics](#)

GRIM-Filter

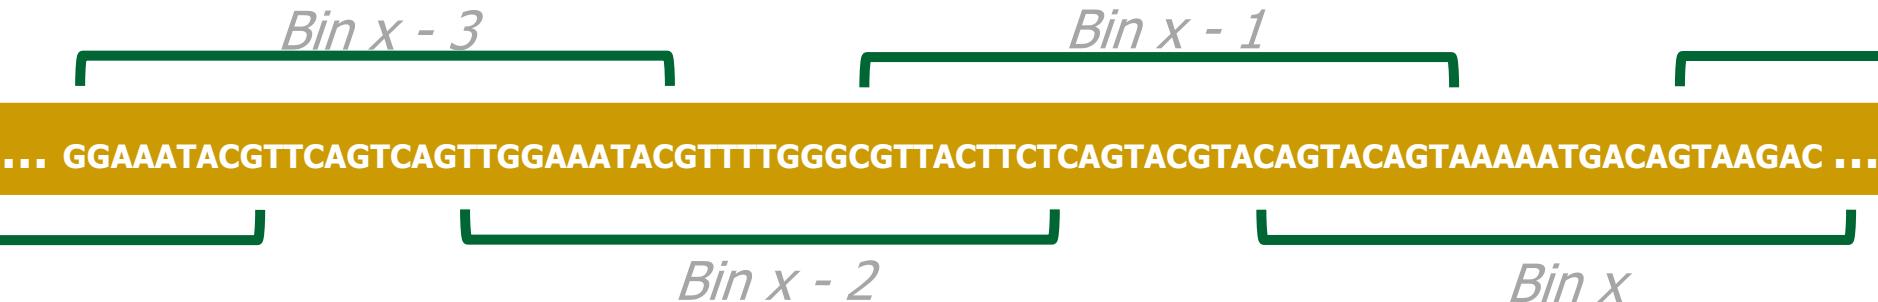
- **Key observation:** FPGA and GPU accelerators are heavily bottlenecked by **data movement**
- **Key idea:** exploit the high memory bandwidth and the logic layer of **3D-stacked memory** to perform **highly-parallel filtering** in the DRAM chip itself
- **GRIM-Filter, an algorithm-hardware co-designed PIM system for pre-alignment filtering**
- **Key results:**
 - GRIM-Filter is 1.8x-3.7x (2.1x on average) **faster than the FastHASH filter** (BMC Genomics'13) across real data sets
 - GRIM-Filter has 5.6x-6.4x (6.0x on average) **lower false accept rate than the FastHASH filter** (BMC Genomics'13) across real data sets

Our Proposal: GRIM-Filter

1. Data Structures: Bins & Bitvectors
2. Checking a Bin
3. Integrating GRIM-Filter into a Mapper

GRIM-Filter: Bins

- We partition the genome into large sequences (**bins**).



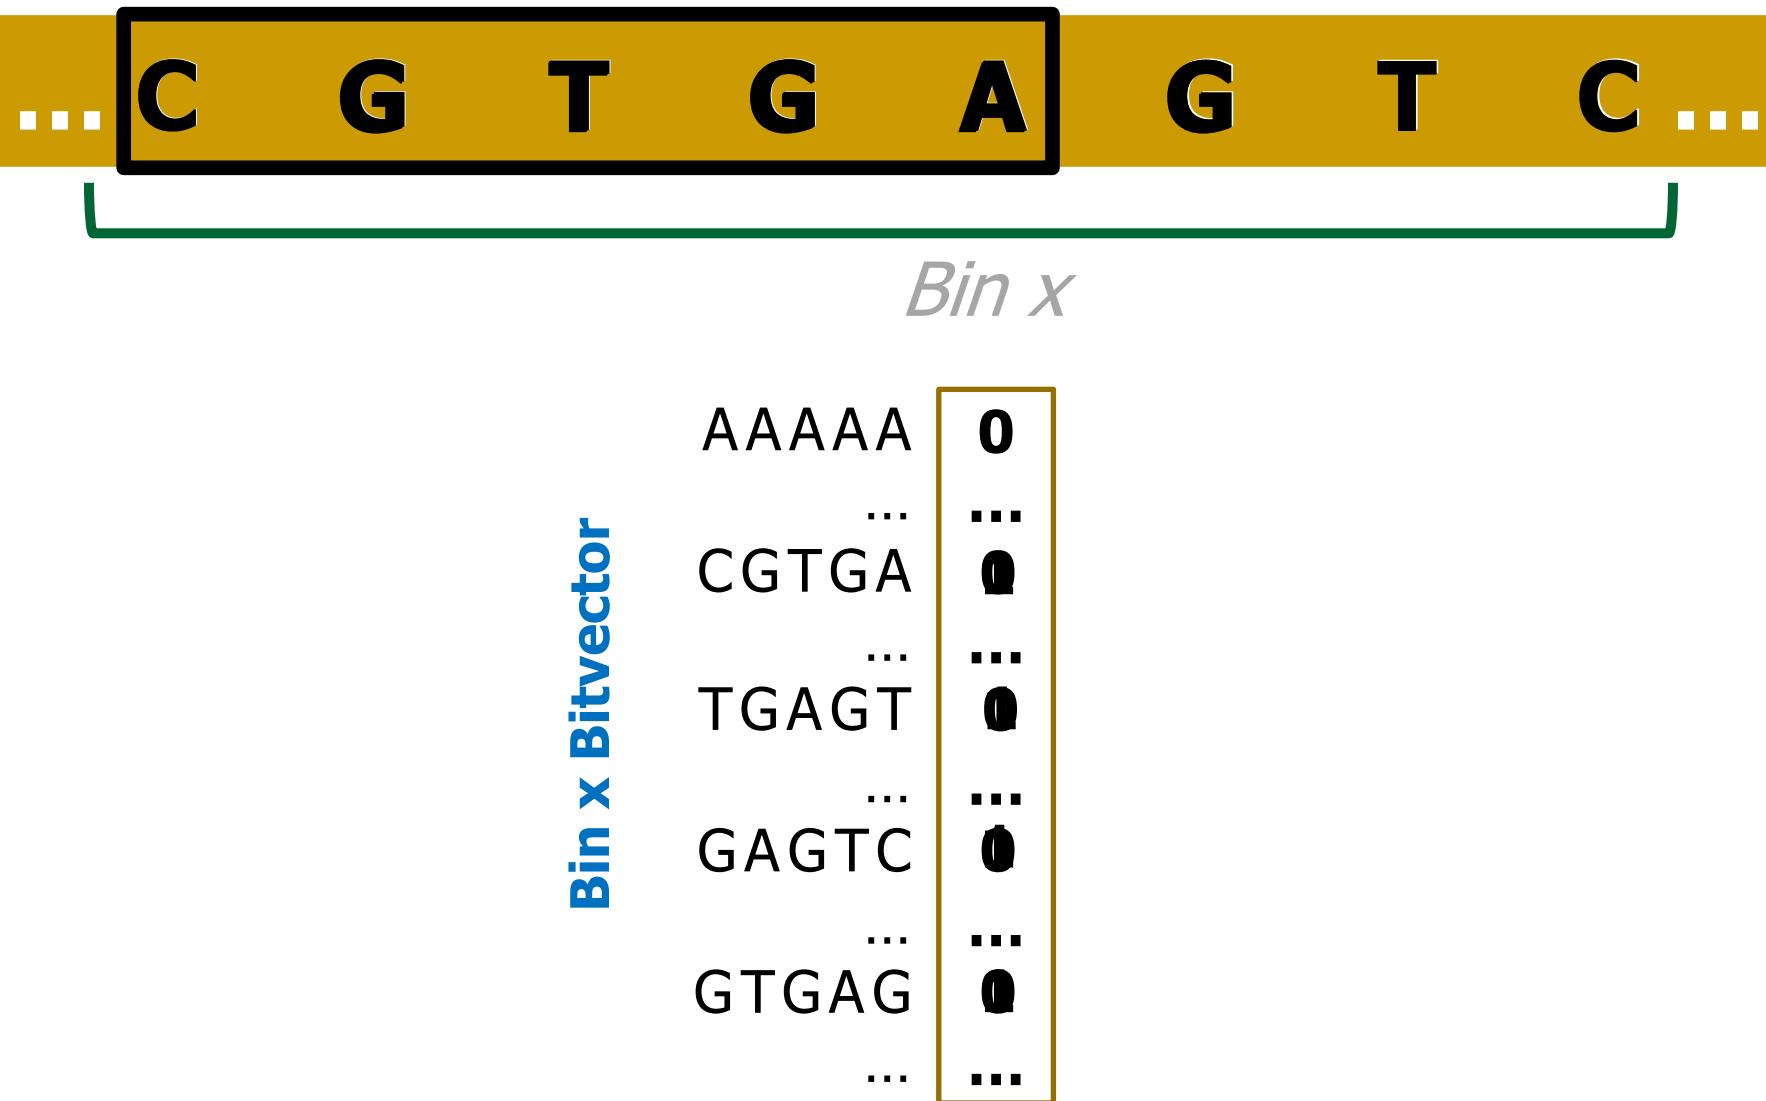
- Represent each bin with a **bitvector** that holds the occurrence of all permutations of a small string (**token**) in the bin
- To account for matches that straddle bins, we employ overlapping bins
 - A read will now always completely fall within a single bin

Bitvector
AAAAAA
AAAAC
AAAAT
...
CCCCC
CCCCT
CCCGG
...
GGGGG

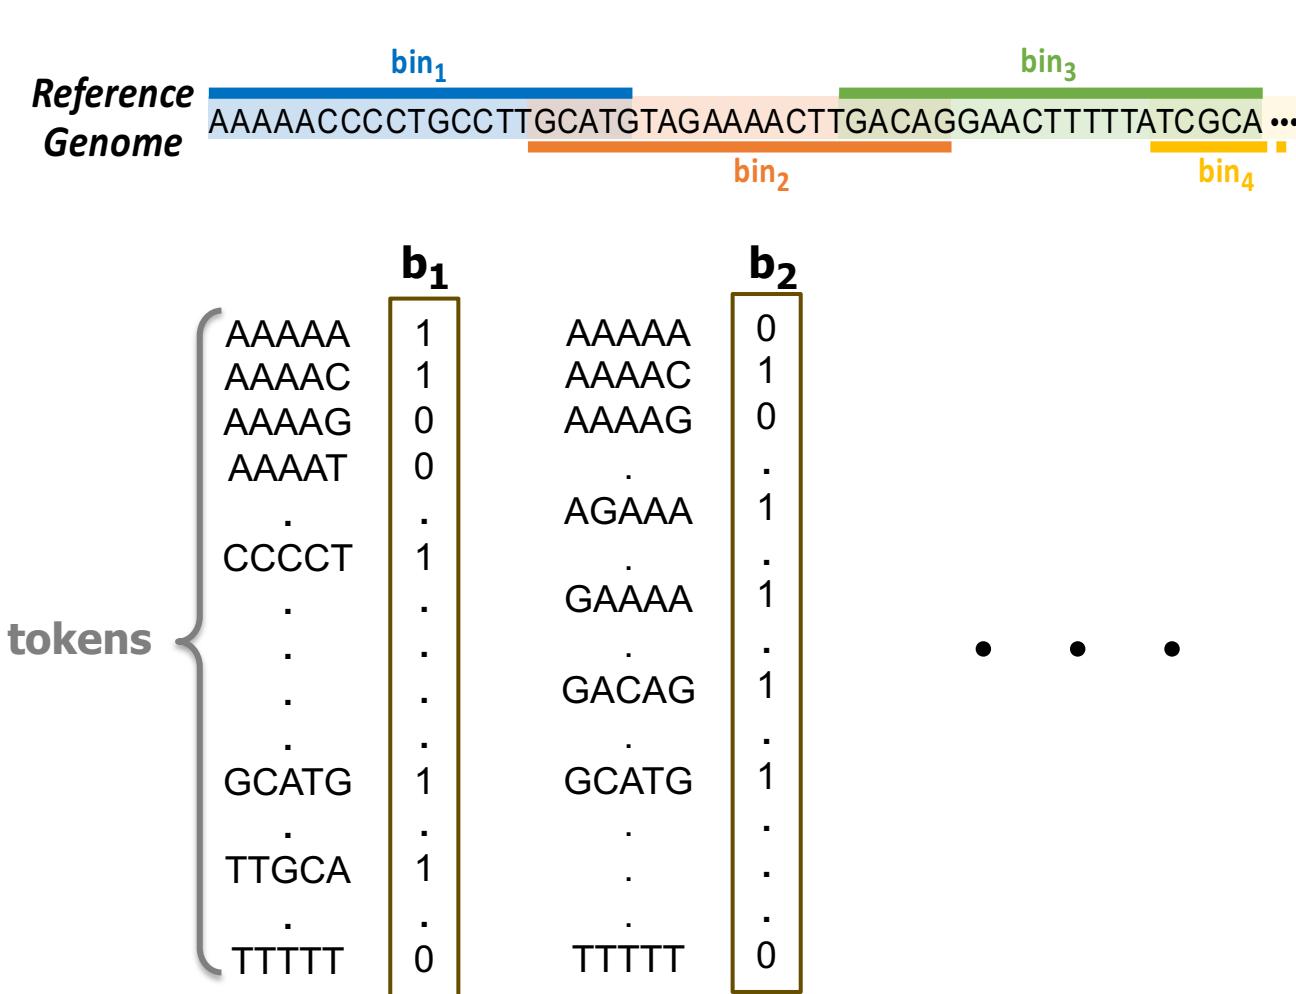
Annotations:

- The entry for "AAAAA" has a green box around the '1' in the bitvector, with an arrow pointing to the text "AAAAAA exists in bin x".
- The entry for "CCCCT" has a pink box around the '0' in the bitvector, with an arrow pointing to the text "CCCT doesn't exist in bin x".

GRIM-Filter: Bitvectors



GRIM-Filter: Bitvectors



Storing all bitvectors requires $4^n * t$ bits in memory, where t = number of bins.

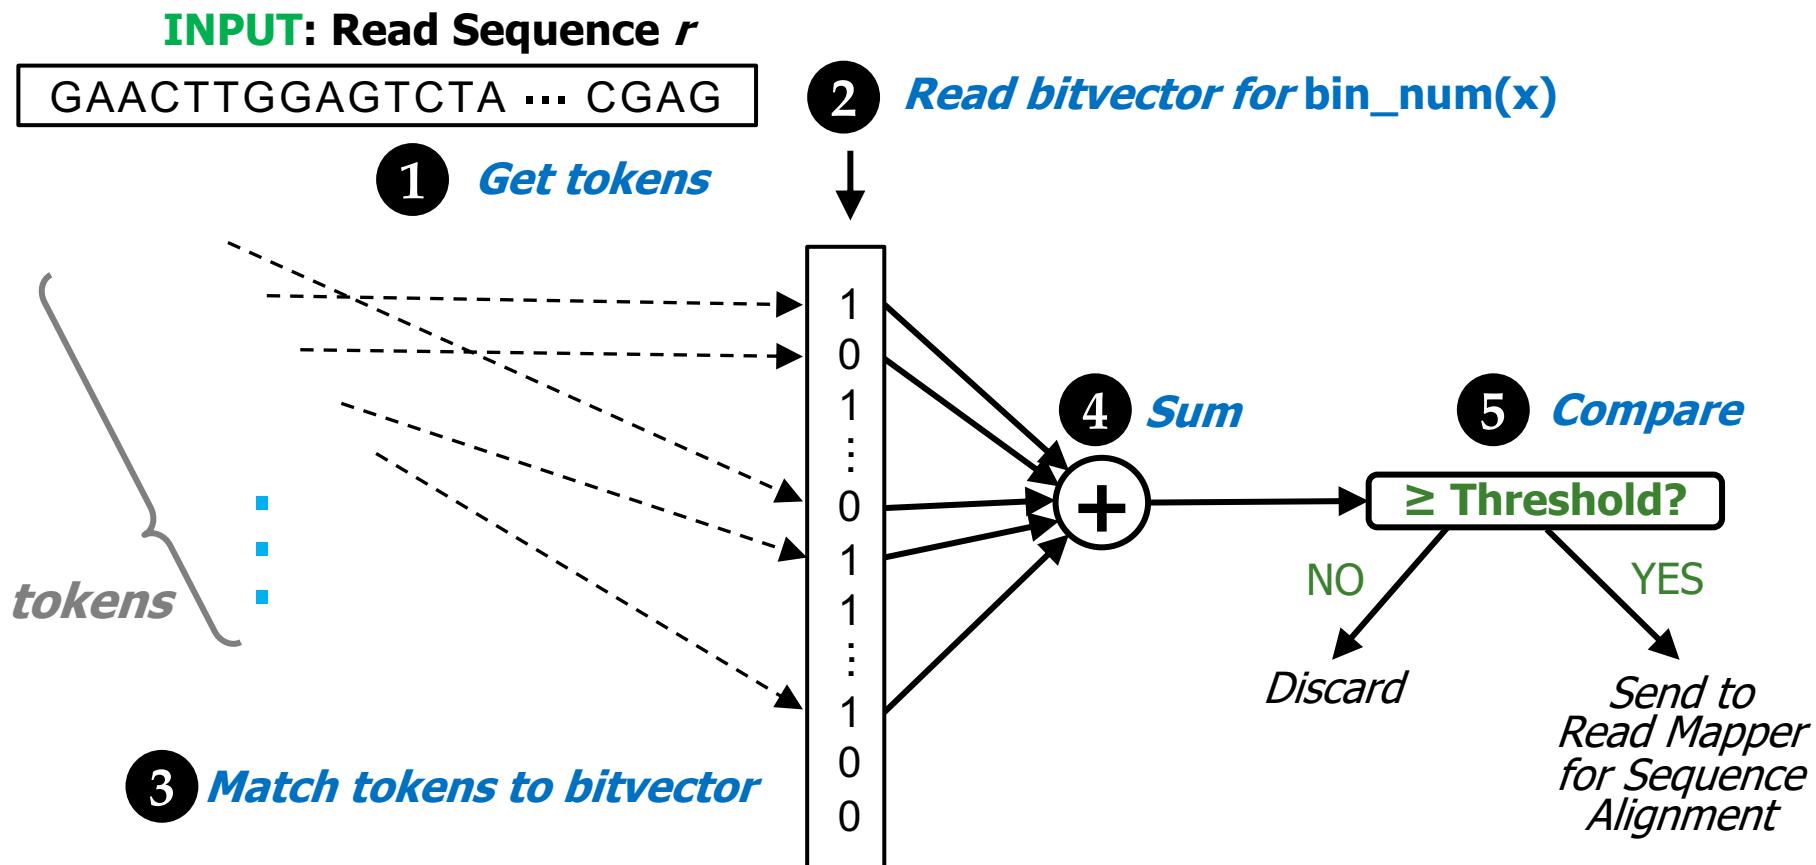
For **bin size** ~200, and **n** = 5, **memory footprint** ~3.8 GB

Our Proposal: GRIM-Filter

1. Data Structures: Bins & Bitvectors
2. Checking a Bin
3. Integrating GRIM-Filter into a Mapper

GRIM-Filter: Checking a Bin

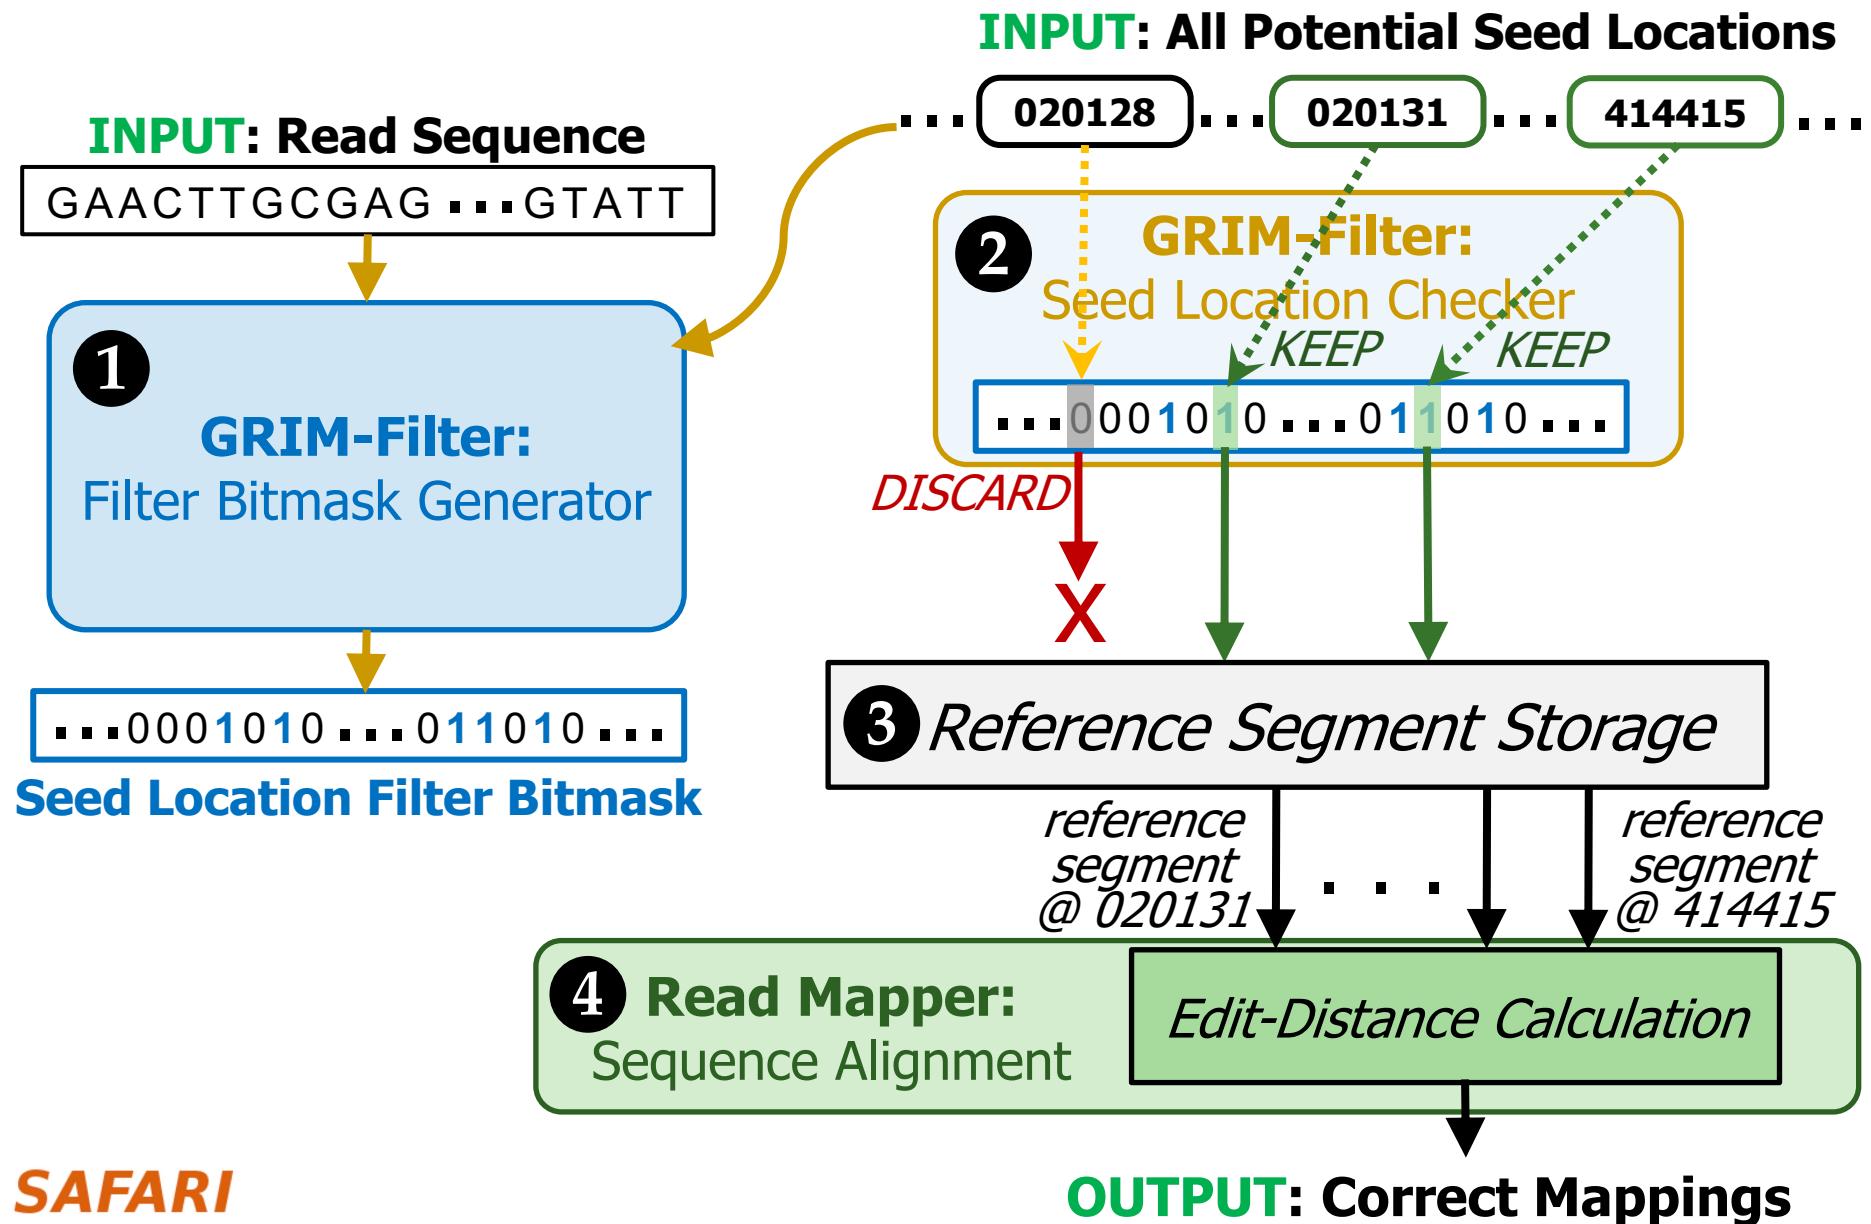
How GRIM-Filter determines whether to **discard** potential match locations in a given bin **prior** to alignment



Our Proposal: GRIM-Filter

1. Data Structures: Bins & Bitvectors
2. Checking a Bin
3. **Integrating GRIM-Filter into a Mapper**

Integrating GRIM-Filter into a Read Mapper



Key Properties of GRIM-Filter

1. Simple Operations:

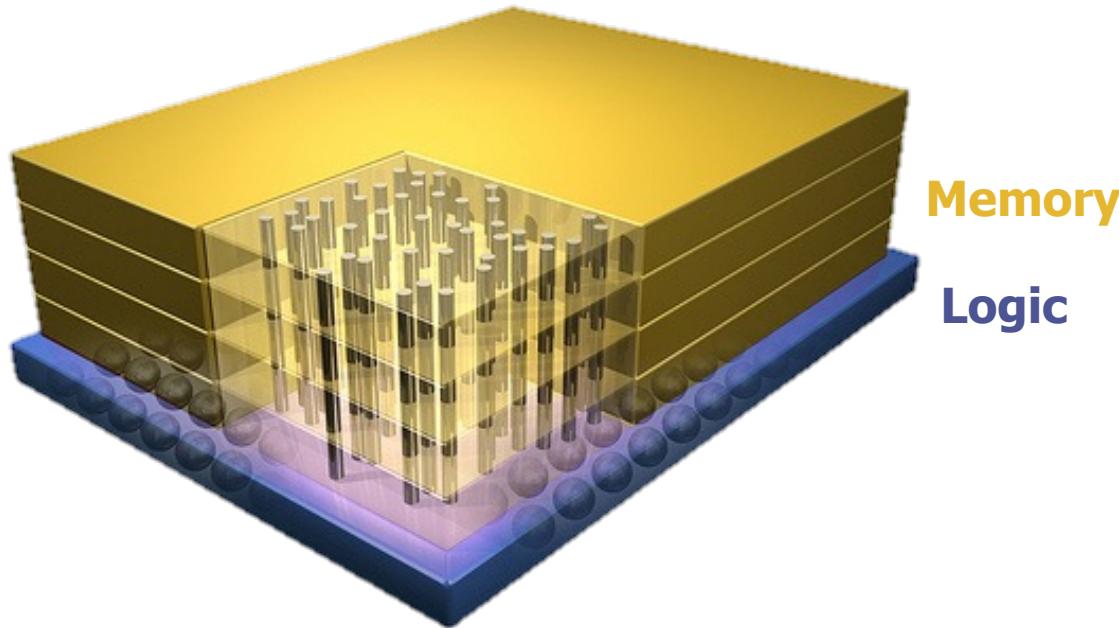
- ❑ To check a given bin, find the **sum** of all bits corresponding to each token in the read
- ❑ **Compare** against threshold to determine whether to align

2. Highly Parallel: Each bin is operated on independently and there are many many bins

3. Memory Bound: Given the frequent accesses to the large bitvectors, we find that GRIM-Filter is memory bound

These properties together make GRIM-Filter a good algorithm to be run in 3D-Stacked DRAM

Opportunity: 3D-Stacked Logic+Memory



Memory
Logic

Other “True 3D” technologies
under development

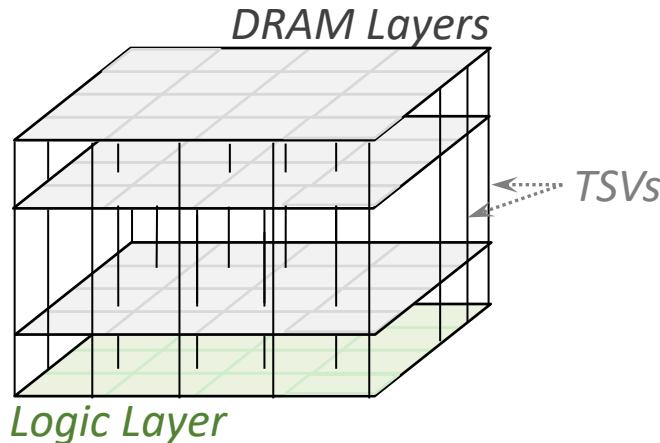
DRAM Landscape (circa 2015)

<i>Segment</i>	<i>DRAM Standards & Architectures</i>
Commodity	DDR3 (2007) [14]; DDR4 (2012) [18]
Low-Power	LPDDR3 (2012) [17]; LPDDR4 (2014) [20]
Graphics	GDDR5 (2009) [15]
Performance	eDRAM [28], [32]; RLDRAM3 (2011) [29]
3D-Stacked	WIO (2011) [16]; WIO2 (2014) [21]; MCDRAM (2015) [13]; HBM (2013) [19]; HMC1.0 (2013) [10]; HMC1.1 (2014) [11]
Academic	SBA/SSA (2010) [38]; Staged Reads (2012) [8]; RAIDR (2012) [27]; SALP (2012) [24]; TL-DRAM (2013) [26]; RowClone (2013) [37]; Half-DRAM (2014) [39]; Row-Buffer Decoupling (2014) [33]; SARP (2014) [6]; AL-DRAM (2015) [25]

Table 1. Landscape of DRAM-based memory

Kim+, "Ramulator: A Flexible and Extensible DRAM Simulator", IEEE CAL 2015.

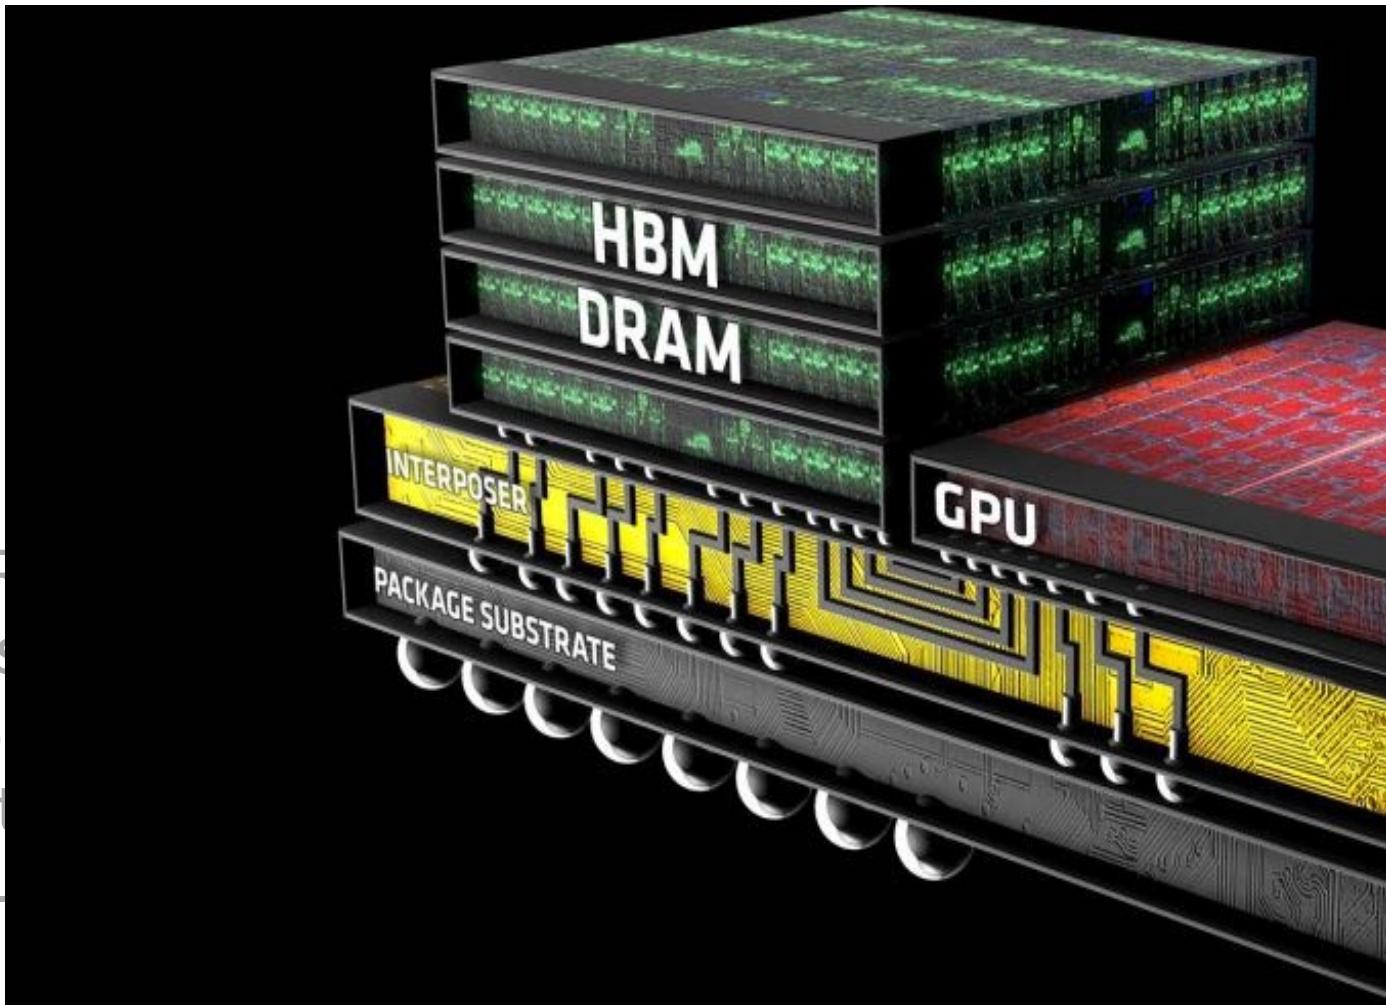
3D-Stacked Memory



- 3D-Stacked DRAM architecture has **extremely high bandwidth** as well as a stacked customizable logic layer
 - Logic Layer enables **Processing-in-Memory**, via high-bandwidth low-latency access to DRAM layers
 - Embed GRIM-Filter operations into **DRAM logic layer** and appropriately distribute bitvectors throughout memory

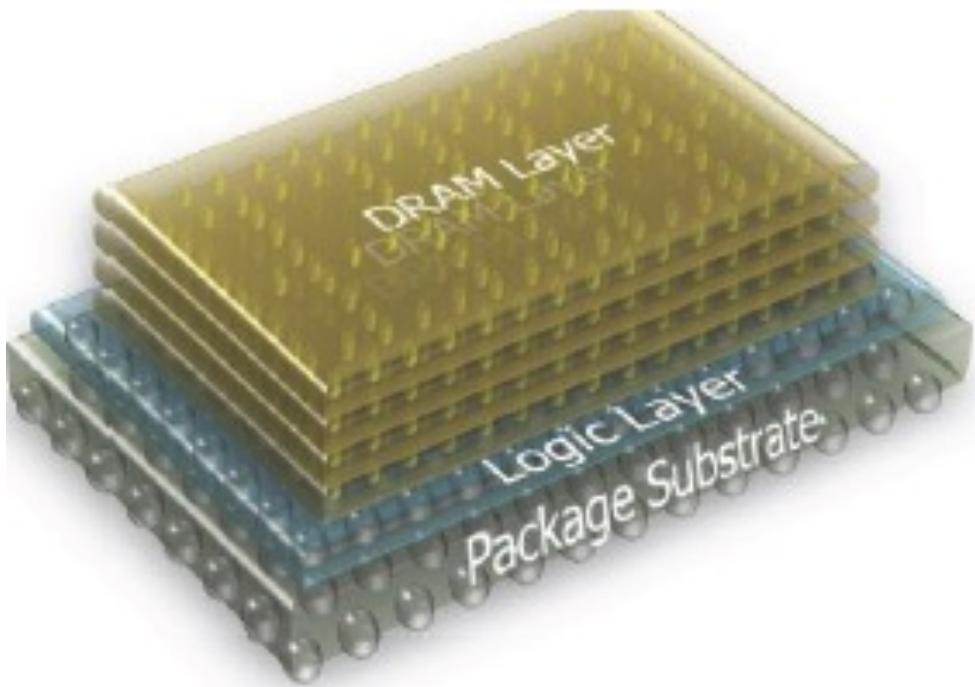
3D-Stacked Memory

- 3D-Stacked DRAM
bandwidth as
 - Logic Layer enables computation to
 - Embed GRIM appropriately



