

# Accelerating Genome Analysis

## A Primer on an Ongoing Journey

Onur Mutlu

[omutlu@gmail.com](mailto:omutlu@gmail.com)

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

January 24, 2018

AACBB Keynote, Vienna

# 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

# Agenda

---


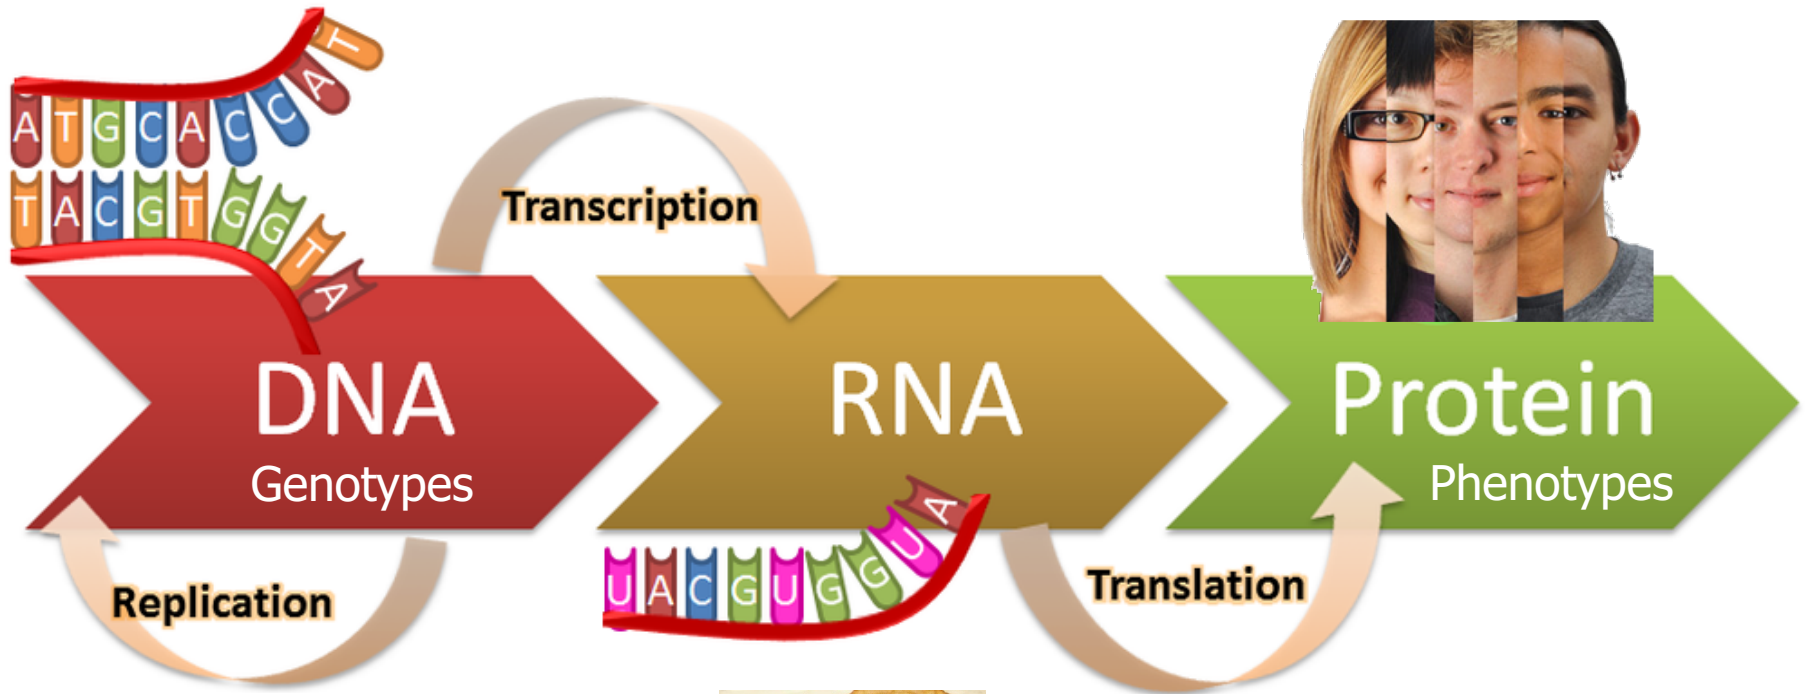
- The Problem: DNA Read Mapping
  - State-of-the-art Read Mapper Design
- Algorithmic Acceleration
  - Exploiting Structure of the Genome
  - Exploiting SIMD Instructions
- Hardware Acceleration
  - Specialized Architectures
  - Processing in Memory
- Future Opportunities: New Sequencing Technologies

# What Is a Genome Made Of?







# The Central Dogma of Molecular Biology



# DNA Under Electron Microscope



human chromosome #12  
from HeLa's cell

# DNA Sequencing

---

- Goal:

- Find the complete sequence of A, C, G, T's in 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)

# Untangling Yarn Balls & DNA Sequencing

---






# Genome Sequencers



Roche/454




AB SOLiD



Illumina MiSeq



Complete Genomics



Illumina HiSeq2000



Pacific Biosciences RS



Oxford Nanopore MinION



Illumina NovaSeq 6000



**SAFARI** Ion Torrent PGM



Ion Torrent Proton

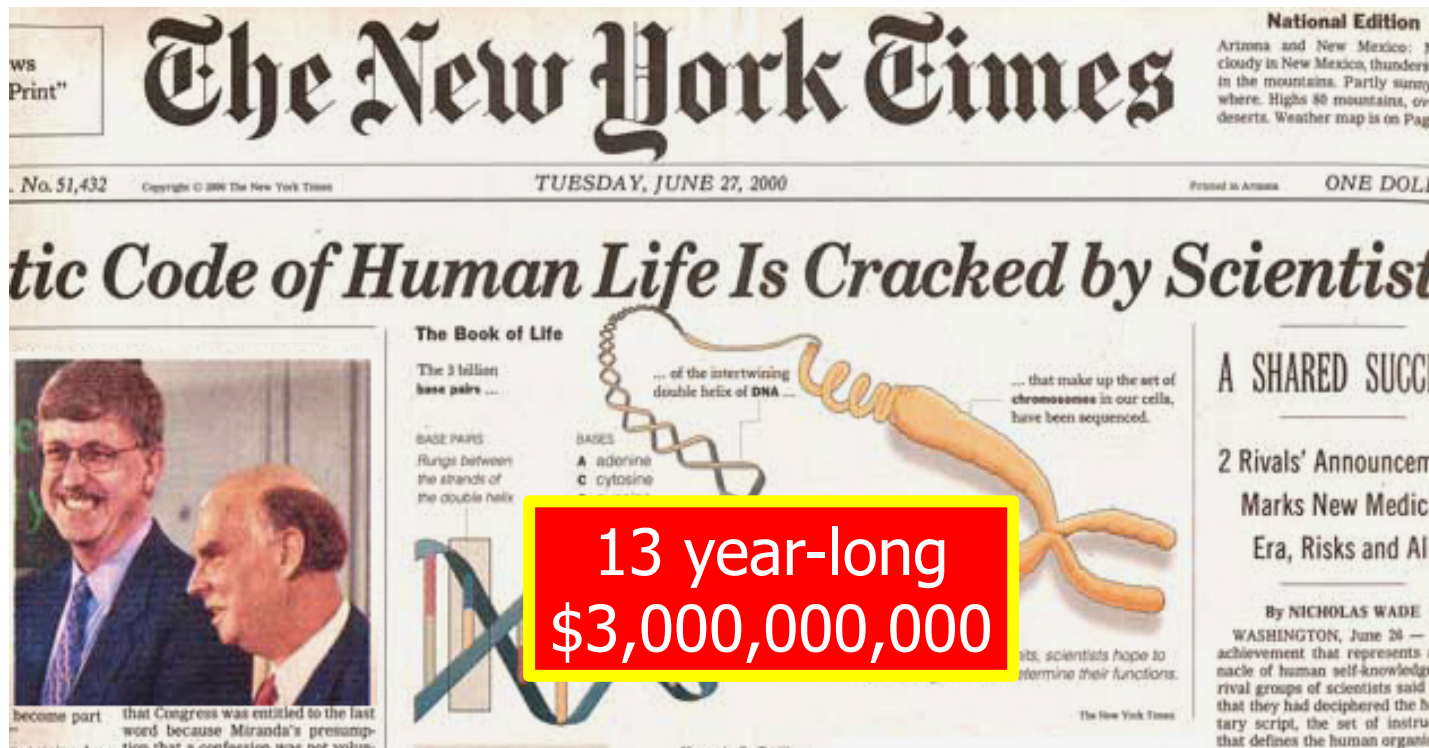


Oxford Nanopore GridION

... 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

# High-Throughput Sequencing (HTS)





# High-Throughput Sequencing (HTS)

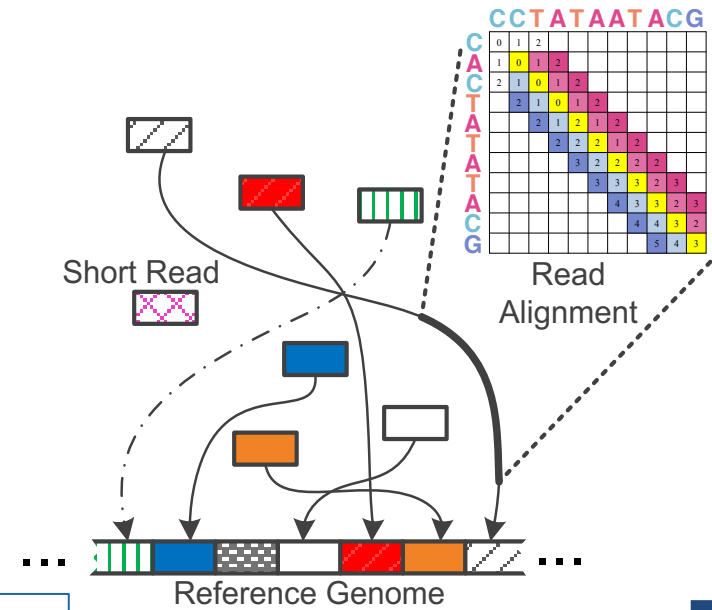



As a workaround, HTS technologies sequence random short DNA fragments (75-300 basepairs long) of copies of the original molecule.

# 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
    - 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



## 1 Sequencing

# Genome Analysis

## 2 Read Mapping

reference: TTTATCGCTTCCATGACGCAG

read1: ATCGCATCC

read2: TATCGCATC

read3: CATCCATGA

read4: CGCTTCCAT


read5: CCATGACGC

read6: TTCCATGAC



## 3 Variant Calling

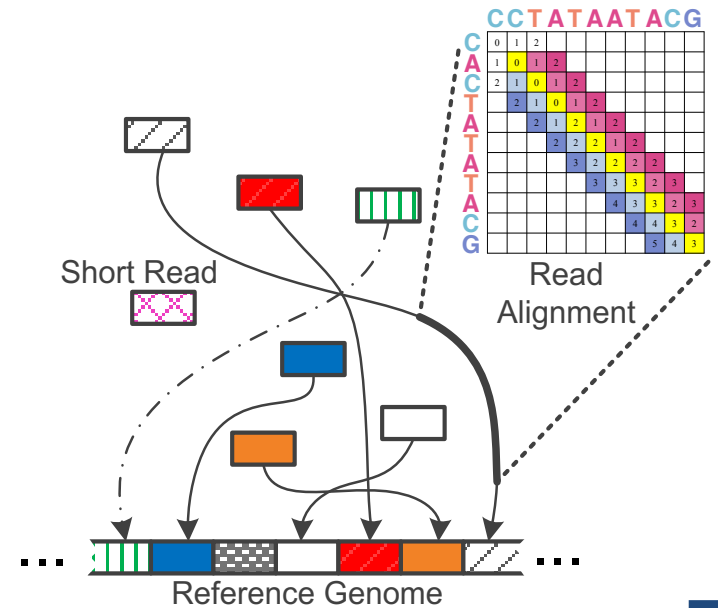
## 4 Scientific Discovery



Billions of Short Reads

ATATATACGTACTAGTACGT  
 TTTAGTACGTACGT  
 ATACGTACTAGTACGT  
 CGCCCCTACGTA  
 ACGTACTAGTACGT  
 TTAGTACGTACGT  
 TACGTACTAAAGTACGT  
 TACGTACTAGTACGT  
 TTTAAACGTA  
 CGTACTAGTACGT  
 GGGAGTACGTACGT