3D-Stacked Memory

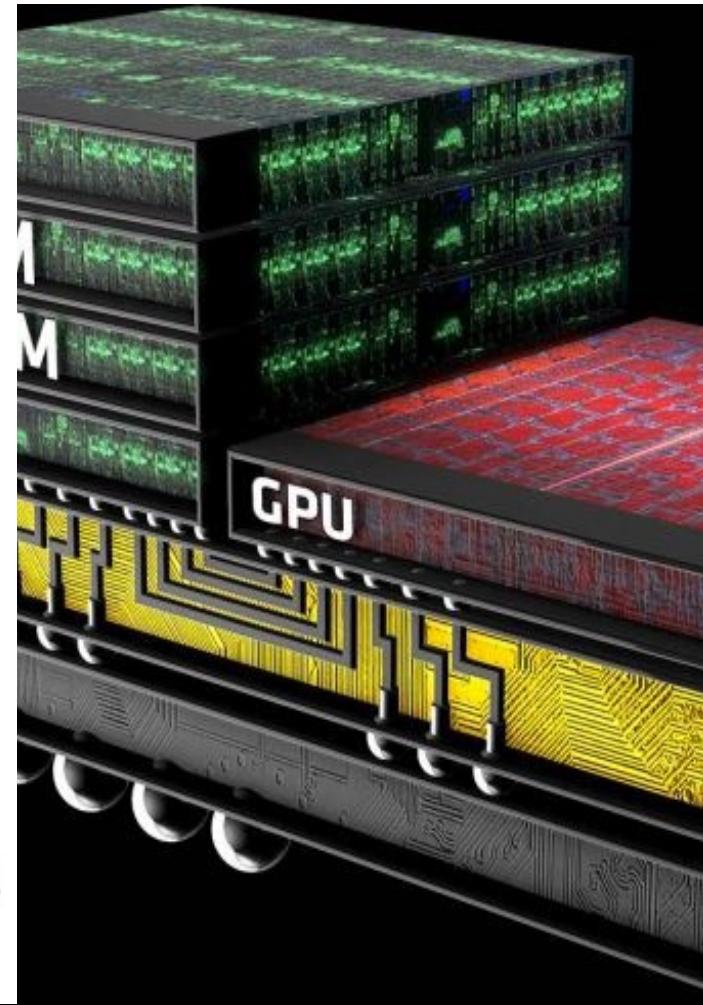
Micron's HMC



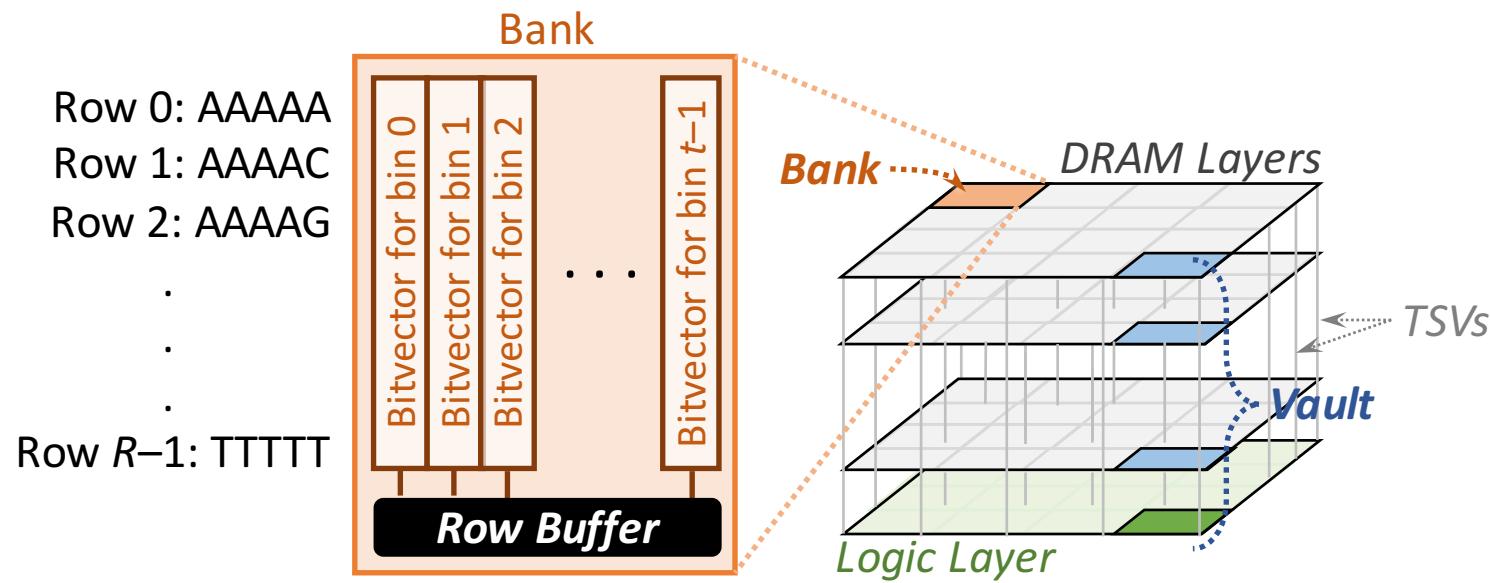
Micron has working demonstration components

http://images.anandtech.com/doc/9266/HBMCar_678x452.jpg

<http://i1-news.softpedia-static.com/images/news2/Micron-and-Samsung-Join-Force-to>Create-Next-Gen-Hybrid-Memory-2.png>

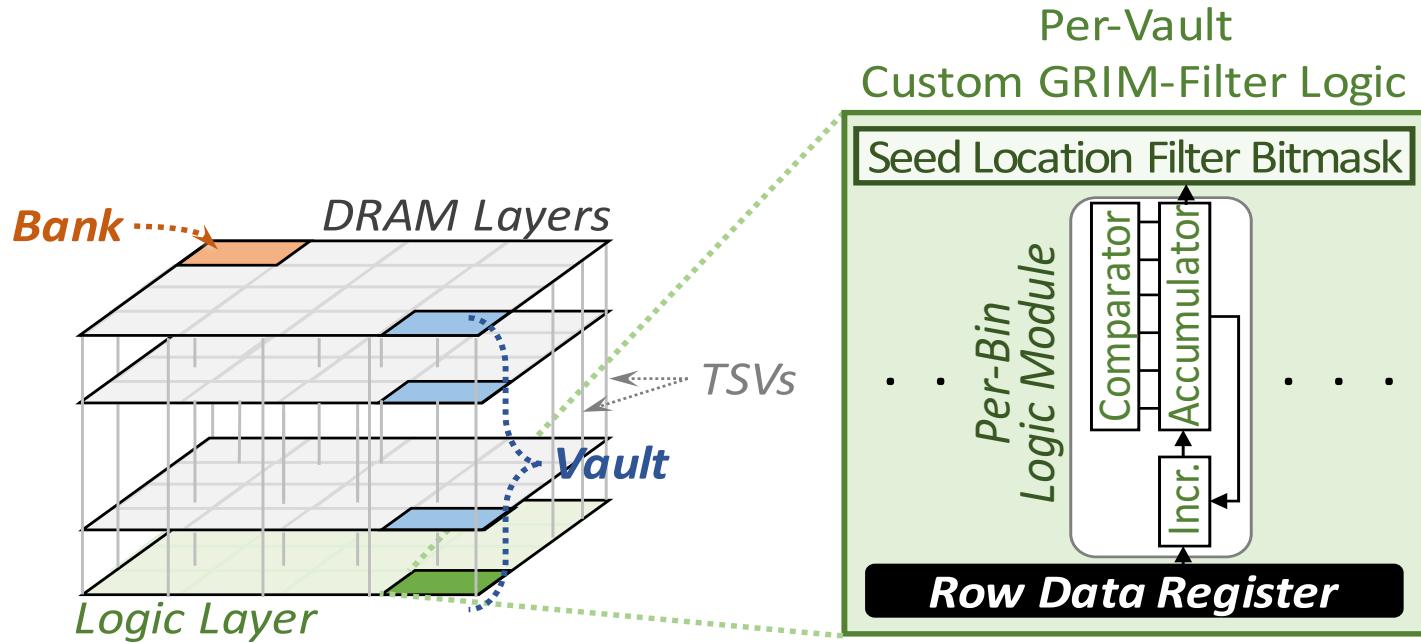


GRIM-Filter in 3D-Stacked DRAM



- Each DRAM layer is organized as an array of **banks**
 - A **bank** is an array of cells with a row buffer to transfer data
- The layout of bitvectors in a bank enables filtering many bins in parallel

GRIM-Filter in 3D-Stacked DRAM



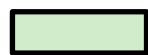
- Customized logic for accumulation and comparison per genome segment
 - Low area overhead, simple implementation
 - For HBM2, we use 4096 incrementer LUTs, 7-bit counters, and comparators in logic layer

Methodology

- Performance simulated using an in-house 3D-Stacked DRAM simulator
- Evaluate 10 real read data sets (From the 1000 Genomes Project)
 - Each data set consists of 4 million reads of length 100
- Evaluate two key metrics
 - Performance
 - False negative rate
 - The fraction of locations that pass the filter but result in a mismatch
- Compare against a state-of-the-art filter, FastHASH [Xin+, BMC Genomics 2013] when using mrFAST, but **GRIM-Filter can be used with ANY read mapper**

GRIM-Filter Performance

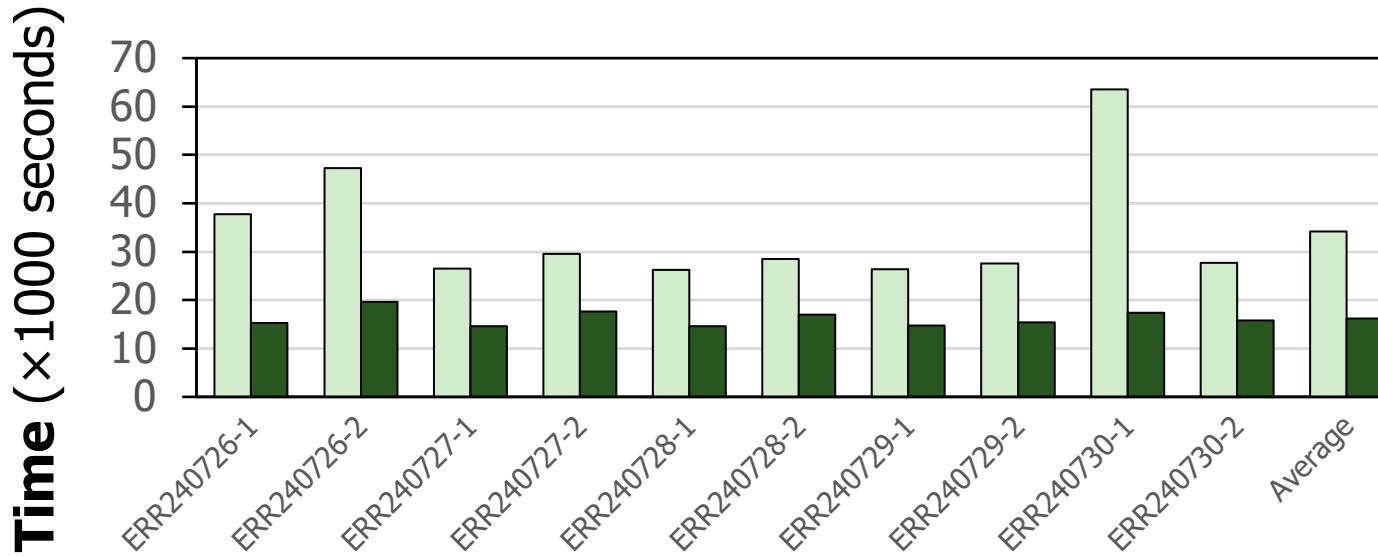
Benchmarks and their Execution Times



FastHASH filter



GRIM-Filter



Sequence Alignment
Error Tolerance (e)
 $e = 0.05$

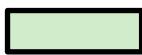
1.8x-3.7x performance benefit across real data sets

2.1x average performance benefit

GRIM-Filter gets performance due to its hardware-software co-design

GRIM-Filter False Negative Rate

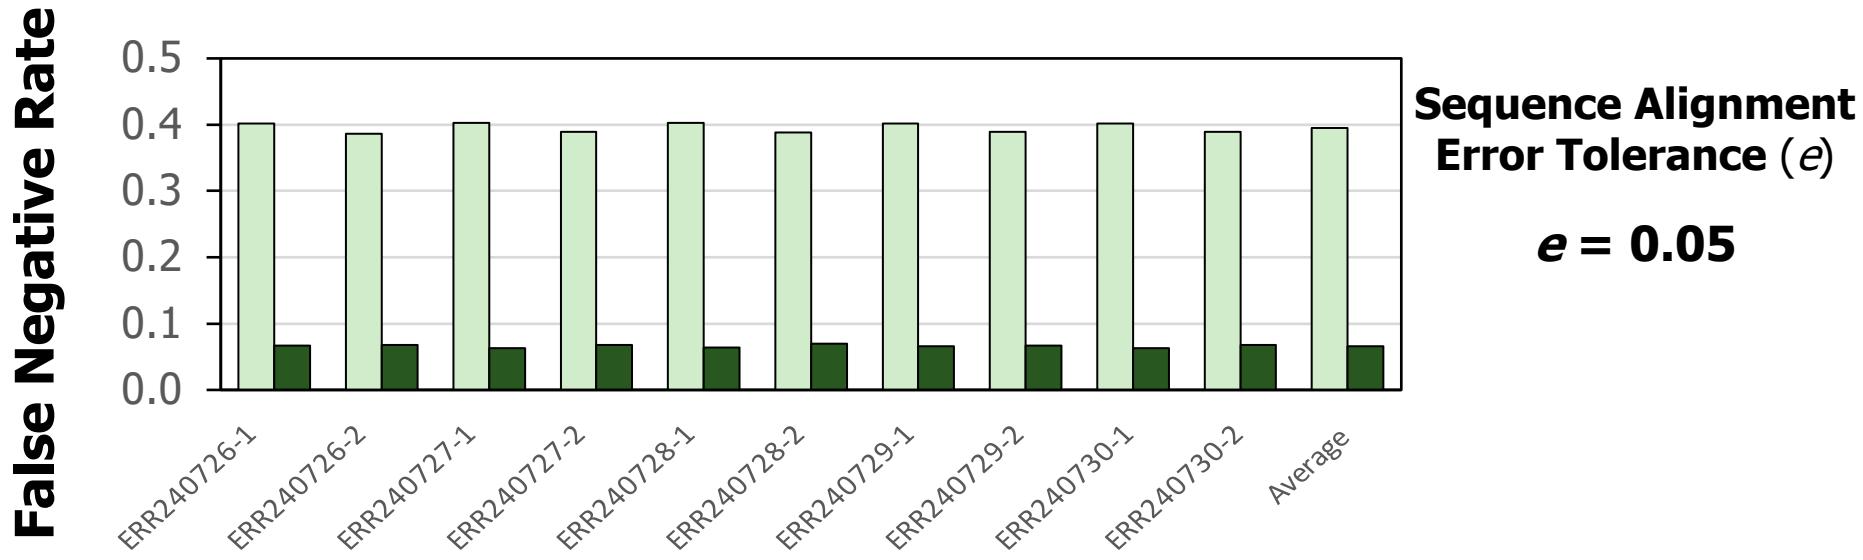
Benchmarks and their False Negative Rates



FastHASH filter



GRIM-Filter



5.6x-6.4x False Negative reduction across real data sets

6.0x average reduction in False Negative Rate

GRIM-Filter utilizes more information available in the read to filter

More on GRIM-Filter

- Jeremie S. Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, and Onur Mutlu,
"GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies"

BMC Genomics, 2018.

Proceedings of the 16th Asia Pacific Bioinformatics Conference (APBC), Yokohama, Japan, January 2018.

[[Slides \(pptx\)](#) ([pdf](#))]

[[Source Code](#)]

[[arxiv.org Version \(pdf\)](#)]

[[Talk Video at AACBB 2019](#)]

Research | [Open Access](#) | Published: 09 May 2018

GRIM-Filter: Fast seed location filtering in DNA read mapping using processing-in-memory technologies

[Jeremie S. Kim](#) , [Damla Senol Cali](#), [Hongyi Xin](#), [Donghyuk Lee](#), [Saugata Ghose](#), [Mohammed Alser](#), [Hasan Hassan](#), [Oguz Ergin](#), [Can Alkan](#)  & [Onur Mutlu](#) 

[BMC Genomics](#) **19**, Article number: 89 (2018) | [Cite this article](#)

4340 Accesses | 39 Citations | 9 Altmetric | [Metrics](#)

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

In-Storage Genome Filtering [ASPLOS 2022]

- Nika Mansouri Ghiasi, Jisung Park, Harun Mustafa, Jeremie Kim, Ataberk Olgun, Arvid Gollwitzer, Damla Senol Cali, Can Firtina, Haiyu Mao, Nour Almadhoun Alserr, Rachata Ausavarungnirun, Nandita Vijaykumar, Mohammed Alser, and Onur Mutlu,

"GenStore: A High-Performance and Energy-Efficient In-Storage Computing System for Genome Sequence Analysis"

Proceedings of the 27th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Virtual, February-March 2022.

[[Lightning Talk Slides \(pptx\)](#) ([pdf](#))]

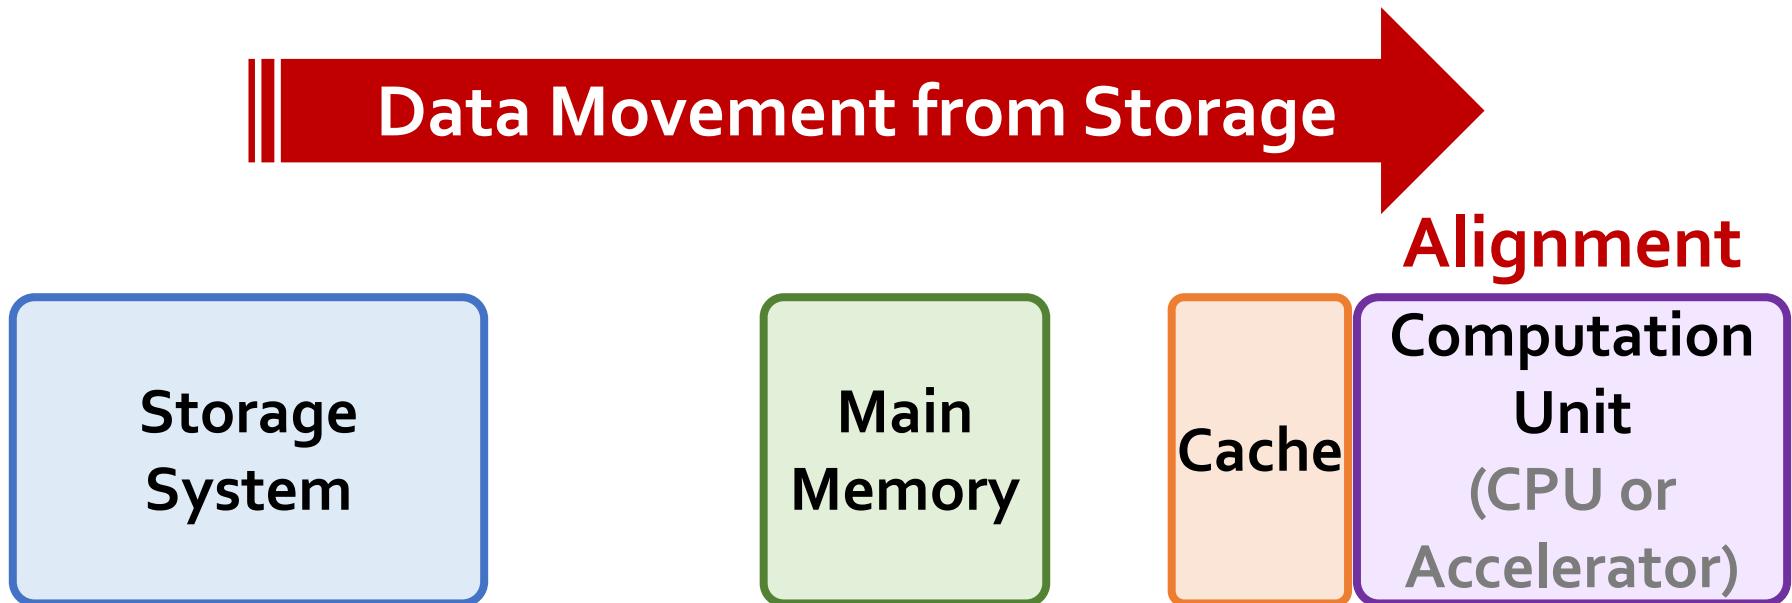
[[Lightning Talk Video](#) (90 seconds)]

GenStore: A High-Performance In-Storage Processing System for Genome Sequence Analysis

Nika Mansouri Ghiasi¹ Jisung Park¹ Harun Mustafa¹ Jeremie Kim¹ Ataberk Olgun¹
Arvid Gollwitzer¹ Damla Senol Cali² Can Firtina¹ Haiyu Mao¹ Nour Almadhoun Alserr¹
Rachata Ausavarungnirun³ Nandita Vijaykumar⁴ Mohammed Alser¹ Onur Mutlu¹

¹ETH Zürich ²Bionano Genomics ³KMUTNB ⁴University of Toronto

Genome Sequence Analysis

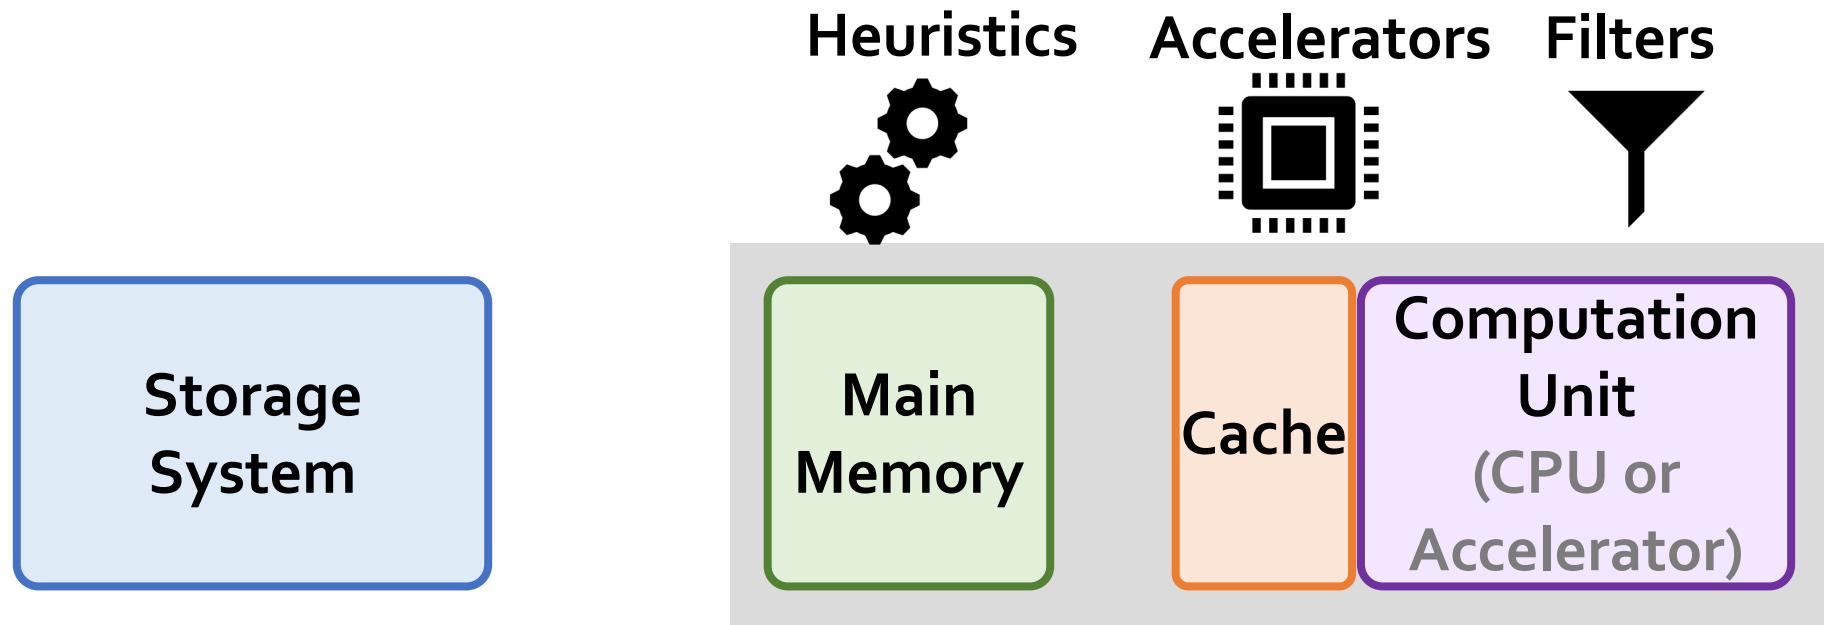


Computation overhead



Data movement overhead

Accelerating Genome Sequence Analysis



Computation overhead



Data movement overhead

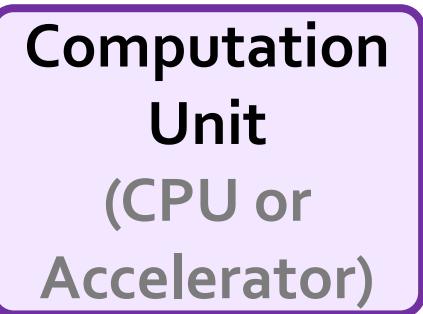
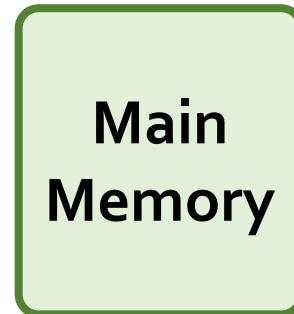
Key Idea



*Filter reads that do **not** require alignment
inside the storage system*



Filtered Reads



Exactly-matching reads

Do not need expensive approximate string matching during alignment

Non-matching reads

Do not have potential matching locations and can skip alignment

Filtering Opportunities

- Sequencing machines produce one of two kinds of reads
 - **Short reads:** highly accurate and short
 - **Long reads:** less accurate and long

Reads that do not require the expensive alignment step:

Exactly-matching reads

Do not need expensive approximate string matching during alignment

- Low sequencing error rates (short reads) combined with
- Low genetic variation

Non-matching reads

Do not have potential matching locations, so they skip alignment

- High sequencing error rates (long reads) or
- High genetic variation (short or long reads)

Challenges



*Filter reads that do **not** require alignment
inside the storage system*

Storage
System

Main
Memory

Cache

Computation
Unit
(CPU or
Accelerator)

Filtered Reads

Read mapping workloads can exhibit different behavior

There are **limited hardware resources**
in the storage system

GenStore



*Filter reads that do **not** require alignment
inside the storage system*

GenStore-Enabled
Storage
System

Main
Memory

Cache

Computation
Unit
(CPU or
Accelerator)



Computation overhead



Data movement overhead

GenStore provides significant speedup (**1.4x - 33.6x**) and
energy reduction (**3.9x – 29.2x**) at low cost

In-Storage Genome Filtering [ASPLOS 2022]

- Nika Mansouri Ghiasi, Jisung Park, Harun Mustafa, Jeremie Kim, Ataberk Olgun, Arvid Gollwitzer, Damla Senol Cali, Can Firtina, Haiyu Mao, Nour Almadhoun Alserr, Rachata Ausavarungnirun, Nandita Vijaykumar, Mohammed Alser, and Onur Mutlu,

"GenStore: A High-Performance and Energy-Efficient In-Storage Computing System for Genome Sequence Analysis"

Proceedings of the 27th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Virtual, February-March 2022.

[[Lightning Talk Slides \(pptx\)](#) ([pdf](#))]

[[Lightning Talk Video](#) (90 seconds)]

GenStore: A High-Performance In-Storage Processing System for Genome Sequence Analysis

Nika Mansouri Ghiasi¹ Jisung Park¹ Harun Mustafa¹ Jeremie Kim¹ Ataberk Olgun¹
Arvid Gollwitzer¹ Damla Senol Cali² Can Firtina¹ Haiyu Mao¹ Nour Almadhoun Alserr¹
Rachata Ausavarungnirun³ Nandita Vijaykumar⁴ Mohammed Alser¹ Onur Mutlu¹

¹ETH Zürich ²Bionano Genomics ³KMUTNB ⁴University of Toronto

PIM Review and Open Problems

A Modern Primer on Processing in Memory

Onur Mutlu^{a,b}, Saugata Ghose^{b,c}, Juan Gómez-Luna^a, Rachata Ausavarungnirun^d

SAFARI Research Group

^a*ETH Zürich*

^b*Carnegie Mellon University*

^c*University of Illinois at Urbana-Champaign*

^d*King Mongkut's University of Technology North Bangkok*

Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, and Rachata Ausavarungnirun,

"A Modern Primer on Processing in Memory"

Invited Book Chapter in Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann, Springer, to be published in 2021.

PIM Review and Open Problems (II)

A Workload and Programming Ease Driven Perspective of Processing-in-Memory

Saugata Ghose[†] Amirali Boroumand[†] Jeremie S. Kim^{†\\$} Juan Gómez-Luna^{\\$} Onur Mutlu^{\\$†}

[†]*Carnegie Mellon University*

^{\\$}*ETH Zürich*

Saugata Ghose, Amirali Boroumand, Jeremie S. Kim, Juan Gomez-Luna, and Onur Mutlu,

"Processing-in-Memory: A Workload-Driven Perspective"

Invited Article in IBM Journal of Research & Development, Special Issue on Hardware for Artificial Intelligence, to appear in November 2019.

[Preliminary arXiv version]

More on Processing-in-Memory

- Onur Mutlu,

"Memory-Centric Computing Systems"

Invited Tutorial at *66th International Electron Devices Meeting (IEDM)*, Virtual, 12 December 2020.

[Slides (pptx) (pdf)]

[Executive Summary Slides (pptx) (pdf)]

[Tutorial Video (1 hour 51 minutes)]

[Executive Summary Video (2 minutes)]

[Abstract and Bio]

[Related Keynote Paper from VLSI-DAT 2020]

[Related Review Paper on Processing in Memory]

<https://www.youtube.com/watch?v=H3sEaINPBOE>



Memory-Centric Computing Systems

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

12 December 2020

IEDM Tutorial



SAFARI

ETH zürich

Carnegie Mellon



0:06 / 1:51:05



IEDM 2020 Tutorial: Memory-Centric Computing Systems, Onur Mutlu, 12 December 2020

1,641 views • Dec 23, 2020

48

0

SHARE

SAVE

...



Onur Mutlu Lectures
13.9K subscribers

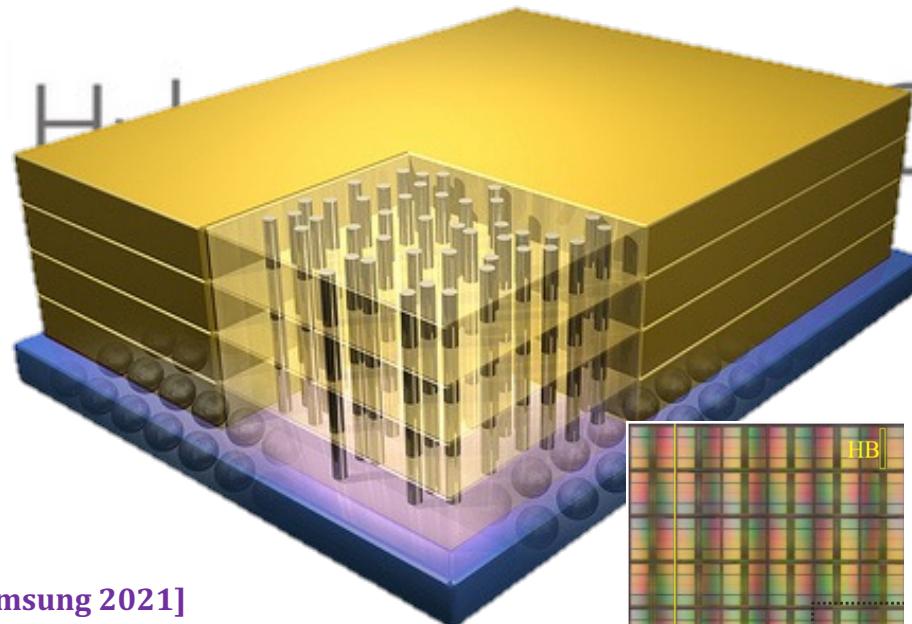
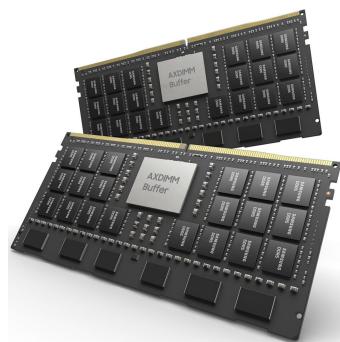
<https://www.youtube.com/watch?v=H3sEaINPBOE>

ANALYTICS

EDIT VIDEO

<https://www.youtube.com/onurmutlulectures>

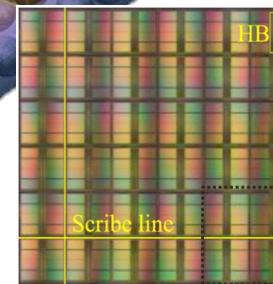
Processing-in-Memory Landscape Today



[Samsung 2021]



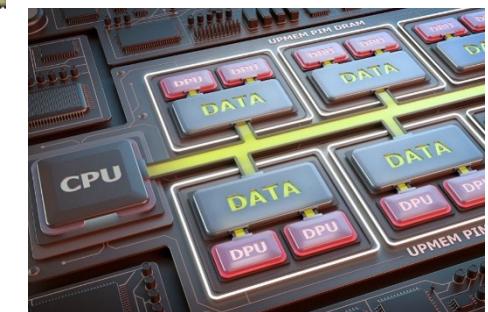
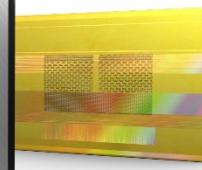
[Alibaba 2022]



[SK Hynix 2022]



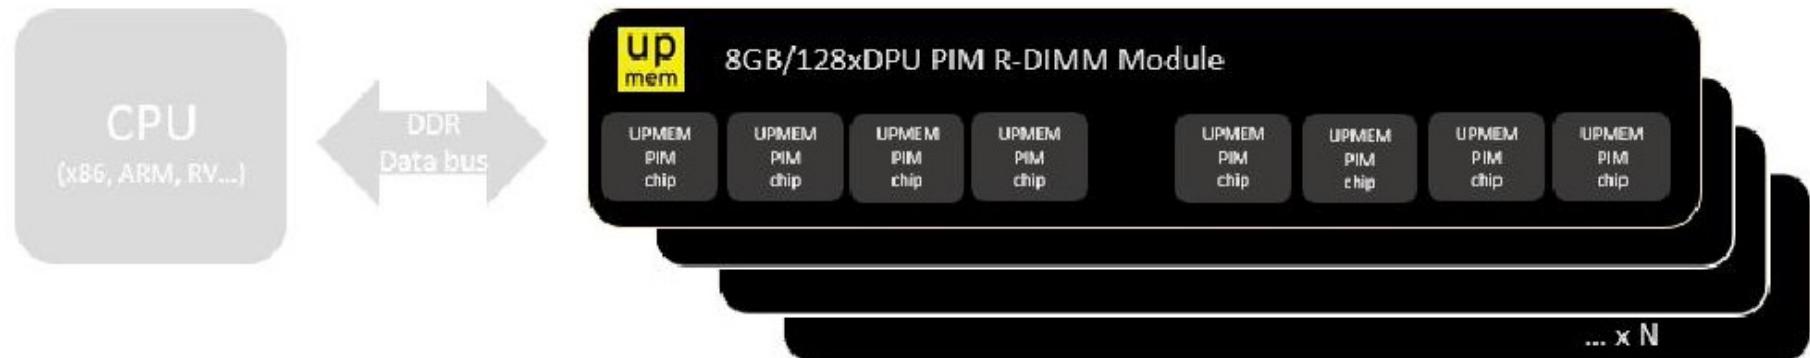
[Samsung 2021]



[UPMEM 2019]

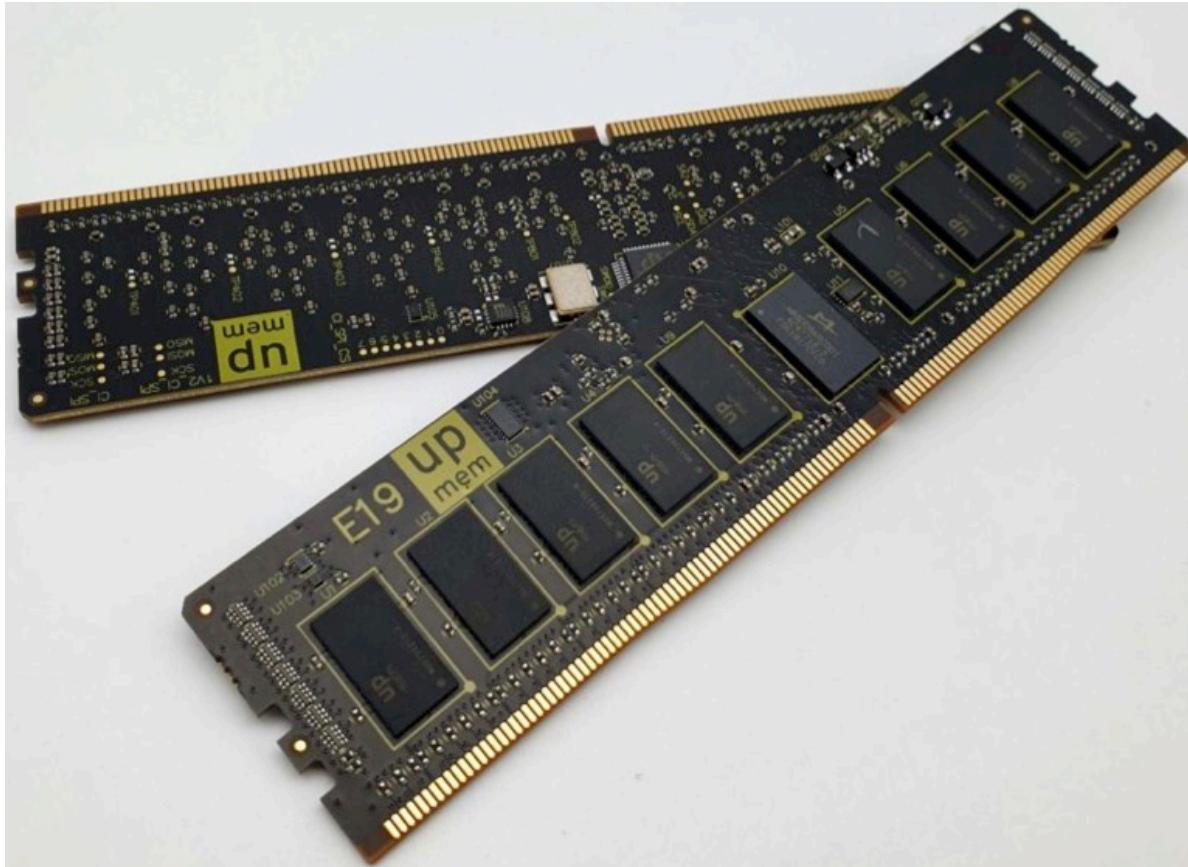
UPMEM Processing-in-DRAM Engine (2019)

- Processing in DRAM Engine
- Includes **standard DIMM modules**, with a **large number of DPU processors** combined with DRAM chips.
- Replaces **standard** DIMMs
 - DDR4 R-DIMM modules
 - 8GB+128 DPUs (16 PIM chips)
 - Standard 2x-nm DRAM process
 - **Large amounts of** compute & memory bandwidth

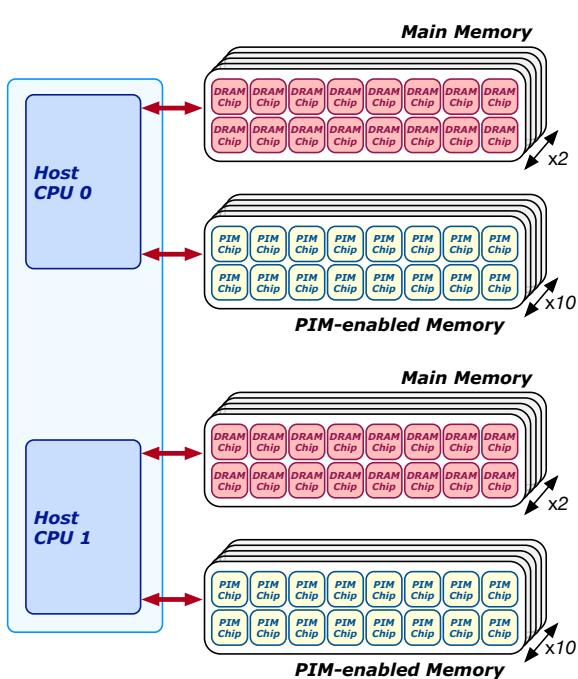


UPMEM Memory Modules

- E19: 8 chips DIMM (1 rank). DPUs @ 267 MHz
- P21: 16 chips DIMM (2 ranks). DPUs @ 350 MHz



2,560-DPU Processing-in-Memory System



Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture

JUAN GÓMEZ-LUNA, ETH Zürich, Switzerland

IZZAT EL HAJI, American University of Beirut, Lebanon

IVAN FERNANDEZ, ETH Zürich, Switzerland and University of Malaga, Spain

CHRISTINA GIANNOULA, ETH Zürich, Switzerland and NTUA, Greece

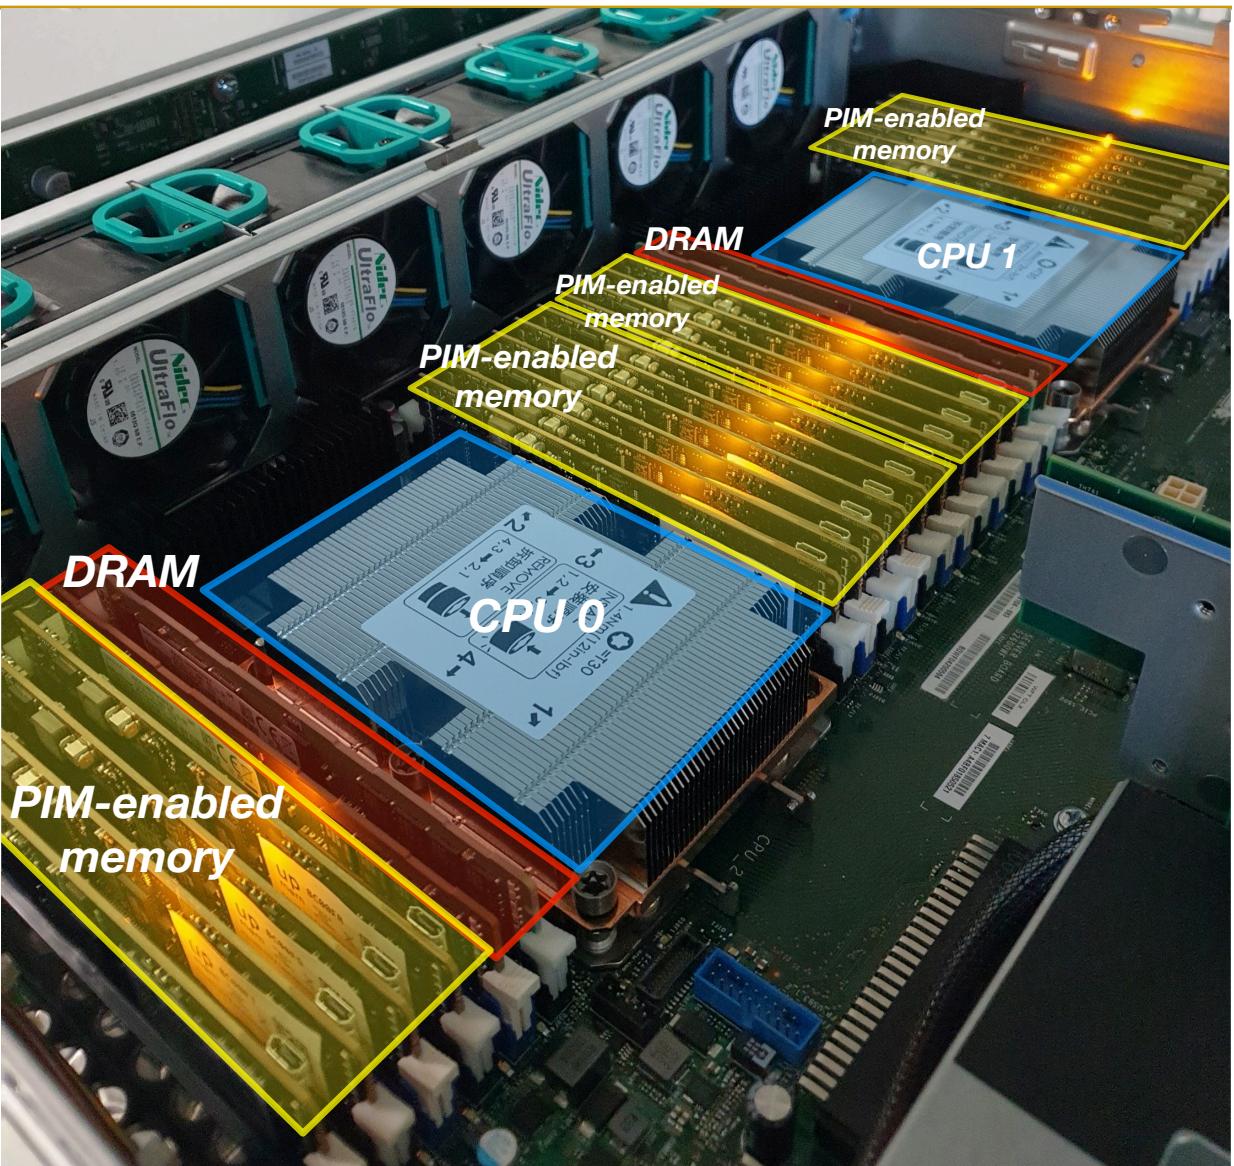
GERALDO F. OLIVEIRA, ETH Zürich, Switzerland

ONUR MUTLU, ETH Zürich, Switzerland

Many modern workloads, such as neural networks, databases, and graph processing, are fundamentally memory-bound. For such workloads, the data movement between main memory and CPU cores imposes a significant overhead in terms of both latency and energy. A major reason is that this communication happens through a narrow bus with high latency and limited bandwidth, and the low data reuse in memory-bound workloads is insufficient to amortize the cost of main memory access. Fundamentally addressing this data movement bottleneck requires a paradigm where the memory system assumes an active role in computing by integrating processing capabilities. This paradigm is known as *processing-in-memory* (PIM).