## 1 Sequencing




## Read Mapping 2

**Bottlenecked in Mapping!!**

Illumina HiSeq4000

300 M

bases/min




on average

2 M

bases/min

(0.6%)


# The Read Mapping Bottleneck



Illumina HiSeq4000

ACGTACGTACGTACGT  
CCCCCCTATATATACGTACTAGTACGT  
CGACTTTAGTACGTACGT  
TATATATACGTACTAGTACGT  
ACGTACGCCCCGTACGTA  
TATATATACGTACTAGTACGT  
GACTTTAGTACGTACGT  
TATATATACGTACTAAAGTACGT  
TATATATACGTACTAGTACGT  
CGTTTTTAAACGTA  
TATATACGTACTAGTACGT  
GACGGGGAGTACGTACGT  
TATATATACGTACTAAAGTACGT

300 Million  
bases/minute




2 Million  
bases/minute

150X slower

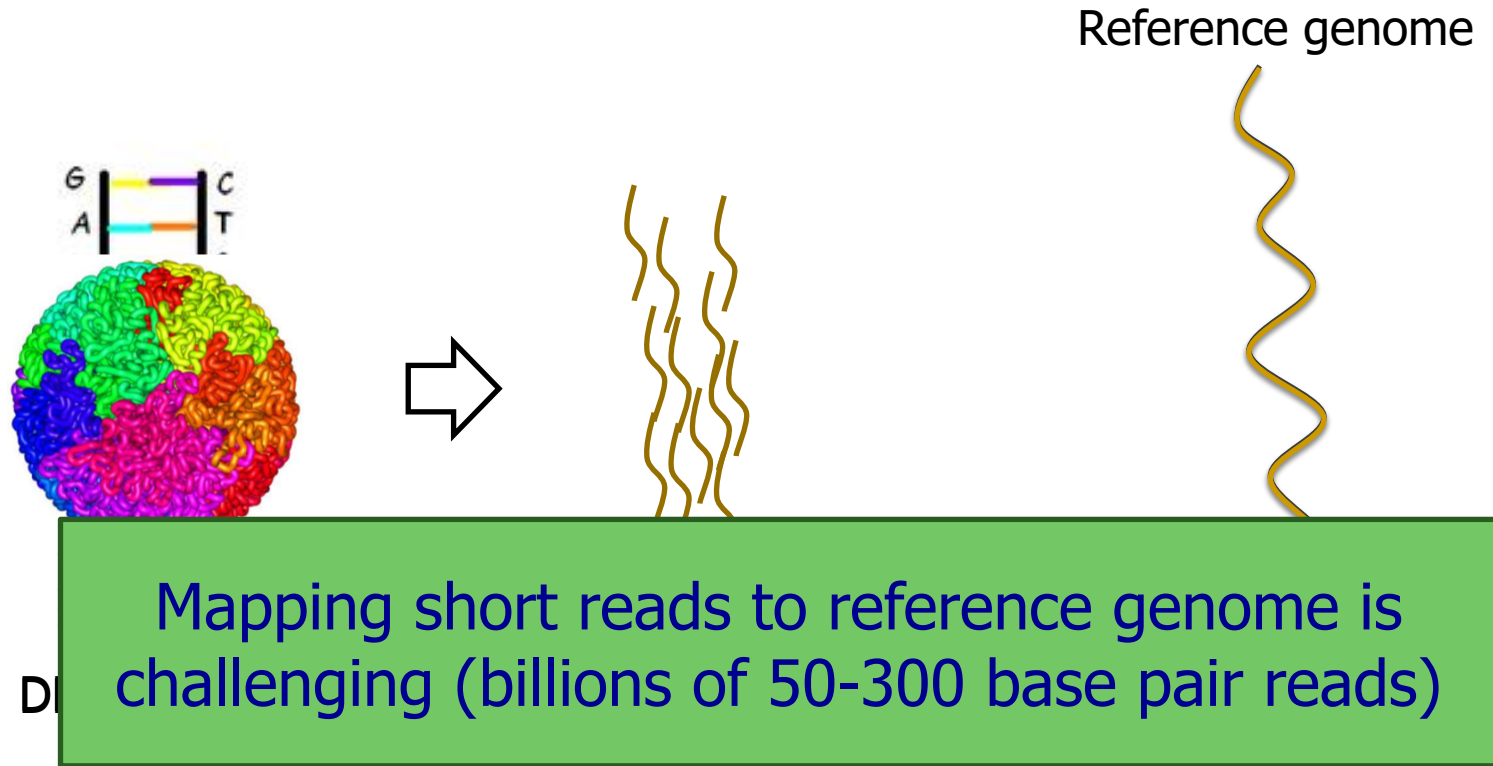
# Read Mapping Execution Time Breakdown

---



# Read Mapping

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



# 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?



# 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.

organization x operation

Ref	o	-	-	r	g	a	n	i	z	a	t	i	o	n
Read	o	p	e	r	-	-	-	-	-	a	t	i	o	n

Ref	o	-	-	r	g	a	n	i	z	a	t	i	o	n
Read	o	p	e	r	-	a	-	-	-	-	t	i	o	n

match
deletion
insertion
mismatch

organization x translation

Ref	o	r	g	a	n	i	z	-	a	t	i	o	n
Read	t	r	-	a	n	-	s	l	a	t	i	o	n

Ref	o	r	g	a	n	-	i	z	a	t	i	o	n
Read	t	r	-	a	n	s	l	-	a	t	i	o	n

Ref	o	r	g	a	n	i	z	a	t	i	o	n
Read	t	r	-	a	n	s	l	a	t	i	o	n

# 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
- Future Opportunities: New Sequencing Technologies

# Read Mapping Algorithms: Two Styles

---

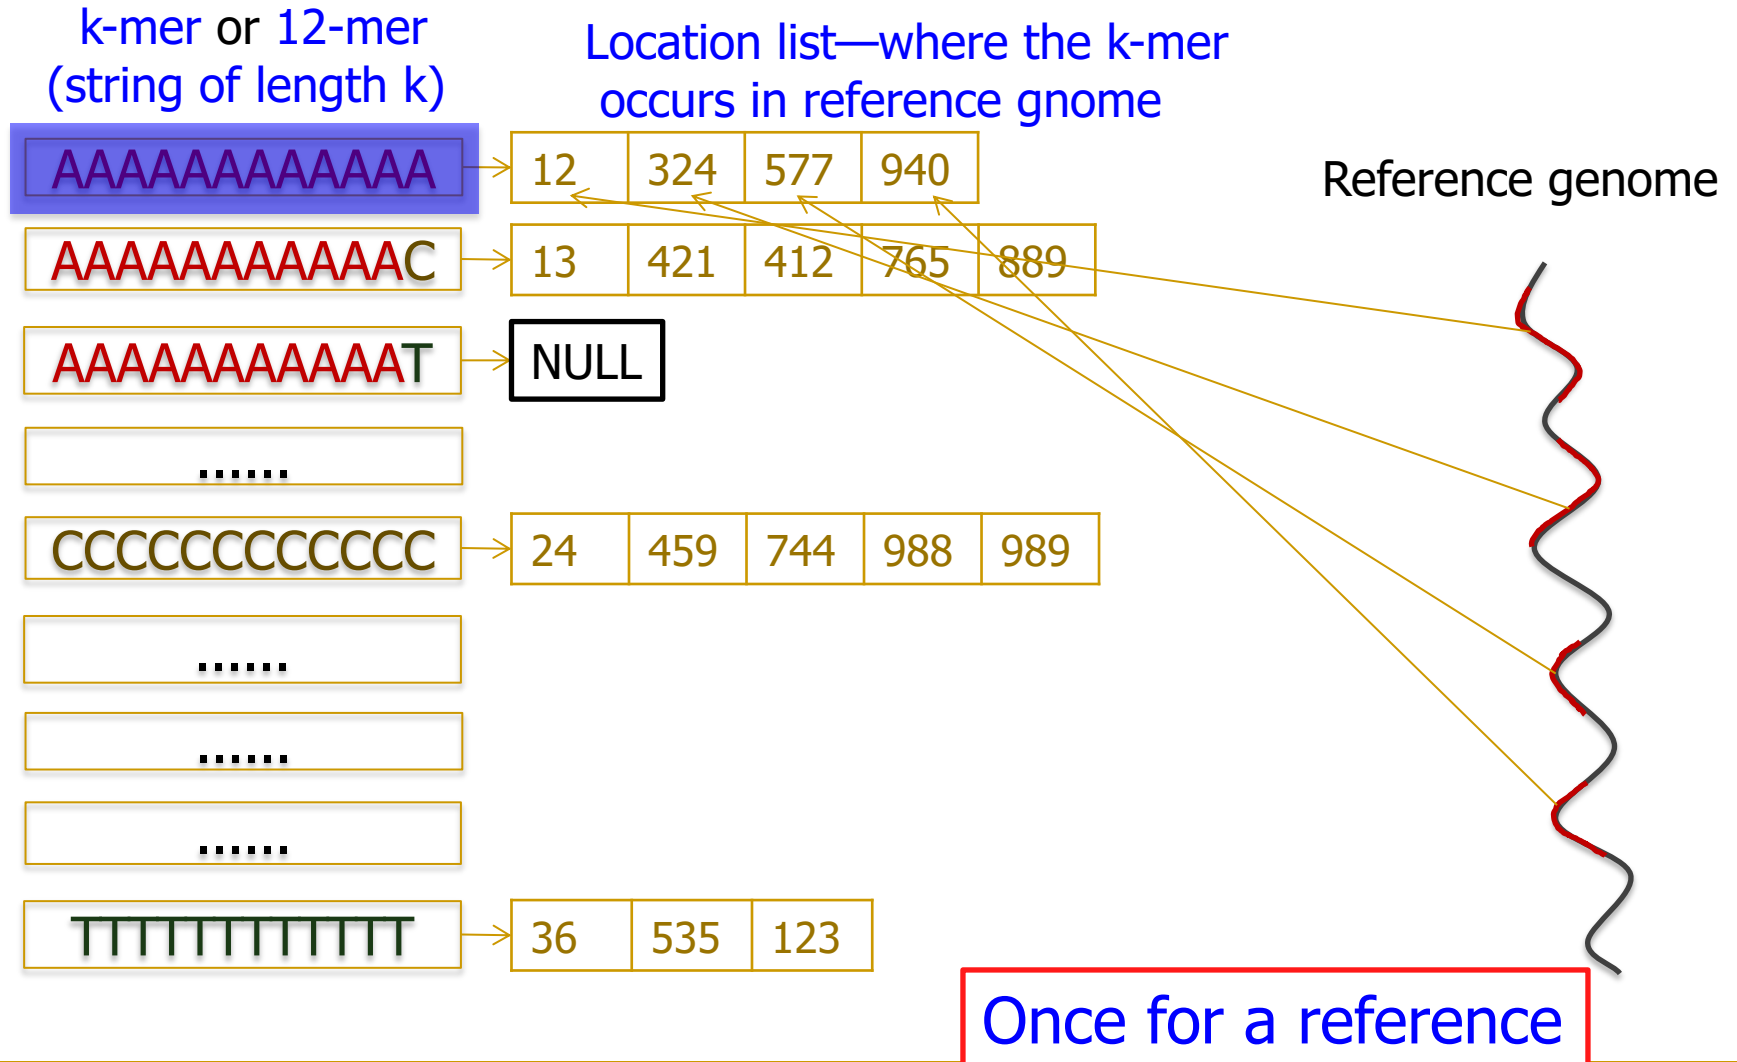
- Hash based seed-and-extend (hash table, suffix array, suffix tree)
  - ❑ Index the “k-mers” in the genome into a hash table (pre-processing)
  - ❑ When searching a read, find the location of a k-mer in the read; then extend through alignment
  - ❑ More sensitive, but slow
  - ❑ Requires large memory; 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]

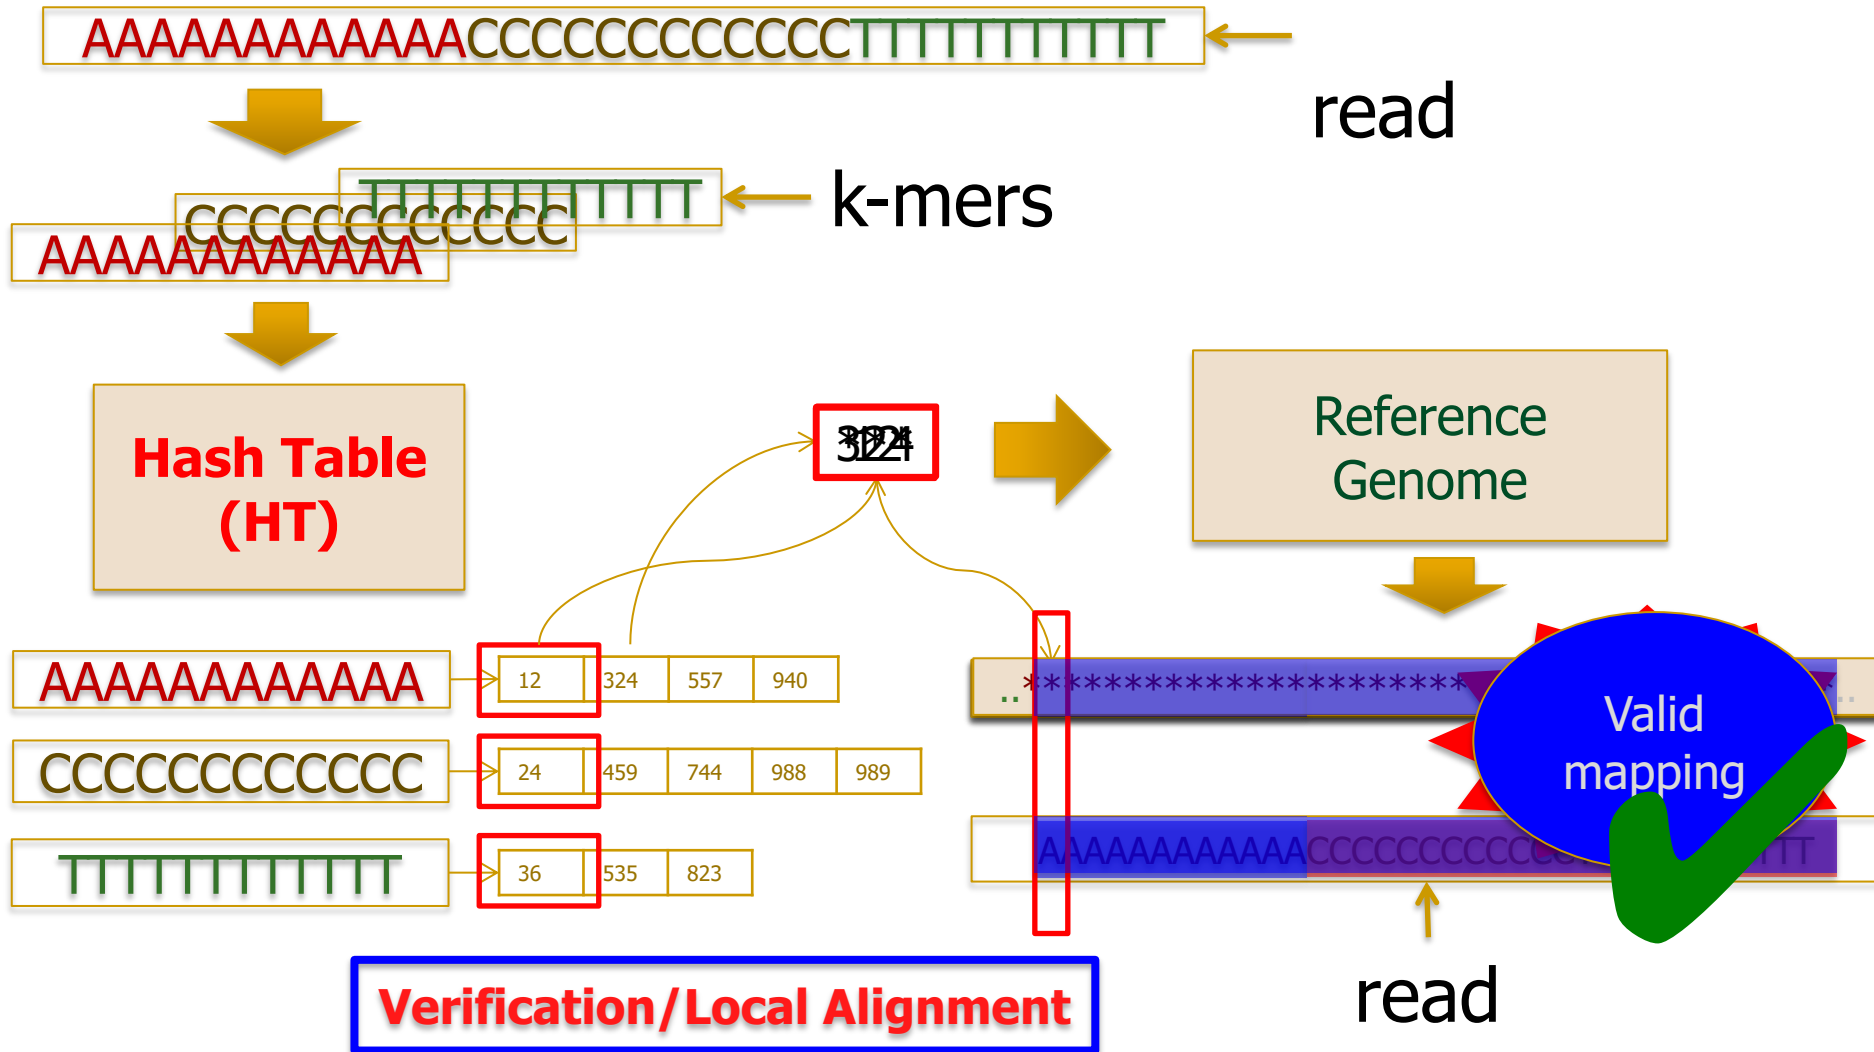


# Hash Table Based Read Mappers

---

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

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



# Advantages of Hash Table Based Mappers

---

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

nature  
genetics

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

---

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

Can Alkan<sup>1,2</sup>, Jeffrey M Kidd<sup>1</sup>, Tomas Marques-Bonet<sup>1,3</sup>, Gozde Aksay<sup>1</sup>, Francesca Antonacci<sup>1</sup>, Fereydoun Hormozdiari<sup>4</sup>, Jacob O Kitzman<sup>1</sup>, Carl Baker<sup>1</sup>, Maika Malig<sup>1</sup>, Onur Mutlu<sup>5</sup>, S Cenk Sahinalp<sup>4</sup>, Richard A Gibbs<sup>6</sup> & Evan E Eichler<sup>1,2</sup>

---


Alkan+, "**Personalized copy number and segmental duplication maps using next-generation sequencing**", Nature Genetics 2009.



# 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

# 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
- Future Opportunities: New Sequencing Technologies

# Reducing the Cost of Verification

---

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

# 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

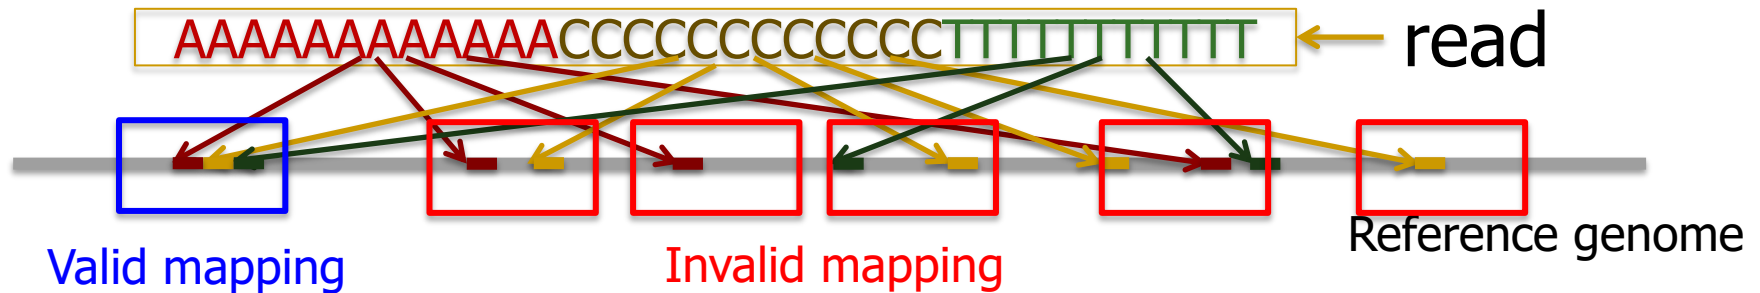
# FastHASH Mechanisms [Xin+, BMC Genomics 2013]

---

- **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

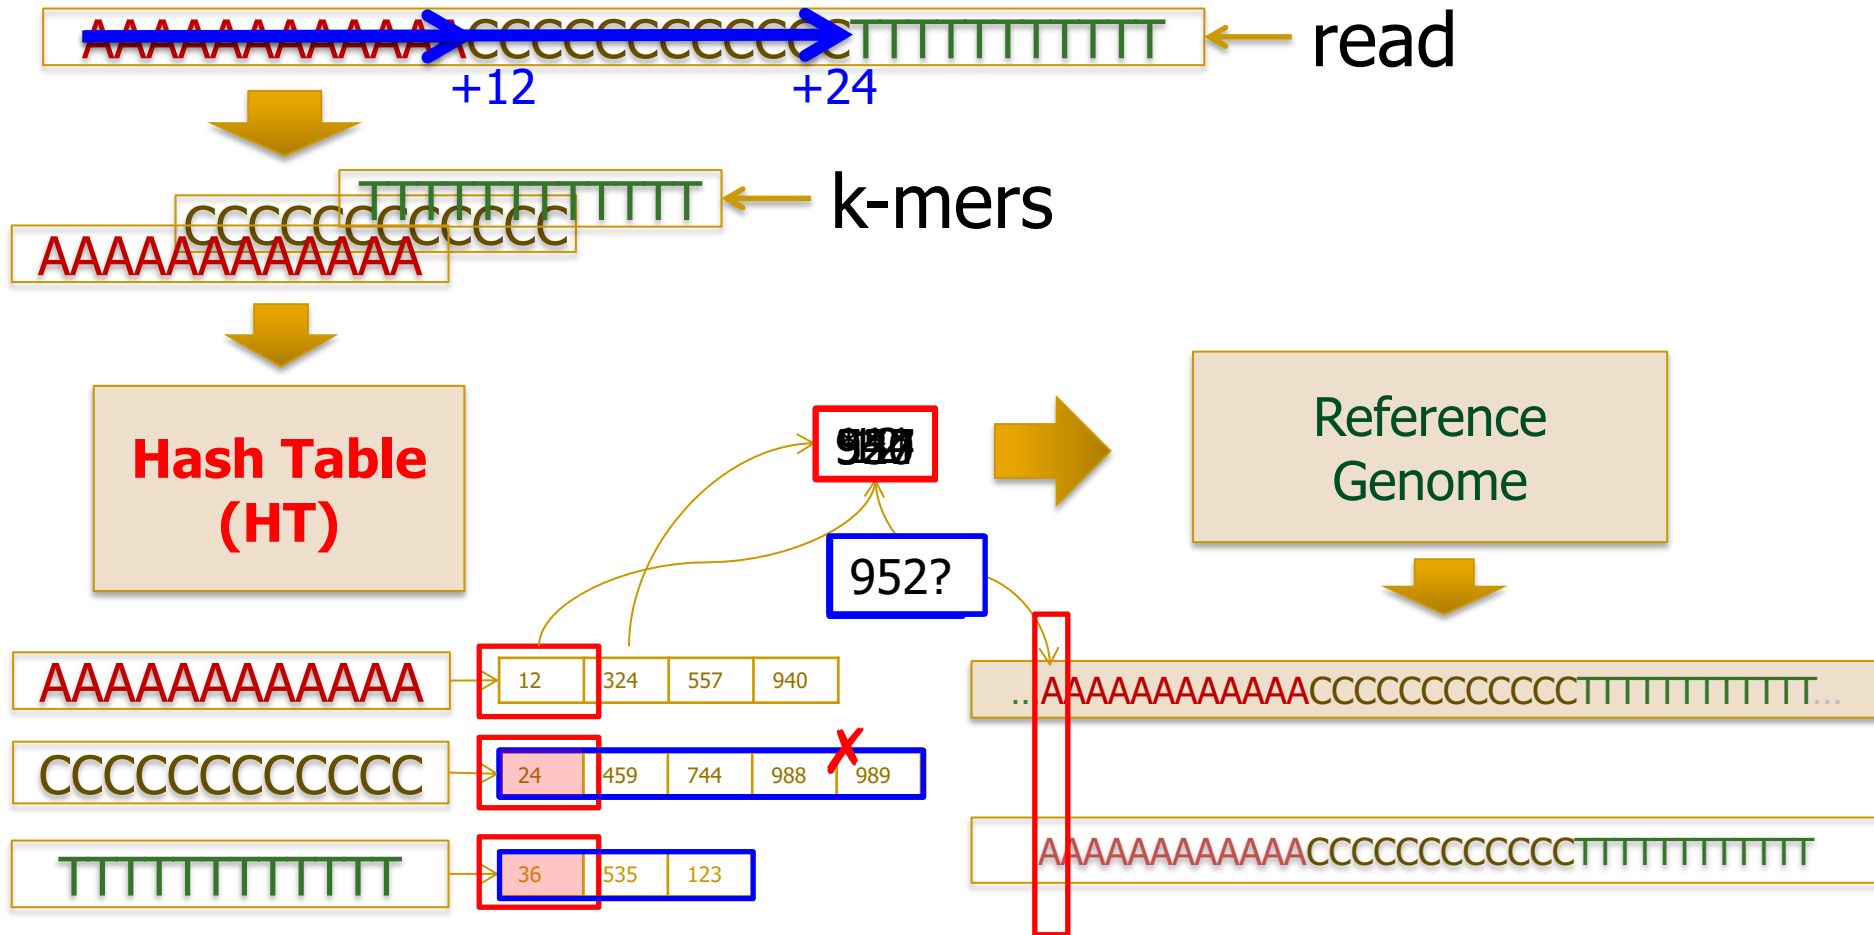
# Adjacency Filtering (AF)