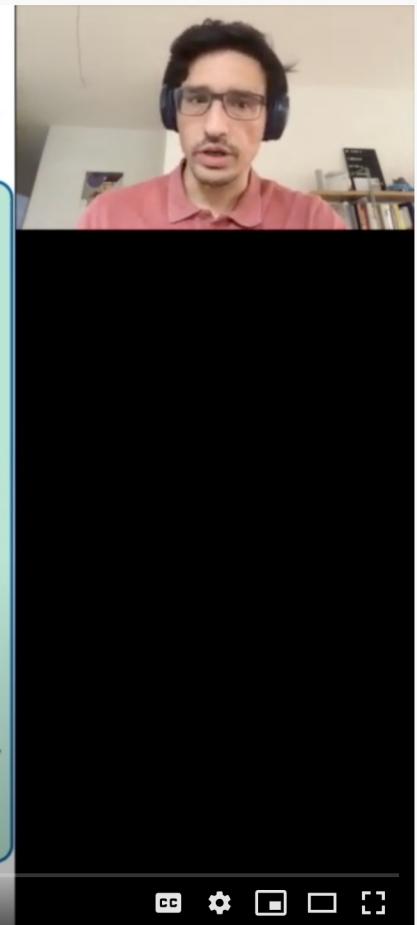
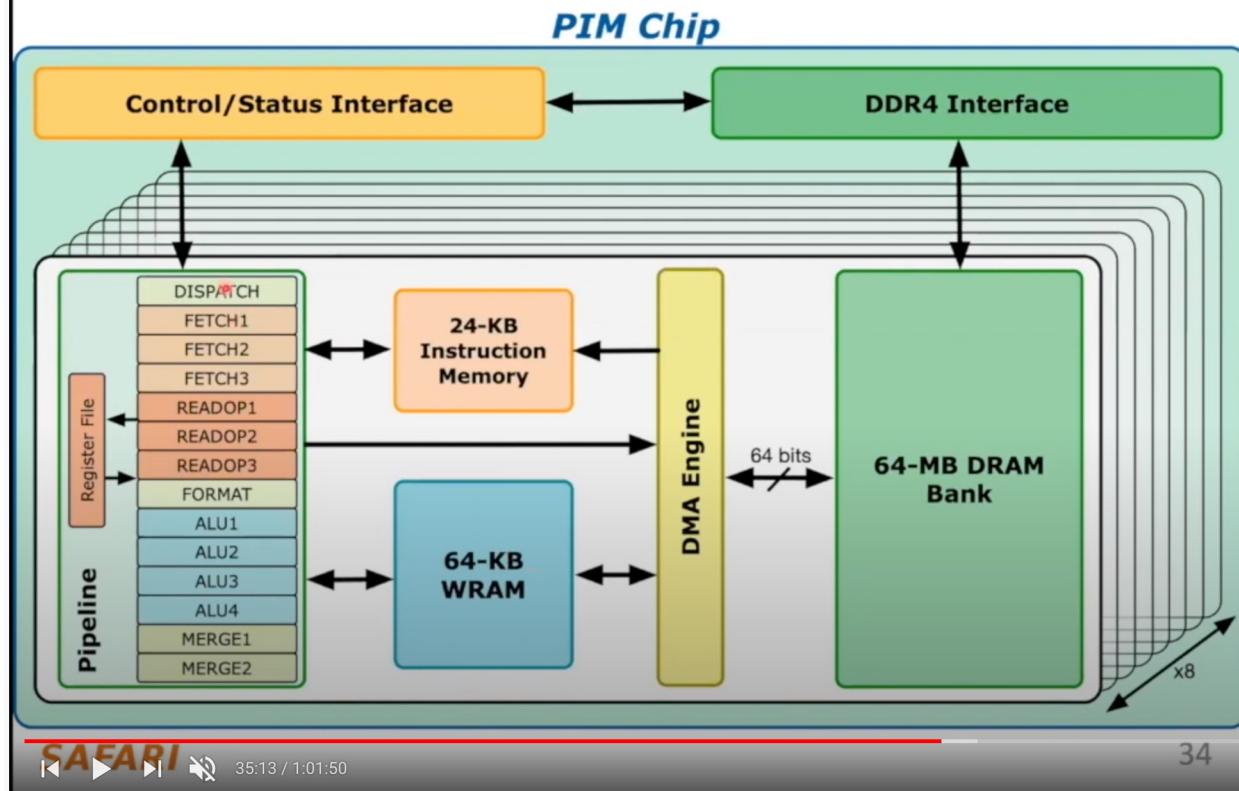
Recent research explores different forms of PIM architectures, motivated by the emergence of new 3D-stacked memory technologies that integrate memory with a logic layer where processing elements can be easily placed. Past works evaluate these architectures in simulation or, at best, with simplified hardware prototypes. In contrast, the UPMEM company has designed and manufactured the first publicly-available real-world PIM architecture. The UPMEM PIM architecture combines traditional DRAM memory arrays with general-purpose in-order cores, called DRAM Processing Units (DPUs), integrated in the same chip.

This paper provides the first comprehensive analysis of the first publicly-available real-world PIM architecture. We make two key contributions. First, we conduct an experimental characterization of the UPMEM-based PIM system using microbenchmarks to assess various architecture limits such as compute throughput and memory bandwidth, yielding new insights. Second, we present PrIM (Processing-In-Memory benchmarks), a benchmark suite of 16 workloads from different application domains (e.g., dense/sparse linear algebra, databases, data analytics, graph processing, neural networks, bioinformatics, image processing), which we identify as memory-bound. We evaluate the performance and scaling characteristics of PrIM benchmarks on the UPMEM PIM architecture, and compare their performance and energy consumption to their state-of-the-art CPU and GPU counterparts. Our extensive evaluation conducted on two real UPMEM-based PIM systems with 640 and 2,556 DPUs provides new insights about suitability of different workloads to the PIM system, programming recommendations for software designers, and suggestions and hints for hardware and architecture designers of future PIM systems.



More on the UPMEM PIM System

DRAM Processing Unit (II)



ETH ZÜRICH HAUPTGEBÄUDE

Computer Architecture - Lecture 12d: Real Processing-in-DRAM with UPMEM (ETH Zürich, Fall 2020)

1,120 views · Oct 31, 2020

1 like · 0 dislikes · SHARE · SAVE · ...



Onur Mutlu Lectures
16.7K subscribers

ANALYTICS

EDIT VIDEO

<https://www.youtube.com/watch?v=Sscy1Wrr22A&list=PL5Q2soXY2Zi9xidyIqBxUz7xRPS-wisBN&index=26>

Experimental Analysis of the UPMEM PIM Engine

Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture

JUAN GÓMEZ-LUNA, ETH Zürich, Switzerland

IZZAT EL HAJJ, American University of Beirut, Lebanon

IVAN FERNANDEZ, ETH Zürich, Switzerland and University of Malaga, Spain

CHRISTINA GIANNOULA, ETH Zürich, Switzerland and NTUA, Greece

GERALDO F. OLIVEIRA, ETH Zürich, Switzerland

ONUR MUTLU, ETH Zürich, Switzerland

Many modern workloads, such as neural networks, databases, and graph processing, are fundamentally memory-bound. For such workloads, the data movement between main memory and CPU cores imposes a significant overhead in terms of both latency and energy. A major reason is that this communication happens through a narrow bus with high latency and limited bandwidth, and the low data reuse in memory-bound workloads is insufficient to amortize the cost of main memory access. Fundamentally addressing this *data movement bottleneck* requires a paradigm where the memory system assumes an active role in computing by integrating processing capabilities. This paradigm is known as *processing-in-memory (PIM)*.

Recent research explores different forms of PIM architectures, motivated by the emergence of new 3D-stacked memory technologies that integrate memory with a logic layer where processing elements can be easily placed. Past works evaluate these architectures in simulation or, at best, with simplified hardware prototypes. In contrast, the UPMEM company has designed and manufactured the first publicly-available real-world PIM architecture. The UPMEM PIM architecture combines traditional DRAM memory arrays with general-purpose in-order cores, called *DRAM Processing Units (DPUs)*, integrated in the same chip.

This paper provides the first comprehensive analysis of the first publicly-available real-world PIM architecture. We make two key contributions. First, we conduct an experimental characterization of the UPMEM-based PIM system using microbenchmarks to assess various architecture limits such as compute throughput and memory bandwidth, yielding new insights. Second, we present *PrIM* (*Processing-In-Memory benchmarks*), a benchmark suite of 16 workloads from different application domains (e.g., dense/sparse linear algebra, databases, data analytics, graph processing, neural networks, bioinformatics, image processing), which we identify as memory-bound. We evaluate the performance and scaling characteristics of PrIM benchmarks on the UPMEM PIM architecture, and compare their performance and energy consumption to their state-of-the-art CPU and GPU counterparts. Our extensive evaluation conducted on two real UPMEM-based PIM systems with 640 and 2,556 DPUs provides new insights about suitability of different workloads to the PIM system, programming recommendations for software designers, and suggestions and hints for hardware and architecture designers of future PIM systems.

Upcoming TECHCON Presentation

■ Dr. Juan Gomez-Luna

- Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware
- Based on two major works
 - <https://arxiv.org/pdf/2105.03814.pdf>
 - <https://arxiv.org/pdf/2207.07886.pdf>



Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-In-Memory Hardware

Year: 2021, Pages: 1-7

DOI Bookmark: [10.1109/IGSC54211.2021.9651614](https://doi.org/10.1109/IGSC54211.2021.9651614)

Authors

Juan Gómez-Luna, ETH Zürich

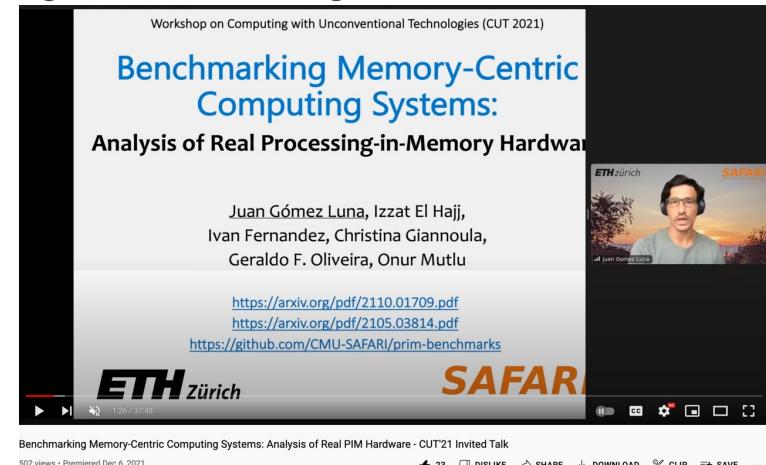
Izzat El Hajj, American University of Beirut

Ivan Fernandez, University of Malaga

Christina Giannoula, National Technical University of Athens

Geraldo F. Oliveira, ETH Zürich

Onur Mutlu, ETH Zürich



Workshop on Computing with Unconventional Technologies (CUT 2021)

Benchmarking Memory-Centric Computing Systems:
Analysis of Real Processing-in-Memory Hardware

Juan Gómez Luna, Izzat El Hajj,
Ivan Fernandez, Christina Giannoula,
Geraldo F. Oliveira, Onur Mutlu

<https://arxiv.org/pdf/2110.01709.pdf>
<https://arxiv.org/pdf/2105.03814.pdf>
<https://github.com/CMU-SAFARI/prim-benchmarks>

ETH Zürich SAFARI

Benchmarking Memory-Centric Computing Systems: Analysis of Real PIM Hardware - CUT'21 Invited Talk

502 views • Premiered Dec 6, 2021

Onur Mutlu Lectures 26.9K subscribers

ANALYTICS EDIT VIDEO

UPMEM PIM System Summary & Analysis

- Juan Gomez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu,

"Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware"

Invited Paper at Workshop on Computing with Unconventional Technologies (CUT), Virtual, October 2021.

[[arXiv version](#)]

[[PrIM Benchmarks Source Code](#)]

[[Slides \(pptx\) \(pdf\)](#)]

[[Talk Video](#) (37 minutes)]

[[Lightning Talk Video](#) (3 minutes)]

Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware

Juan Gómez-Luna

ETH Zürich

Izzat El Hajj

*American University
of Beirut*

Ivan Fernandez

*University
of Malaga*

Christina Giannoula

*National Technical
University of Athens*

Geraldo F. Oliveira

ETH Zürich

Onur Mutlu

ETH Zürich

PrIM Benchmarks: Application Domains

Domain	Benchmark	Short name
Dense linear algebra	Vector Addition	VA
	Matrix-Vector Multiply	GEMV
Sparse linear algebra	Sparse Matrix-Vector Multiply	SpMV
Databases	Select	SEL
	Unique	UNI
Data analytics	Binary Search	BS
	Time Series Analysis	TS
Graph processing	Breadth-First Search	BFS
Neural networks	Multilayer Perceptron	MLP
Bioinformatics	Needleman-Wunsch	NW
Image processing	Image histogram (short)	HST-S
	Image histogram (large)	HST-L
Parallel primitives	Reduction	RED
	Prefix sum (scan-scan-add)	SCAN-SSA
	Prefix sum (reduce-scan-scan)	SCAN-RSS
	Matrix transposition	TRNS

PrIM Benchmarks are Open Source

- All microbenchmarks, benchmarks, and scripts
- <https://github.com/CMU-SAFARI/prim-benchmarks>

The screenshot shows the GitHub repository page for 'CMU-SAFARI / prim-benchmarks'. The repository has 2 stars, 1 fork, and 1 issue. The README.md file is the main document viewed. It contains a brief introduction to PrIM (Processing-In-Memory Benchmarks), mentioning it's the first benchmark suite for UPMEM PIM architecture. It highlights the repository's purpose of evaluating, analyzing, and characterizing the UPMEM PIM architecture. The README also notes that PrIM provides workloads for programming, architecture, and system researchers, and includes microbenchmarks for assessing architecture limits.

CMU-SAFARI / prim-benchmarks

Code Issues Pull requests Actions Projects Wiki Security Insights Settings

main prim-benchmarks / README.md Go to file ...

Juan Gomez Luna PrIM -- first commit Latest commit 3de4b49 9 days ago History

1 contributor

168 lines (132 sloc) 5.79 KB Raw Blame

PrIM (Processing-In-Memory Benchmarks)

PrIM is the first benchmark suite for a real-world processing-in-memory (PIM) architecture. PrIM is developed to evaluate, analyze, and characterize the first publicly-available real-world processing-in-memory (PIM) architecture, the [UPMEM](#) PIM architecture. The UPMEM PIM architecture combines traditional DRAM memory arrays with general-purpose in-order cores, called DRAM Processing Units (DPUs), integrated in the same chip.

PrIM provides a common set of workloads to evaluate the UPMEM PIM architecture with and can be useful for programming, architecture and system researchers all alike to improve multiple aspects of future PIM hardware and software. The workloads have different characteristics, exhibiting heterogeneity in their memory access patterns, operations and data types, and communication patterns. This repository also contains baseline CPU and GPU implementations of PrIM benchmarks for comparison purposes.

PrIM also includes a set of microbenchmarks can be used to assess various architecture limits such as compute throughput and memory bandwidth.

Understanding a Modern PIM Architecture

Benchmarking a New Paradigm: Experimental Analysis and Characterization of a Real Processing-in-Memory System

JUAN GÓMEZ-LUNA¹, IZZAT EL HAJJ², IVAN FERNANDEZ^{1,3}, CHRISTINA GIANNOULA^{1,4},
GERALDO F. OLIVEIRA¹, AND ONUR MUTLU¹

¹ETH Zürich

²American University of Beirut

³University of Malaga

⁴National Technical University of Athens

Corresponding author: Juan Gómez-Luna (e-mail: juang@ethz.ch).

<https://arxiv.org/pdf/2105.03814.pdf>

<https://github.com/CMU-SAFARI/prim-benchmarks>

Understanding a Modern PIM Architecture

**Understanding a Modern Processing-in-Memory Architecture:
Benchmarking and Experimental Characterization**

Juan Gómez Luna, Izzat El Hajj,
Ivan Fernandez, Christina Giannoula,
Geraldo F. Oliveira, Onur Mutlu

<https://arxiv.org/pdf/2105.03814.pdf>
<https://github.com/CMU-SAFARI/prim-benchmarks>

ETH Zürich SAFARI zoom

2:26 / 2:57:10

SAFARI Live Seminar: Understanding a Modern Processing-in-Memory Architecture

2,579 views • Streamed live on Jul 12, 2021

93 likes 0 dislikes SHARE SAVE ...



Onur Mutlu Lectures
18.7K subscribers

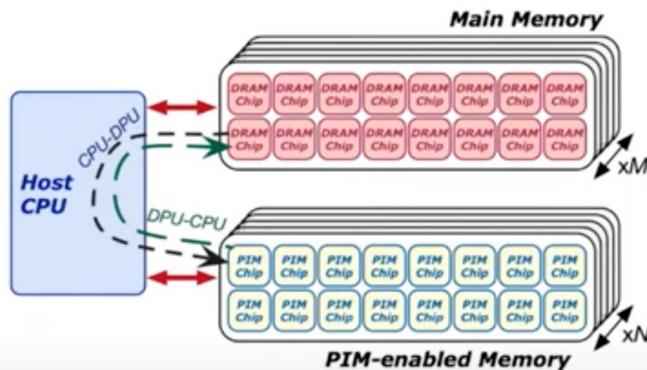
SUBSCRIBED



More on Analysis of the UPMEM PIM Engine

Inter-DPU Communication

- There is no direct communication channel between DPUs



- Inter-DPU communication takes place via the host CPU using CPU-DPU and DPU-CPU transfers
- Example communication patterns:
 - Merging of partial results to obtain the final result
 - Only DPU-CPU transfers
 - Redistribution of intermediate results for further computation
 - DPU-CPU transfers and CPU-DPU transfers



SAFARI Live Seminar: Understanding a Modern Processing-in-Memory Architecture

1,868 views • Streamed live on Jul 12, 2021

81 likes 0 dislikes SHARE SAVE ...



Onur Mutlu Lectures
17.6K subscribers

Talk Title: Understanding a Modern Processing-in-Memory Architecture: Benchmarking and Experimental Characterization
Dr. Juan Gómez-Luna, SAFARI Research Group, D-ITET, ETH Zurich

ANALYTICS EDIT VIDEO

More on Analysis of the UPMEM PIM Engine

Data Movement in Computing Systems

- Data movement dominates performance and is a major system energy bottleneck
- Total system energy: data movement accounts for
 - 62% in consumer applications*,
 - 40% in scientific applications*,
 - 35% in mobile applications*

The diagram illustrates the data movement architecture within a System-on-Chip (SoC). It shows a central SoC boundary (dashed line) containing a CPU (blue), GPU (green), and L2 cache (orange). Bidirectional arrows connect the CPU and GPU to the L2 cache. Outside the SoC, a large green rectangle represents DRAM. A thick yellow double-headed arrow connects the SoC's L2 cache to the DRAM. Below the SoC, four peripheral components are connected to the L2 cache via bidirectional arrows: Video Encoder, Video Decoder, Audio, and Display Engine. The entire diagram is labeled "Data Movement" at the top.

* Boroumand et al., "Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks," ASPLOS 2018
* Kestor et al., "Quantifying the Energy Cost of Data Movement in Scientific Applications," IISWC 2013
* Pandian and Wu, "Quantifying the energy cost of data movement for emerging smart phone workloads on mobile platforms," IISWC 2014

SAFARI

3

▶ ▶ 🔍 2:27 / 21:28

▶ CC HD □ □

Understanding a Modern Processing-in-Memory Arch: Benchmarking & Experimental Characterization; 21m

3,482 views • Premiered Jul 25, 2021

1 38 0 SHARE SAVE ...

Onur Mutlu Lectures
17.9K subscribers

ANALYTICS

EDIT VIDEO

https://www.youtube.com/watch?v=Pp9jSU2b9oM&list=PL5Q2soXY2Zi8_VVChACnON4sfh2bJ5IrD&index=159

More on PRIM Benchmarks

- Juan Gomez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu,
"Benchmarking a New Paradigm: An Experimental Analysis of a Real Processing-in-Memory Architecture"

Preprint in [arXiv](#), 9 May 2021.

[[arXiv preprint](#)]

[[PrIM Benchmarks Source Code](#)]

[[Slides \(pptx\)](#) ([pdf](#))]

[[Long Talk Slides \(pptx\)](#) ([pdf](#))]

[[Short Talk Slides \(pptx\)](#) ([pdf](#))]

[[SAFARI Live Seminar Slides \(pptx\)](#) ([pdf](#))]

[[SAFARI Live Seminar Video](#) (2 hrs 57 mins)]

[[Lightning Talk Video](#) (3 minutes)]

UPMEM PIM System Summary & Analysis

- Juan Gomez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu,

"Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware"

Invited Paper at Workshop on Computing with Unconventional Technologies (CUT), Virtual, October 2021.

[[arXiv version](#)]

[[PrIM Benchmarks Source Code](#)]

[[Slides \(pptx\) \(pdf\)](#)]

[[Talk Video](#) (37 minutes)]

[[Lightning Talk Video](#) (3 minutes)]

Benchmarking Memory-Centric Computing Systems: Analysis of Real Processing-in-Memory Hardware

Juan Gómez-Luna

ETH Zürich

Izzat El Hajj

*American University
of Beirut*

Ivan Fernandez

*University
of Malaga*

Christina Giannoula

*National Technical
University of Athens*

Geraldo F. Oliveira

ETH Zürich

Onur Mutlu

ETH Zürich

ML Training on a Real PIM System

Machine Learning Training on a Real Processing-in-Memory System

Juan Gómez-Luna¹ Yuxin Guo¹ Sylvan Brocard² Julien Legriel²
Remy Cimadomo² Geraldo F. Oliveira¹ Gagandeep Singh¹ Onur Mutlu¹

¹ETH Zürich ²UPMEM

An Experimental Evaluation of Machine Learning Training on a Real Processing-in-Memory System

Juan Gómez-Luna¹ Yuxin Guo¹ Sylvan Brocard² Julien Legriel²
Remy Cimadomo² Geraldo F. Oliveira¹ Gagandeep Singh¹ Onur Mutlu¹

¹ETH Zürich ²UPMEM

Short version: <https://arxiv.org/pdf/2206.06022.pdf>

Long version: <https://arxiv.org/pdf/2207.07886.pdf>

<https://www.youtube.com/watch?v=qeuKNs5XI3g&t=11226s>

AIM (PIM Sequence Alignment Framework)

Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez-Luna,
Onur Mutlu, Izzat El Hajj

"A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems"

arXiv, 2022

[\[Source code\]](#)

A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems

Safaa Diab¹, Amir Nassereldine¹, Mohammed Alser², Juan Gómez Luna², Onur Mutlu², Izzat El Hajj¹

¹*American University of Beirut, Lebanon* ²*ETH Zürich, Switzerland*

Connecting Basecalling and Read Mapping in PIM

Haiyu Mao, Mohammed Alser, Mohammad Sadrosadati, Can Firtina, Akanksha Baranwal, Damla Senol Cali, Aditya Manglik, Nour Almadhoun Alserr, and Onur Mutlu,

"GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping"

Proceedings of the *55th International Symposium on Microarchitecture (MICRO)*, Chicago, Illinois, October 2022.

GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping

Haiyu Mao¹ Mohammed Alser¹ Mohammad Sadrosadati¹ Can Firtina¹ Akanksha Baranwal¹
Damla Senol Cali² Aditya Manglik¹ Nour Almadhoun Alserr¹ Onur Mutlu¹

¹ETH Zürich

²Bionano Genomics

Agenda

- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

Newer Genome Sequencing Technologies

Nanopore sequencing technology and tools for genome assembly: computational analysis of the current state, bottlenecks and future directions

Damla Senol Cali ✉, Jeremie S Kim, Saugata Ghose, Can Alkan, Onur Mutlu

Briefings in Bioinformatics, bby017, <https://doi.org/10.1093/bib/bby017>

Published: 02 April 2018 Article history ▾



Oxford Nanopore MinION

Senol Cali+, “[Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions](#),” *Briefings in Bioinformatics*, 2018.

[Open arxiv.org version] [Slides (pptx) (pdf)] [Talk Video at AACBB 2019]

New Applications: Graph Genomes

- Damla Senol Cali, Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo, Can Firtina, Meryem Banu Cavlak, Jeremie Kim, Nika MansouriGhiasi, Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu,
"SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping"

Proceedings of the 49th International Symposium on Computer Architecture (ISCA), New York, June 2022.

[[arXiv version](#)]

SeGram: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping

Damla Senol Cali¹ Konstantinos Kanellopoulos² Joël Lindegger² Zülal Bingöl³
Gurpreet S. Kalsi⁴ Ziyi Zuo⁵ Can Firtina² Meryem Banu Cavlak² Jeremie Kim²
Nika Mansouri Ghiasi² Gagandeep Singh² Juan Gómez-Luna² Nour Almadhoun Alserr²
Mohammed Alser² Sreenivas Subramoney⁴ Can Alkan³ Saugata Ghose⁶ Onur Mutlu²

¹Bionano Genomics ²ETH Zürich ³Bilkent University ⁴Intel Labs

⁵Carnegie Mellon University ⁶University of Illinois Urbana-Champaign

New Applications: Ref Genome Updates

RESEARCH

AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes

Jeremie S. Kim¹, Can Firtina¹, Meryem Banu Cavlak², Damla Senol Cali³, Nastaran Hajinazar^{1,4},
Mohammed Alser¹, Can Alkan² and Onur Mutlu^{1,2,3*}

https://people.inf.ethz.ch/omutlu/pub/AirLift_genome-remapper_arxiv21.pdf

Remapping Reads Between References

- Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Nastaran Hajinazar, Mohammed Alser, Can Alkan, and Onur Mutlu,
["AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes"](#)

Preprint in [arXiv](#) and [bioRxiv](#), 2021.

[[bioRxiv preprint](#)]

[[arXiv preprint](#)]

[[AirLift Source Code and Data](#)]

METHOD

AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes

Jeremie S. Kim^{1†}, Can Firtina^{1†}, Meryem Banu Cavlak², Damla Senol Cali³, Nastaran Hajinazar^{1,4},
Mohammed Alser¹, Can Alkan² and Onur Mutlu^{1,2,3*}

Mapping Constant Regions Between References

- Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Can Alkan, and Onur Mutlu,

["FastRemap: A Tool for Quickly Remapping Reads between Genome Assemblies"](#)

Bioinformatics, btac554.

[[FastRemap Source Code](#)]

FastRemap: A Tool for Quickly Remapping Reads between Genome Assemblies

Jeremie S. Kim¹

Can Firtina¹

Meryem Banu Cavlak¹

Damla Senol Cali^{2,3}

Can Alkan⁴

Onur Mutlu^{1,2,4}

¹*ETH Zürich*

²*Carnegie Mellon University*

³*Bionano Genomics*

⁴*Bilkent University*

Newer Genome Sequencing Technologies

Nanopore sequencing technology and tools for genome assembly: computational analysis of the current state, bottlenecks and future directions

Damla Senol Cali ✉, Jeremie S Kim, Saugata Ghose, Can Alkan, Onur Mutlu

Briefings in Bioinformatics, bby017, <https://doi.org/10.1093/bib/bby017>

Published: 02 April 2018 Article history ▾



Oxford Nanopore MinION

Senol Cali+, “[Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions](#),” *Briefings in Bioinformatics*, 2018.

[Open arxiv.org version] [Slides (pptx) (pdf)] [Talk Video at AACBB 2019]

Recall: High-Throughput Sequencing

- Massively parallel sequencing technology
 - Illumina, Roche 454, Ion Torrent, SOLID...
- Small DNA fragments are first amplified and then sequenced in parallel, leading to
 - High throughput
 - High speed
 - Low cost
 - Short reads
 - Amplification step limits the read length since too short or too long fragments are not amplified well.
- Sequencing is done by either reading optical signals as each base is added, or by detecting hydrogen ions instead of light, leading to:
 - Low error rates (relatively)
 - Reads lack information about their order and which part of genome they are originated from

Nanopore Sequencing Technology

- **Nanopore sequencing** is an emerging and a promising single-molecule DNA sequencing technology

- First nanopore sequencing device, **MinION**, made commercially available by **Oxford Nanopore Technologies (ONT)** in **May 2014**.
 - Inexpensive
 - Long read length (> 882K bp)
 - Portable: Pocket-sized
 - Produces data in real-time

Nanopore Sequencing Technology



an emerging and a promising
sequencing technology
read length → Longer read length

- First nanopore sequencing device, **MinION**, made commercially available by **Oxford Nanopore Technologies (ONT)** in **May 2014**.
 - Inexpensive
 - Long read length (> 882K bp)
 - Portable: Pocket-sized
 - Produces data in real-time



Oxford Nanopore Sequencers



MinION Mk1B



MinION Mk1C



GridION Mk1



PromethION 24/48

	MinION Mk1B	MinION Mk1C	GridION Mk1	PromethION 24	PromethION 48
Read length	> 2Mb	> 2Mb	> 2Mb	> 2Mb	> 2Mb
Yield per flow cell	50 Gb	50 Gb	50 Gb	220 Gb	220 Gb
Number of flow cells per device	1	1	5	24	48
Yield per device	<50 Gb	<50 Gb	<250 Gb	<5.2 Tb	<10.5 Tb
Starting price	\$1,000	\$4,990	\$49,995	\$195,455	\$327,455

Illumina Sequencers

illumina®



iSeq 100

MiniSeq

MiSeq

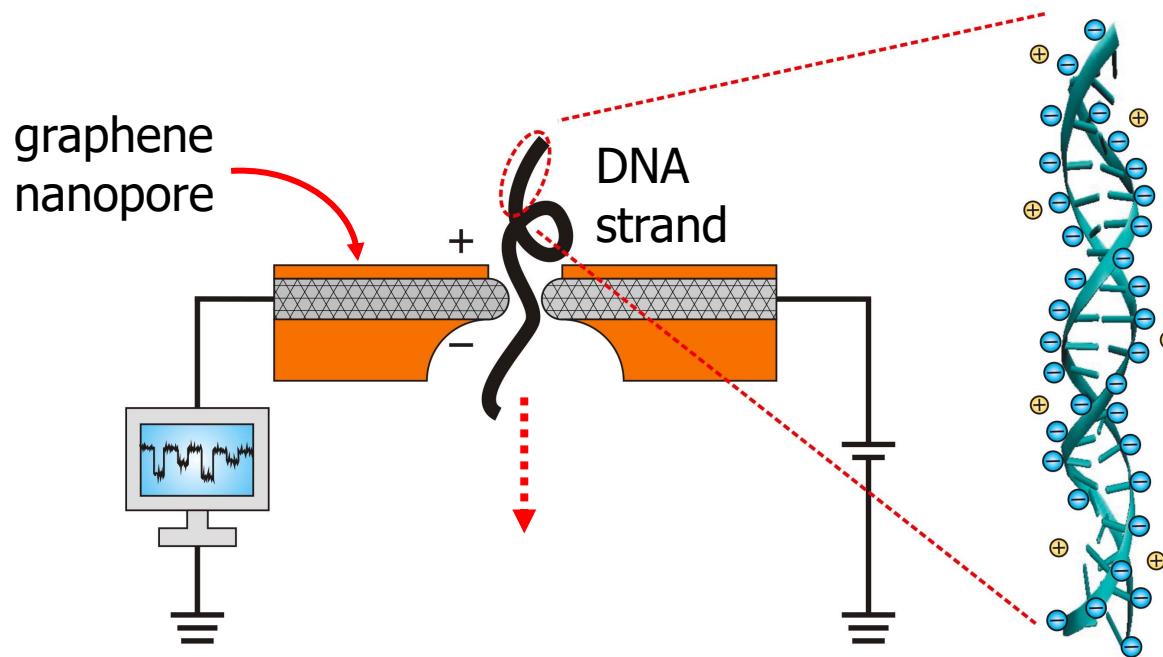
NextSeq 550

NextSeq 2000

NovaSeq 6000

Run time	9.5–19 hrs	4–24 hrs	4–55 hrs	12–30 hrs	24–48 hrs	13–44 hrs
Max. reads per run	4 million	25 million	25 million	400 million	1 billion	20 billion
Max. read length	2 × 150 bp	2 × 150 bp	2 × 300 bp	2 × 150 bp	2 × 150 bp	2 × 250
Max. output	1.2 Gb	7.5 Gb	15 Gb	120 Gb	300 Gb	6000 Gb
Estimated price	\$19,900	\$49,500	\$128,000	\$275,000	\$335,000	\$985,000

How Does Nanopore Sequencing Work?



- **Nanopore** is a nano-scale hole (<20nm).
- In nanopore sequencers, an **ionic current** passes through the nanopores
- When the DNA strand passes through the nanopore, the sequencer measures the the **change in current**
- This change is used to identify the bases in the strand with the help of **different electrochemical structures** of the different bases

Advantages of Nanopore Sequencing

Nanopores:

- Do *not* require any labeling of the DNA or nucleotide for detection during sequencing
- Rely on the electronic or chemical structure of the different nucleotides for identification
- Allow sequencing **very long reads**, and
- Provide **portability, low cost, and high throughput.**

Challenges of Nanopore Sequencing

- One major drawback: **high error rates**
- Nanopore sequence analysis tools have a critical role to:
 - overcome high error rates
 - take better advantage of the technology
- **Faster tools** are critically needed to:
 - Take better advantage of the **real-time data production** capability of nanopore sequencing
 - Enable **fast, real-time data analysis**

Nanopore Genome Assembly Pipeline

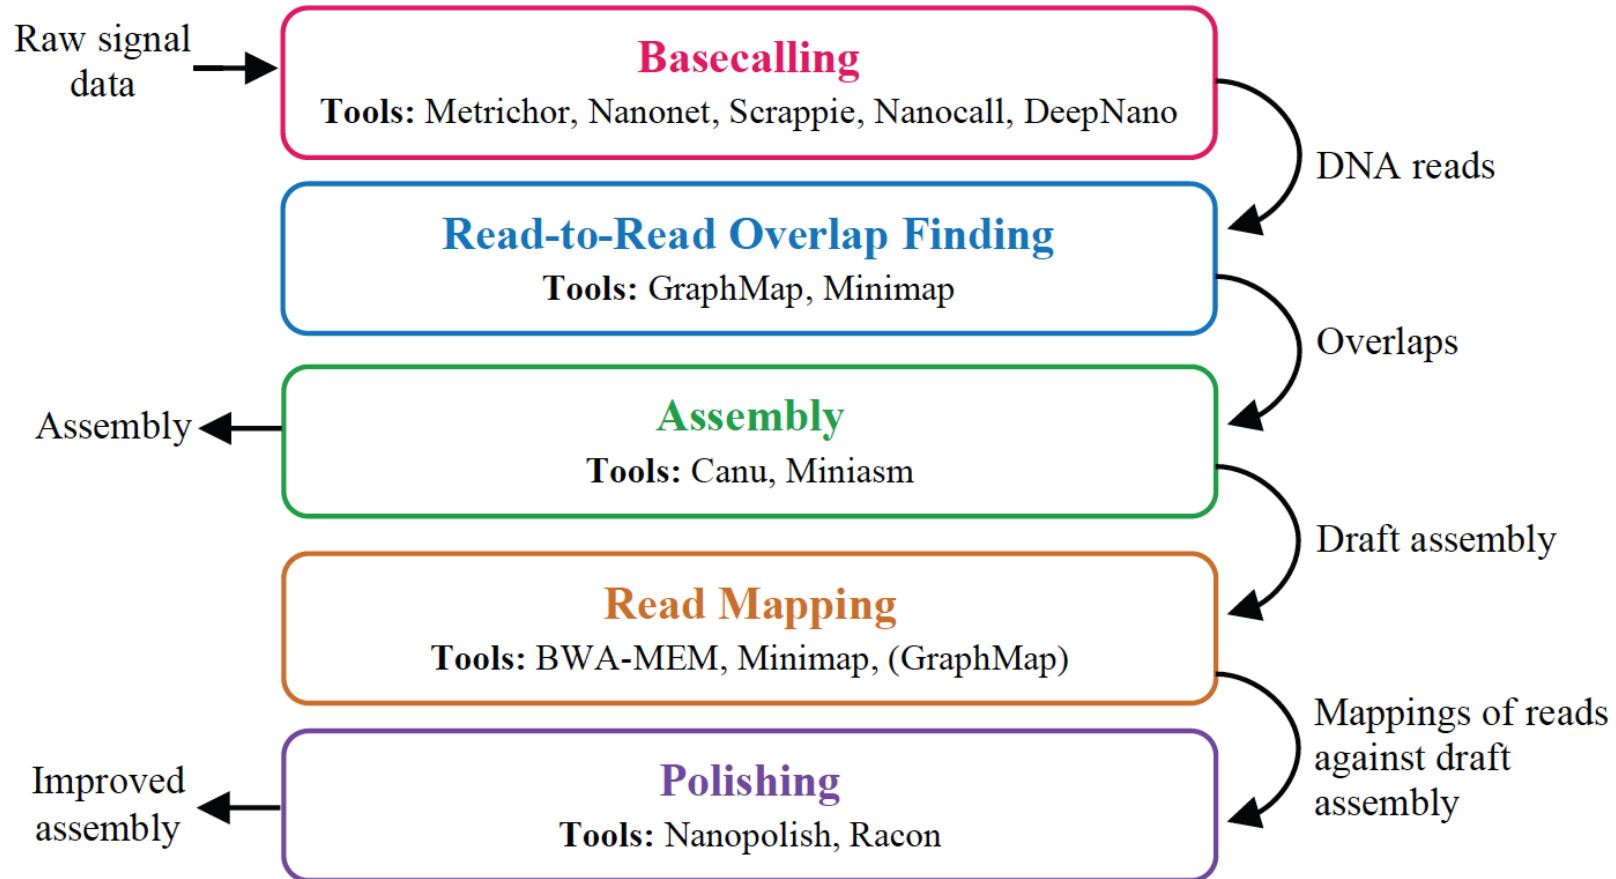


Figure 1. The analyzed genome assembly pipeline using nanopore sequence data, with its five steps and the associated tools for each step.