- **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
  - 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



# Cheap K-mer Selection (CKS)

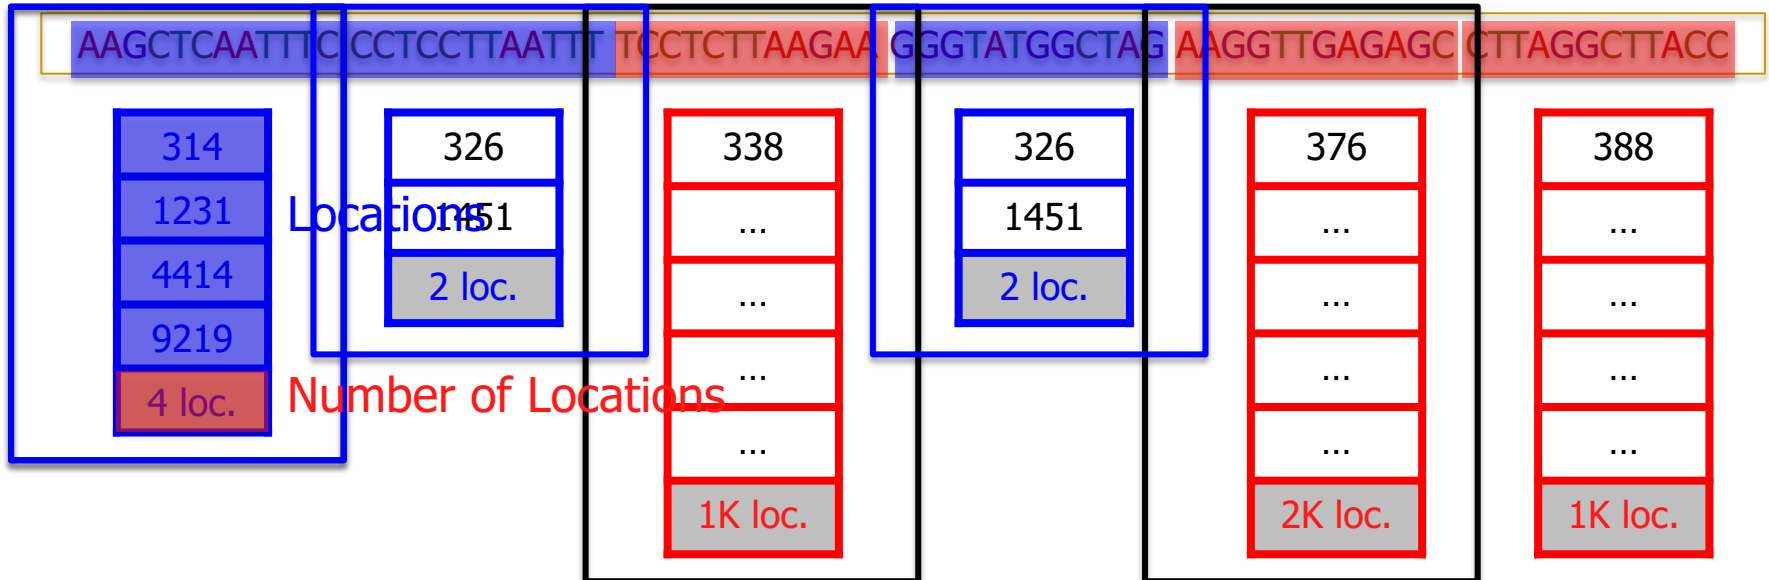
---

- **Goal:** Reduce the number of potential mappings
- **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 fewest locations to verify

# Cheap K-mer Selection

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

read



Expensive 3 k-mers

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

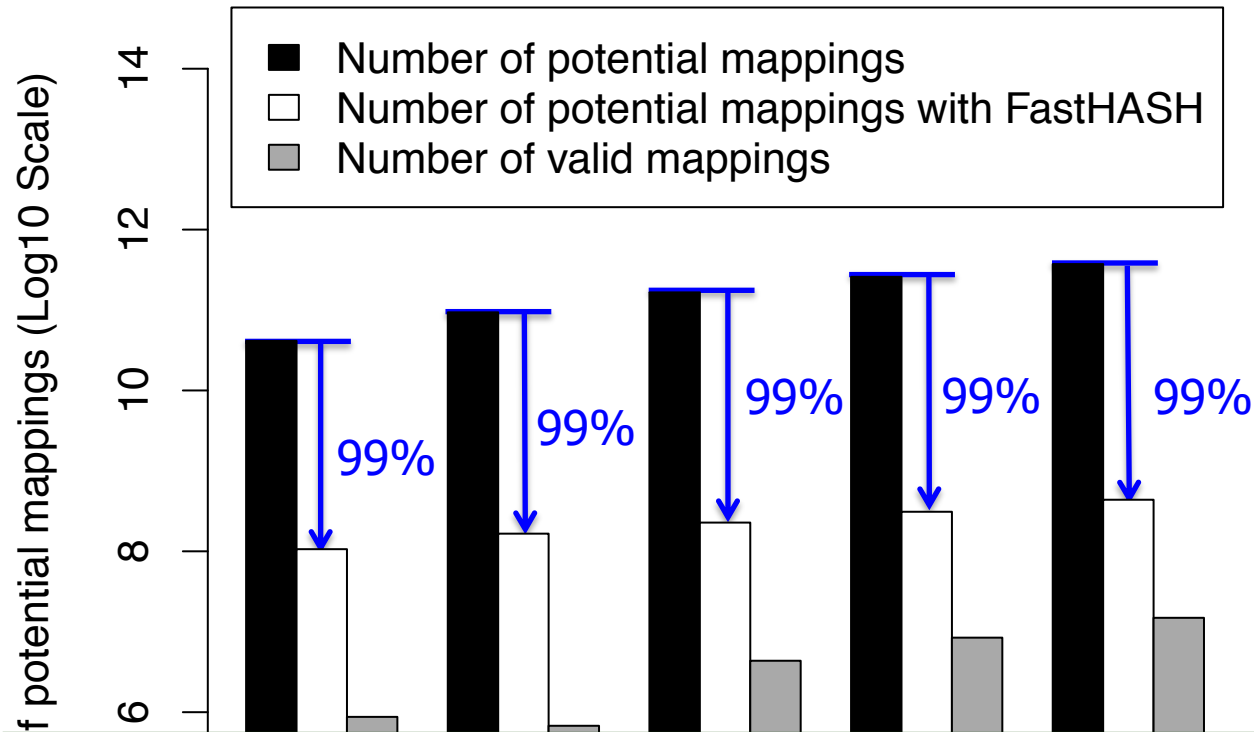
# FastHASH Speedup



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 Conclusion

---

- Problem: Existing read mappers perform poorly in mapping billions of short reads to the reference genome, in the presence of errors
  - Observation: Most of the verification calculations are unnecessary → filter them out
  - Key Idea: To reduce the cost of unnecessary verification
    - 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 Xin<sup>1</sup>, Donghyuk Lee<sup>1</sup>, Farhad Hormozdiari<sup>2</sup>, Samihan Yedkar<sup>1</sup>, Onur Mutlu<sup>1\*</sup>, Can Alkan<sup>3\*</sup>

*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
  - Future Opportunities: New Sequencing Technologies
-



# An Example: 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<sup>1,\*</sup>, John Greth<sup>2</sup>, John Emmons<sup>2</sup>, Gennady Pekhimenko<sup>1</sup>,  
Carl Kingsford<sup>3</sup>, Can Alkan<sup>4,\*</sup> and Onur Mutlu<sup>2,\*</sup>

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.
- Key idea:
  - Compute “Shifted Hamming Distance”: **AND of  $2E$  Hamming Distances of two strings**, to identify 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**

# 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

# 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
- Future Opportunities: New Sequencing Technologies

# Location Filtering

---

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

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

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

# Ideal Filtering Algorithm

---

**Minimal False  
Accept Rate**

**Maximal True  
Reject Rate**

Filter out all  
incorrect mappings

**Zero False  
Reject Rate**

**Faster Than  
Mapper**

Do not filter out any  
correct mappings

# Alignment vs. Pre-alignment (Filtering)

Needleman-Wunsch

C T A T A A T A C G


	0	1	2							
A	1	0	1	2						

GateKeeper

C T A T A A T A C G

A		1	1	0						

- Independent vectors can be processed in parallel using hardware technologies




$$dp[i][j] = 1 + \max \begin{cases} dp[i][j-1] & // \text{Insert.} \\ dp[i-1][j] & // \text{Del.} \\ dp[i-1][j-1] & // \text{Subs.} \end{cases}$$

Each cell depends on three pre-computed cells!

$$dp[i][j] = \begin{cases} 0 & \text{if } X[i] = Y[j] \\ 1 & \text{if } X[i] \neq Y[j] \end{cases}$$

No data dependencies!