Senol Cali+, “**Nanopore Sequencing Technology and Tools for Genome Assembly**” Briefings in Bioinformatics, 2018.

Nanopore Genome Assembly Tools (I)

Table 12. Accuracy analysis results for the full pipeline with a focus on the last two steps.

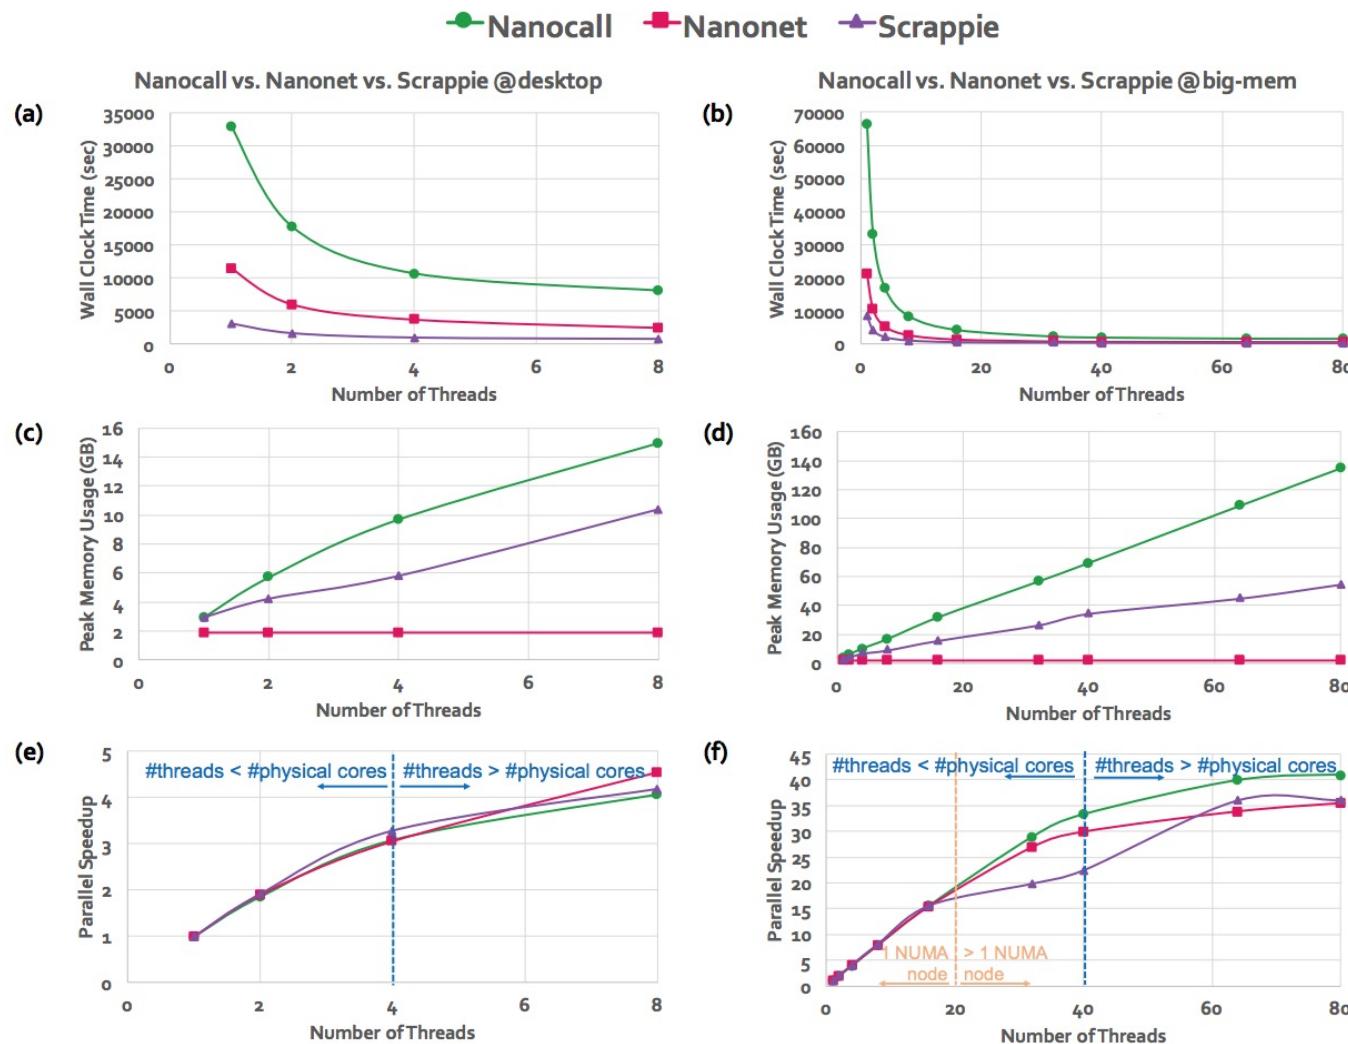
						Number of Bases	Number of Contigs	Identity (%)	Coverage (%)	Number of Mismatches	Number of Indels
1	Metricor	+	—	+	Canu	+	BWA-MEM	+	Nanopolish	4,683,072	1 99.48 99.93 8,198 15,581
2	Metricor	+	Minimap	+	Miniasm	+	BWA-MEM	+	Nanopolish	4,540,352	1 92.33 96.31 162,884 182,965
3	Metricor	+	GraphMap	+	Miniasm	+	BWA-MEM	+	Nanopolish	4,637,916	2 92.38 95.80 159,206 180,603
4	Metricor	+	—	+	Canu	+	BWA-MEM	+	Racon	4,650,502	1 98.46 100.00 18,036 51,842
5	Metricor	+	—	+	Canu	+	Minimap	+	Racon	4,648,710	1 98.45 100.00 17,906 52,168
6	Metricor	+	Minimap	+	Miniasm	+	BWA-MEM	+	Racon	4,598,267	1 97.70 99.91 24,014 82,906
7	Metricor	+	Minimap	+	Miniasm	+	Minimap	+	Racon	4,600,109	1 97.78 100.00 23,339 79,721
8	Nanonet	+	—	+	Canu	+	BWA-MEM	+	Racon	4,622,285	1 98.48 100.00 16,872 52,509
9	Nanonet	+	—	+	Canu	+	Minimap	+	Racon	4,620,597	1 98.49 100.00 16,874 52,232
10	Nanonet	+	Minimap	+	Miniasm	+	BWA-MEM	+	Racon	4,593,402	1 98.01 99.97 20,322 72,284
11	Nanonet	+	Minimap	+	Miniasm	+	Minimap	+	Racon	4,592,907	1 98.04 100.00 20,170 70,705
12	Scrapie	+	—	+	Canu	+	BWA-MEM	+	Racon	4,673,871	1 98.40 99.98 13,583 60,612
13	Scrapie	+	—	+	Canu	+	Minimap	+	Racon	4,673,606	1 98.40 99.98 13,798 60,423
14	Scrapie	+	Minimap	+	Miniasm	+	BWA-MEM	+	Racon	5,157,041	8 97.87 99.80 18,085 78,492
15	Scrapie	+	Minimap	+	Miniasm	+	Minimap	+	Racon	5,156,375	8 97.87 99.94 17,922 77,807
16	Nanocall	+	—	+	Canu	+	BWA-MEM	+	Racon	1,383,851	86 93.49 28.82 19,057 65,244
17	Nanocall	+	—	+	Canu	+	Minimap	+	Racon	1,367,834	86 94.43 28.74 15,610 55,275
18	Nanocall	+	Minimap	+	Miniasm	+	BWA-MEM	+	Racon	4,707,961	5 90.75 97.11 91,502 347,005
19	Nanocall	+	Minimap	+	Miniasm	+	Minimap	+	Racon	4,673,069	5 92.23 97.10 72,646 291,918
20	DeepNano	+	—	+	Canu	+	BWA-MEM	+	Racon	7,429,290	106 96.46 99.24 27,811 102,682
21	DeepNano	+	—	+	Canu	+	Minimap	+	Racon	7,404,454	106 96.03 99.21 34,023 110,640
22	DeepNano	+	Minimap	+	Miniasm	+	BWA-MEM	+	Racon	4,566,253	1 96.76 99.86 25,791 125,386
23	DeepNano	+	Minimap	+	Miniasm	+	Minimap	+	Racon	4,571,810	1 96.90 99.97 24,994 119,519

Nanopore Genome Assembly Tools (II)

Table 13. Performance analysis results for the full pipeline with a focus on the last two steps.

				Step 4: Read Mapper			Step 5: Polisher		
				Wall Clock Time (h:m:s)	CPU Time (h:m:s)	Memory Usage (GB)	Wall Clock Time (h:m:s)	CPU Time (h:m:s)	Memory Usage (GB)
1	Metricchor	+ —	+ Canu	+ BWA-MEM	+ Nanopolish	24:43	15:47:21	5.26	5:51:00 191:18:52 13.38
2	Metricchor	+ Minimap	+ Miniasm	+ BWA-MEM	+ Nanopolish	12:33	7:50:54	3.75	122:52:00 4458:36:10 31.36
3	Metricchor	+ GraphMap	+ Miniasm	+ BWA-MEM	+ Nanopolish	12:47	7:57:58	3.60	129:46:00 4799:03:51 31.31
4	Metricchor	+ —	+ Canu	+ BWA-MEM	+ Racon	24:20	15:43:40	6.60	14:44 9:09:22 8.11
5	Metricchor	+ —	+ Canu	+ Minimap	+ Racon	3	1:35	0.26	15:12 9:45:33 14.55
6	Metricchor	+ Minimap	+ Miniasm	+ BWA-MEM	+ Racon	12:10	7:48:10	5.19	15:43 9:33:39 9.98
7	Metricchor	+ Minimap	+ Miniasm	+ Minimap	+ Racon	3	1:24	0.26	20:28 8:57:40 18.24
8	Nanonet	+ —	+ Canu	+ BWA-MEM	+ Racon	9:08	5:53:18	4.84	6:33 4:02:10 4.47
9	Nanonet	+ —	+ Canu	+ Minimap	+ Racon	2	54	0.26	6:45 4:17:26 7.93
10	Nanonet	+ Minimap	+ Miniasm	+ BWA-MEM	+ Racon	4:40	2:58:02	3.88	7:08 4:19:30 5.35
11	Nanonet	+ Minimap	+ Miniasm	+ Minimap	+ Racon	2	46	0.26	7:01 4:18:48 9.53
12	Scrappie	+ —	+ Canu	+ BWA-MEM	+ Racon	33:41	21:11:06	8.66	13:32 8:24:44 7.58
13	Scrappie	+ —	+ Canu	+ Minimap	+ Racon	3	1:39	0.27	18:45 7:43:17 13.20
14	Scrappie	+ Minimap	+ Miniasm	+ BWA-MEM	+ Racon	22:41	14:31:00	6.08	14:37 8:53:59 9.50
15	Scrappie	+ Minimap	+ Miniasm	+ Minimap	+ Racon	3	1:27	0.27	15:10 9:02:45 12.72
16	Nanocall	+ —	+ Canu	+ BWA-MEM	+ Racon	4:52	3:01:15	3.80	11:07 3:26:52 5.63
17	Nanocall	+ —	+ Canu	+ Minimap	+ Racon	3	1:16	0.22	7:28 2:50:35 3.62
18	Nanocall	+ Minimap	+ Miniasm	+ BWA-MEM	+ Racon	16:06	10:27:20	5.06	18:56 11:32:45 11.47
19	Nanocall	+ Minimap	+ Miniasm	+ Minimap	+ Racon	4	1:18	0.26	11:49 7:08:59 10.98
20	DeepNano	+ —	+ Canu	+ BWA-MEM	+ Racon	17:36	11:30:20	4.43	12:48 7:13:04 8.88
21	DeepNano	+ —	+ Canu	+ Minimap	+ Racon	3	1:24	0.28	11:39 6:55:01 3.73
22	DeepNano	+ Minimap	+ Miniasm	+ BWA-MEM	+ Racon	8:15	5:22:29	4.11	14:16 8:34:32 10.30
23	DeepNano	+ Minimap	+ Miniasm	+ Minimap	+ Racon	3	1:10	0.26	12:29 7:55:32 17.11

Nanopore Genome Assembly Tools (III)



More on Nanopore Sequencing & Tools

Nanopore sequencing technology and tools for genome assembly: computational analysis of the current state, bottlenecks and future directions

Damla Senol Cali ✉, Jeremie S Kim, Saugata Ghose, Can Alkan, Onur Mutlu

Briefings in Bioinformatics, bby017, <https://doi.org/10.1093/bib/bby017>

Published: 02 April 2018 Article history ▾



BiB



arXiv

Senol Cali+, “**Nanopore Sequencing Technology and Tools for Genome Assembly: Computational Analysis of the Current State, Bottlenecks and Future Directions**,” *Briefings in Bioinformatics*, 2018.
[Preliminary [arxiv.org version](#)]

Why Do We Care? An Example from 2020

200 Oxford Nanopore sequencers have left UK for China, to support rapid, near-sample coronavirus sequencing for outbreak surveillance

Fri 31st January 2020

Following extensive support of, and collaboration with, public health professionals in China, Oxford Nanopore has shipped an additional 200 MinION sequencers and related consumables to China. These will be used to support the ongoing surveillance of the current coronavirus outbreak, adding to a large number of the devices already installed in the country.



Each MinION sequencer is approximately the size of a stapler, and can provide rapid sequence information about the coronavirus.



700Kg of Oxford Nanopore sequencers and consumables are on their way for use by Chinese scientists in understanding the current coronavirus outbreak.

Sequencing of COVID-19

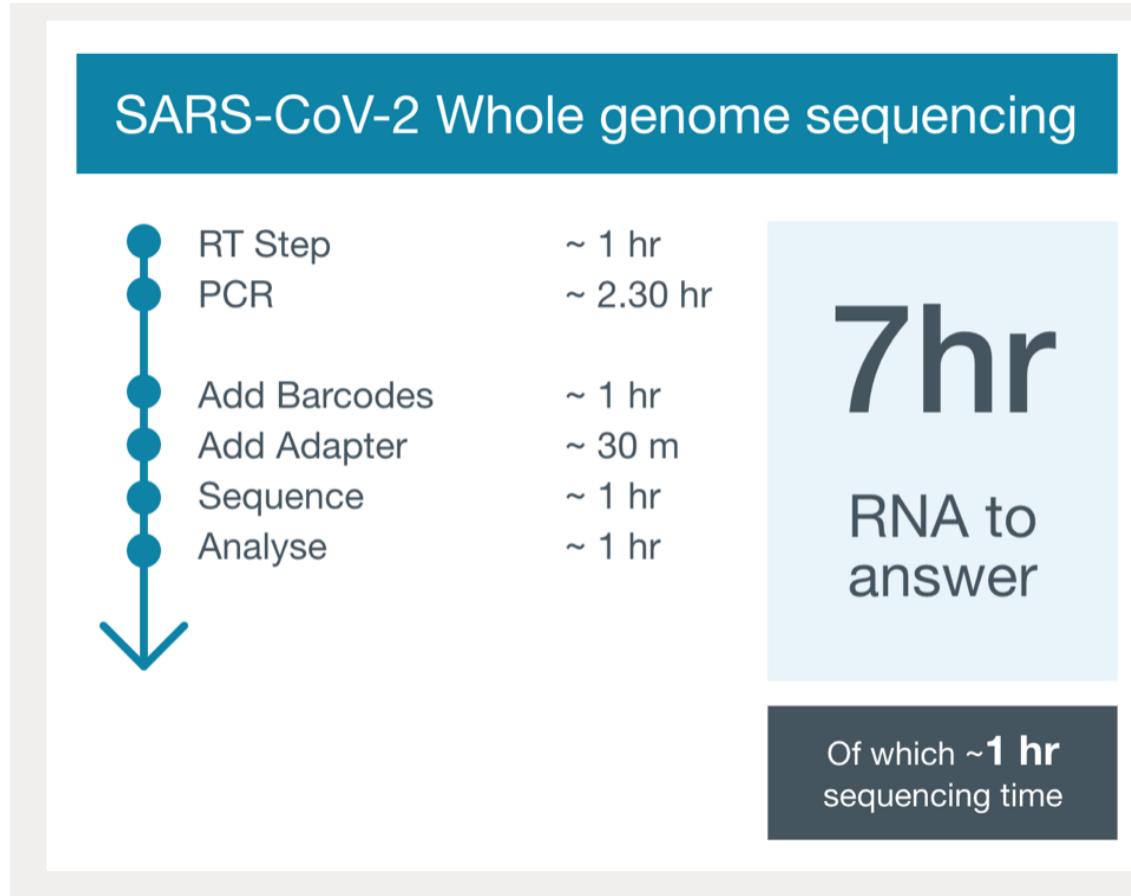
■ **Whole genome sequencing (WGS) and sequence data analysis are important**

- To detect the virus from a human sample such as saliva, Bronchoalveolar fluid etc.
- To understand the sources and modes of transmission of the virus
- To discover the genomic characteristics of the virus, and compare with better-known viruses (e.g., 02-03 SARS epidemic)
- To design and evaluate the diagnostic tests and deep-dive studies

■ **Two key areas of COVID-19 genomic research**

- To sequence the genome of the virus itself, COVID-19, in order to track the mutations in the virus.
- To explore the genes of infected patients. This analysis can be used to understand why some people get more severe symptoms than others, as well as, help with the development of new treatments in the future.

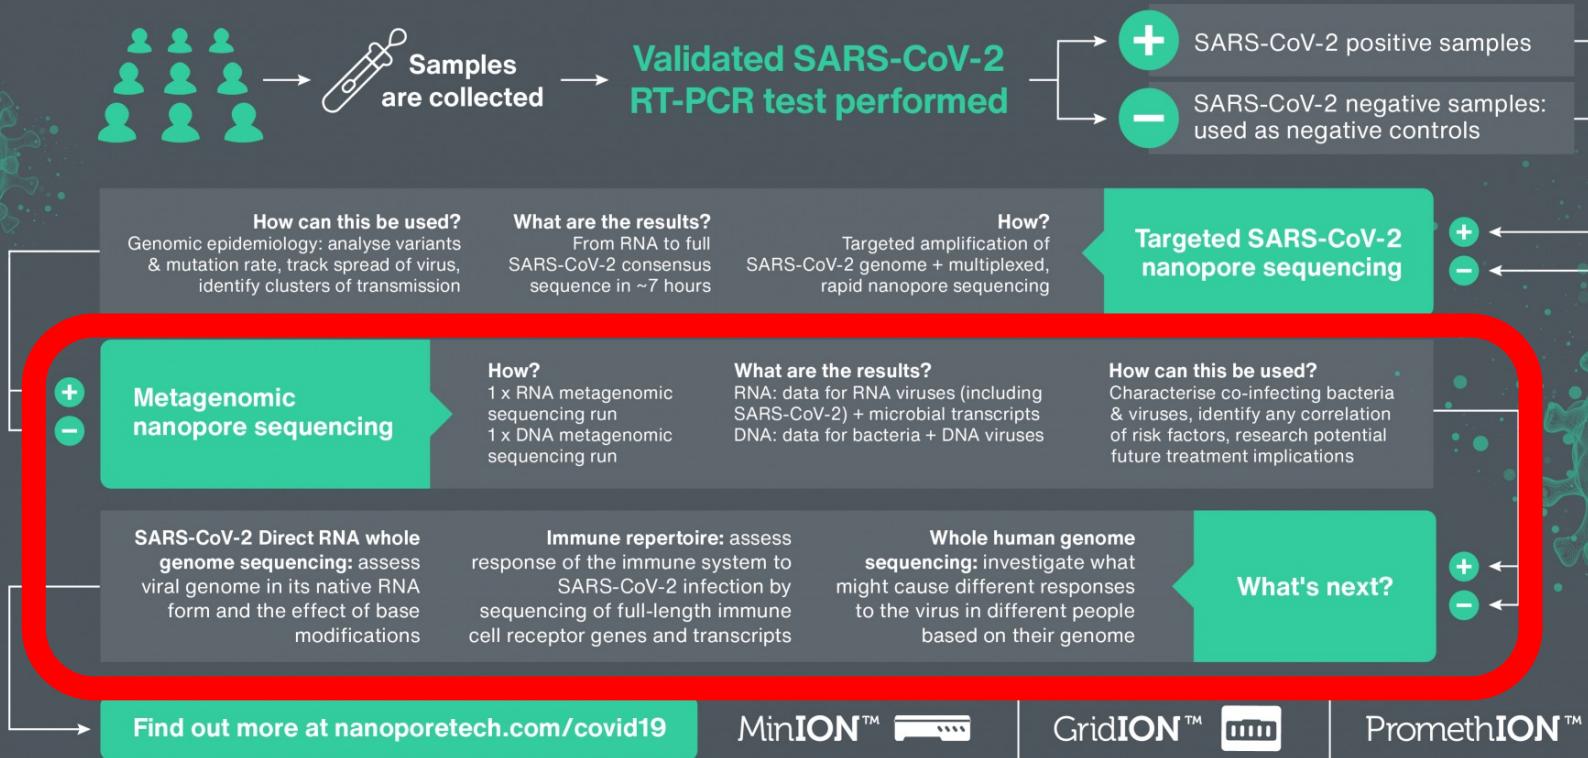
COVID-19 Nanopore Sequencing (I)



- From ONT (<https://nanoporetech.com/covid-19/overview>)

COVID-19 Nanopore Sequencing (II)

How are scientists using nanopore sequencing to research COVID-19?



Oxford Nanopore Technologies, the Wheel icon, GridION, PromethION and MinION are registered trademarks of Oxford Nanopore Technologies in various countries. © 2020 Oxford Nanopore Technologies. All rights reserved. Oxford Nanopore Technologies' products are currently for research use only IG_1061[EN]_V1_03April2020

- From ONT (<https://nanoporetech.com/covid-19/overview>)

A Bright Future for Intelligent Genome Analysis

Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu
[“Accelerating Genome Analysis: A Primer on an Ongoing Journey”](#) IEEE Micro, August 2020.



MinION from ONT

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Sept.-Oct. 2020, pp. 65-75, vol. 40
DOI Bookmark: [10.1109/MM.2020.3013728](https://doi.org/10.1109/MM.2020.3013728)

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41
DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)



SmidgION from ONT

Agenda

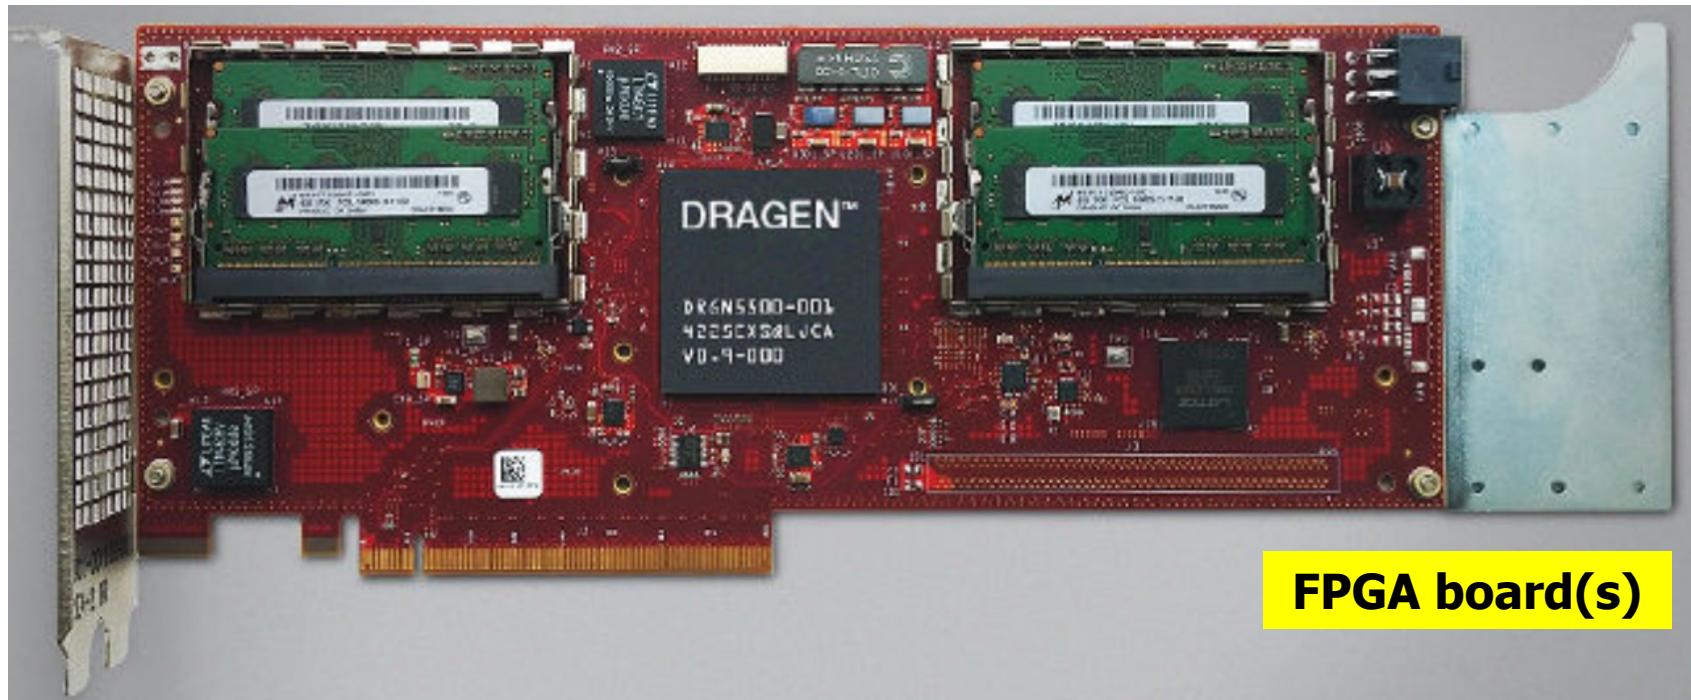
- The Problem: DNA Read Mapping
 - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
 - Exploiting Structure of the Genome
 - Exploiting SIMD Instructions
- Hardware Acceleration
 - Specialized Architectures
 - Processing in Memory & Storage
- Future Opportunities: New Technologies & Applications

Conclusion

Things Are Happening In Industry

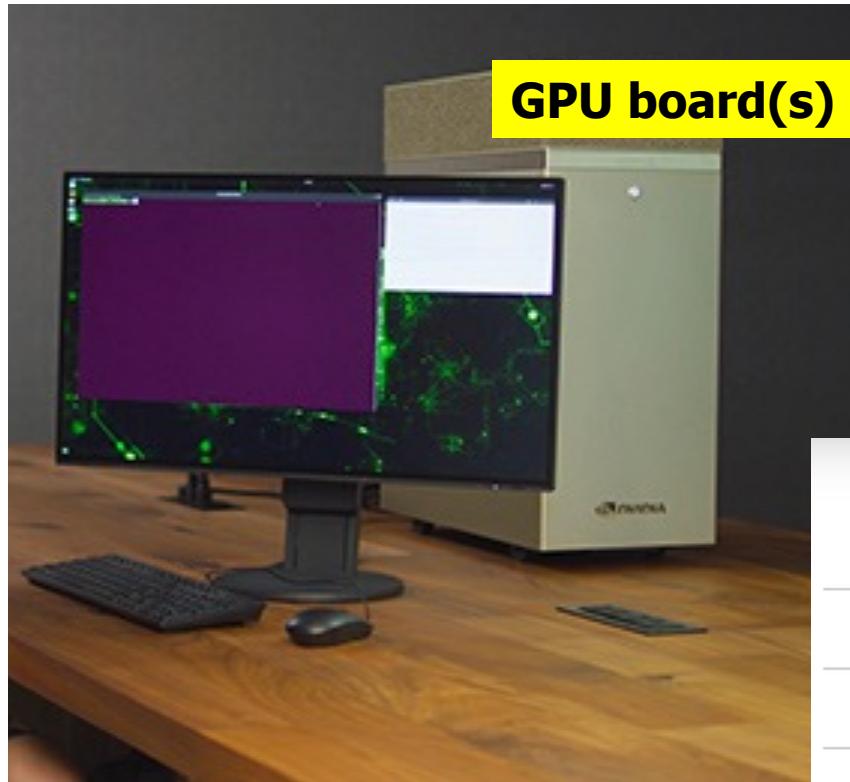
Illumina DRAGEN Bio-IT Platform (2018)

- Processes whole genome at 30x coverage in ~25 minutes with hardware support for data compression

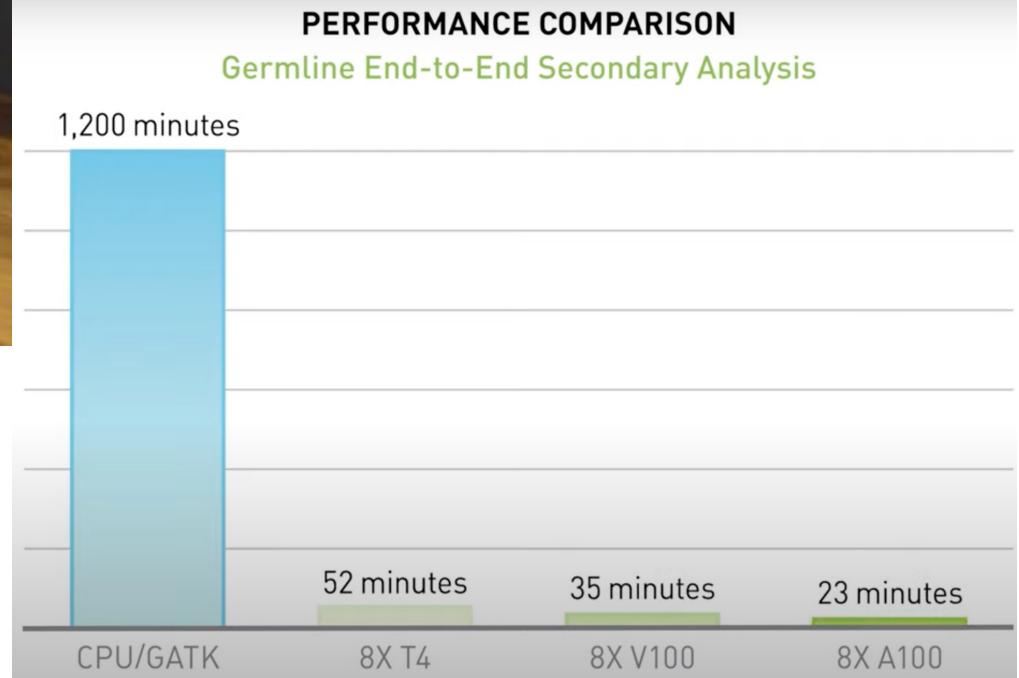


emea.illumina.com/products/by-type/informatics-products/dragen-bio-it-platform.html
emea.illumina.com/company/news-center/press-releases/2018/2349147.html

NVIDIA Clara Parabricks (2020)



A University of Michigan startup in 2018 joined NVIDIA in 2020

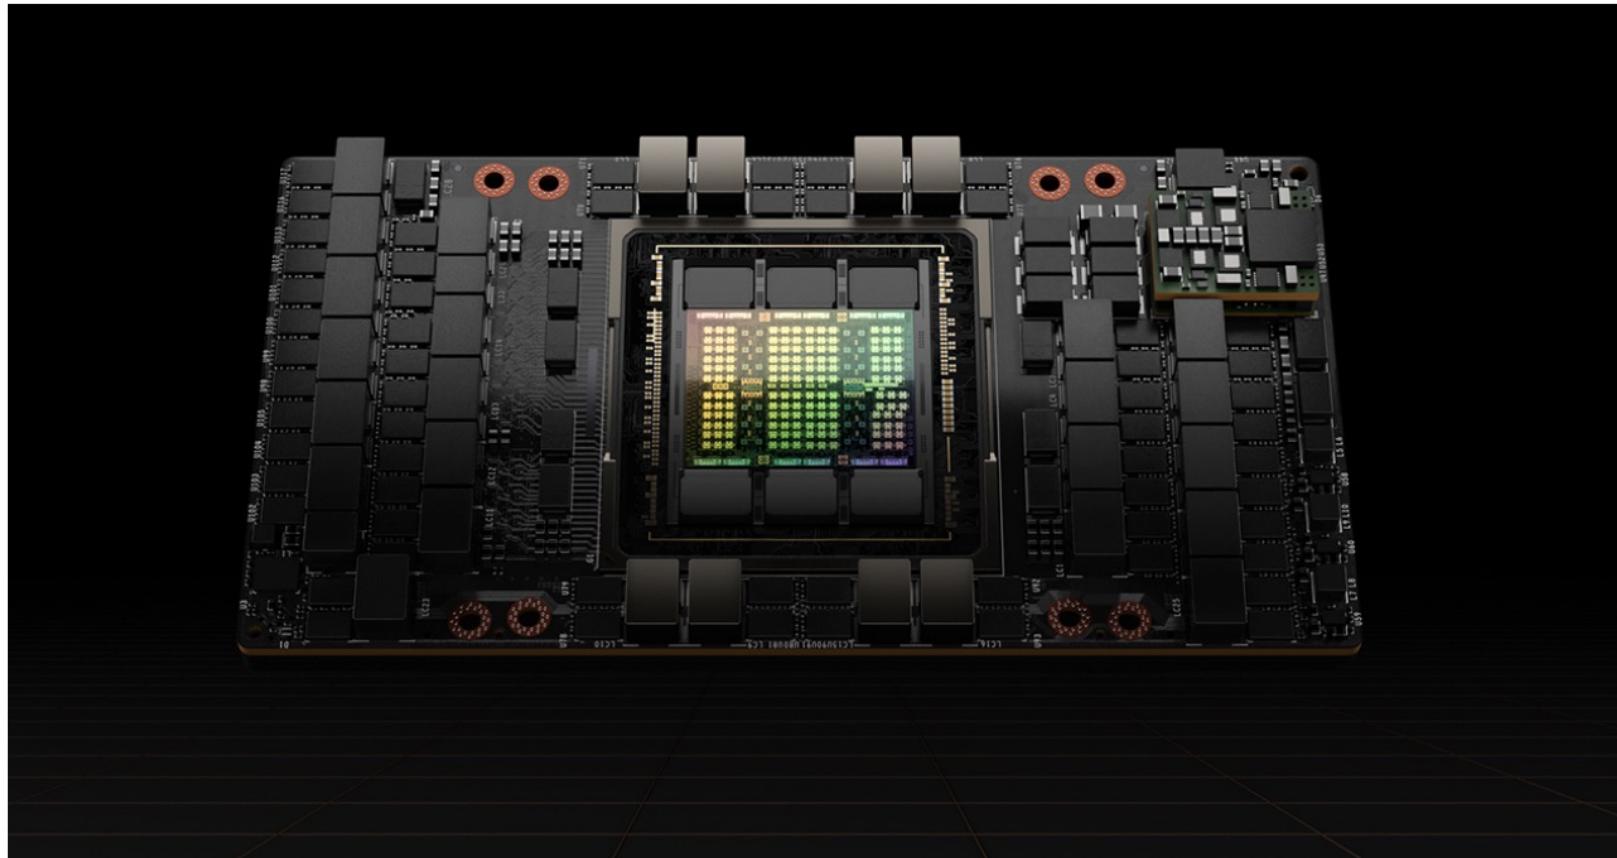


NVIDIA Hopper DPX Instructions (2022)

NVIDIA Hopper GPU Architecture Accelerates Dynamic Programming Up to 40x Using New DPX Instructions

Dynamic programming algorithms are used in healthcare, robotics, quantum computing, data science and more.

March 22, 2022 by DION HARRIS



Recall Our Dream (from 2007)

- An embedded device that can perform comprehensive genome analysis in real time (within a minute)
- Still a long ways to go
 - Energy efficiency
 - Performance (latency)
 - Security & privacy
 - **Huge memory bottleneck**

Conclusion

- System design for bioinformatics is a critical problem
 - It has large scientific, medical, societal, personal implications
- This talk is about accelerating a key step in bioinformatics:
genome sequence analysis
 - In particular, read mapping
- We covered various recent ideas to accelerate read mapping
 - My personal journey since September 2006
- Many future opportunities exist
 - Especially with new sequencing technologies
 - Especially with new applications and use cases

A Bright Future for Intelligent Genome Analysis

Mohammed Alser, Zülal Bingöl, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu
[“Accelerating Genome Analysis: A Primer on an Ongoing Journey”](#) IEEE Micro, August 2020.



MinION from ONT

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Sept.-Oct. 2020, pp. 65-75, vol. 40
DOI Bookmark: [10.1109/MM.2020.3013728](https://doi.org/10.1109/MM.2020.3013728)

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41
DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)



SmidgION from ONT

Resources & Acknowledgments

Accelerating Genome Analysis: Overview

- Mohammed Alser, Zulal Bingol, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, and Onur Mutlu,

"Accelerating Genome Analysis: A Primer on an Ongoing Journey"

IEEE Micro (IEEE MICRO), Vol. 40, No. 5, pages 65-75, September/October 2020.

[[Slides \(pptx\)\(pdf\)](#)]

[[Talk Video \(1 hour 2 minutes\)](#)]

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Mohammed Alser
ETH Zürich

Zülal Bingöl
Bilkent University

Damla Senol Cali
Carnegie Mellon University

Jeremie Kim
ETH Zurich and Carnegie Mellon University

Saugata Ghose
University of Illinois at Urbana–Champaign and Carnegie Mellon University

Can Alkan
Bilkent University

Onur Mutlu
ETH Zurich, Carnegie Mellon University, and Bilkent University

PIM Review and Open Problems

A Modern Primer on Processing in Memory

Onur Mutlu^{a,b}, Saugata Ghose^{b,c}, Juan Gómez-Luna^a, Rachata Ausavarungnirun^d

SAFARI Research Group

^a*ETH Zürich*

^b*Carnegie Mellon University*

^c*University of Illinois at Urbana-Champaign*

^d*King Mongkut's University of Technology North Bangkok*

Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, and Rachata Ausavarungnirun,

"A Modern Primer on Processing in Memory"

Invited Book Chapter in Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann, Springer, to be published in 2021.

PIM Review and Open Problems (II)

A Workload and Programming Ease Driven Perspective of Processing-in-Memory

Saugata Ghose[†] Amirali Boroumand[†] Jeremie S. Kim^{†\\$} Juan Gómez-Luna^{\\$} Onur Mutlu^{\\$†}

[†]*Carnegie Mellon University*

^{\\$}*ETH Zürich*

Saugata Ghose, Amirali Boroumand, Jeremie S. Kim, Juan Gomez-Luna, and Onur Mutlu,

"Processing-in-Memory: A Workload-Driven Perspective"

Invited Article in IBM Journal of Research & Development, Special Issue on Hardware for Artificial Intelligence, to appear in November 2019.

[Preliminary arXiv version]

More on Memory-Centric System Design

- Onur Mutlu,

"Memory-Centric Computing Systems"

Invited Tutorial at *66th International Electron Devices Meeting (IEDM)*, Virtual, 12 December 2020.

[Slides (pptx) (pdf)]

[Executive Summary Slides (pptx) (pdf)]

[Tutorial Video (1 hour 51 minutes)]

[Executive Summary Video (2 minutes)]

[Abstract and Bio]

[Related Keynote Paper from VLSI-DAT 2020]

[Related Review Paper on Processing in Memory]

<https://www.youtube.com/watch?v=H3sEaINPBOE>



Memory-Centric Computing Systems

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

12 December 2020

IEDM Tutorial



SAFARI

ETH zürich

Carnegie Mellon



0:06 / 1:51:05



IEDM 2020 Tutorial: Memory-Centric Computing Systems, Onur Mutlu, 12 December 2020

1,641 views • Dec 23, 2020

48

0

SHARE

SAVE

...



Onur Mutlu Lectures
13.9K subscribers

<https://www.youtube.com/watch?v=H3sEaINPBOE>

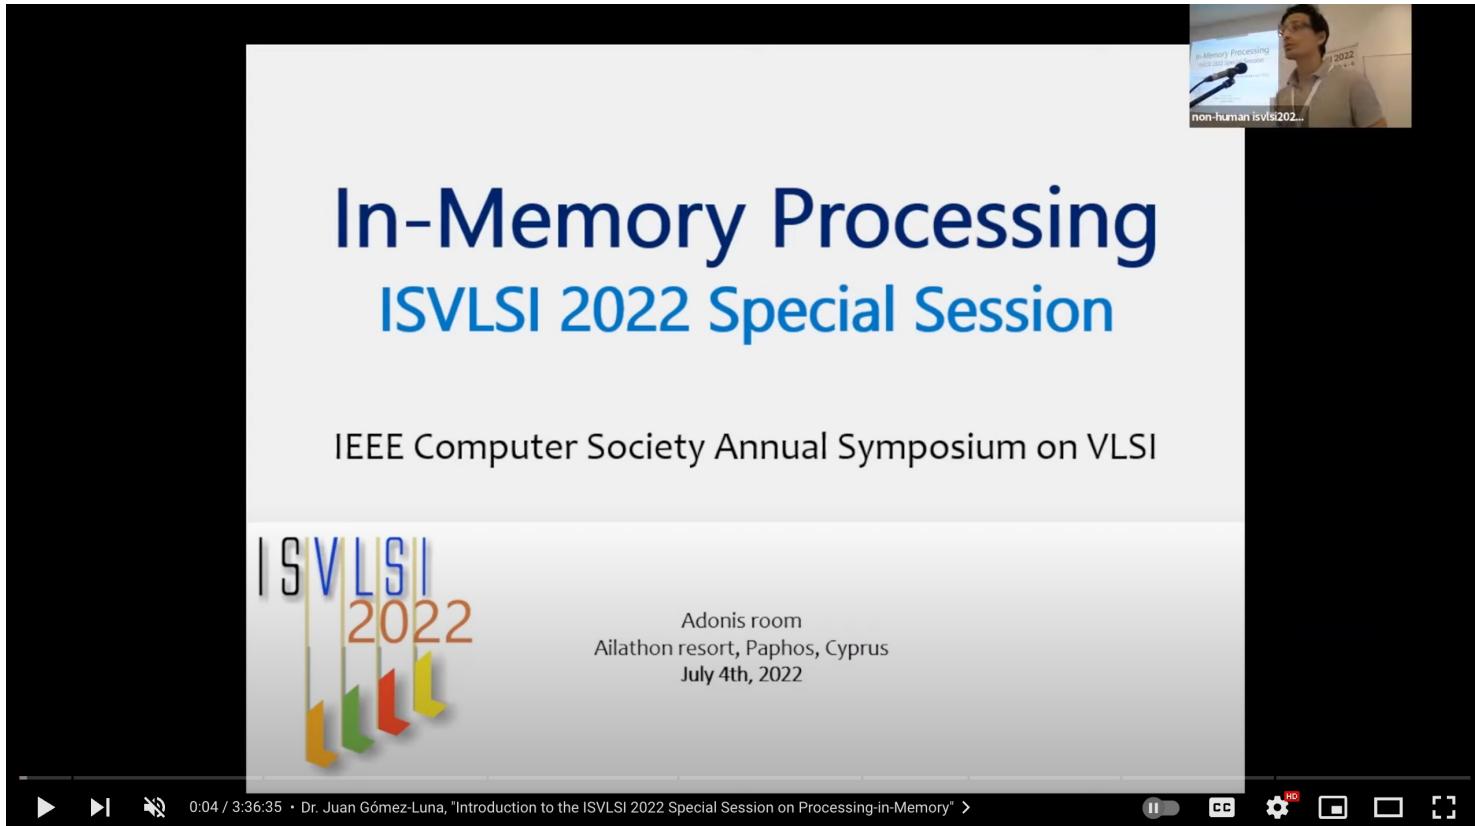
ANALYTICS

EDIT VIDEO

<https://www.youtube.com/onurmutlulectures>

Special Research Sessions & Courses

- Special Session at ISVLSI 2022: 9 cutting-edge talks



ISVLSI 2022 Special Session on Processing-in-Memory

1,286 views • Premiered Aug 9, 2022

61 DISLIKE SHARE DOWNLOAD CLIP SAVE ...



Onur Mutlu Lectures
26.9K subscribers

ANALYTICS

EDIT VIDEO

Overview Readings (II)

Gagandeep Singh, Mohammed Alser, Damla Senol Cali, Dionysios Diamantopoulos, Juan Gomez-Luna, Henk Corporaal, Onur Mutlu,

"FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications"

IEEE Micro, 2021.

[[Source Code](#)]



Home / Magazines / IEEE Micro / 2021.04

IEEE Micro

FPGA-Based Near-Memory Acceleration of Modern Data-Intensive Applications

July-Aug. 2021, pp. 39-48, vol. 41

DOI Bookmark: [10.1109/MM.2021.3088396](https://doi.org/10.1109/MM.2021.3088396)

Authors

Gagandeep Singh, ETH Zürich, Zürich, Switzerland

Mohammed Alser, ETH Zürich, Zürich, Switzerland

Damla Senol Cali, Carnegie Mellon University, Pittsburgh, PA, USA

Dionysios Diamantopoulos, Zürich Lab, IBM Research Europe, Rüschlikon, Switzerland

Juan Gomez-Luna, ETH Zürich, Zürich, Switzerland

Henk Corporaal, Eindhoven University of Technology, Eindhoven, The Netherlands

Onur Mutlu, ETH Zürich, Zürich, Switzerland

Overview Readings (III)

Mohammed Alser, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu

"From Molecules to Genomic Variations: Intelligent Algorithms and Architectures for Intelligent Genome Analysis"

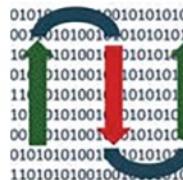
Computational and Structural Biotechnology Journal, 2022

[[Source code](#)]



ELSEVIER

01010... 0010101010010
00101010010010101011
10101010010101010111
01010100101010101010
11010100101010101010
10101010010101010111
00101010010101010111
01010101001010101010



COMPUTATIONAL
AND STRUCTURAL
BIOTECHNOLOGY
JOURNAL

journal homepage: www.elsevier.com/locate/csbj



Review

From molecules to genomic variations: Accelerating genome analysis via intelligent algorithms and architectures



Mohammed Alser *, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu *

ETH Zurich, Gloriastrasse 35, 8092 Zürich, Switzerland

Detailed Lectures on Genome Analysis

- Computer Architecture, Fall 2020, Lecture 3a
 - **Introduction to Genome Sequence Analysis** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=CrRb32v7SJc&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=5>
- Computer Architecture, Fall 2020, Lecture 8
 - **Intelligent Genome Analysis** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=ygmQpdDTL7o&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=14>
- Computer Architecture, Fall 2020, Lecture 9a
 - **GenASM: Approx. String Matching Accelerator** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=XoLpzmN-Pas&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=15>
- Accelerating Genomics Project Course, Fall 2020, Lecture 1
 - **Accelerating Genomics** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=rgjl8ZyLsAg&list=PL5Q2soXY2Zi9E2bBVAgCqLgwiDRQDTyId>

Genomics (Spring 2022)

■ Spring 2022 Edition:

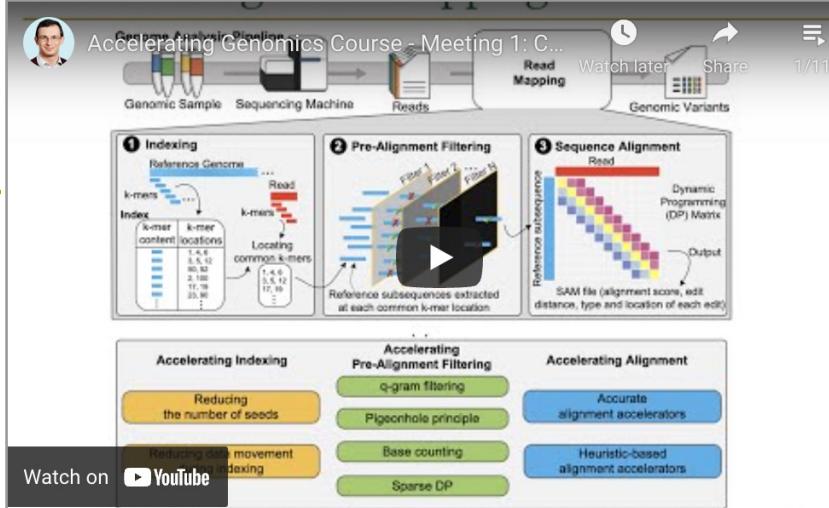
- ❑ https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=bioinformatics

■ Youtube Livestream:

- ❑ https://www.youtube.com/watch?v=DEL5A_Y3TI&list=PL5Q2soXY2Zi8NrPDgOR1yRU_Cxxjw-u18

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ Genomics lectures
- ❑ Hands-on research exploration
- ❑ Many research readings



Spring 2022 Meetings/Schedule

Week	Date	Livestream	Meeting	Learning Materials	Assignments
W1	11.3 Fri.	YouTube Live	M1: P&S Accelerating Genomics Course Introduction & Project Proposals PDF PPT	Required Materials Recommended Materials	
W2	18.3 Fri.	YouTube Live	M2: Introduction to Sequencing PDF PPT		
W3	25.3 Fri.	YouTube Premiere	M3: Read Mapping PDF PPT		
W4	01.04 Fri.	YouTube Premiere	M4: GateKeeper PDF PPT		
W5	08.04 Fri.	YouTube Premiere	M5: MAGNET & Shouji PDF PPT		
W6	15.4 Fri.	YouTube Premiere	M6: SneakySnake PDF PPT		
W7	29.4 Fri.	YouTube Premiere	M7: GenStore PDF PPT		
W8	06.05 Fri.	YouTube Premiere	M8: GRIM-Filter PDF PPT		
W9	13.05 Fri.	YouTube Premiere	M9: Genome Assembly PDF PPT		
W10	20.05 Fri.	YouTube Live	M10: Genomic Data Sharing Under Differential Privacy PDF PPT		
W11	10.06 Fri.	YouTube Premiere	M11: Accelerating Genome Sequence Analysis PDF PPT		

Genomics (Fall 2021)

■ Fall 2021 Edition:

- ❑ https://safari.ethz.ch/projects_and_seminars/fall2021/doku.php?id=bioinformatics

■ Youtube Livestream:

- ❑ <https://www.youtube.com/watch?v=MnogTeMjY8k&list=PL5Q2soXY2Zi8sngH-TrNZnDhDkPq55J9J>

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ Genomics lectures
- ❑ Hands-on research exploration
- ❑ Many research readings

Mobile Genomics Course - Meeting 1: Course...

Understanding **genetic variations**

Predicting the presence and relative abundances of **microbes** in a sample

Watch on YouTube

Rapid surveillance of **disease outbreaks**

Developing **personalized medicine**

Fall 2021 Meetings/Schedule

Week	Date	Livestream	Meeting	Learning Materials	Assignments
W1	5.10 Tue.	YouTube Live	M1: P&S Accelerating Genomics Course Introduction & Project Proposals PDF PPT YouTube Video	Required Materials Recommended Materials	
W2	20.10 Wed.	YouTube Live	M2: Introduction to Sequencing PDF PPT		
W3	27.10 Wed.	YouTube Live	M3: Read Mapping PDF PPT		
W4	3.11 Wed.	YouTube Live	M4: GateKeeper PDF PPT		
W5	10.11 Wed.	YouTube Live	M5: MAGNET & Shouji PDF PPT		
W6	17.11 Wed.		M6.1: SneakySnake PDF PPT Video		
			M6.2: GRIM-Filter PDF PPT YouTube Video		
W7	24.11 Wed.		M7: GenASM PDF PPT YouTube Video		
W8	01.12 Wed.	YouTube Live	M8: Genome Assembly PDF PPT		
W9	13.12 Mon.	YouTube Live	M9: GRIM-Filter PDF PPT		
W10	15.12 Wed.	YouTube Live	M10: Genomic Data Sharing Under Differential Privacy PDF PPT		

Comp Arch (Fall'21)

Fall 2021 Edition:

- <https://safari.ethz.ch/architecture/fall2021/doku.php?id=schedule>

Fall 2020 Edition:

- <https://safari.ethz.ch/architecture/fall2020/doku.php?id=schedule>

Youtube Livestream (2021):

- https://www.youtube.com/watch?v=4yfkM_5EFg_o&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILKTOF

Youtube Livestream (2020):

- <https://www.youtube.com/watch?v=c3mPdZA-Fmc&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN>

Master's level course

- Taken by Bachelor's/Masters/PhD students
- Cutting-edge research topics + fundamentals in Computer Architecture
- 5 Simulator-based Lab Assignments
- Potential research exploration
- Many research readings

<https://www.youtube.com/onurmutlulectures>



- Lectures/Schedule
- Lecture Buzzwords
- Readings
- HWs
- Labs
- Exams
- Related Courses
- Tutorials

- Computer Architecture FS20: Course Webpage
- Computer Architecture FS20: Lecture Videos
- Digitaltechnik SS21: Course Webpage
- Digitaltechnik SS21: Lecture Videos
- Moodle
- HotCRP
- Verilog Practice Website (HDLBits)

Lecture Video Playlist on YouTube

Livestream Lecture Playlist

Computer Architecture Processing-in-Memory System

<https://arxiv.org/pdf/2105.03814.pdf>

Recorded Lecture Playlist

TESLA Full Self-Driving Computer (2019)

ML accelerator: 260 mm², 6 billion transistors, 600 GFLOPS GPU, 12 ARM 2.2 GHz CPUs.

Two redundant chips for better safety.

<https://arxiv.org/pdf/1902.06236.pdf>

Fall 2021 Lectures & Schedule

Week	Date	Livestream	Lecture	Readings	Lab	HW
W1	30.09 Thu.	YouTube Live	L1: Introduction and Basics (PDF) (PPT)	Required Mentioned	Lab 1 Out	HW 0 Out
	01.10 Fri.	YouTube Live	L2: Trends, Tradeoffs and Design Fundamentals (PDF) (PPT)	Required Mentioned		
W2	07.10 Thu.	YouTube Live	L3a: Memory Systems: Challenges and Opportunities (PDF) (PPT)	Described Suggested		HW 1 Out
			L3b: Course Info & Logistics (PDF) (PPT)			
			L3c: Memory Performance Attacks (PDF) (PPT)	Described Suggested		
	08.10 Fri.	YouTube Live	L4a: Memory Performance Attacks (PDF) (PPT)	Described Suggested	Lab 2 Out	
			L4b: Data Retention and Memory Refresh (PDF) (PPT)	Described Suggested		
			L4c: RowHammer (PDF) (PPT)	Described Suggested		

DDCA (Spring 2022)

Spring 2022 Edition:

- https://safari.ethz.ch/digitaltechnik/spring2022/do_ku.php?id=schedule

Spring 2021 Edition:

- https://safari.ethz.ch/digitaltechnik/spring2021/do_ku.php?id=schedule

Youtube Livestream (Spring 2022):

- <https://www.youtube.com/watch?v=cpXdE3HwvK0&list=PL5Q2soXY2Zi97Ya5DEUpMpO2bbAoaG7c6>

Youtube Livestream (Spring 2021):

- https://www.youtube.com/watch?v=LbC0EZY8yw4&list=PL5Q2soXY2Zi_uej3aY39YB5pfW4SJ7LIN

Bachelor's course

- 2nd semester at ETH Zurich
- Rigorous introduction into "How Computers Work"
- Digital Design/Logic
- Computer Architecture
- 10 FPGA Lab Assignments


Digital Design and Computer Architecture - Spring 2021

Trace: - schedule

Home Announcements Materials Resources

- Lectures/Schedule
- Lecture Buzzwords
- Readings
- Optional HWs
- Labs
- Extra Assignments
- Exams
- Technical Docs

- Computer Architecture (CMU) SS15: Lecture Videos
- Computer Architecture (CMU) SS15: Course Website
- Digitaltechnik SS18: Lecture Videos
- Digitaltechnik SS18: Course Website
- Digitaltechnik SS19: Lecture Videos
- Digitaltechnik SS19: Course Website
- Digitaltechnik SS20: Lecture Videos
- Digitaltechnik SS20: Course Website
- Moodle

Lecture Video Playlist on YouTube

Livestream Lecture Playlist

Onur Mutlu, Digital Design and Computer Architecture Today

Computing landscape is very different from 10-20 years ago

Applications and technology both demand novel architectures



Every component and its interfaces, as well as entire system designs are being re-examined

Watch on YouTube Recorded Lecture Playlist

Digital Design & Computer Architecture: Lect...

ANSWER

Digital Design & Computer Architecture: Lect...

How Computers Work (from the ground up)

Watch on YouTube

Spring 2021 Lectures/Schedule

Week	Date	Livestream	Lecture	Readings	Lab	HW
W1	25.02 Thu.	YouTube Live	L1: Introduction and Basics (PDF) (PPT)	Required Suggested Mentioned		
	26.02 Fri.	YouTube Live	L2a: Tradeoffs, Metrics, Mindset (PDF) (PPT)	Required		
			L2b: Mysteries in Computer Architecture (PDF) (PPT)	Required Mentioned		
W2	04.03 Thu.	YouTube Live	L3a: Mysteries in Computer Architecture II (PDF) (PPT)	Required Suggested Mentioned		

Seminar in Comp Arch (Spring & Fall)

■ Spring 2022 Edition:

- https://safari.ethz.ch/architecture_seminar/spring2022/doku.php?id=schedule

■ Fall 2021 Edition:

- https://safari.ethz.ch/architecture_seminar/fall2021/doku.php?id=schedule

■ Youtube Livestream (Spring 2022):

- https://www.youtube.com/watch?v=rS9UPk509AQ&list=PL5Q2soXY2Zi_hxizriwKmFHgcoe2Q8-m0

■ Youtube Livestream (Fall 2021):

- https://www.youtube.com/watch?v=4TcP297mdsI&list=PL5Q2soXY2Zi_7UBNmC9B8Yr5JSwTG9yH4

■ Critical analysis course

- Taken by Bachelor's/Masters/PhD students
- Cutting-edge research topics + fundamentals in Computer Architecture
- 20+ research papers, presentations, analyses

The screenshot shows a website for the Seminar in Computer Architecture - Spring 2022. At the top right is a navigation bar with 'Search', 'Recent Changes', 'Media Manager', and a user icon. Below it is a 'Lecture Video Playlist on YouTube' section with a thumbnail for a video titled 'Exponential Growth of Neural Networks'. The video stats show 1800x more compute in just 2 years, with a graph comparing training compute (PFLOPs-days) and model memory requirement (GB) from 2018 to 2020+ across various models like T5, T-NLG, Megatron-LM, GPT-2, BERT Large, BERT Base, and MSFT-1T. A 'Watch on YouTube' button and a source link are at the bottom.

Week	Date	Livestream	Lecture	Readings	Assignments
W1	24.02 Thu.	YouTube Live	L1a: Course Logistics (PDF) (PPT) L1b: Introduction and Basics (PDF) (PPT) L1c: Architectural Design Fundamentals (PDF) (PPT)	Suggested	
W2	03.03 Thu.	YouTube Live	L2: Memory-Centric Computing (PDF) (PPT)	Suggested	
W3	10.03 Thu.	YouTube Live	L3: Memory-Centric Computing II (PDF) (PPT)	Suggested	
W4	17.03 Thu.	YouTube Live	L4: Memory-Centric Computing III (PDF) (PPT)	Suggested	
W5	24.03 Thu.	YouTube Live	L5: Accelerating Genome Analysis (PDF) (PPT)	Suggested	
W6	31.03 Thu.	YouTube Live	L6a: Rethinking Virtual Memory I (PDF) (PPT) L6b: Rethinking Virtual Memory II (PDF) (PPT)	Suggested	
W7	07.04 Thu.	YouTube Live	S1.1: A Logic-in-Memory Computer; <i>IEEE Trans. Comput.</i> , 1970 (PDF) (PPT)	Suggested	

PIM Course (Spring 2022)

■ Spring 2022 Edition:

- ❑ https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=processing_in_memory

■ Youtube Livestream:

- ❑ <https://www.youtube.com/watch?v=9e4Chnwdovo&list=PL5Q2soXY2Zi-841fUYYUK9EsXKhQKRPyX>

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ Processing-in-Memory lectures
- ❑ Hands-on research exploration
- ❑ Many research readings

PIM Review and Open Problem
Processing in Memory Course, Meeting 1, Ex...
Watch later Share 1/13

A Modern Primer on Processing in Memory

Onur Mutlu^{a,b}, Saugata Ghose^{b,c}, Juan Gómez-Luna^a, Rachata Ausavarungnirun^d
SAFARI Research Group
^aCarrieys University
^cUniversity of Illinois Urbana-Champaign
^dKing Mongkut's University of Technology North Bangkok

Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, and Rachata Ausavarungnirun,
"A Modern Primer on Processing in Memory"
Invited Book Chapter in *Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann*, Springer, to be published in 2021.

Watch on YouTube <https://arxiv.org/pdf/1903.03988.pdf> 108

Spring 2022 Meetings/Schedule

Week	Date	Livestream	Meeting	Learning Materials	Assignments
W1	10.03 Thu.	YouTube Live	M1: P&S PIM Course Presentation (PDF) (PPT)	Required Materials Recommended Materials	HW 0 Out
W2	15.03 Tue. 17.03 Thu.	YouTube Premiere	Hands-on Project Proposals M2: Real-world PIM: UPMEM PIM (PDF) (PPT)		
W3	24.03 Thu.	YouTube Live	M3: Real-world PIM: Microbenchmarking of UPMEM PIM (PDF) (PPT)		
W4	31.03 Thu.	YouTube Live	M4: Real-world PIM: Samsung HBM-PIM (PDF) (PPT)		
W5	07.04 Thu.	YouTube Live	M5: How to Evaluate Data Movement Bottlenecks (PDF) (PPT)		
W6	14.04 Thu.	YouTube Live	M6: Real-world PIM: SK Hynix AIM (PDF) (PPT)		
W7	21.04 Thu.	YouTube Premiere	M7: Programming PIM Architectures (PDF) (PPT)		
W8	28.04 Thu.	YouTube Premiere	M8: Benchmarking and Workload Suitability on PIM (PDF) (PPT)		
W9	05.05 Thu.	YouTube Premiere	M9: Real-world PIM: Samsung AxDIMM (PDF) (PPT)		
W10	12.05 Thu.	YouTube Premiere	M10: Real-world PIM: Alibaba HB-PNM (PDF) (PPT)		
W11	19.05 Thu.	YouTube Live	M11: SpMV on a Real PIM Architecture (PDF) (PPT)		
W12	26.05 Thu.	YouTube Live	M12: End-to-End Framework for Processing-using-Memory (PDF) (PPT)		
W13	02.06 Thu.	YouTube Live	M13: Bit-Serial SIMD Processing using DRAM (PDF) (PPT)		
W14	09.06 Thu.	YouTube Live	M14: Analyzing and Mitigating ML Inference Bottlenecks (PDF) (PPT)		
W15	15.06 Thu.	YouTube Live	M15: In-Memory HTAP Databases with HW/SW Co-design (PDF) (PPT)		
W16	23.06 Thu.	YouTube Live	M16: In-Storage Processing for Genome Analysis (PDF) (PPT)		
W17	18.07 Mon.	YouTube Premiere	M17: How to Enable the Adoption of PIM? (PDF) (PPT)		
W18	09.08 Tue.	YouTube Premiere	SS1: ISVLSI 2022 Special Session on PIM (PDF & PPT)		

Hetero. Systems (Spring'22)

■ Spring 2022 Edition:

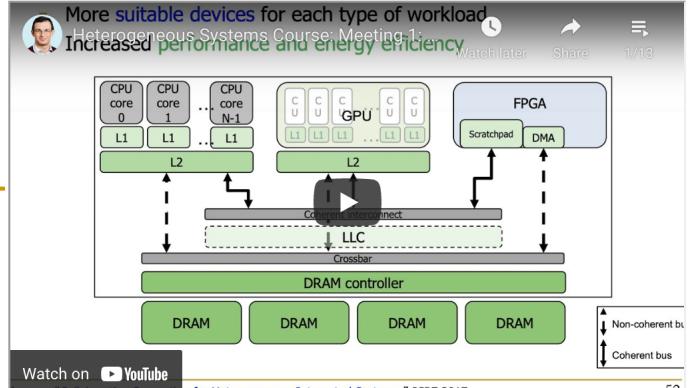
- ❑ https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=heterogeneous_systems

■ Youtube Livestream:

- ❑ <https://www.youtube.com/watch?v=oFO5fTrgFIY&list=PL5Q2soXY2Zi9XrgXR38IMFTjmY6h7Gzm>

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ GPU and Parallelism lectures
- ❑ Hands-on research exploration
- ❑ Many research readings



Watch on [YouTube](#)
hang+, "Collaborative Computing for Heterogeneous Integrated Systems," ICPE 2017.

52

Spring 2022 Meetings/Schedule

Week	Date	Livestream	Meeting	Learning Materials	Assignments
W1	15.03 Tue.	YouTube Premiere	M1: P&S Course Presentation PDF PPT	Required Materials Recommended Materials	HW 0 Out
W2	22.03 Tue.	YouTube Premiere	M2: SIMD Processing and GPUs PDF PPT		
W3	29.03 Tue.	YouTube Premiere	M3: GPU Software Hierarchy PDF PPT		
W4	05.04 Tue.	YouTube Premiere	M4: GPU Memory Hierarchy PDF PPT		
W5	12.04 Tue.	YouTube Premiere	M5: GPU Performance Considerations PDF PPT		
W6	19.04 Tue.	YouTube Premiere	M6: Parallel Patterns: Reduction PDF PPT		
W7	26.04 Tue.	YouTube Premiere	M7: Parallel Patterns: Histogram PDF PPT		
W8	03.05 Tue.	YouTube Premiere	M8: Parallel Patterns: Convolution PDF PPT		
W9	10.05 Tue.	YouTube Premiere	M9: Parallel Patterns: Prefix Sum (Scan) PDF PPT		
W10	17.05 Tue.	YouTube Premiere	M10: Parallel Patterns: Sparse Matrices PDF PPT		
W11	24.05 Tue.	YouTube Premiere	M11: Parallel Patterns: Graph Search PDF PPT		
W12	01.06 Wed.	YouTube Premiere	M12: Parallel Patterns: Merge Sort PDF PPT		
W13	07.06 Tue.	YouTube Premiere	M13: Dynamic Parallelism PDF PPT		
W14	15.06 Wed.	YouTube Premiere	M14: Collaborative Computing PDF PPT		
W15	24.06 Fri.	YouTube Premiere	M15: GPU Acceleration of Genome Sequence Alignment PDF PPT		
W16	14.07 Thu.	YouTube Premiere	M16: Accelerating Agent-based Simulations PDF ODP		

HW/SW Co-Design (Spring 2022)

■ Spring 2022 Edition:

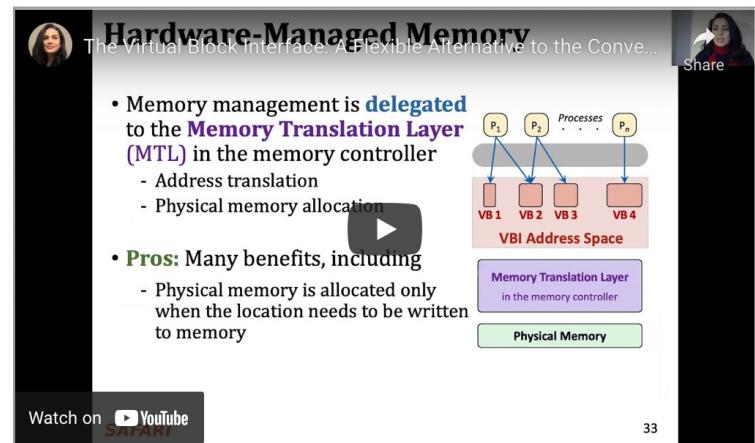
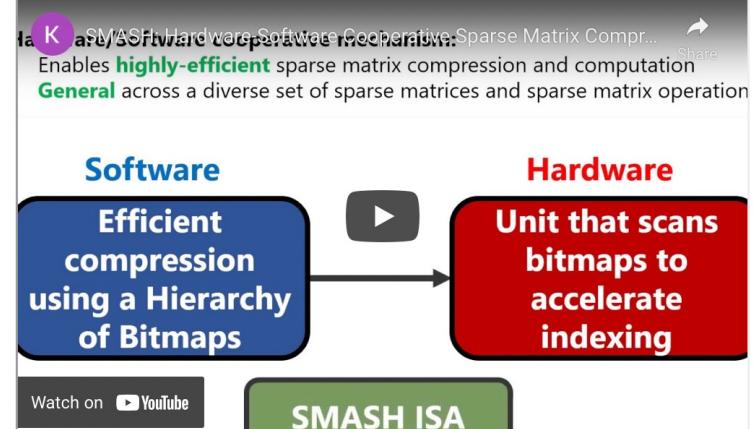
- ❑ https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=hw_sw_co_design

■ Youtube Livestream:

- ❑ <https://youtube.com/playlist?list=PL5Q2soXY2Zi8nH7un3ghD2nutKWWdk-NK>

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ HW/SW co-design lectures
- ❑ Hands-on research exploration
- ❑ Many research readings



2022 Meetings/Schedule (Tentative)

Week	Date	Livestream	Meeting	Materials	Assignments
W0	16.03	YouTube Live	Intro to HW/SW Co-Design (PPTX) (PDF)	Required	HW 0 Out
W1	23.03		Project selection	Required	
W2	30.03	YouTube Live	Virtual Memory (I) (PPTX) (PDF)		
W3	13.04	YouTube Live	Virtual Memory (II) (PPTX) (PDF)		

SSD Course (Spring 2022)

■ Spring 2022 Edition:

- ❑ https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=modern_ssds

■ Youtube Livestream:

- ❑ <https://www.youtube.com/watch?v=q4rm71DsY4&list=PL5Q2soXY2Zi8vabcse1kL22DEcgMI2RAq>

■ Project course

- ❑ Taken by Bachelor's/Master's students
- ❑ SSD Basics and Advanced Topics
- ❑ Hands-on research exploration
- ❑ Many research readings

P&S Modern SSDs

Basics of NAND Flash-Based SSDs

Dr. Jisung Park
Prof. Onur Mutlu
ETH Zürich
Spring 2022
25 March 2021

Modern Solid-State Drives (SSDs) Course - Meeting 2: Basics of NAND Flash-Based SSDs (Spring 2022)

807 views • Streamed live on Mar 25, 2022

16 DISLIKE SHARE DOWNLOAD CLIP SAVE

P&S Modern SSDs

Introduction to MQSim

Rakesh Nadig
Dr. Jisung Park
Prof. Onur Mutlu
ETH Zürich
Spring 2022
8th April 2022

Modern Solid-State Drives (SSDs) Course - Meeting 4: Introduction to MQSim (Spring 2022)

310 views • Streamed live on Apr 8, 2022

17 DISLIKE SHARE DOWNLOAD CLIP SAVE

Funding Acknowledgments

- Alibaba, AMD, ASML, [Google](#), Facebook, [Hi-Silicon](#), HP Labs, [Huawei](#), IBM, [Intel](#), [Microsoft](#), Nvidia, Oracle, Qualcomm, Rambus, Samsung, Seagate, [VMware](#), [Xilinx](#)
- NSF
- NIH
- GSRC
- [SRC](#)
- CyLab
- EFCL

Acknowledgments



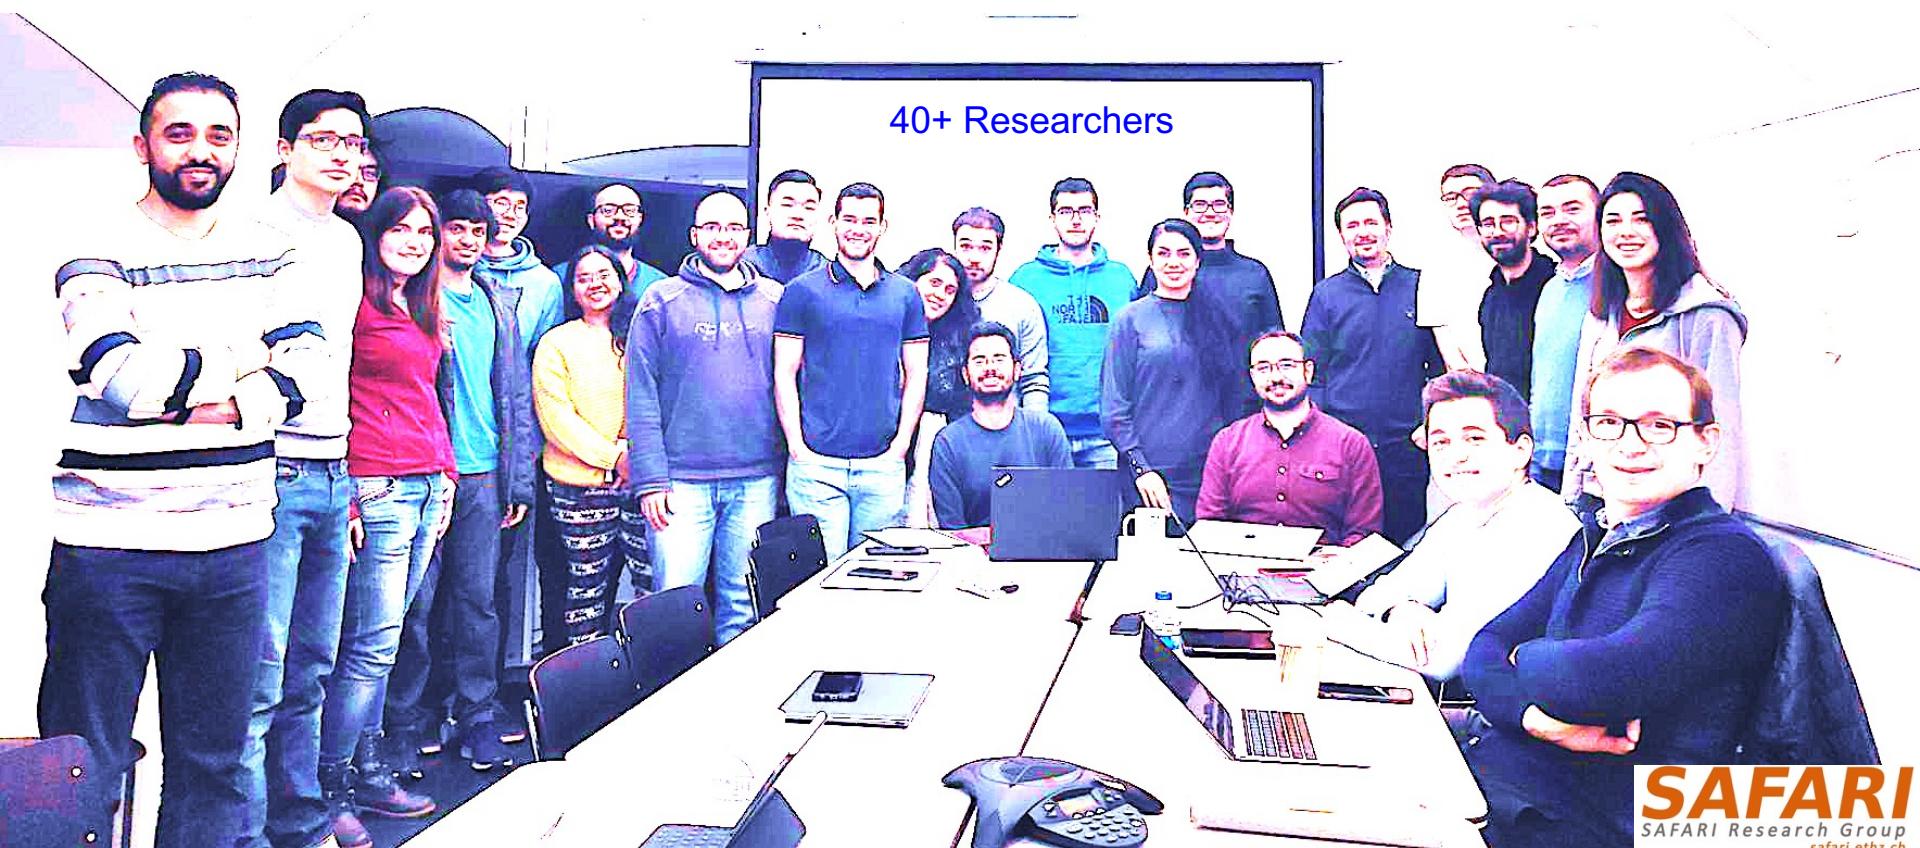
Think BIG, Aim HIGH!

<https://safari.ethz.ch>

Onur Mutlu's SAFARI Research Group

Computer architecture, HW/SW, systems, bioinformatics, security, memory

<https://safari.ethz.ch/safari-newsletter-january-2021/>



SAFARI
SAFARI Research Group
safari.ethz.ch

Think BIG, Aim HIGH!

SAFARI

<https://safari.ethz.ch>

SAFARI Newsletter April 2020 Edition

- <https://safari.ethz.ch/safari-newsletter-april-2020/>



Dear SAFARI friends,

2019 and the first three months of 2020 have been very positive eventful times for SAFARI.

SAFARI Newsletter January 2021 Edition

- <https://safari.ethz.ch/safari-newsletter-january-2021/>



Newsletter
January 2021

*Think Big, Aim High, and
Have a Wonderful 2021!*



Dear SAFARI friends,

Happy New Year! We are excited to share our group highlights with you in this second edition of the SAFARI newsletter (You can find the first edition from April 2020 [here](#)). 2020 has

SAFARI Newsletter December 2021 Edition

- <https://safari.ethz.ch/safari-newsletter-december-2021/>



Think Big, Aim High



View in your browser

December 2021



Referenced Papers, Talks, Artifacts

- All are available at

<https://people.inf.ethz.ch/omutlu/projects.htm>

<https://www.youtube.com/onurmutlulectures>

<https://github.com/CMU-SAFARI/>

Open Source Tools: SAFARI GitHub



SAFARI Research Group at ETH Zurich and Carnegie Mellon University

Site for source code and tools distribution from SAFARI Research Group at ETH Zurich and Carnegie Mellon University.

📍 ETH Zurich and Carnegie Mellon U... 🌐 <https://safari.ethz.ch/> 📩 omutlu@gmail.com

Overview

Repositories 71

Projects

Packages

Teams 1

People 44

Settings

Pinned

Customize pins

💻 **ramulator** Public

A Fast and Extensible DRAM Simulator, with built-in support for modeling many different DRAM technologies including DDRx, LPDDRx, GDDRx, WIOx, HBMx, and various academic proposals. Described in the...