# Our Solution: GateKeeper



- 1 High throughput DNA sequencing (HTS) technologies
- 2 Read Pre-Alignment Filtering  
Fast & Low False Positive Rate
- 3 Read Alignment  
Slow & Zero False Positives



# GateKeeper Walkthrough

Generate 2E+1  
masks

Amend random zeros:  
101 → 111 & 1001 → 1111

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

Query :GAGAGAGATATTTAGTGTTGCAGCACTACAACACAAAAGAGGACCAACTTACGTGTCTAAAAGGGGGAACATTGTTGGGCCGGA

Reference :GAGAGAGATAGTTAGTGTTGCAGCCACTACAACACAAAAGAGGACCAACTTACGTGTCTAAAAGGGGAGACATTGTTGGGCCGG

Hamming Mask :0000000000100000000000111111101111000111011010110111111110001000001111011010010101

1-Deletion Mask :111111111111001111101111100011000000000000000

2-Deletion Mask :00000000101101110011111111111111101111000111011010110111111111000100010011101101001010

3-Deletion Mask :111111111110110110011011101101100110010010011111111111110010110011001011011101101111

1-Insertion Mask :11111111111011111011111101101100100100111111111111100101100110010101110111011110

2-Insertion Mask :000000100111110011111111110010001101010100110101111111111111011100111111000111101100

3-Insertion Mask :11111110111011001100111111111101011011111100110010110111111110111011111010111001000

--- Masks after amendment ---

Hamming Mask :000000000010000000000011111111111110001111111101111111111111111000100000111111111111111

1-Deletion Mask :111111111111111111111111100011000000000000000

2-Deletion Mask :000000001111111111111111111111111111000111111111111111111000100011111111111111110

3-Deletion Mask :11111111111111111111111111111111111000111

1-Insertion Mask :1111111111111111111111111111111111110001110

2-Insertion Mask :00000011111111111111111111111111111100011100

3-Insertion Mask :111000

AND Mask :00000000001000000000001000100000000000000

Needleman-Wunsch Alignment :  
GAGAGAGATATTTAGTGTTGCAG-CACTACAACACAAAAGAGGACCAACTTACGTGTCTAAAAGGGGGAACATTGTTGGGCCGG  
||||| : :  
GAGAGAGATAGTTAGTGTTGCAGCCACTACAACACAAAAGAGGACCAACTTACGTGTCTAAAAGGGGAGACATTGTTGGGCCGG

# GateKeeper Walkthrough (cont'd)

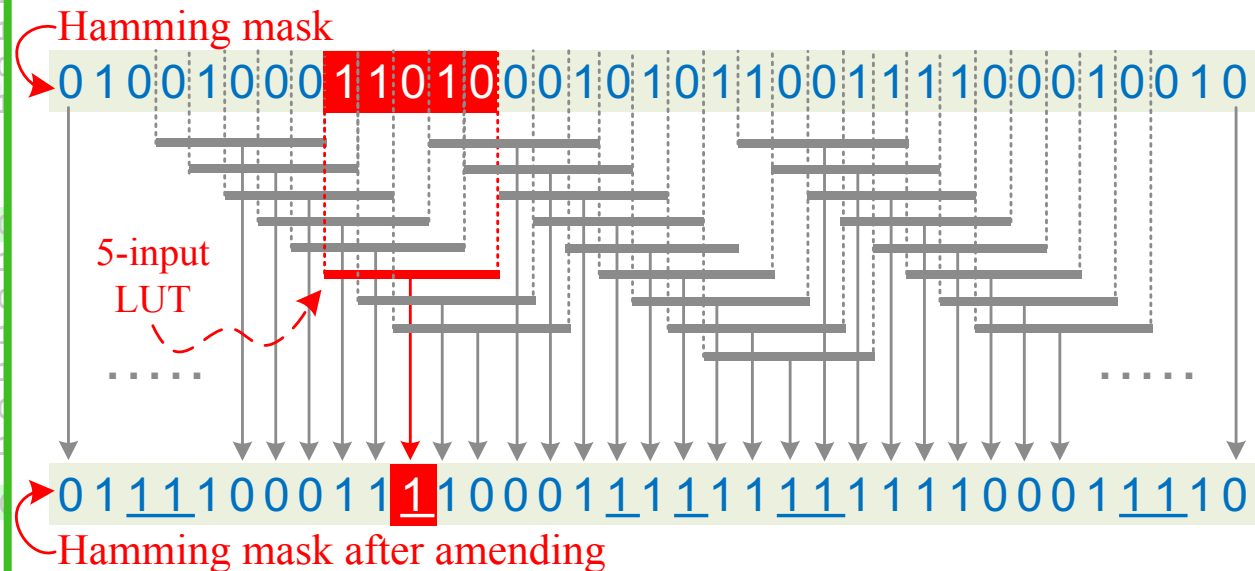
Generate  $2E+1$  masks

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

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

- $E$  right-shift registers (length=ReadLength)
- $E$  left-shift registers (length=ReadLength)
- $(2E+1) * (\text{ReadLength})$  2-XOR operations.

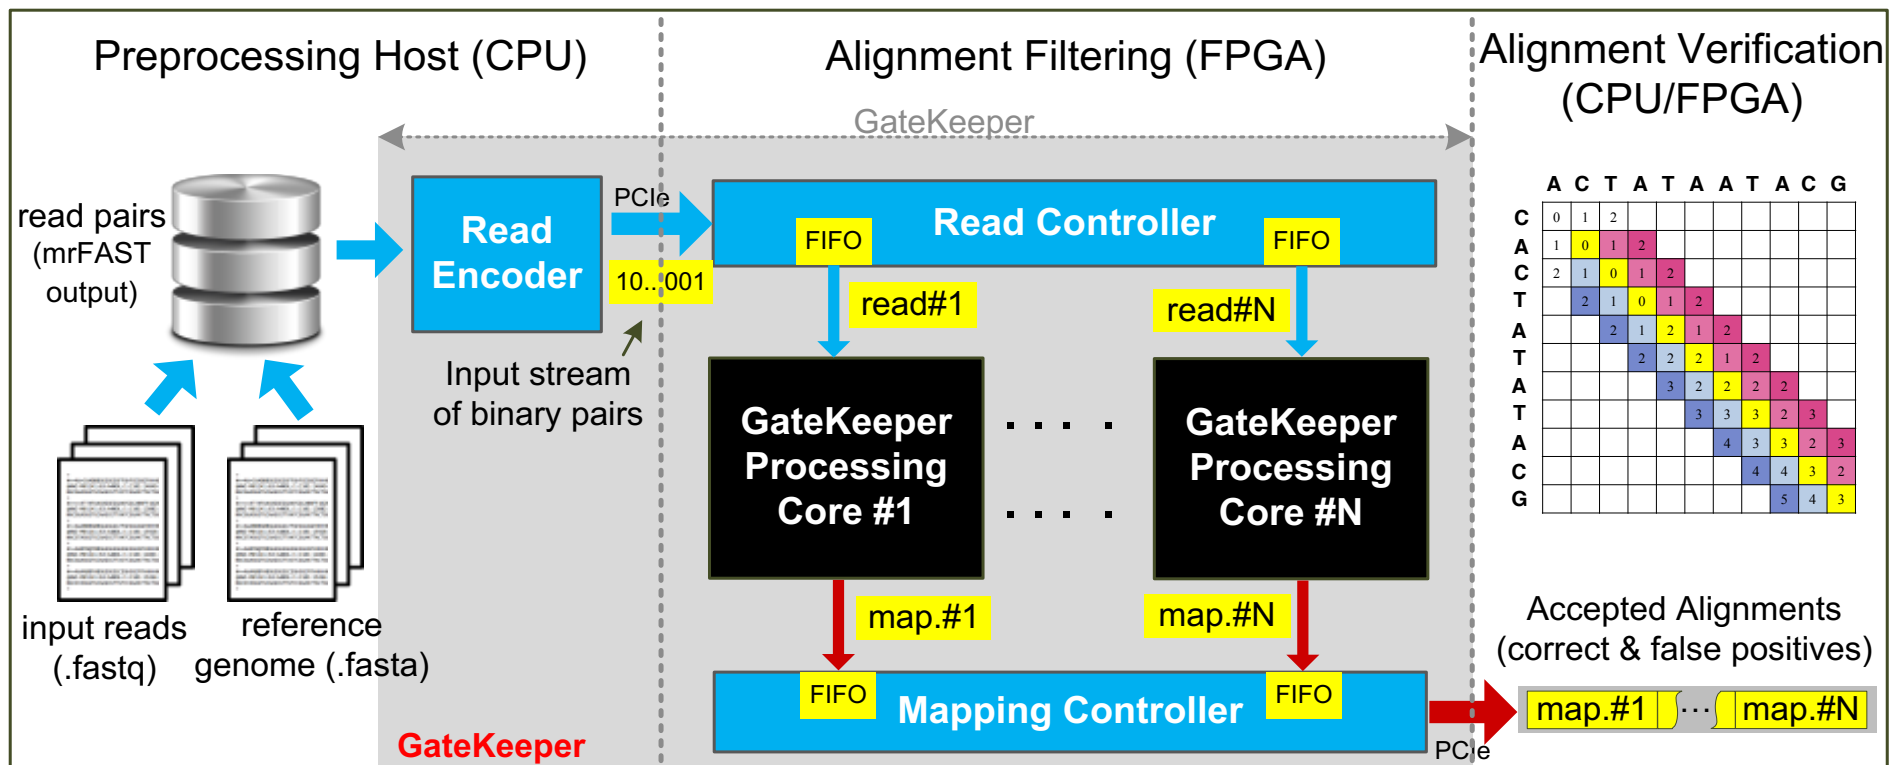
- $(2E) * (\text{ReadLength})$  2-AND operations.
- $(\text{ReadLength}/4)$  5-input LUT.
- $\log_2 \text{ReadLength}$ -bit counter.



- $(2E+1) * (\text{ReadLength})$  5-input LUT.

# 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



# 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 @  $\sim 2$  MHz
- 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](https://github.com/BilkentCompGen/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

# More on GateKeeper

---

- Download and test for yourself  
<https://github.com/BilkentCompGen/GateKeeper>

Alser+, "**GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping**", Bioinformatics, 2017.

---

*Sequence analysis*

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

Mohammed Alser<sup>1,\*</sup>, Hasan Hassan<sup>2</sup>, Hongyi Xin<sup>3</sup>, Oğuz Ergin<sup>2</sup>, Onur Mutlu<sup>4,\*</sup>, and Can Alkan<sup>1,\*</sup>

# Next Talk: MAGNET

- 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.

[illegible]



# 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
- Future Opportunities: New Sequencing Technologies

# Read Mapping & Filtering

---

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

# Hash Tables in Read Mapping

Read Sequence (100 bp)



**Matching...**

**~~Mismatch.~~**


False  
Negative

Hash Table

37    140  
894   1203  
1564

Reference Genome

Filter



# Read Mapping & Filtering in Memory

---

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

# 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

AAAAA	1	<u>AAAAA</u> exists in bin x
AAAAC	0	
AAAAT	1	
...	...	
CCCCC	1	
<b>CCCCT</b>	0	<b><u>CCCCT</u></b> doesn't exist in bin x
CCCCG	0	
...	...	
GGGGG	1	

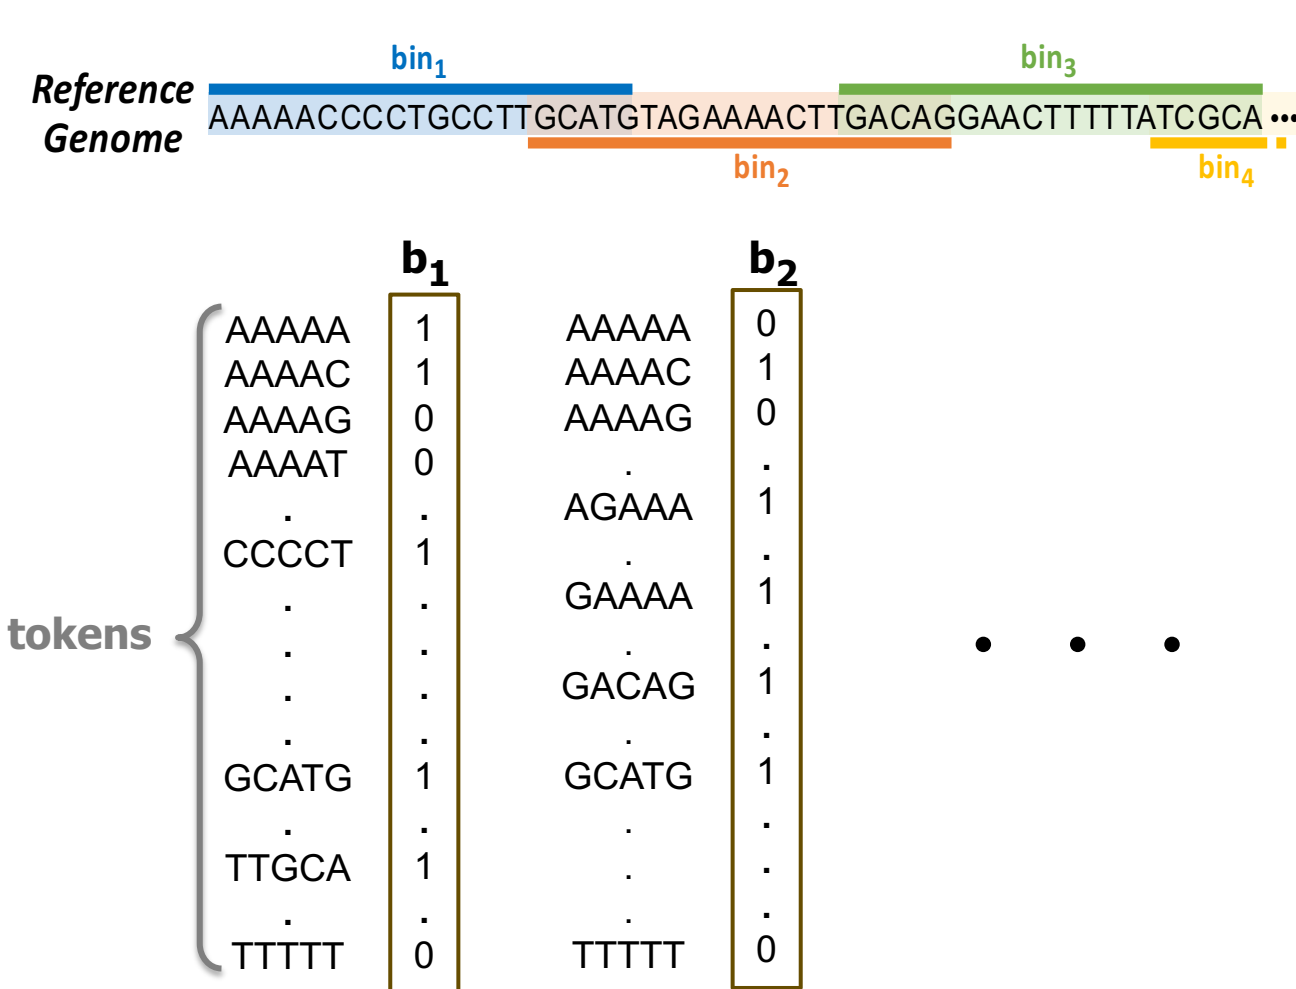
# GRIM-Filter: Bitvectors



Bin x Bitvector

AAAAA	0
...	...
CGTGA	0
...	...
TGAGT	0
...	...
GAGTC	0
...	...
GTGAG	0
...	...