● C++ ⭐ 311 🏷 161

💻 **prim-benchmarks** Public

PrIM (Processing-In-Memory benchmarks) is the first benchmark suite for a real-world processing-in-memory (PIM) architecture. PrIM is developed to evaluate, analyze, and characterize the first publ...

● C ⭐ 53 🏷 21

💻 **DAMOV** Public

DAMOV is a benchmark suite and a methodical framework targeting the study of data movement bottlenecks in modern applications. It is intended to study new architectures, such as near-data processin...

● C++ ⭐ 26 🏷 4

💻 **SneakySnake** Public

SneakySnake is the first and the only pre-alignment filtering algorithm that works efficiently and fast on modern CPU, FPGA, and GPU architectures. It greatly (by more than two orders of magnitude...

● VHDL ⭐ 41 🏷 8

💻 **MQSim** Public

MQSim is a fast and accurate simulator modeling the performance of modern multi-queue (MQ) SSDs as well as traditional SATA based SSDs. MQSim faithfully models new high-bandwidth protocol implement...

● C++ ⭐ 146 🏷 93

💻 **rowhammer** Public

Source code for testing the Row Hammer error mechanism in DRAM devices. Described in the ISCA 2014 paper by Kim et al. at http://users.ece.cmu.edu/~omutlu/pub/dram-row-hammer_isca14.pdf.

● C ⭐ 189 🏷 41

<https://github.com/CMU-SAFARI/>

Accelerating Genome Analysis

A Primer on an Ongoing Journey

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

6 September 2022

Barcelona Supercomputing Center

SAFARI

ETH zürich

Carnegie Mellon

Some Recent Papers

Connecting Basecalling and Read Mapping in PIM

- Haiyu Mao, Mohammed Alser, Mohammad Sadrosadati, Can Firtina, Akanksha Baranwal, Damla Senol Cali, Aditya Manglik, Nour Almadhoun Alserr, and Onur Mutlu,

["GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping"](#)

Proceedings of the 55th International Symposium on Microarchitecture (MICRO), Chicago, Illinois, October 2022.

GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping

Haiyu Mao¹ Mohammed Alser¹ Mohammad Sadrosadati¹ Can Firtina¹ Akanksha Baranwal¹
Damla Senol Cali² Aditya Manglik¹ Nour Almadhoun Alserr¹ Onur Mutlu¹

¹*ETH Zürich*

²*Bionano Genomics*

Finding Approximate Seed Matches

- Can Firtina, Jisung Park, Mohammed Alser, Jeremie S. Kim, Damla Senol Cali, Taha Shahroodi, Nika Mansouri-Ghiasi, Gagandeep Singh, Konstantinos Kanellopoulos, Can Alkan, and Onur Mutlu,
["BLEND: A Fast, Memory-Efficient, and Accurate Mechanism to Find Fuzzy Seed Matches"](#)
Preprint in [arXiv](#), 2021.
[\[arXiv preprint\]](#)
[\[BLEND Source Code and Data\]](#)

BLEND: A Fast, Memory-Efficient, and Accurate Mechanism to Find Fuzzy Seed Matches

Can Firtina¹ Jisung Park¹ Mohammed Alser¹ Jeremie S. Kim¹ Damla Senol Cali²
Taha Shahroodi³ Nika Mansouri-Ghiasi¹ Gagandeep Singh¹ Konstantinos Kanellopoulos¹
Can Alkan⁴ Onur Mutlu¹

¹*ETH Zurich*

²*Bionano Genomics*

³*TU Delft*

⁴*Bilkent University*

Hardware Acceleration for pHMMs

- Can Firtina, Kamlesh Pillai, Gurpreet S. Kalsi, Bharathwaj Suresh, Damla Senol Cali, Jeremie S. Kim, Taha Shahroodi, Meryem Banu Cavlak, Joel Lindegger, Mohammed Alser, Juan Gómez-Luna, Sreenivas Subramoney, and Onur Mutlu,
["ApHMM: A Profile Hidden Markov Model Acceleration Framework for Genome Analysis"](#)

Preprint in [arXiv](#), 2022.

[[Source Code](#)]

ApHMM: A Profile Hidden Markov Model Acceleration Framework for Genome Analysis

Can Firtina¹ Kamlesh Pillai² Gurpreet S. Kalsi² Bharathwaj Suresh² Damla Senol Cali³
Jeremie S. Kim¹ Taha Shahroodi⁴ Meryem Banu Cavlak¹ Joel Lindegger¹ Mohammed Alser¹
Juan Gómez Luna¹ Sreenivas Subramoney² Onur Mutlu¹

¹*ETH Zurich* ²*Intel Labs* ³*Bionano Genomics* ⁴*TU Delft*

Remapping Reads Between References

- Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Nastaran Hajinazar, Mohammed Alser, Can Alkan, and Onur Mutlu,
["AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes"](#)
Preprint in arXiv and bioRxiv, 2021.
[[bioRxiv preprint](#)]
[[arXiv preprint](#)]
[[AirLift Source Code and Data](#)]

METHOD

AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes

Jeremie S. Kim^{1†}, Can Firtina^{1†}, Meryem Banu Cavlak², Damla Senol Cali³, Nastaran Hajinazar^{1,4},
Mohammed Alser¹, Can Alkan² and Onur Mutlu^{1,2,3*}

Mapping Constant Regions Between References

- Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Can Alkan, and Onur Mutlu,
["FastRemap: A Tool for Quickly Remapping Reads between Genome Assemblies"](#)
Bioinformatics, btac554.
[\[FastRemap Source Code\]](#)

FastRemap: A Tool for Quickly Remapping Reads between Genome Assemblies

Jeremie S. Kim¹

Can Firtina¹

Meryem Banu Cavlak¹

Damla Senol Cali^{2,3}

Can Alkan⁴

Onur Mutlu^{1,2,4}

¹ETH Zürich

²Carnegie Mellon University

³Bionano Genomics

⁴Bilkent University

COVIDHunter

Mohammed Alser, Jeremie S. Kim, Nour Almadhoun Alserr, Stefan W. Tell,
Onur Mutlu

["COVIDHunter: COVID-19 Pandemic Wave Prediction and Mitigation via Seasonality Aware Modeling"](#)

Frontiers in Public Health 2022

[\[Source Code\]](#)



Frontiers in **Public Health**

ORIGINAL RESEARCH

published: 17 June 2022

doi: 10.3389/fpubh.2022.877621

COVIDHunter: COVID-19 Pandemic Wave Prediction and Mitigation via Seasonality Aware Modeling

Mohammed Alser, Jeremie S. Kim, Nour Almadhoun Alserr, Stefan W. Tell and Onur Mutlu*

Department of Information Technology and Electrical Engineering (D-ITET), ETH Zurich, Zurich, Switzerland

Packaging Omics Methods

Mohammed Alser, Sharon Waymost, Ram Ayyala, Brendan Lawlor, Richard J. Abdill,
Neha Rajkumar, Nathan LaPierre, Jaqueline Brito, Andre M. Ribeiro-dos-Santos, Can
Firtina, Nour Almadhoun, Varuni Sarwal, Eleazar Eskin, Qiyang Hu, Derek Strong,
Byoung-Do (BD)Kim, Malak S. Abedalthagafi, Onur Mutlu, Serghei Mangul

"Packaging, containerization, and virtualization of computational omics methods:
Advances, challenges, and opportunities"

arrXiv 2022

Packaging, containerization, and virtualization of computational omics methods: Advances, challenges, and opportunities

Mohammed Alser¹, Sharon Waymost², Ram Ayyala^{3,4}, Brendan Lawlor⁵, Richard J. Abdill⁶,
Neha Rajkumar⁷, Nathan LaPierre², Jaqueline Brito⁴, André M. Ribeiro-dos-Santos⁸, Can Firtina¹,
Nour Almadhoun¹, Varuni Sarwal², Eleazar Eskin^{2,9,10}, Qiyang Hu¹¹, Derek Strong¹²,
Byoung-Do (BD) Kim¹², Malak S. Abedalthagafi^{13,14,15*}, Onur Mutlu^{1,*}, Serghei Mangul^{4,*}

Demeter (HD Food Microbiome Profiling)

Taha Shahroodi, Mahdi Zahedi, Can Firtina, Mohammed Alser, Stephan Wong,
Onur Mutlu, Said Hamdioui

[“Demeter: A Fast and Energy-Efficient Food Profiler using Hyperdimensional Computing in Memory”](#)

IEEE Access, 2022



 RESEARCH ARTICLE

Demeter: A Fast and Energy-Efficient Food Profiler Using Hyperdimensional Computing in Memory

TAHA SHAHROODI^{ID1}, MAHDI ZAHEDI^{ID1}, CAN FIRGINA², MOHAMMED ALSER^{ID2},
STEPHAN WONG¹, (Senior Member, IEEE), ONUR MUTLU^{ID2}, (Fellow, IEEE),
AND SAID HAMDIOUI^{ID1}, (Senior Member, IEEE)

¹Q&CE Department, EEMCS Faculty, Delft University of Technology (TU Delft), 2628 CD Delft, The Netherlands

²SAFARI Research Group, D-ITET, ETH Zürich, 8092 Zürich, Switzerland

AIM (PIM Sequence Alignment Framework)

Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez-Luna,
Onur Mutlu, Izzat El Hajj

["A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems"](#)

arXiv, 2022

[\[Source code\]](#)

A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems

Safaa Diab¹, Amir Nassereldine¹, Mohammed Alser², Juan Gómez Luna², Onur Mutlu², Izzat El Hajj¹

¹*American University of Beirut, Lebanon* ²*ETH Zürich, Switzerland*

Scrooge

Joël Lindegger, Damla Senol Cali, Mohammed Alser, Juan Gómez-Luna,
Nika Mansouri Ghiasi, Onur Mutlu

["Scrooge: A Fast and Memory-Frugal Genomic Sequence Aligner for CPUs, GPUs, and ASICs"](#)

arXiv, 2022

[\[Source code\]](#)

Bioinformatics
doi.10.1093/bioinformatics/xxxxxx
Advance Access Publication Date: Day Month Year
Original paper

OXFORD

Genome analysis

Scrooge: A Fast and Memory-Frugal Genomic Sequence Aligner for CPUs, GPUs, and ASICs

**Joël Lindegger^{1,*}, Damla Senol Cali², Mohammed Alser¹,
Juan Gómez-Luna¹, Nika Mansouri Ghiasi¹ and Onur Mutlu^{1,*}**

¹Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich 8006, Switzerland and

²Bionano Genomics, San Diego, CA 92121, USA.

Intelligent Genome Analysis

Mohammed Alser, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao,
Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu

[“From Molecules to Genomic Variations: Intelligent Algorithms and Architectures for Intelligent Genome Analysis”](#)

Computational and Structural Biotechnology Journal, 2022

[\[Source code\]](#)

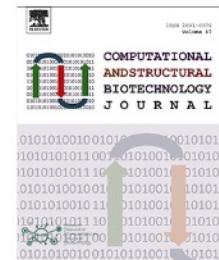


ELSEVIER

010100000000000000
001001000000000000
100010100000000000
010010100000000000
110010100000000000
010010100000000000
110010100000000000
100010100000000000
000010100000000000
010101000000000000
110101010000000000

COMPUTATIONAL
AND STRUCTURAL
BIOTECHNOLOGY
JOURNAL

journal homepage: www.elsevier.com/locate/csbj



Review

From molecules to genomic variations: Accelerating genome analysis via intelligent algorithms and architectures



Mohammed Alser *, Joel Lindegger, Can Firtina, Nour Almadhoun, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, Onur Mutlu *

ETH Zurich, Gloriastrasse 35, 8092 Zürich, Switzerland

Pairwise Sequence Alignment using PIM

- Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez Luna, Onur Mutlu, and Izzat El Hajj,
["High-throughput Pairwise Alignment with the Wavefront Algorithm using Processing-in-Memory"](#)
Preprint in [arXiv](#), 2022.

High-throughput Pairwise Alignment with the Wavefront Algorithm using Processing-in-Memory

Safaa Diab¹, Amir Nassereldine¹, Mohammed Alser², Juan Gómez Luna², Onur Mutlu², Izzat El Hajj¹

¹*American University of Beirut, Lebanon* ²*ETH Zürich, Switzerland*

Backup Slides for Further Info

Detailed Lectures on PIM (I)

- Computer Architecture, Fall 2020, Lecture 6
 - **Computation in Memory** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=oGcZAGwfEUE&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=12>
- Computer Architecture, Fall 2020, Lecture 7
 - **Near-Data Processing** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=j2GIigqn1Qw&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=13>
- Computer Architecture, Fall 2020, Lecture 11a
 - **Memory Controllers** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=TeG773OgiMQ&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=20>
- Computer Architecture, Fall 2020, Lecture 12d
 - **Real Processing-in-DRAM with UPMEM** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=Sscy1Wrr22A&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=25>

Detailed Lectures on PIM (II)

- Computer Architecture, Fall 2020, Lecture 15
 - **Emerging Memory Technologies** (ETH Zürich, Fall 2020)
 - https://www.youtube.com/watch?v=AIE1rD9G_YU&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=28
- Computer Architecture, Fall 2020, Lecture 16a
 - **Opportunities & Challenges of Emerging Memory Technologies** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=pmLszWGmMGQ&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=29>
- Computer Architecture, Fall 2020, Guest Lecture
 - **In-Memory Computing: Memory Devices & Applications** (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=wNmqQHiEZNk&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=41>

Genome Analysis



NO machine can read the
entire content of a genome



```
>CCTCCTCAGTGCCACCCAGCCCCTGGCAGCTCCAAACAGGGCTCTTATTAAAACACCCGTCCCTGCCCTGGAGTGAGGTGTCAAG  
GACCTAAACTAAAAAAAAAAAAAGAAAAAGAAAAAGAAAAAGAATTAAAATTAAGTAATTCTTGAAAAAAAACTAATTCTAAGCTTCTT  
CATGTCAAGGACCTAACATGTCTAACAGCACTTTTGACCATTATTTGGATCTGAAAGAAAATCAAGAATAATGAAGGACTGATACATTG  
GAAGAGGAGAGTCAGGACCTACAGAAAAAAAAGAAAAAGAAAAGAAAAAGAATTAAAATTAAGTAATTCTTGAAAAAAA  
ACTAATTCTAAGCTTCTTCATGTCAAGGACCTAACATGTCTGTGCAGGTCTTGCATTCCCTGTCAAAAGAAAAAGAATTAAAATTT  
AAGTAATTCTTGAAAAAAAACTAATTCTAACAGCTTCTCATGTCAAGGACCTAACATGTCAAGGCCAAGAGTTGCAAAAAAAAAGAAAAA  
GAAAAGAAAAAGAATTAAAATTTAGTAATTCTTGAAAAAAAACTAATTCTAACAGCTTCTCATGTCAAGGACCTAACATGTAGCCAGAATGG  
TTGTGGGATGGGAGCCTCTGGACCGACCAGGTAGCTCTTTCCACACTGTAGTCTCAAAGCTTCTCATGTGGTTCTGTAGTGAAA  
AAAAAAAAGAAAAAGAAAAAGAAAAAGAATTAAAATTTAAGTAATTCTTGAAAAAAAACTAATTCTAACAGCTTTCTCATGTCAAGGAC  
TAATGTAGCTACTGAACGTTATCTAGGGAAAGATTGAAGGGAGCTAACAGGTCAACACACCACCTCCAGAAAGCTTCTCA.....
```

Genome Analysis



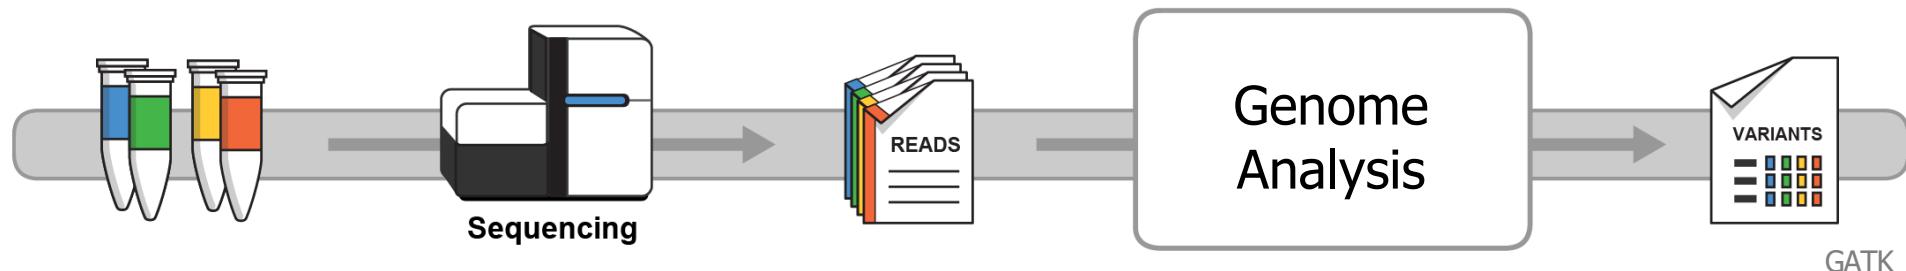
NO machine can read the *entire* content of a genome



Why?

```
>CCTGACCATGGAAGACTAAAGTAAAGGAAATTTAGCTTCAAGCTTCAAGGGACCTAATGTAGCTACTGAACGTTATCTAGGGGAAAGATTGAAGGGAGCTAAGGTCAACACACCACCTCCAGAAAGCTTCA.....  
CAAG  
TTCTT  
CATTG  
AAAAA  
ATTT  
AAAAA  
ATGG  
GAAA
```

Genome Sequencer is a Chopper



CCCCCTATATATAACGTACTAGTACGT
ACGACTTTAGTACGTACGT
TATATATAACGTACTAGTACGT
ACGTACG CCCCTACGTA
TATATATAACGTACTAGTACGT
ACGACTTTAGTACGTACGT
TATATATAACGTACTAAAGTACGT
TATATATAACGTACTAGTACGT
ACG TTTTAAAAACGTA
TATATATAACGTACTAGTACGT
ACGACGGGAGTACGTACGT



1×10^{12} bases*



44 hours*



<1000 \$

* NovaSeq 6000

Oxford Nanopore Sequencers



MinION Mk1B



MinION Mk1C



GridION Mk1



PromethION 24/48

	MinION Mk1B	MinION Mk1C	GridION Mk1	PromethION 24	PromethION 48
Read length	> 2Mb	> 2Mb	> 2Mb	> 2Mb	> 2Mb
Yield per flow cell	50 Gb	50 Gb	50 Gb	220 Gb	220 Gb
Number of flow cells per device	1	1	5	24	48
Yield per device	<50 Gb	<50 Gb	<250 Gb	<5.2 Tb	<10.5 Tb
Starting price	\$1,000	\$4,990	\$49,995	\$195,455	\$327,455

Illumina Sequencers

illumina®



iSeq 100

MiniSeq

MiSeq

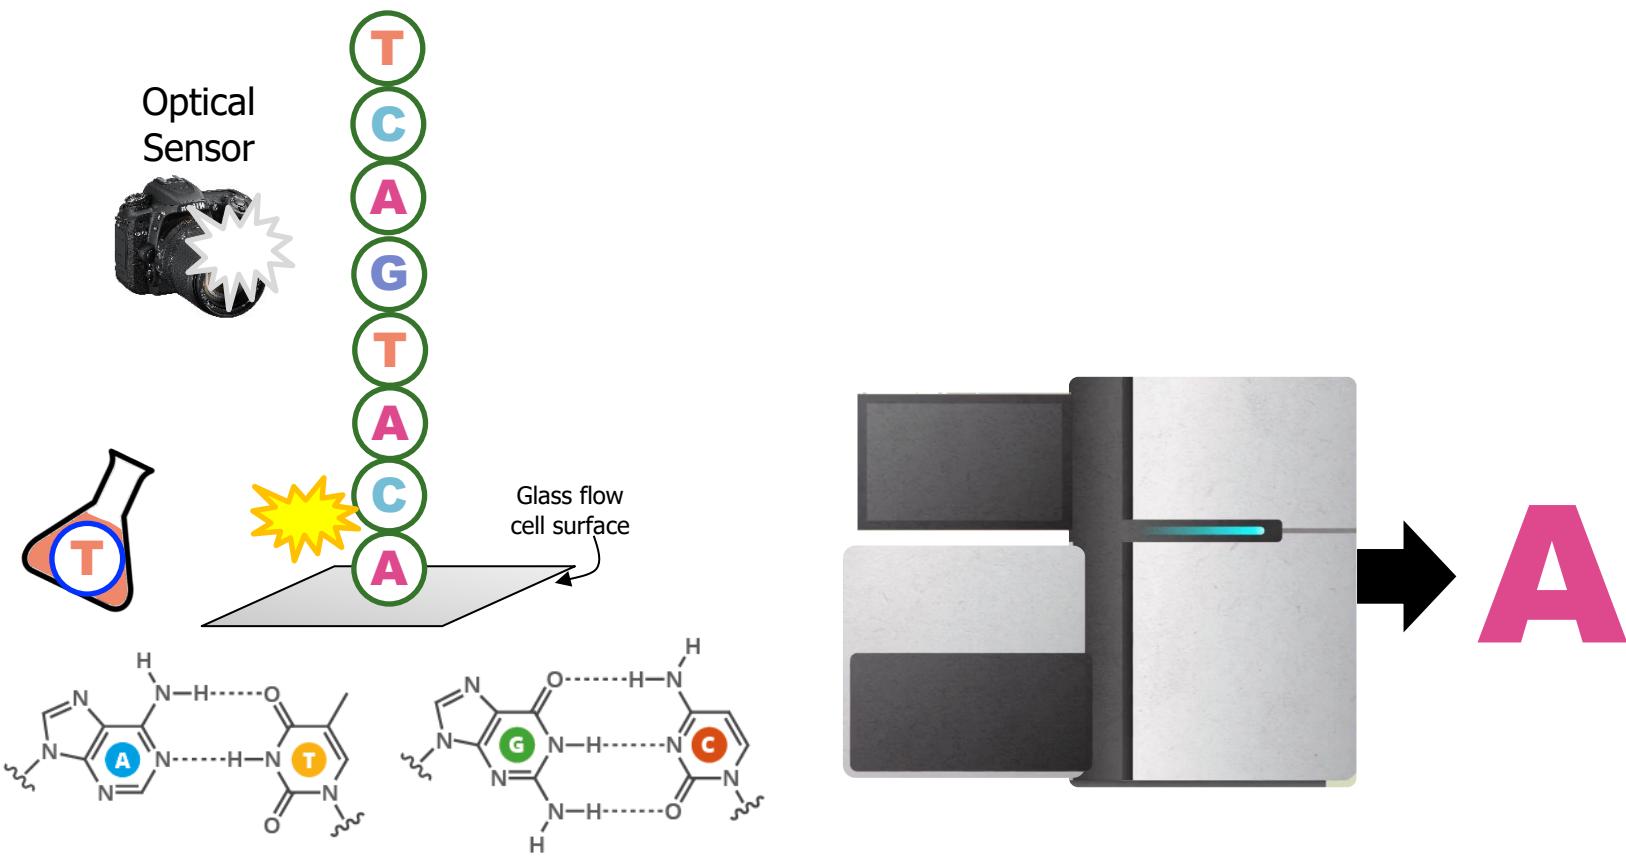
NextSeq 550

NextSeq 2000

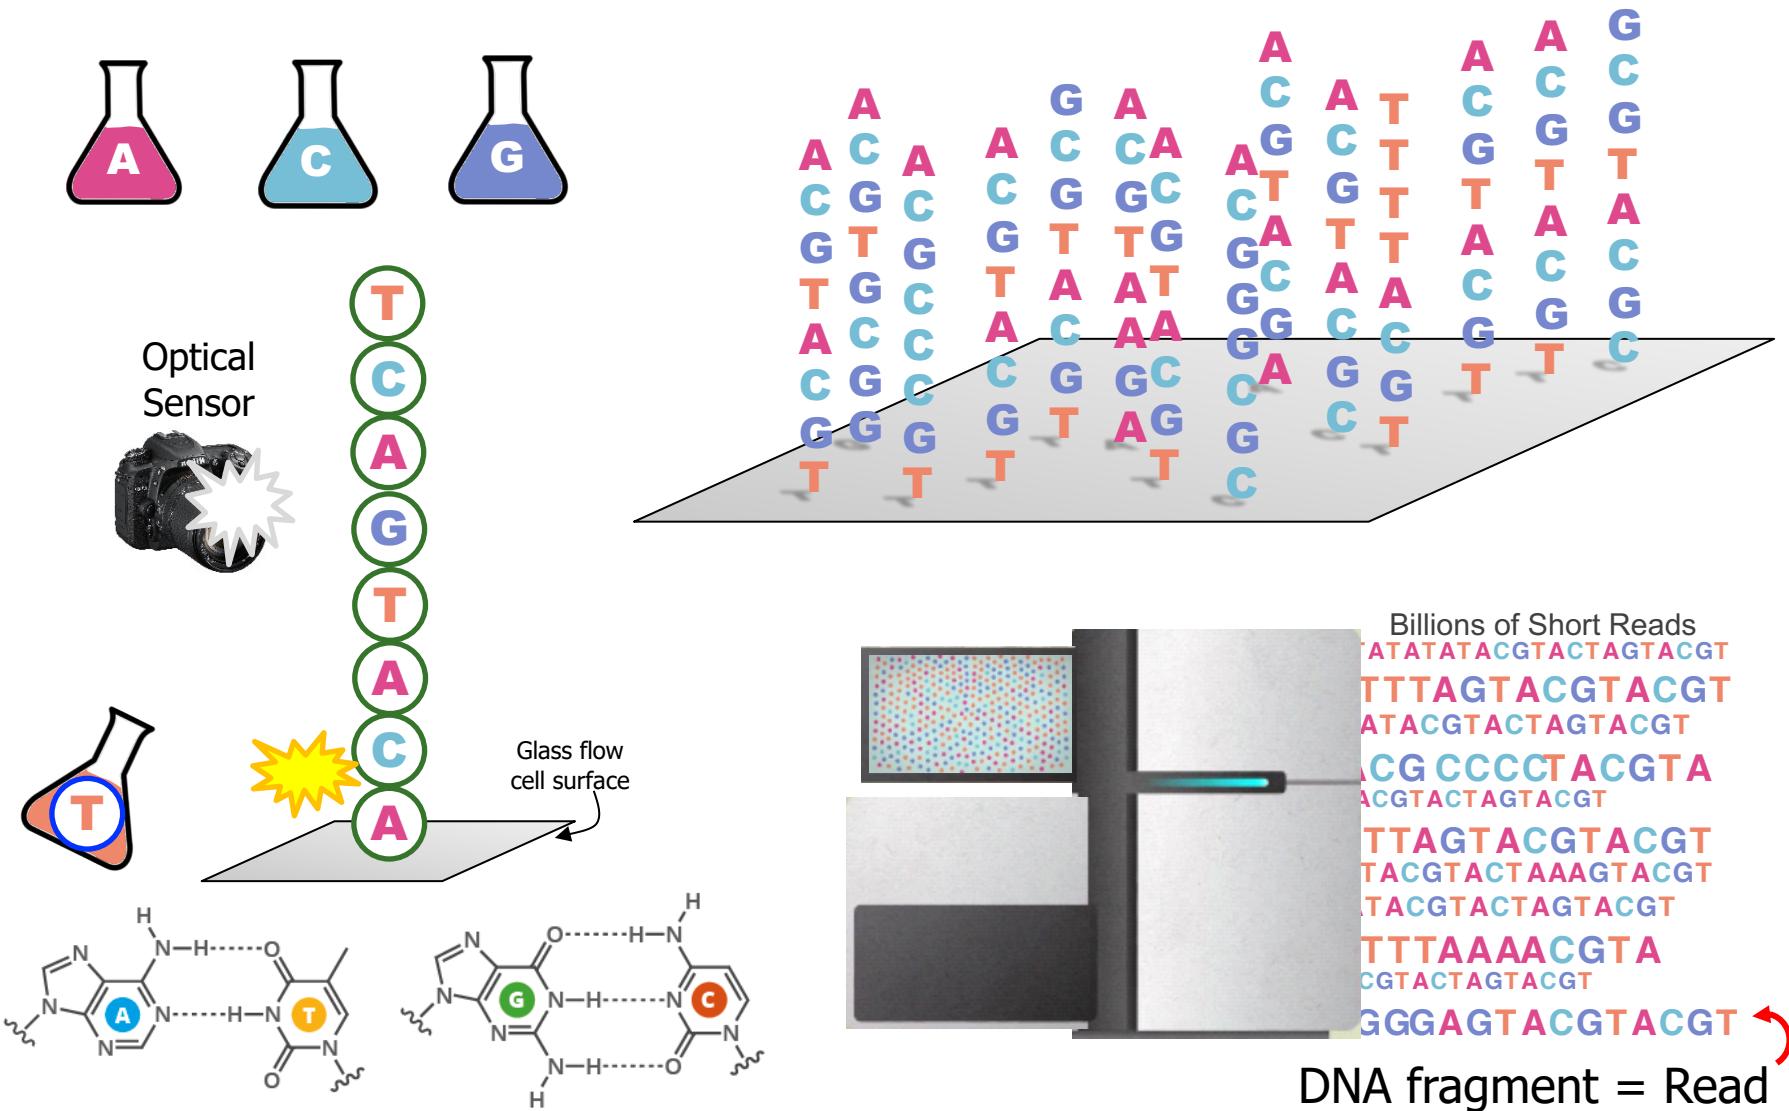
NovaSeq 6000

Run time	9.5–19 hrs	4–24 hrs	4–55 hrs	12–30 hrs	24–48 hrs	13–44 hrs
Max. reads per run	4 million	25 million	25 million	400 million	1 billion	20 billion
Max. read length	2 × 150 bp	2 × 150 bp	2 × 300 bp	2 × 150 bp	2 × 150 bp	2 × 250
Max. output	1.2 Gb	7.5 Gb	15 Gb	120 Gb	300 Gb	6000 Gb
Estimated price	\$19,900	\$49,500	\$128,000	\$275,000	\$335,000	\$985,000

How Does Illumina Machine Work?



How Does Illumina Machine Work?

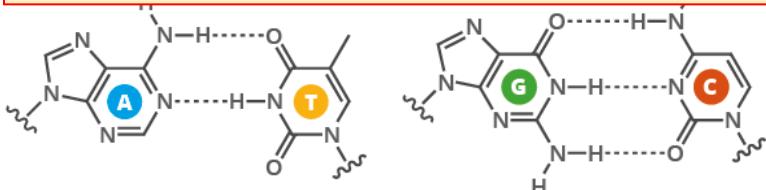


How Does Illumina Machine Work?



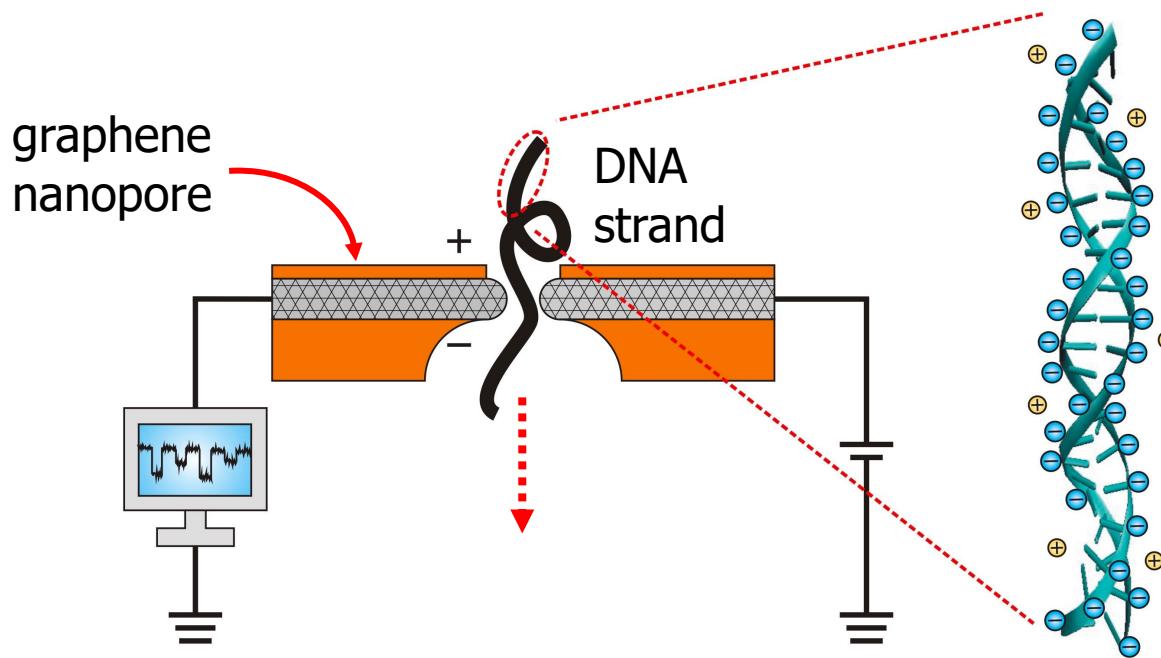
Check Illumina virtual tour:

<https://emea.illumina.com/systems/sequencing-platforms/iseq/tour.html>



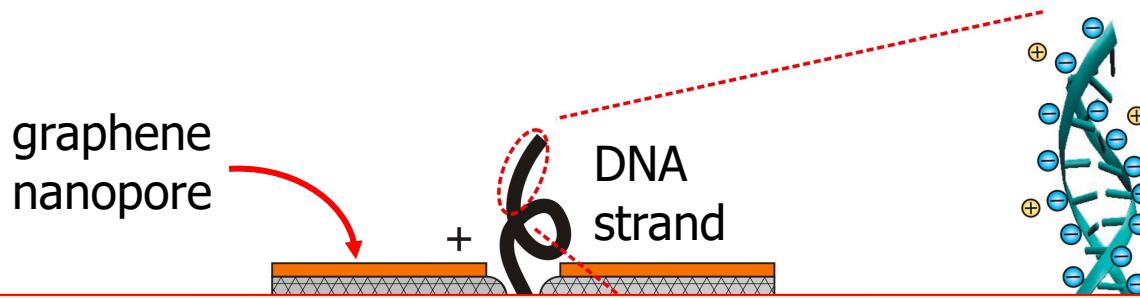
DNA fragment = Read

How Does Nanopore Machine Work?



- **Nanopore** is a nano-scale hole (<20nm).
- In nanopore sequencers, an **ionic current** passes through the nanopores
- When the DNA strand passes through the nanopore, the sequencer measures the the **change in current**
- This change is used to identify the bases in the strand with the help of **different electrochemical structures** of the different bases

How Does Nanopore Machine Work?



Check Nanopore virtual tour:

<https://nanoporetech.com/resource-centre/minion-video>

- measures the the **change in current**
- This change is used to identify the bases in the strand with the help of **different electrochemical structures** of the different bases

Solving the Puzzle



Reference genome



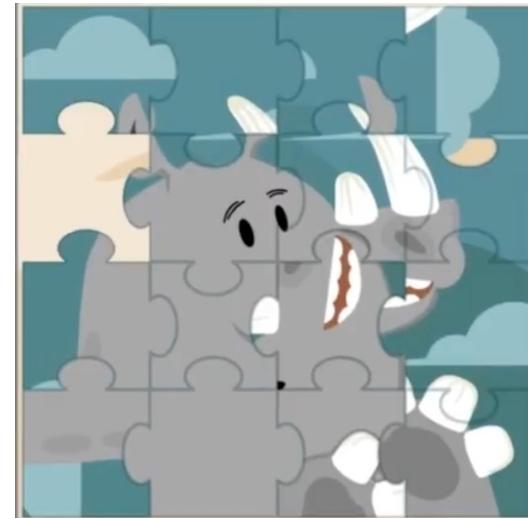
<https://www.pacb.com/smrt-science/smrt-sequencing/hifi-reads-for-highly-accurate-long-read-sequencing/>

HTS Sequencing Output

Small pieces of a puzzle
short reads (Illumina)



Large pieces of a puzzle
long reads (ONT & PacBio)



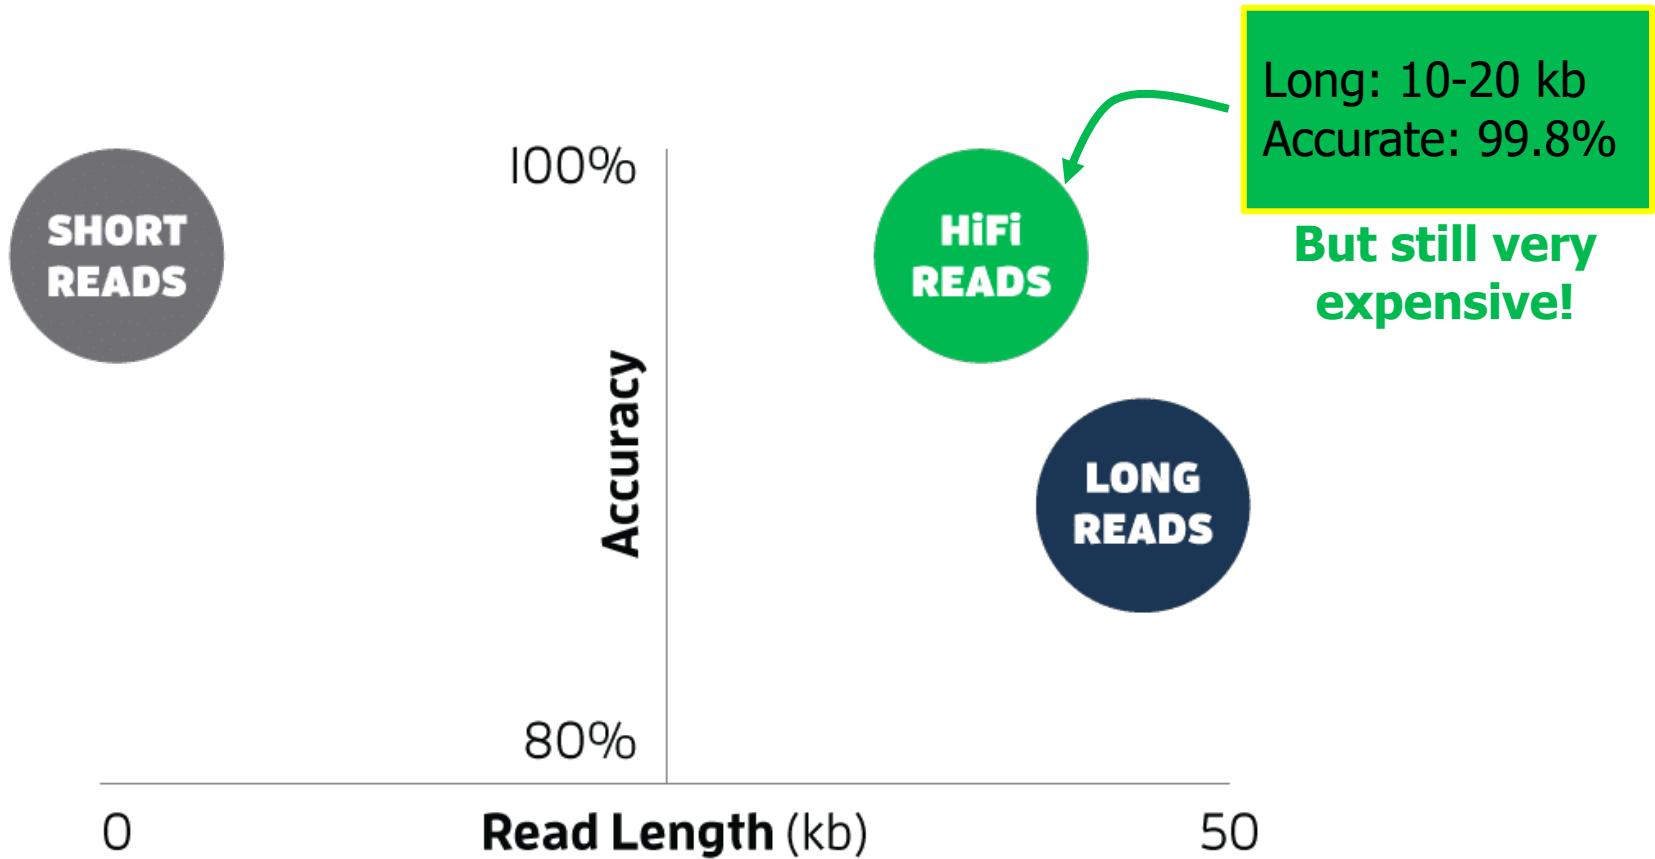
Which sequencing technology is the best?

- 100-300 bp
- low error rate (~0.1%)

- 500-2M bp
- high error rate (~15%)

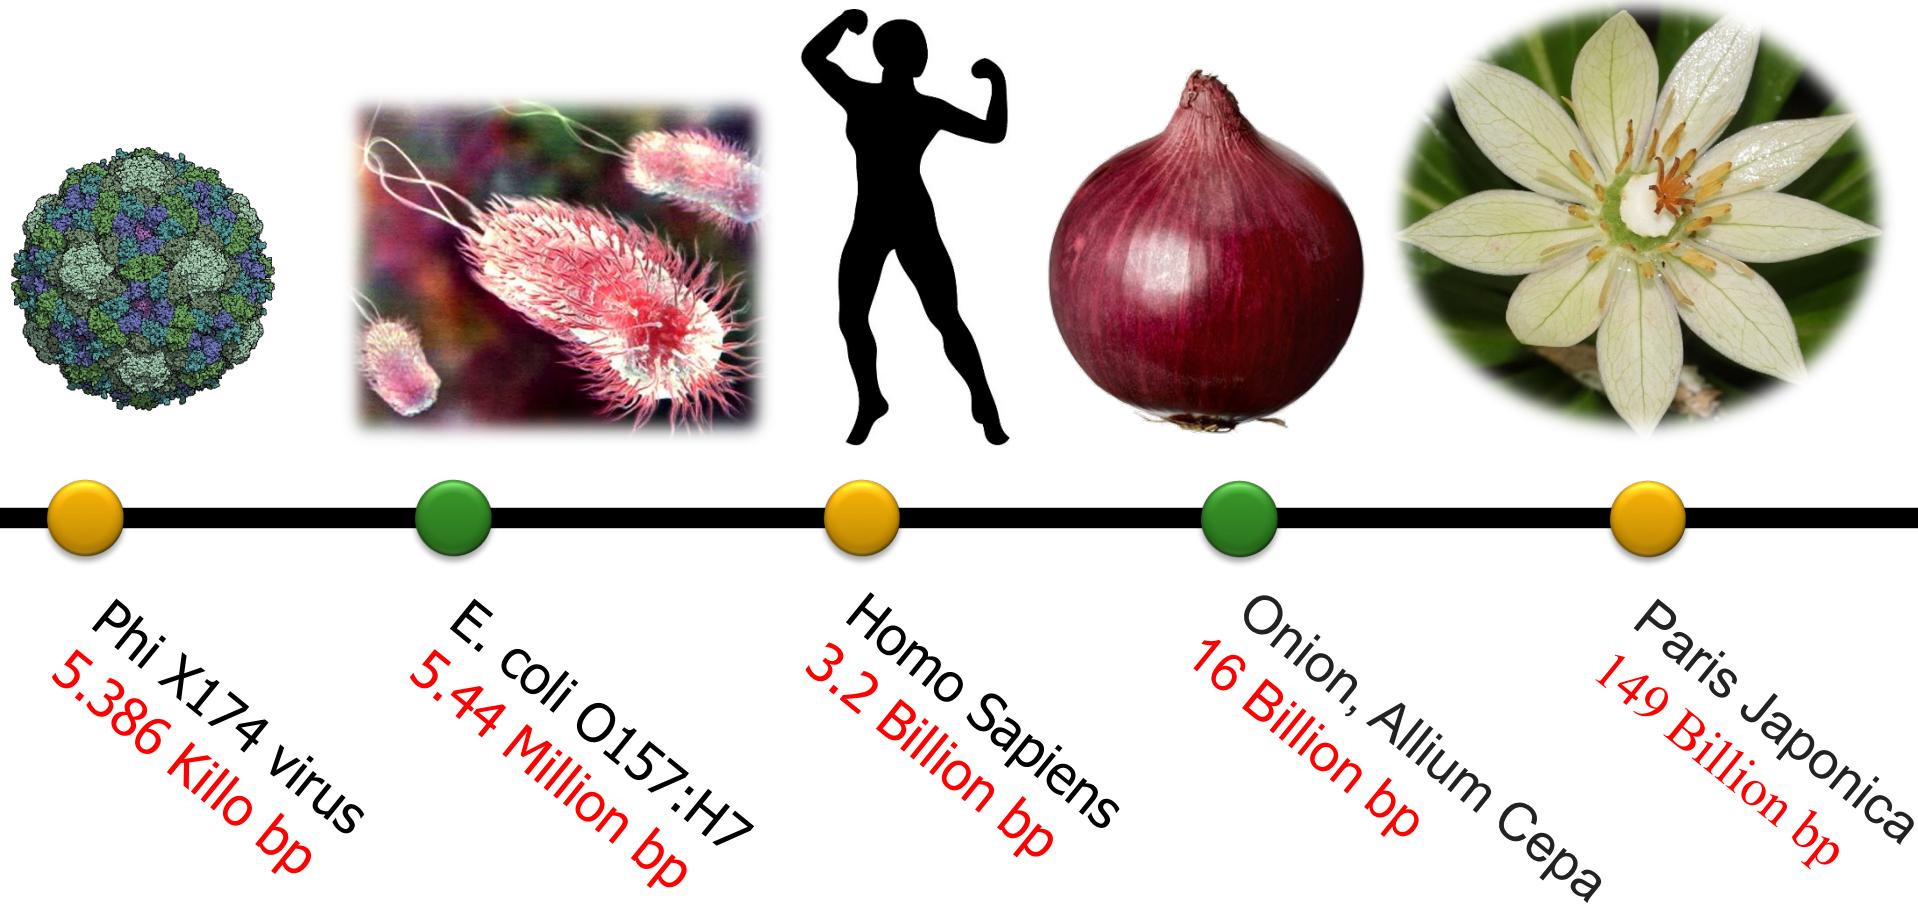
<https://www.pacb.com/smrt-science/smrt-sequencing/hifi-reads-for-highly-accurate-long-read-sequencing/>

HiFi Reads (PacBio)



Wenger+, "Accurate circular consensus long-read sequencing improves variant detection and assembly of a human genome", *Nature Biotechnology*, 2019

How Long is DNA?



Cracking the 1st Human Genome Sequence

- **1990-2003:** The Human Genome Project (HGP) provides a complete and accurate sequence of all **DNA base pairs** that make up the human genome and finds 20,000 to 25,000 human genes.



A
C
G
T

3.2×10^9
bases



13 years



$>3 \times 10^9$ \$

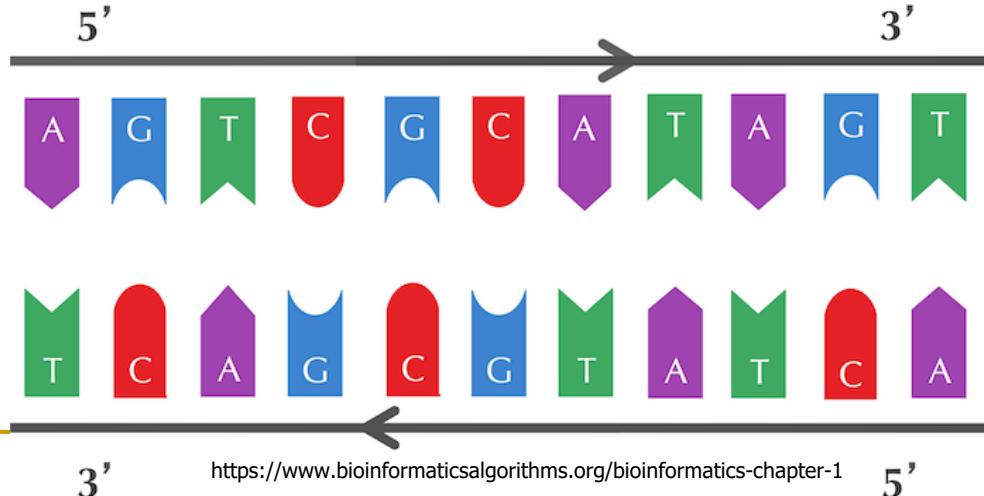
Obtaining the Human Reference Genome

- **GRCh38.p13**
 - Description: Genome Reference Consortium Human Build 38 patch release 13 (GRCh38.p13)
 - Organism name: Homo sapiens (human)
 - Date: 2019/02/28
 - 3,099,706,404 bases
 - Compressed .fna file (964.9 MB)
 - https://www.ncbi.nlm.nih.gov/assembly/GCF_000001405.39

>NC_000001.11 Homo sapiens chromosome 1, GRCh38.p13 Primary Assembly

Challenges in Read Mapping

- Need to find many **mappings** of **each read**
- Need to **tolerate variances/sequencing errors** in each read
- Need to **map** each read **very fast** (i.e., performance is important, life critical in some cases)
- Need to **map** reads to both **forward and reverse strands**



Revisiting the Puzzle



<https://www.pacb.com/smrt-science/smrt-sequencing/hifi-reads-for-highly-accurate-long-read-sequencing/>

Reference Genome Bias

nature genetics

Letter | Open Access | Published: 19 November 2018

Assembly of a pan-genome from deep sequencing of 910 humans of African descent

Rachel M. Sherman , Juliet Forman, [...] Steven L. Salzberg 

Nature Genetics 51, 30–35(2019) | Cite this article

“African pan-genome contains ~10% more DNA bases than the current human reference genome”

Time to Change the Reference Genome

The screenshot shows the header of the *Genome Biology* website. The title "Genome Biology" is in a yellow-bordered box. Below the header, there are navigation links: Home, About, Articles, and Submission Guidelines. The main content area displays the article details.

Home About Articles Submission Guidelines

Opinion | **Open Access** | Published: 09 August 2019

Is it time to change the reference genome?

Sara Ballouz, Alexander Dobin & Jesse A. Gillis [✉](#)

Genome Biology **20**, Article number: 159 (2019) | [Cite this article](#)

12k Accesses | 11 Citations | 45 Altmetric | [Metrics](#)

“Switching to a consensus reference would offer important advantages over the continued use of the current reference with few disadvantages”

MAGNET (AACBB 2018, TIR 2017)

- Key observation: the use of **AND operation** to check if a zero (match) exists in a column introduces filtering inaccuracy.
 - Key Idea: count the **consecutive zeros** in each mask and select the longest in a divide-and-conquer approach.
 - **MAGNET** is **17x to 105x more accurate** than GateKeeper and SHD.

The figure displays a sequence of DNA bases (A, T, C, G) and their binary representation. The sequence starts with 15 As, followed by 15 Ts, 15 Cs, and 15 Gs. Red vertical lines mark positions 2, 4, 1, and 3. Yellow circles labeled 2, 4, 1, and 3 are placed above these lines. Below the sequence, binary digits 0 and 1 are shown for each position, with green boxes highlighting specific patterns.

MAGNET Walkthrough



Find the longest segment of consecutive zeros

Exclude the errors from the search space

Divide the problem into two subproblems and repeat

What if we got a new version
of the reference genome?

AirLift

- **Key observation:** Reference genomes are updated frequently. Repeating *read mapping* is a computationally expensive workload.
- **Key idea:** Update the mapping results of only affected reads depending on how a region in the old reference relates to another region in the new reference.
- **Key results:**
 - reduces number of reads that needs to be re-mapped to new reference by up to 99%
 - reduces overall runtime to re-map reads by 6.94x, 208x, and 16.4x for large (human), medium (*C. elegans*), and small (yeast) reference genomes

Clustering the Reference Genome Regions

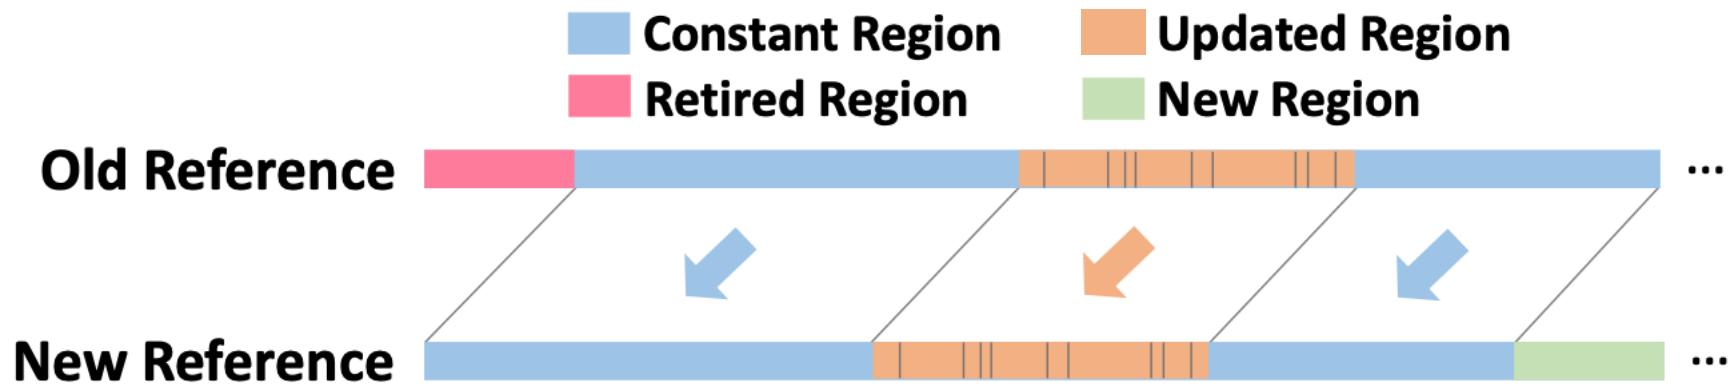


Fig. 2. Reference Genome Regions.

More Details on AirLift

arXiv.org > q-bio > arXiv:1912.08735

Search...

Help | Advanced

Quantitative Biology > Genomics

[Submitted on 18 Dec 2019]

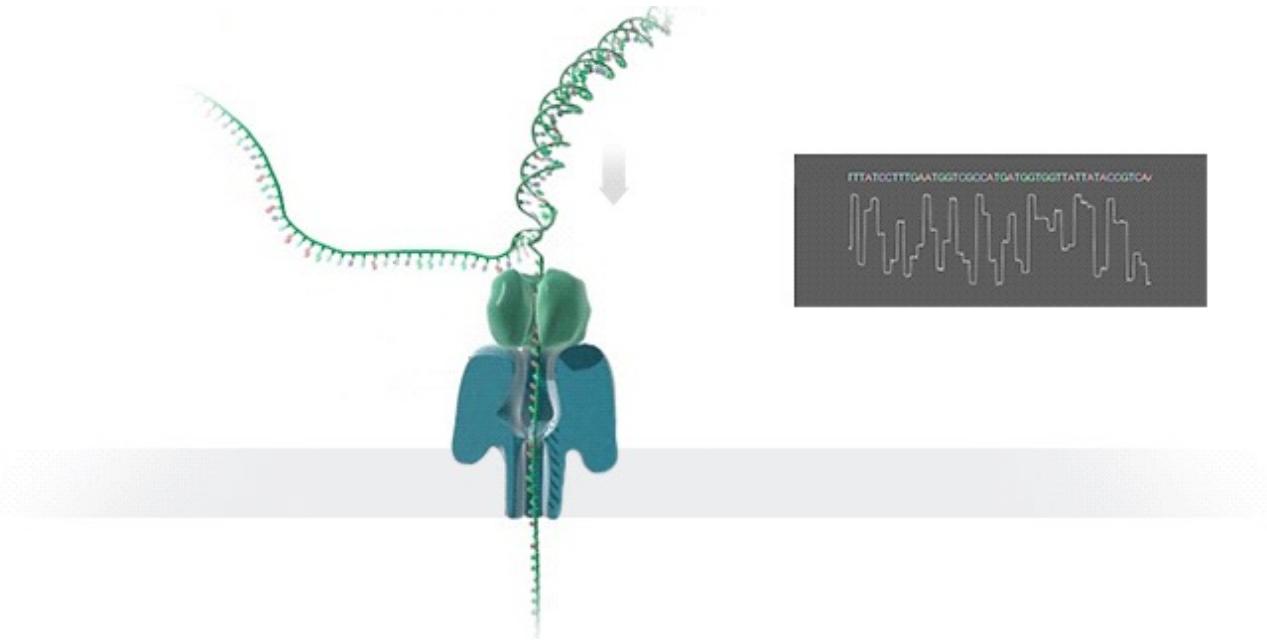
AirLift: A Fast and Comprehensive Technique for Translating Alignments between Reference Genomes

Jeremie S. Kim, Can Firtina, Damla Senol Cali, Mohammed Alser, Nastaran Hajinazar, Can Alkan, Onur Mutlu

GitHub: <https://github.com/CMU-SAFARI/AirLift>

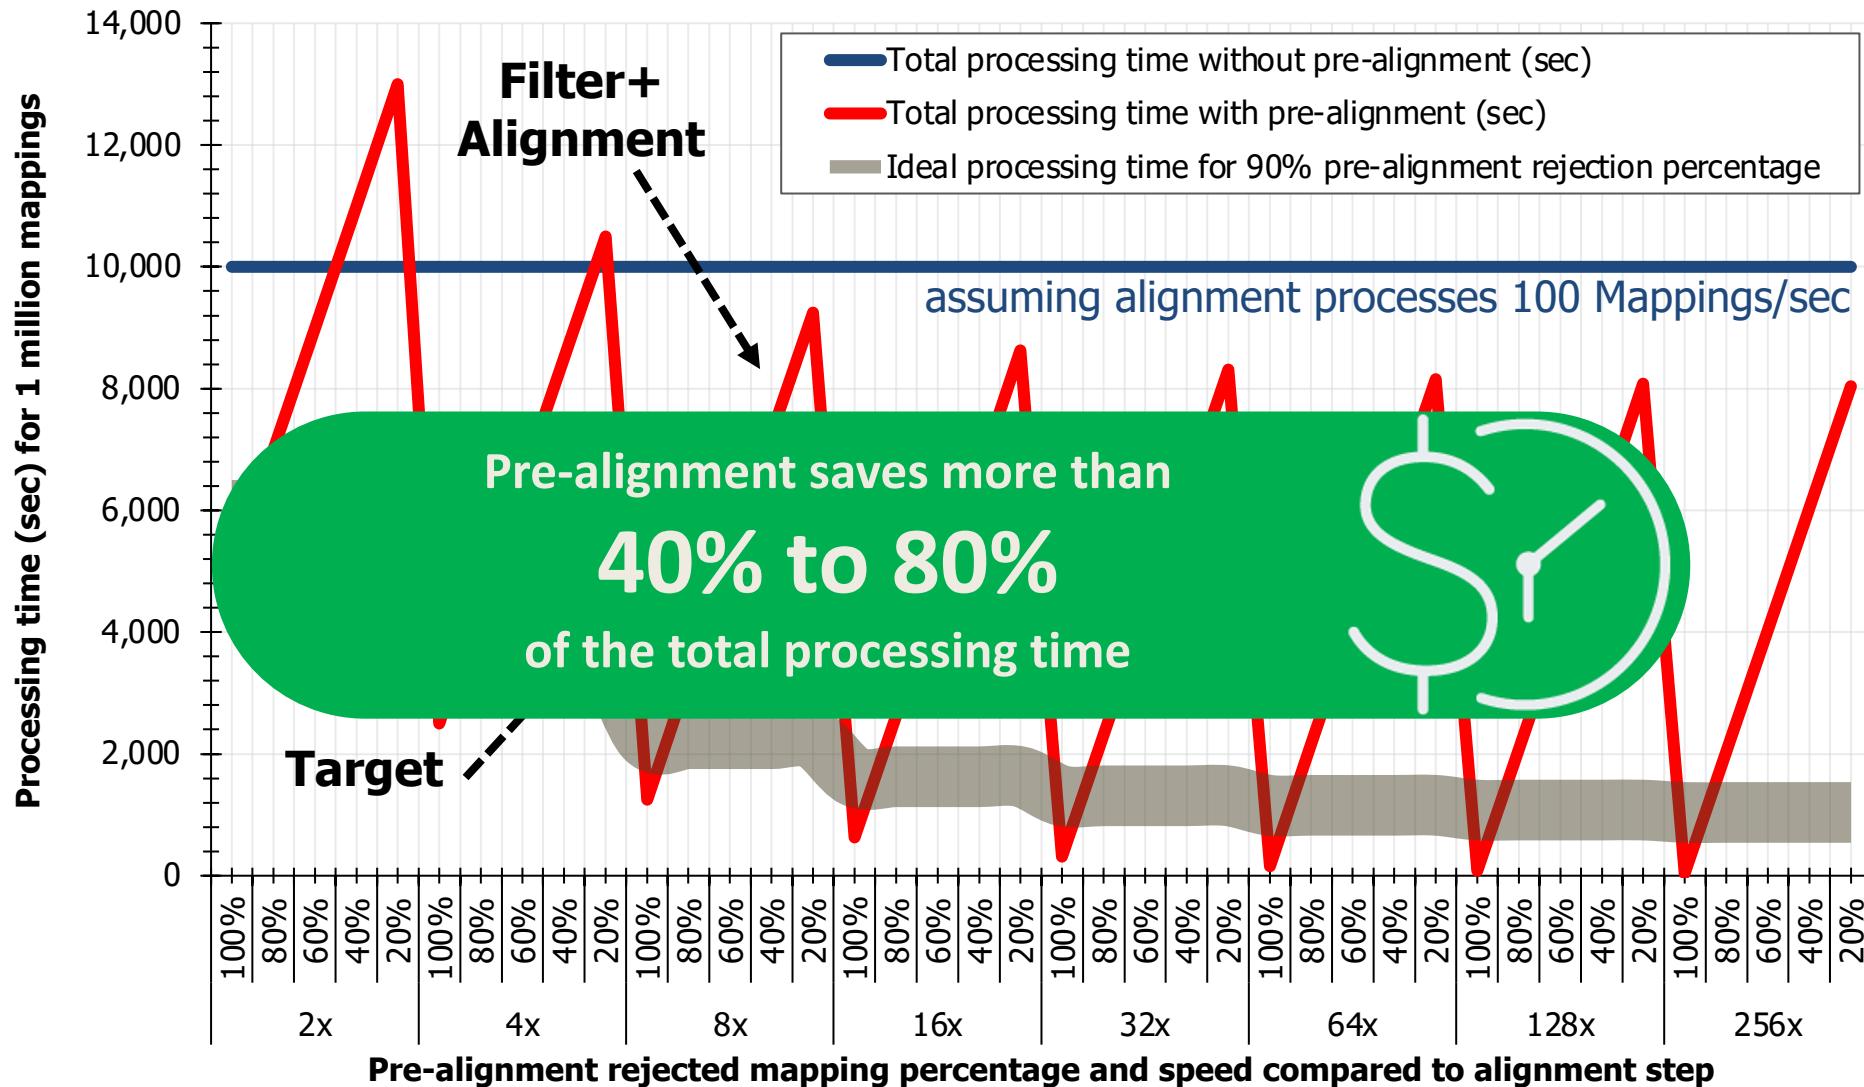
Kim+, "AirLift: A Fast and Comprehensive Technique for Translating Alignments between Reference Genomes", arXiv, 2020

Nanopore Sequencing



- **Nanopore** is a nano-scale hole
- In nanopore sequencers, an **ionic current** passes through the nanopores
- When the DNA strand passes through the nanopore, the sequencer measures the the **change in current**
- This change is used to identify the bases in the strand with the help of **different electrochemical structures** of the different bases

The Effect of Pre-Alignment (Theoretically)



Aside: In-Memory Graph Processing

- Large graphs are everywhere (circa 2015)



36 Million
Wikipedia Pages



1.4 Billion
Facebook Users

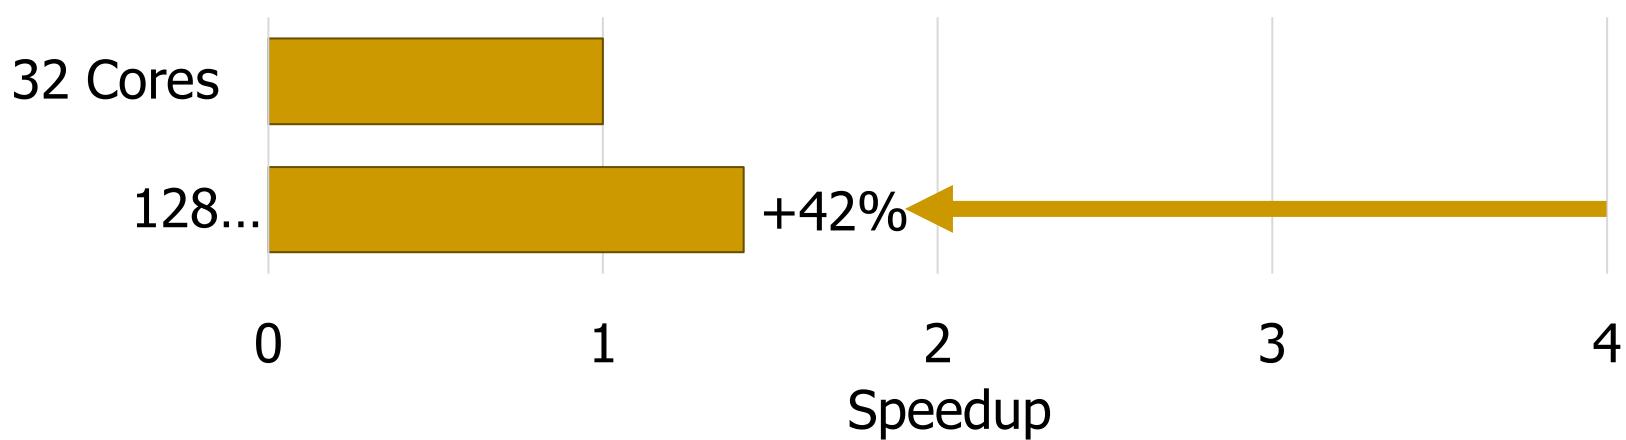


300 Million
Twitter Users



30 Billion
Instagram Photos

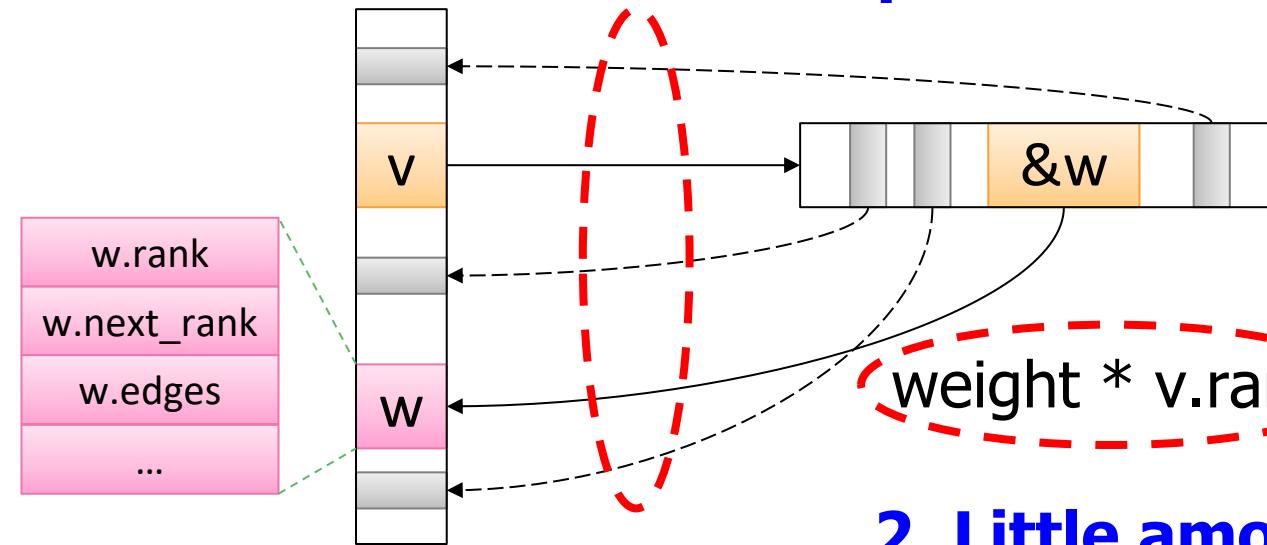
- Scalable large-scale graph processing is challenging



Key Bottlenecks in Graph Processing

```
for (v: graph.vertices) {  
    for (w: v.successors) {  
        w.next_rank += weight * v.rank;  
    }  
}
```

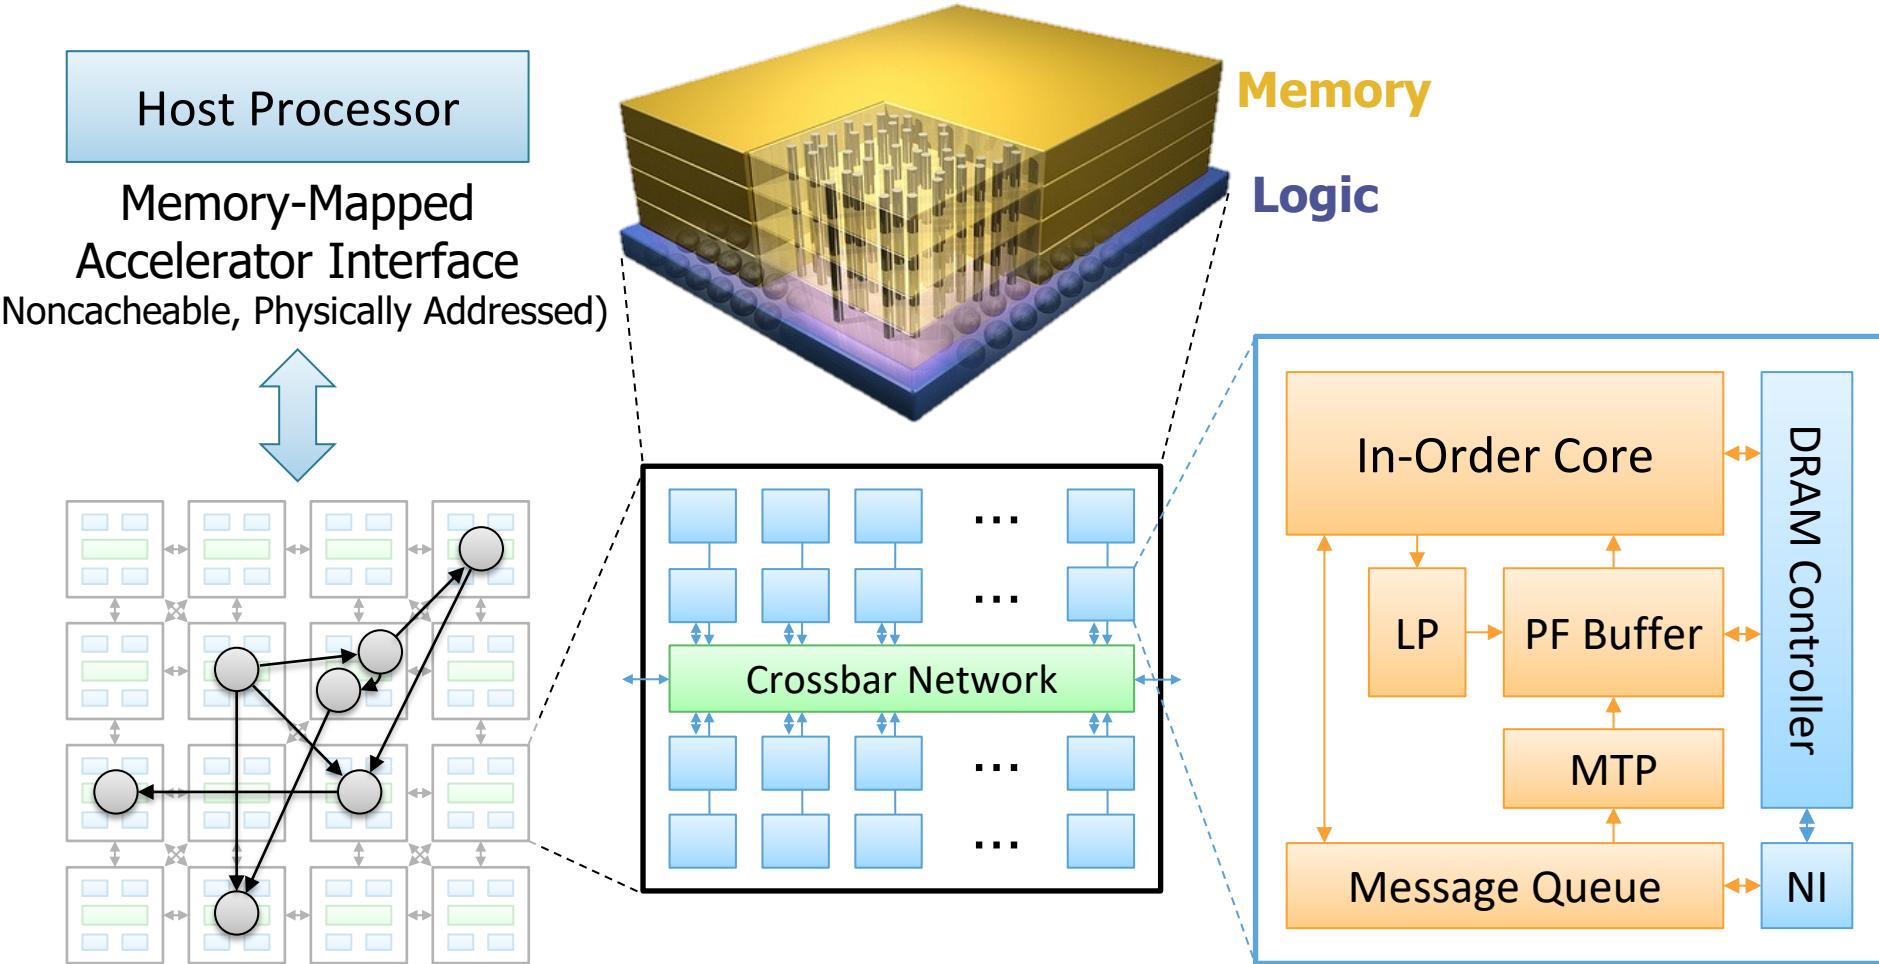
1. Frequent random memory accesses



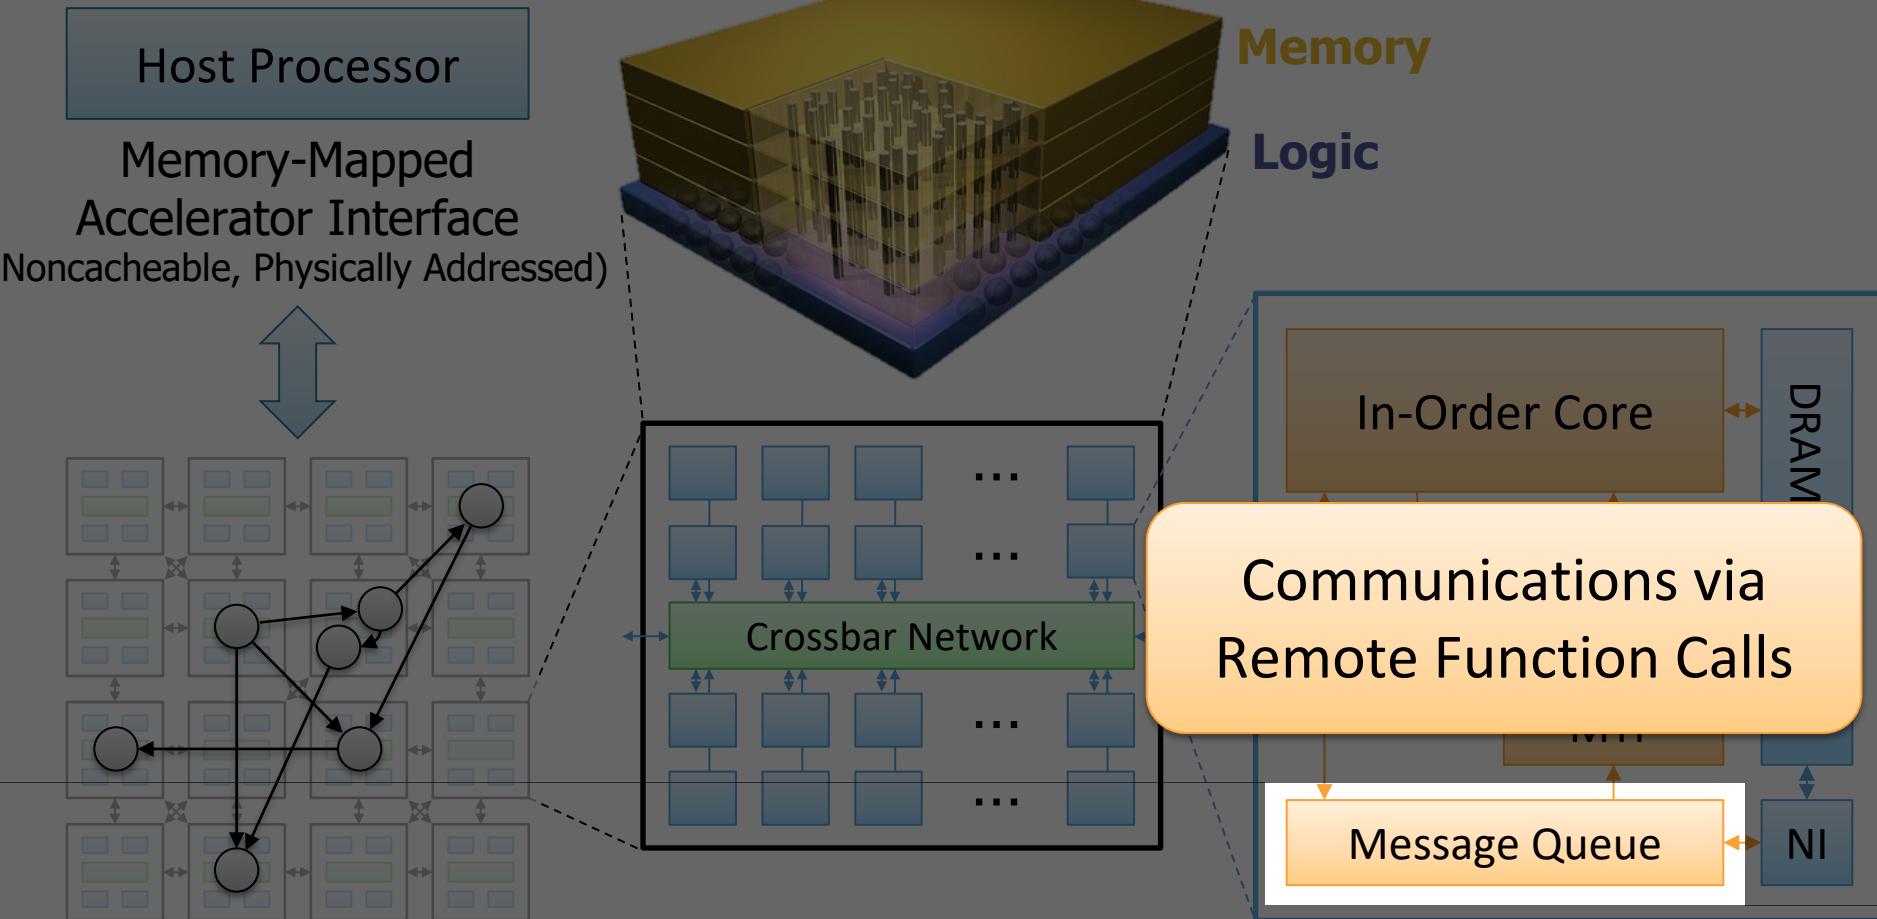
2. Little amount of computation

Tesseract System for Graph Processing

Interconnected set of 3D-stacked memory+logic chips with simple cores

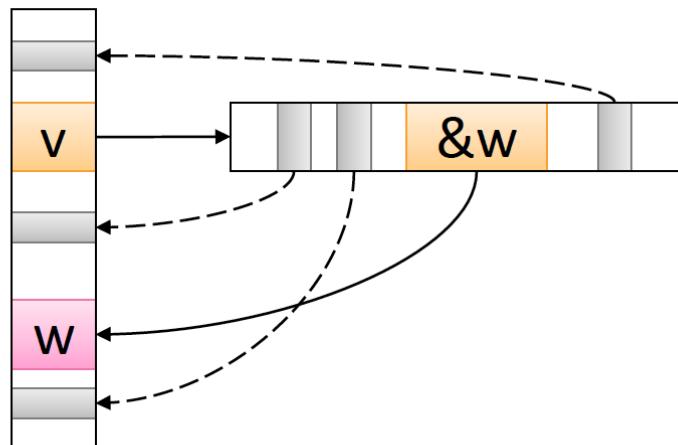


Tesseract System for Graph Processing



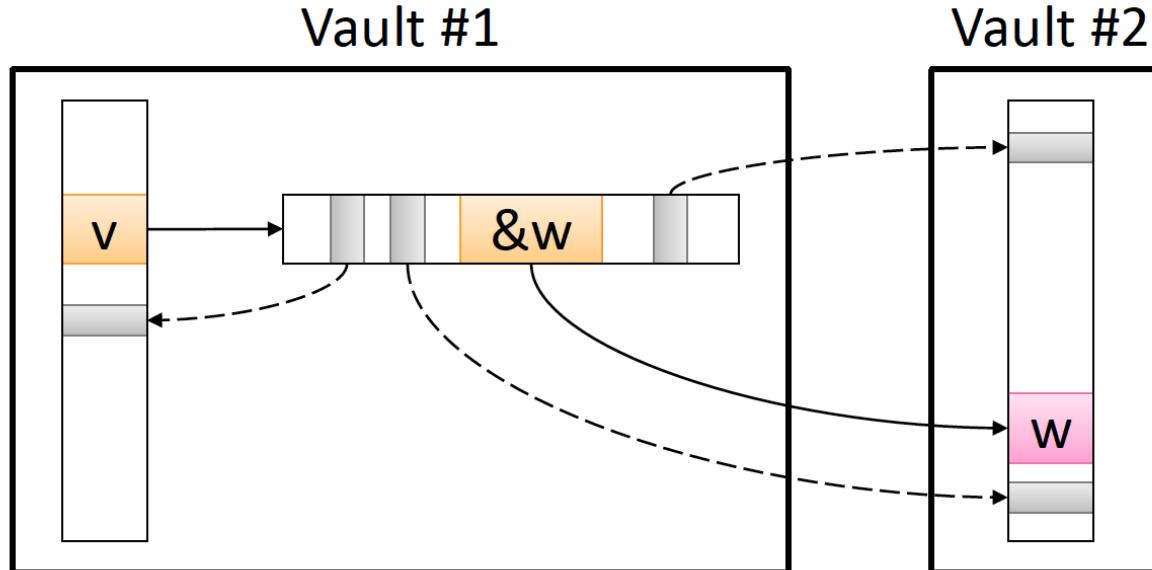
Communications In Tesseract (I)

```
for (v: graph.vertices) {  
    for (w: v.successors) {  
        w.next_rank += weight * v.rank;  
    }  
}
```