# GRIM-Filter: Bitvectors



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

For **bin size**  $\sim 200$ , and **n** = 5, **memory footprint**  $\sim 3.8$  GB

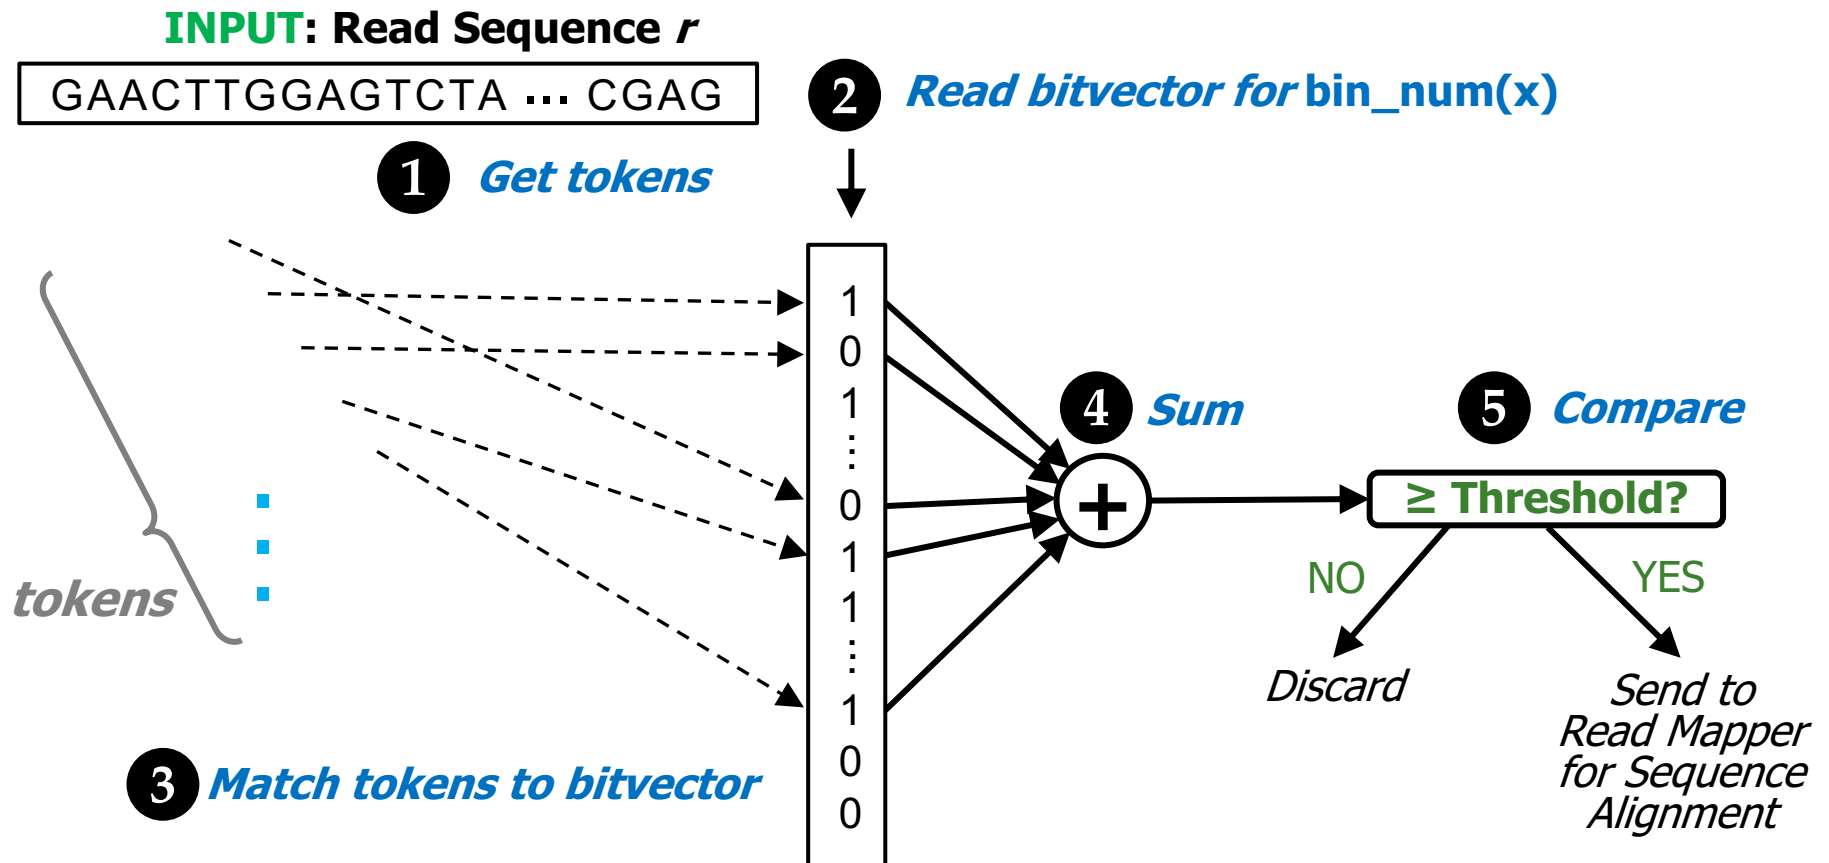


# Our Proposal: GRIM-Filter

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



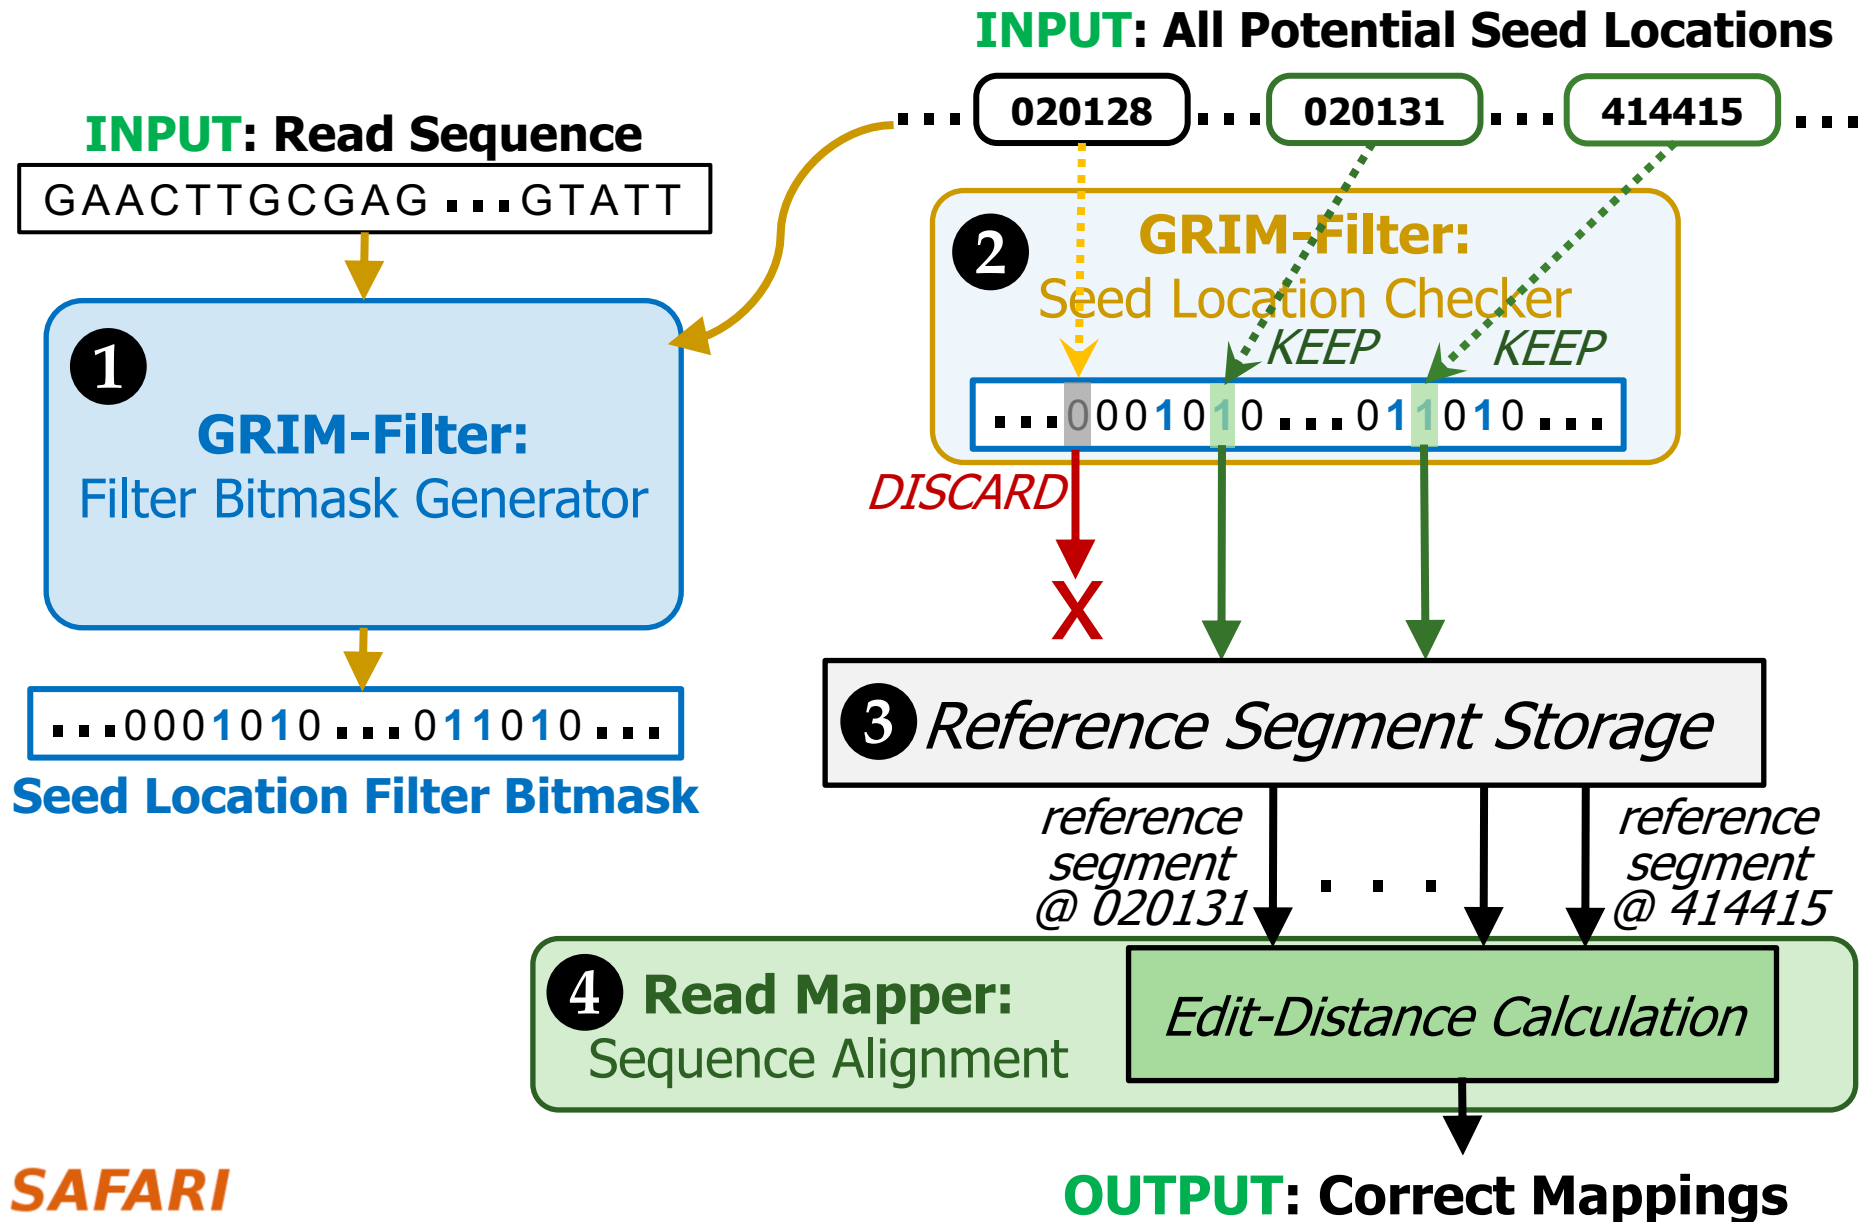
# Our Proposal: GRIM-Filter

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

# 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**


# 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 e computation t
  - Embed GRIM- appropriately






# 3D-Stacked Memory

## Micron's HMC




Micron has working demonstration components

[http://images.anandtech.com/doci/9266/HBMCa\\_678x452.jpg](http://images.anandtech.com/doci/9266/HBMCa_678x452.jpg)



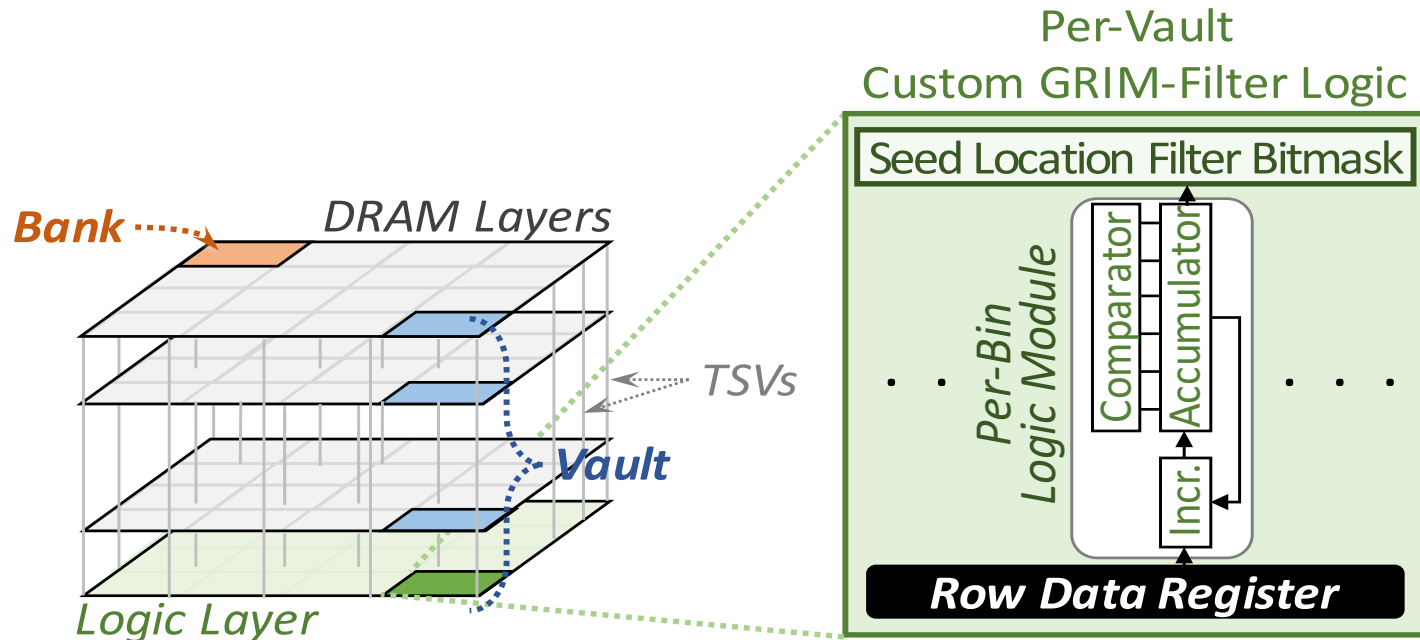
<http://i1-news.softpedia-static.com/images/news2/Micron-and-Samsung-Join-Force-to-Create-Next-Gen-Hybrid-Memory-2.png>

# 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



**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

False Negative Rate



Sequence Alignment  
Error Tolerance ( $e$ )

$e = 0.05$

**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"**  
to appear in ***BMC Genomics***, 2018.  
*Proceedings of the 16th Asia Pacific Bioinformatics Conference (APBC)*,  
Yokohama, Japan, January 2018.  
[arxiv.org Version \(pdf\)](#)

## GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies

Jeremie S. Kim<sup>1,6\*</sup>, Damla Senol Cali<sup>1</sup>, Hongyi Xin<sup>2</sup>, Donghyuk Lee<sup>3</sup>, Saugata Ghose<sup>1</sup>,  
Mohammed Alser<sup>4</sup>, Hasan Hassan<sup>6</sup>, Oguz Ergin<sup>5</sup>, Can Alkan<sup>\*4</sup>, and Onur Mutlu<sup>\*6,1</sup>