Communications In Tesseract (II)

```
for (v: graph.vertices) {  
    for (w: v.successors) {  
        w.next_rank += weight * v.rank;  
    }  
}
```

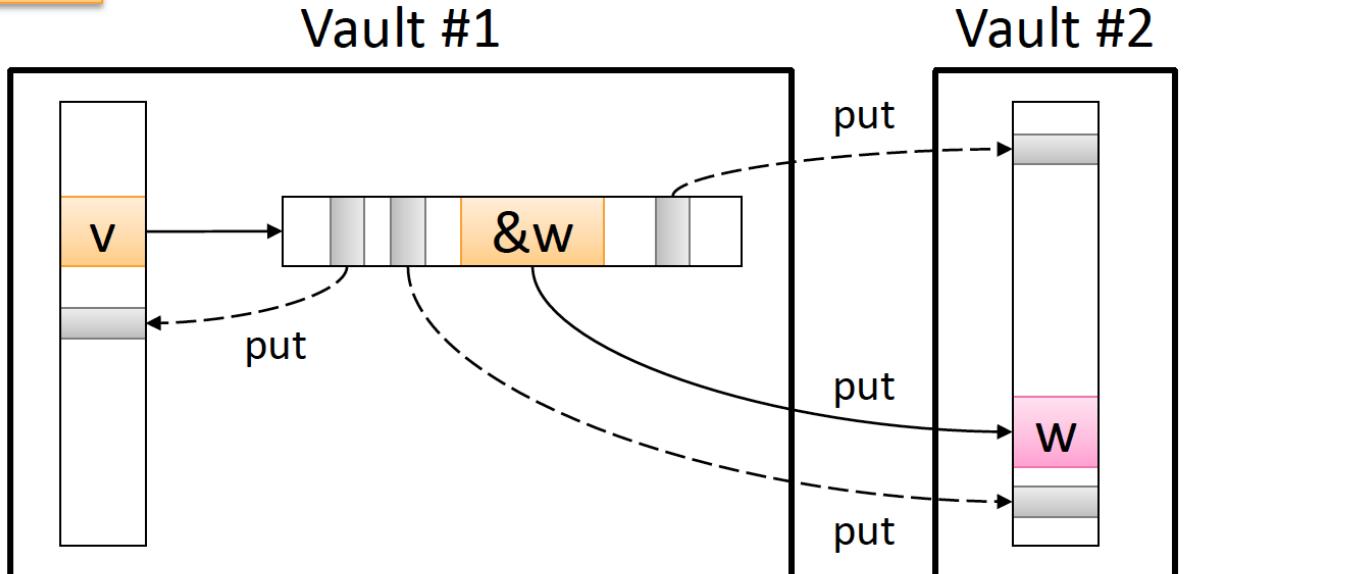


Communications In Tesseract (III)

```
for (v: graph.vertices) {  
    for (w: v.successors) {  
        put(w.id, function() { w.next_rank += weight * v.rank; });  
    }  
}  
barrier();
```

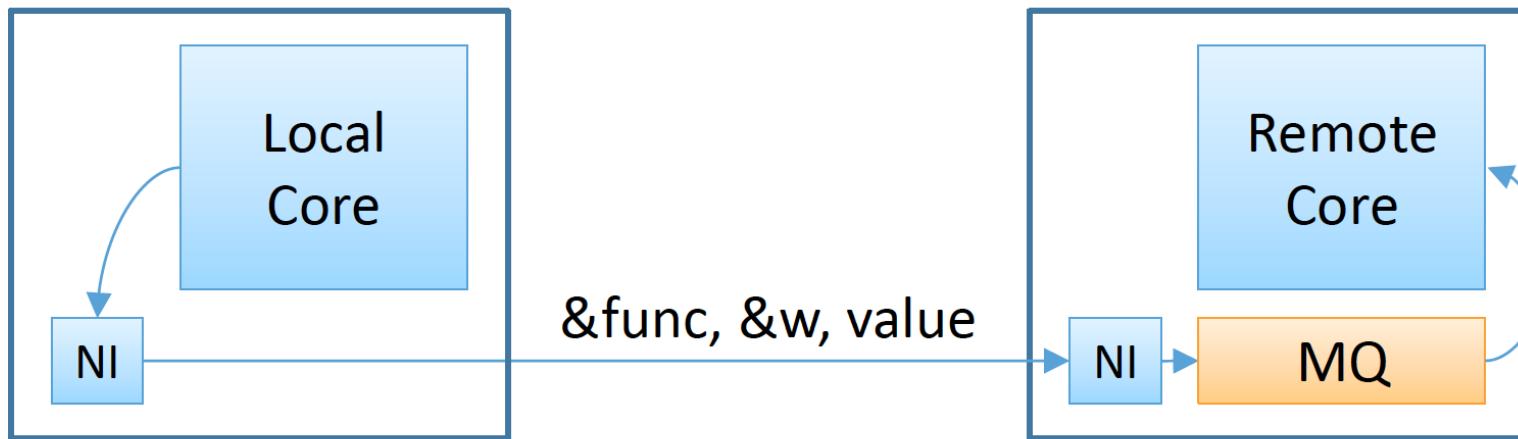
Non-blocking Remote Function Call

Can be **delayed**
until the nearest barrier



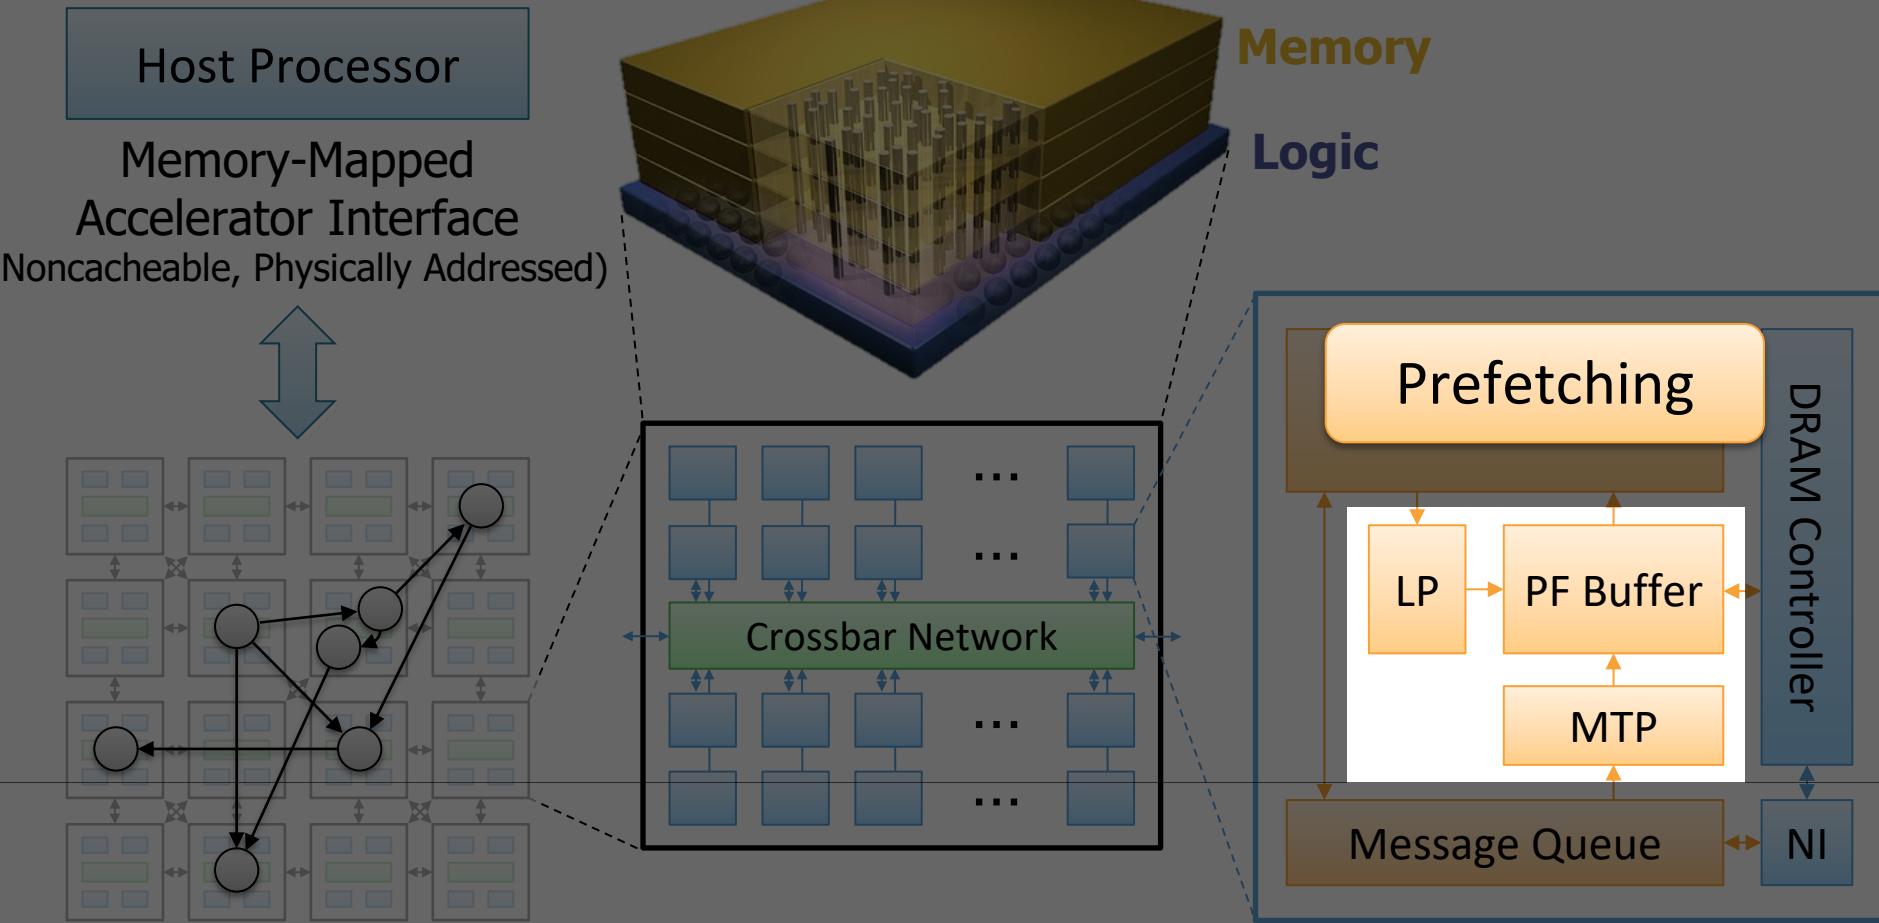
Remote Function Call (Non-Blocking)

1. Send function address & args to the remote core
2. Store the incoming message to the message queue
3. Flush the message queue when it is full or a synchronization barrier is reached



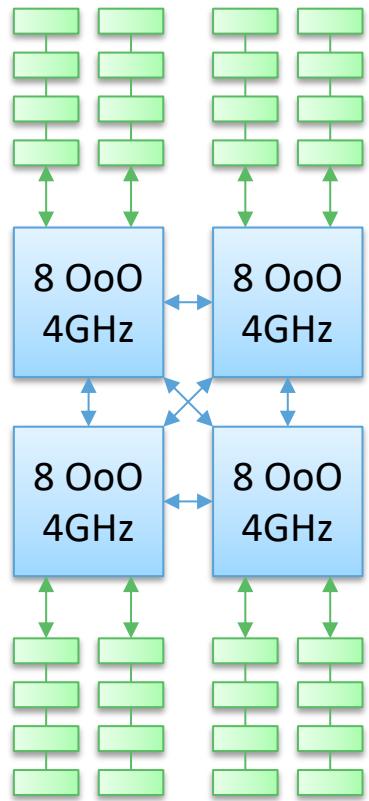
```
put(w.id, function() { w.next_rank += value; })
```

Tesseract System for Graph Processing

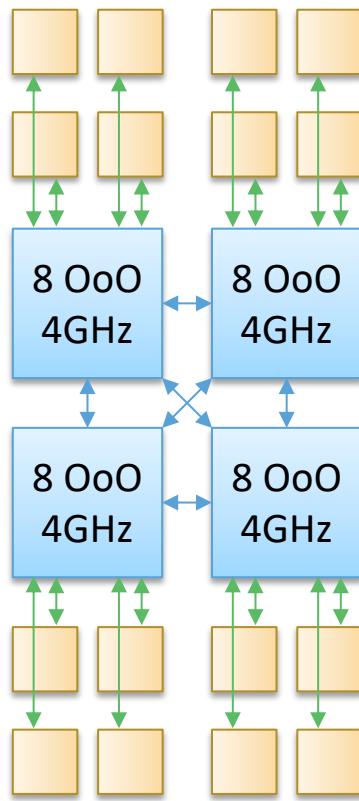


Evaluated Systems

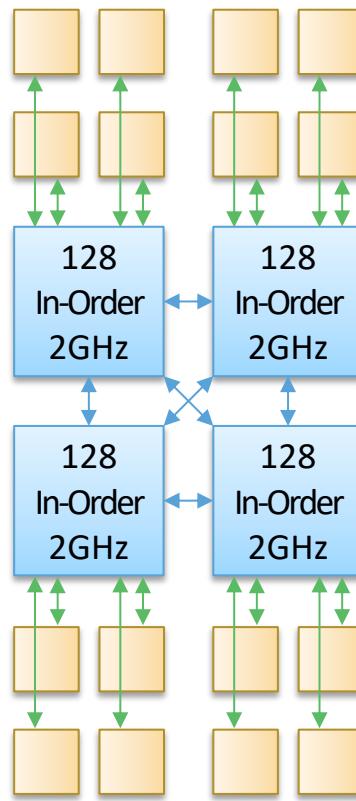
DDR3-OoO



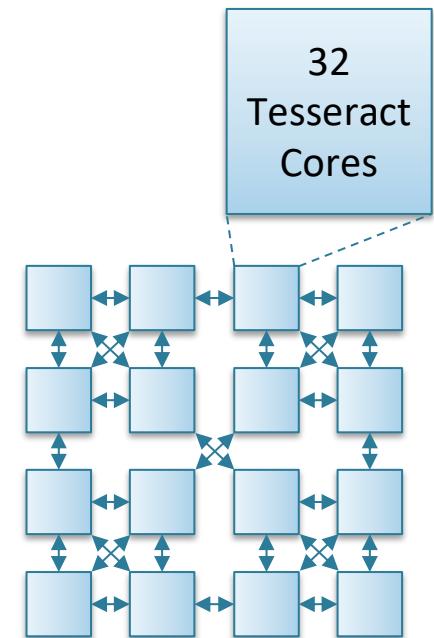
HMC-OoO



HMC-MC



Tesseract



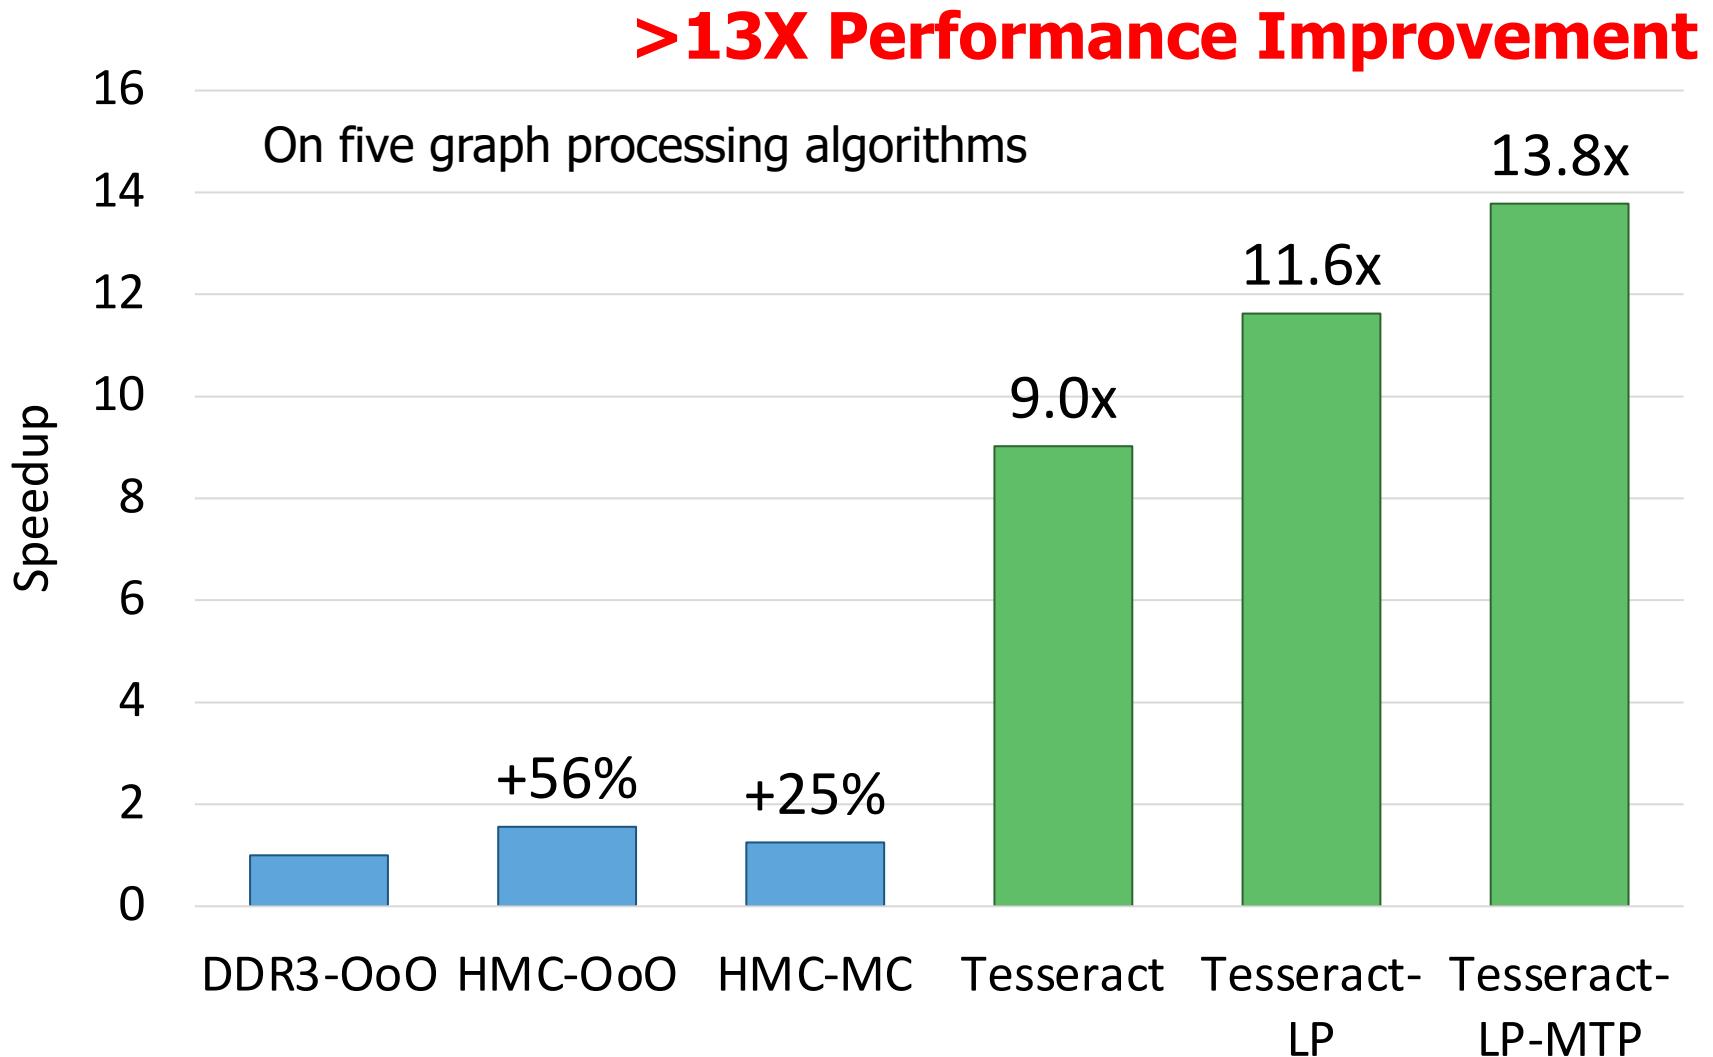
102.4GB/s

640GB/s

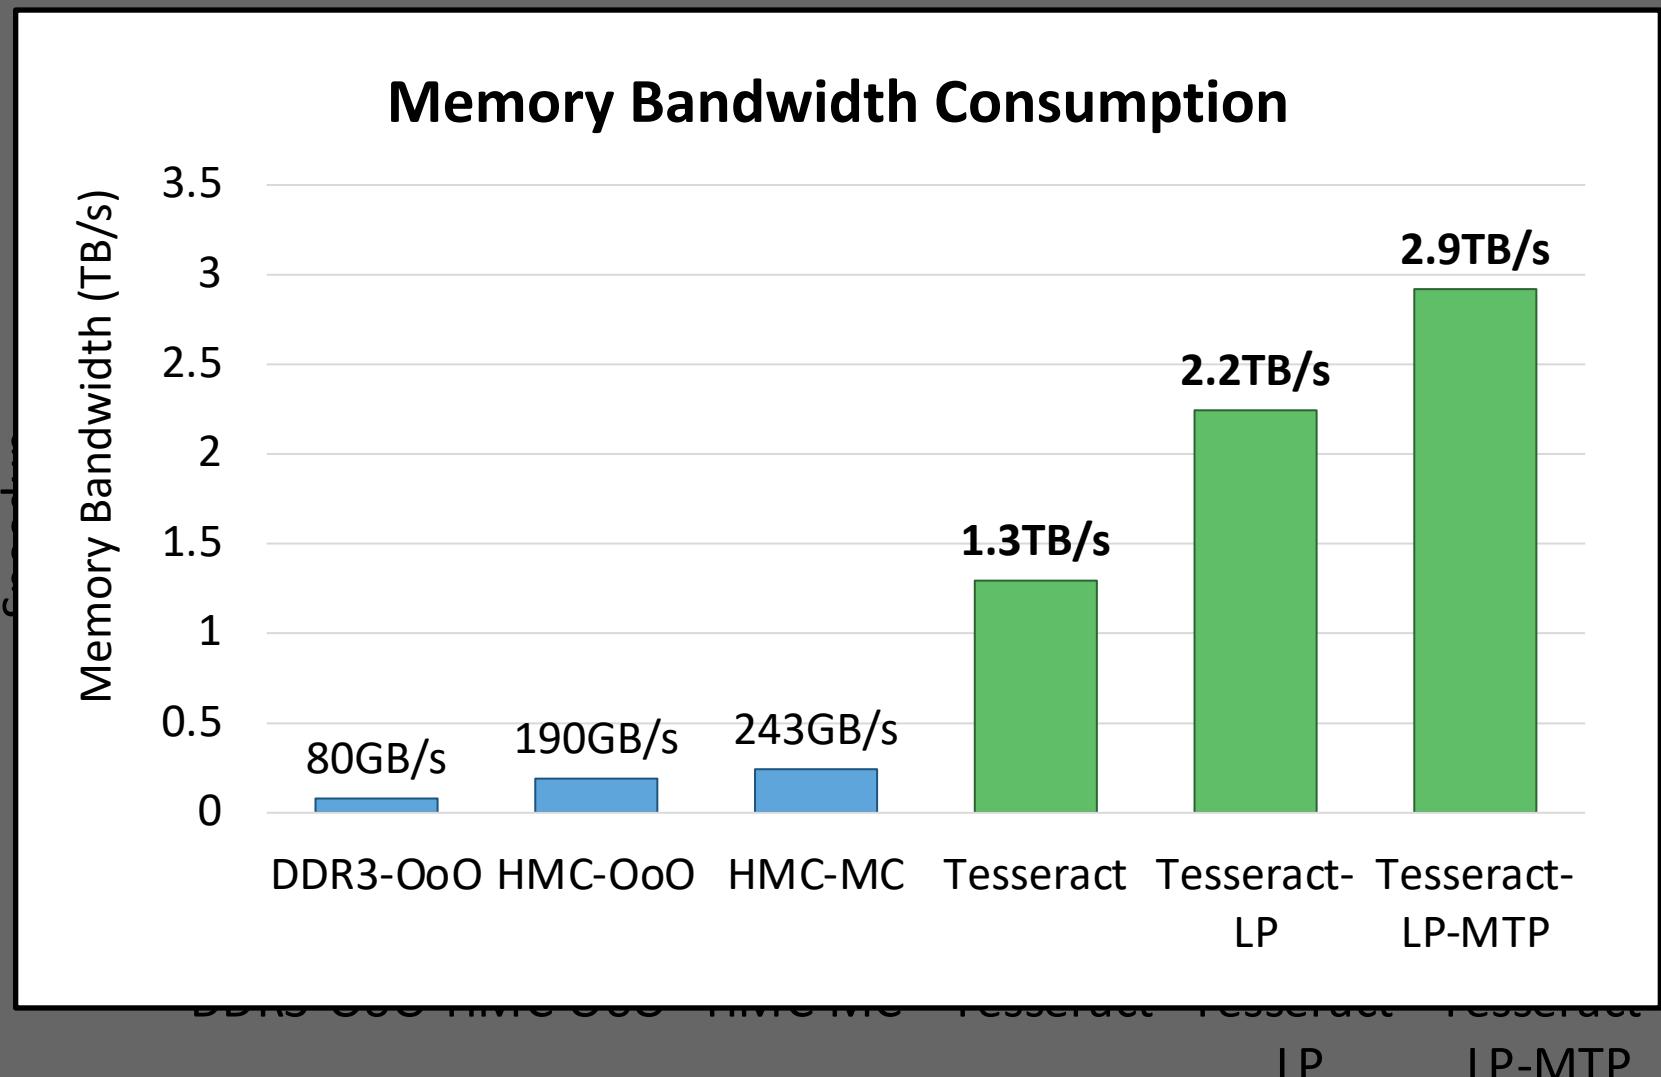
640GB/s

8TB/s

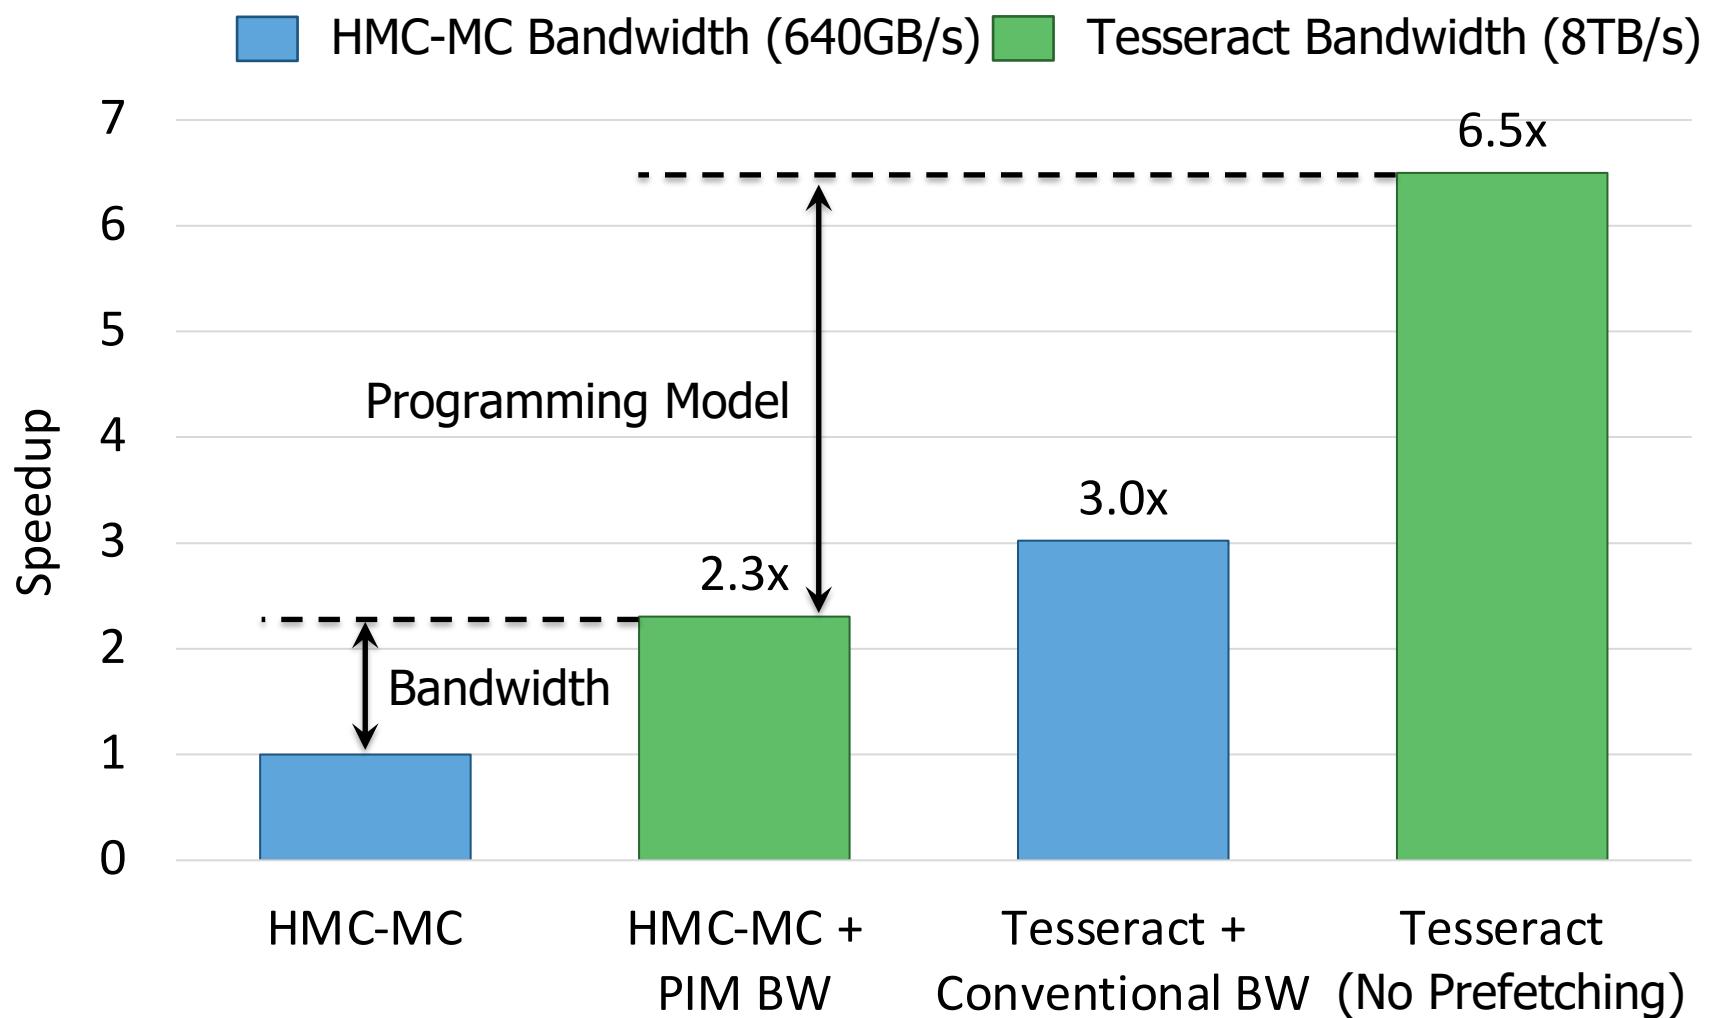
Tesseract Graph Processing Performance



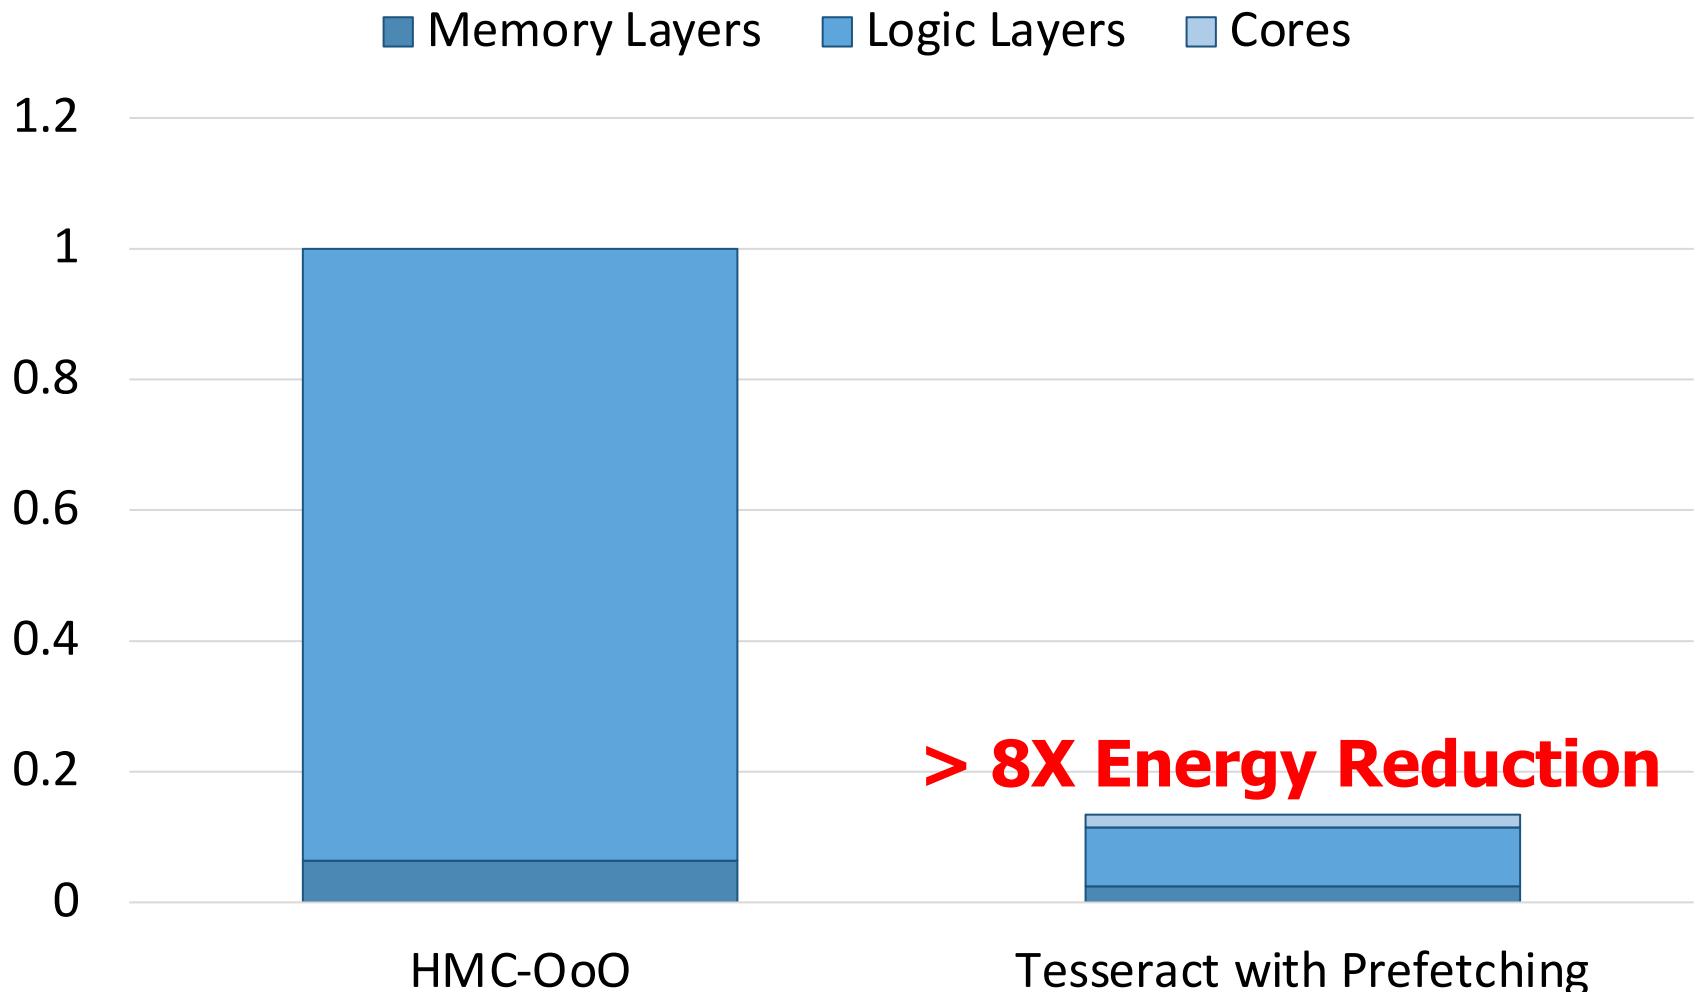
Tesseract Graph Processing Performance



Effect of Bandwidth & Programming Model



Tesseract Graph Processing System Energy



More on Tesseract

- Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyoung Choi,

"A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing"

Proceedings of the 42nd International Symposium on Computer Architecture (ISCA), Portland, OR, June 2015.
[Slides (pdf)] [Lightning Session Slides (pdf)]

A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing

Junwhan Ahn Sungpack Hong[§] Sungjoo Yoo Onur Mutlu[†] Kiyoung Choi

junwhan@snu.ac.kr, sungpack.hong@oracle.com, sungjoo.yoo@gmail.com, onur@cmu.edu, kchoi@snu.ac.kr

Seoul National University

[§]Oracle Labs

[†]Carnegie Mellon University

End of Backup Slides