# 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
- Future Opportunities: New Sequencing Technologies



# 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
  - No amplification → Less limit on 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

# 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



# 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 **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:

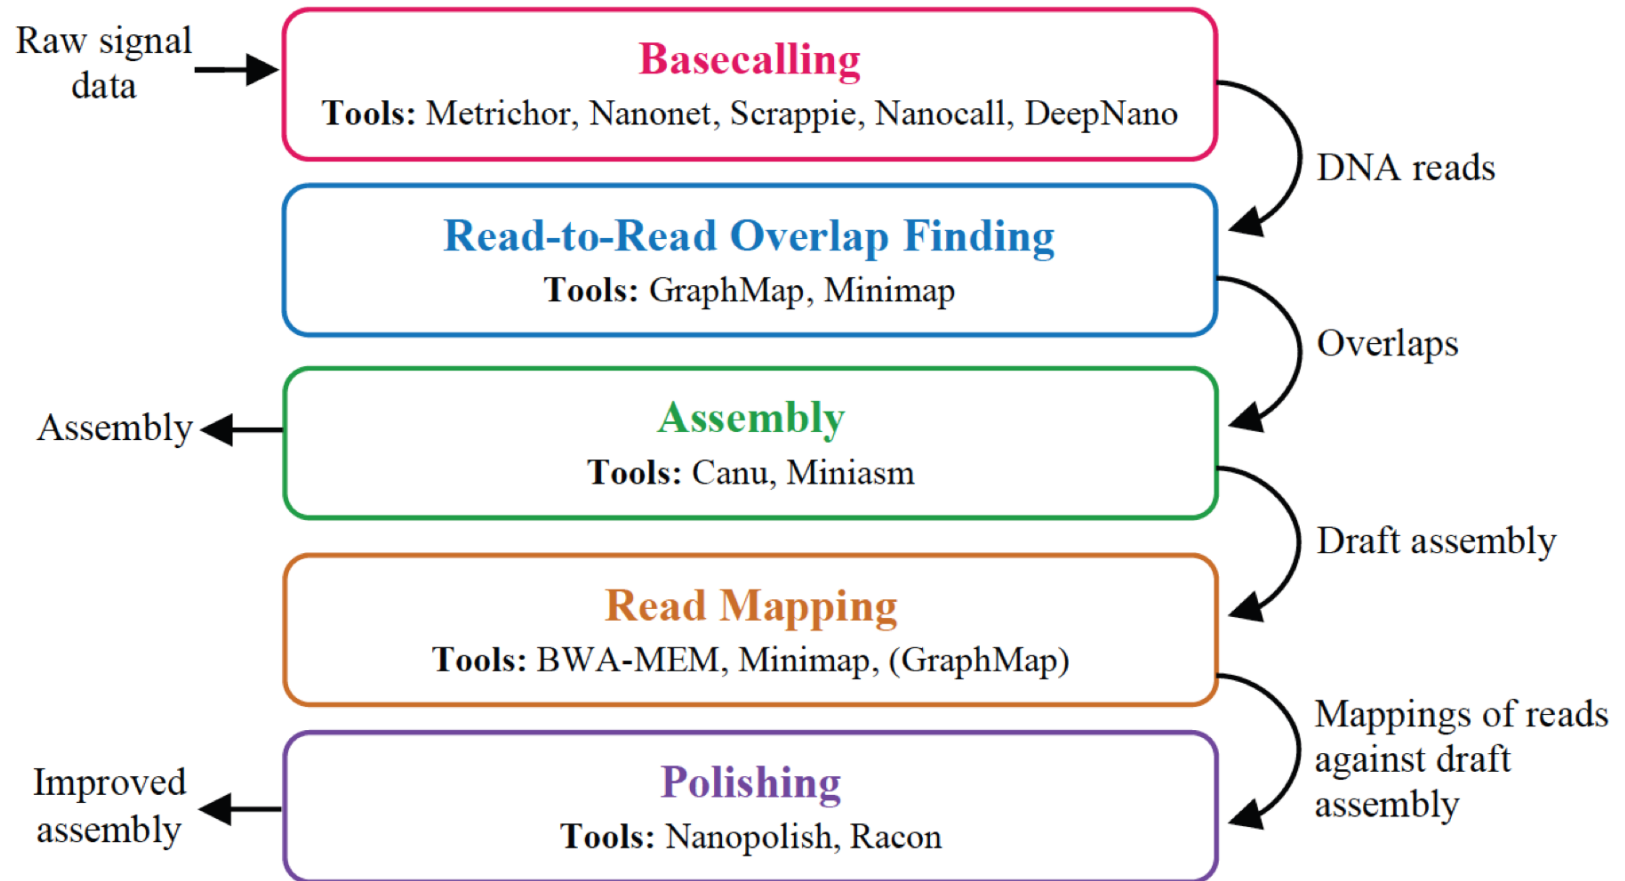
- Do *not* require any labeling of the DNA or nucleotide for detection during sequencing
- Rely on the electronic or chemical structure of the different nucleotides for identification
- Allow sequencing **very long reads**, and
- Provide **portability, low cost, and high throughput**.

# 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 MinION
  - 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.**

# 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<sup>1,\*</sup>, Jeremie Kim<sup>1,3</sup>, Saugata Ghose<sup>1</sup>, Can Alkan<sup>2\*</sup>  
and Onur Mutlu<sup>3,1\*</sup>**

<sup>1</sup>Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA

<sup>2</sup>Department of Computer Engineering, Bilkent University, Bilkent, Ankara, Turkey

<sup>3</sup>Department of Computer Science, Systems Group, ETH Zürich, Zürich, Switzerland

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

[[Preliminary arxiv.org version](#)]



# 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
- Future Opportunities: New Sequencing Technologies

# 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**

# Acknowledgments

---

- Prof. Can Alkan, Bilkent University
- Many students
  - Mohammed Alser, Damla Senol Cali, Jeremie Kim
  - Hasan Hassan
  - Hongyi Xin
  - ...

# Accelerating Genome Analysis

## A Primer on an Ongoing Journey

Onur Mutlu


[omutlu@gmail.com](mailto:omutlu@gmail.com)

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

January 24, 2018

AACBB Keynote, Vienna

# High-Throughput Sequencing



# Nanopore Sequencing

---

- **Basecalling** translates the raw signal output of the nanopore sequencer into bases (A, C, G, T) to generate DNA reads.
  - 1) The raw current signal is divided into discrete blocks (events).
  - 2) Each event is decoded into a most-likely set of bases.
- **Deletions** are the dominant error of nanopore sequencing.
  - In the ideal case, each consecutive event should differ by one base. However, in practice, this is not the case because of the **non-stable speed of the translocation**.
  - Determining the correct length of the **homopolymers** (*i.e.*, repeating stretches of one kind of base, *e.g.*, AAAAAAA) is challenging.

---

# The Importance of Genome Analysis?

Helps, for example, to answer the following 3 questions:

69–92% of the respondents in these studies had positive attitudes towards genomics research and donating their DNA samples.

**Public involvement in pharmacogenomics research: a national survey on public attitudes towards pharmacogenomics research and the willingness to donate DNA samples to a DNA bank in Japan**

Eriko Kobayashi · Nobunori Satoh

**Attitudes and perceptions of patients towards methods of establishing a DNA biobank**

Pulley · Margaret M. Brace · Gordon R. Bernard · Masys

**Genetic research participation in a young adult community sample**

Carla L. Storr · Flora Or · William W. Eaton · Nicholas Jalongo

30 May 2007 / Accepted: 3 July 2007 / Published online: 25 October 2007  
© Science+Business Media B.V. 2007

Received: 2005-11-18 / Accepted: 2005-11-18 / Published: 2005-11-18  
© The Author 2005. Published by Oxford University Press on behalf of the European Public Health Association  
doi:10.1093/eurpub/cki198 Advance Access published on October 5, 2005

© 2008 Wiley-Liss, Inc. American Journal of Medical Genetics Part A 146A:1696–1706 (2008)

**Miscellaneous**

**Genetic research and donation of tissue samples to biobanks. What do potential sample donors in the Swedish general public think?**

Åsa Kettis-Lindblad<sup>1</sup>, Lena Ring<sup>1,2</sup>, Eva Viberth<sup>1</sup>, Mats G. Hansson<sup>3</sup>

**Relationship Between Public Attitudes Toward Genomic Studies Related to Medicine and Their Level of Genomic Literacy in Japan**

Izumi Ishiyama,<sup>1</sup> Akiko Nagai,<sup>1</sup> Kaori Muto,<sup>2</sup> Akiko Tamakoshi,<sup>3</sup> Minori Kokado,<sup>4</sup> Kyoko Mimura,<sup>5</sup> Tetsuro Tanzawa,<sup>6</sup> and Zentarō Yamagata<sup>1\*</sup>




# Pairwise sequence alignment

## Distribution of 116 Blast Hits on the Query Sequence

BLAST

A57075 tensin - chicken (fragment) gi|63805|emb|CAA79215.1| (..S= 492 E=1e-137



Question #1: If I give you a gene sequence, tell me which of the billions of known sequences is most similar to it.

# CODIS: Combined DNA Index System

---

- FBI's program of support for criminal justice.
- CODIS defines 13 human DNA regions (loci) to be stored in the database for personal identification purposes.
- Stored 14.5 million DNA profiles (for offenders, arrestees ..)
- As of September 2016, CODIS has produced over 346,880 hits assisting in more than 332,776 investigations.



<https://www.fbi.gov/services/laboratory/biometric-analysis/codis/ndis-statistics>

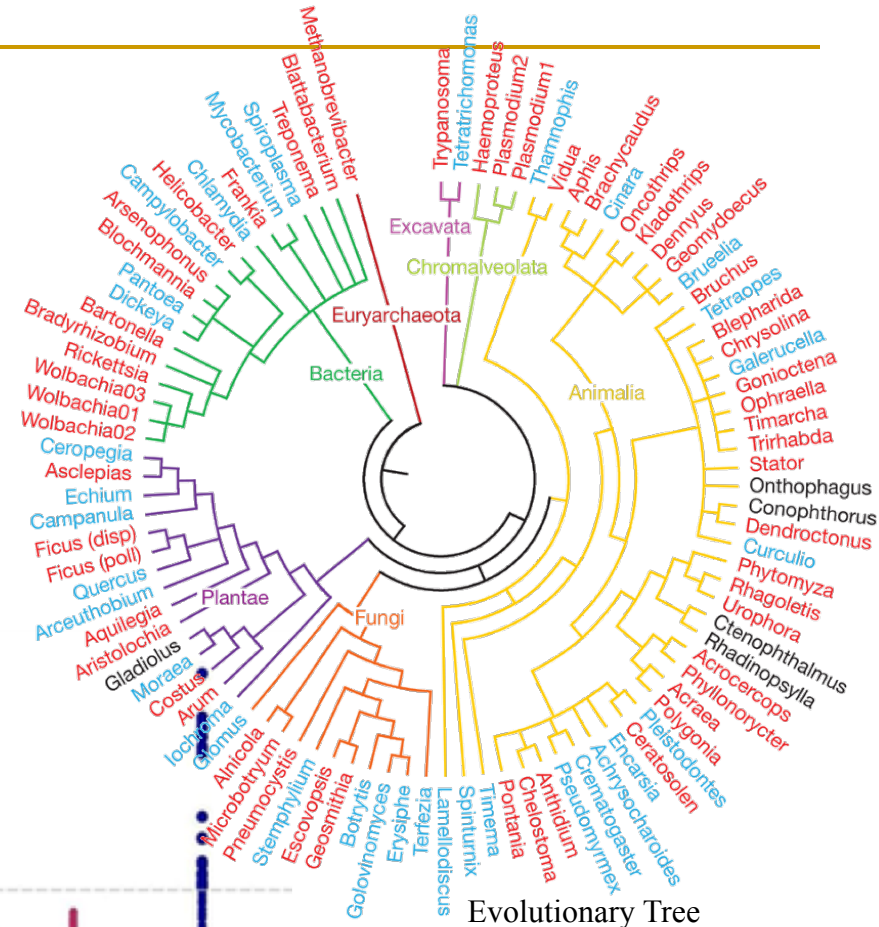

# Multiple sequence alignment

PHDHtm			-----MMMMMMMMMMMMMMMMMMMM-----	
16082665	<i>T acid</i>	10	----MASDRKSEGFQSGAGLIRYFEEEEIKGPALDPKLVVYMGIAVAIIVEIAKIFWFP---	(55)
13541150	<i>T volc</i>	10	----MASDRKSEGFQSGAGLIRYFEEEEIKGPALDPKLVVYIGIAVAIMVELAKIFWFP---	(55)
RFAC01077	<i>F acid</i>	13	-MTSMAKDNQNFQSGAGLIRYFNEEEIKGPAIDPKLIITYIGIAMGVIVELAKVFWFPV---	(58)
15791336	<i>H NRC1</i>	10	----MSSGQNSGGLMSSSAGLVRYFDSEDSNALQIDPRSVVAVGAFFGLVLVLLAQFFA----	(53)
RAG22196	<i>A fulg</i>	14	MAKAPKKGAKTPPLMSSSAGIMRYFEE-EKTQIKVSPKTI LAAGIVTGVLI IILNAYYGLWP-	(68)
RPO01000	<i>P abys</i>	9	-----MAKEKTTL PPTGAGLMRFFDE-DTRAIKITPKGAVALTLILIIIFEIILHVVGPRIFG	(56)
RPH01741	<i>P hori</i>	9	-----MAKEKTTL PPTGAGLMRFFDE-DTRAIKITPKGAIALVLILIIIFEIILHVVGPRIFG	(56)
AE000914	<i>M ther</i>	10	----MAKKDKKTLPPSGAGLVRYFEE-ETKGEKLTPEQVVVMSIILAVFCLVLRFSG----	(52)
RMJ09857	<i>M jann</i>	9	-----MSKREESTGLATSAGLIRYMDE-TFSKIRVKPEHVIGVTVAFVIIIEAILTYGRFL---	(53)
15920503	<i>S toko</i>	13	-MPSSKKKKSTVPLASMAGLIRYYEE-ENEKIKISP KLLIIISIMVAGVIVASILIPPP--	(58)
AE006662	<i>S solf</i>	11	-MPSSKKKKSTVPM SMAGLIRYYEE-ENEKVKISP KIVIGASLALTIIIVIVITKLF-----	(55)
RPK02491	<i>P aero</i>	12	--MARRRKYEGLNPFVAAGLIKFS EEGELEKIKLTPRAAVVISLAIIGLLIAINLLLPPL--	(58)
RAP00437	<i>A pern</i>	13	-MSVRRRRRERRATPVTAAGLLSFYEE-YEGKIKISPTIVVGAAILVSAVVAABHIFLPAVP-	(59)
5803165	<i>H sapi</i>	49	-----SAGTGGMWRFYTE-DSPGLKVGVPVFLVMSLLFIASVFMLHIWGTKYTRS	(96)
13324684	<i>M musc</i>	49	-----SAGTGGMWRFYTE-DSPGLKVGVPVFLVMSLLFIAAVFMLHIWGTKYTRS	(96)
6002114	<i>D mela</i>	53	-----GAGTGGMWRFYTD-DSPGIRKVGVPVFLVMSLLFIASVFMLHIWGTKYNRS	(100)
14574310	<i>C eleg</i>	32	-----GGNNGGLWRFYTE-DSTGLKIGVPVFLVMSLVFIASVFVLHIWGTKFTRS	(81)
10697176	<i>Y lipo</i>	41	-----GGSSSTMLKLYTD-ESQGLKVDPVVVMVLSLGFIFSVVALHLLAKVSTK	(91)
6320857	<i>S cere</i>	40	-----GGSSSSILKLYTD-EANGFRVDSLVLFLSVGFIFSVIALHLLTKFTHI	(88)
6320932	<i>S cere</i>	33	-----TNSNNSILKIYSD-EATGLRVDPLVLFLAVGFIFSVVALHVISKVAGK	(82)

Question #2: If I give you a bunch of sequences, tell me where they are the same and where they are different.

# Phylogenetic tree

- Reveals the genomic variants that cause diseases.
- Helps understanding the evolutionary relationships among various species.



[http://www.nature.com/nature/journal/v465/n7300/fig\\_tab/nature09113\\_F3.html](http://www.nature.com/nature/journal/v465/n7300/fig_tab/nature09113_F3.html)



# The genetic similarity between species

---



Human ~ Human  
99.9%



Human ~ Chimpanzee  
96%



Human ~ Cat  
90%





Human ~ Cow  
80%



Human ~ Banana  
50-60%

**Question #3: Given a bunch of short sequences, Can you identify the approximate species cluster for genomically unknown organisms (bacteria)?**



uncleaned de Bruijn graph


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

---

# ANALYZING THE PROBLEM

# Read Mapping

## 1 High Throughput Sequencing



## 2 Alignment Verification

e.g. Smith-Waterman alignment algorithm

	0	A	C	G	T	A	T	G	C
0	0	0	0	0	0	0	0	0	0
A	0	2	0	0	0	2	0	0	0
C	0	0	4	2	1	0	1	0	2
G	0	0	2	6	4	3	2	3	1
A	0	2	1	4	5	6	4	3	2
A	0	2	1	3	3	7	5	4	3
C	0	2	4	2	2	5	6	4	6
C	0	0	2	3	1	4	4	5	6

- We want ultra fast and accurate alignment.
- Detection of genomic variation.

## 3 Read Mapping



Reference Genome  
(>3 gigabases)




# Key Observations:

---

- Alignment Verification → **90%** of mapper's execution time.
- **>98%** of candidate locations have high dissimilarity with a given read.

Cheng *et al*, *BMC bioinformatics* (2015)  
Xin *et al*, *BMC genomics* (2013)

# Read Mappers Timeline



# Filters, Alignment, and Mappers

Year	Purpose	Architecture	Platform	alignments#/1sec
2015	Filter	Shifted Hamming Distance	Intel SSE	<b>3x faster</b> 3583
	Alignment	Myers's bit-vector [45]	Intel SSE	409
	Alignment	Smith-Waterman [40]	Intel SSE	38
	Mapper			16
2014	M			13
	A			1
				60
				17
2013				4
				15
	Mapper	BWT-FM	FPGA(Virtex6)	1092
	Mapper	BWT-FM	GPU	17
2012	Mapper	Hash-Based (BFAST)	FPGA(Virtex6)	35
			FPGA(Virtex4)	
	Alignment	Smith-Waterman	GPU	131
			Cell BE	16
			CPU	41

ideal filter → fast & accurate to  
compensate the computation  
overhead

