

# Fast SIMD-Based Chunking Algorithm

Yehonatan Dude, Michael Hirsch, Yair Toaff



## PSC 2019

# *Outline*

1. Background
2. Chunking Problem
3. Traditional Solutions
4. Our Solution
5. Future Work



# *Background - Deduplication*

**Deduplication is a technique for eliminating duplicate copies of repeating data.**



# Deduplication process in a nutshell

1. Divide into chunks
2. Calculate the chunks' hashes
3. Store chunks uniquely

Object 1



Object 2





Object 1



Object 2







# *Background - Chunking Methods*

## **How to chunk the input data**

1. Simple - fixed size.
2. Content aware - files, objects, applications.
3. Content sensitive - rolling hash.



The Opera ghost really existed. He was not, as was long believed of the artists, the superstition of the managers, or a product of the brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge. Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to

say, of a spectral shade.

The Opera ghost really existed. He was not, as was long believed of the artists, the superstition of the managers, or a product of the brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge. Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to

say, of a spectral shade.

The Opera ghost really existed.

He was not, as was long believed, a creature of the imagination of the artists, the superstition of the managers, or a product of the *joy* and impressionable brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge.

Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to say, of a spectral shade.

The Opera ghost really existed.

He was not, as was long believed, a creature of the imagination of the artists, the superstition of the managers, or a product of the *absurd* and impressionable brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge.

Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to say, of a spectral shade.

The Opera ghost really existed. He was **not**, **as** was long believed, a creature of the imagination of the artists, the superstition of the managers, or a **p**roduct of the *joy* and impressionable brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge. Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to say, **of** a spectral shade.

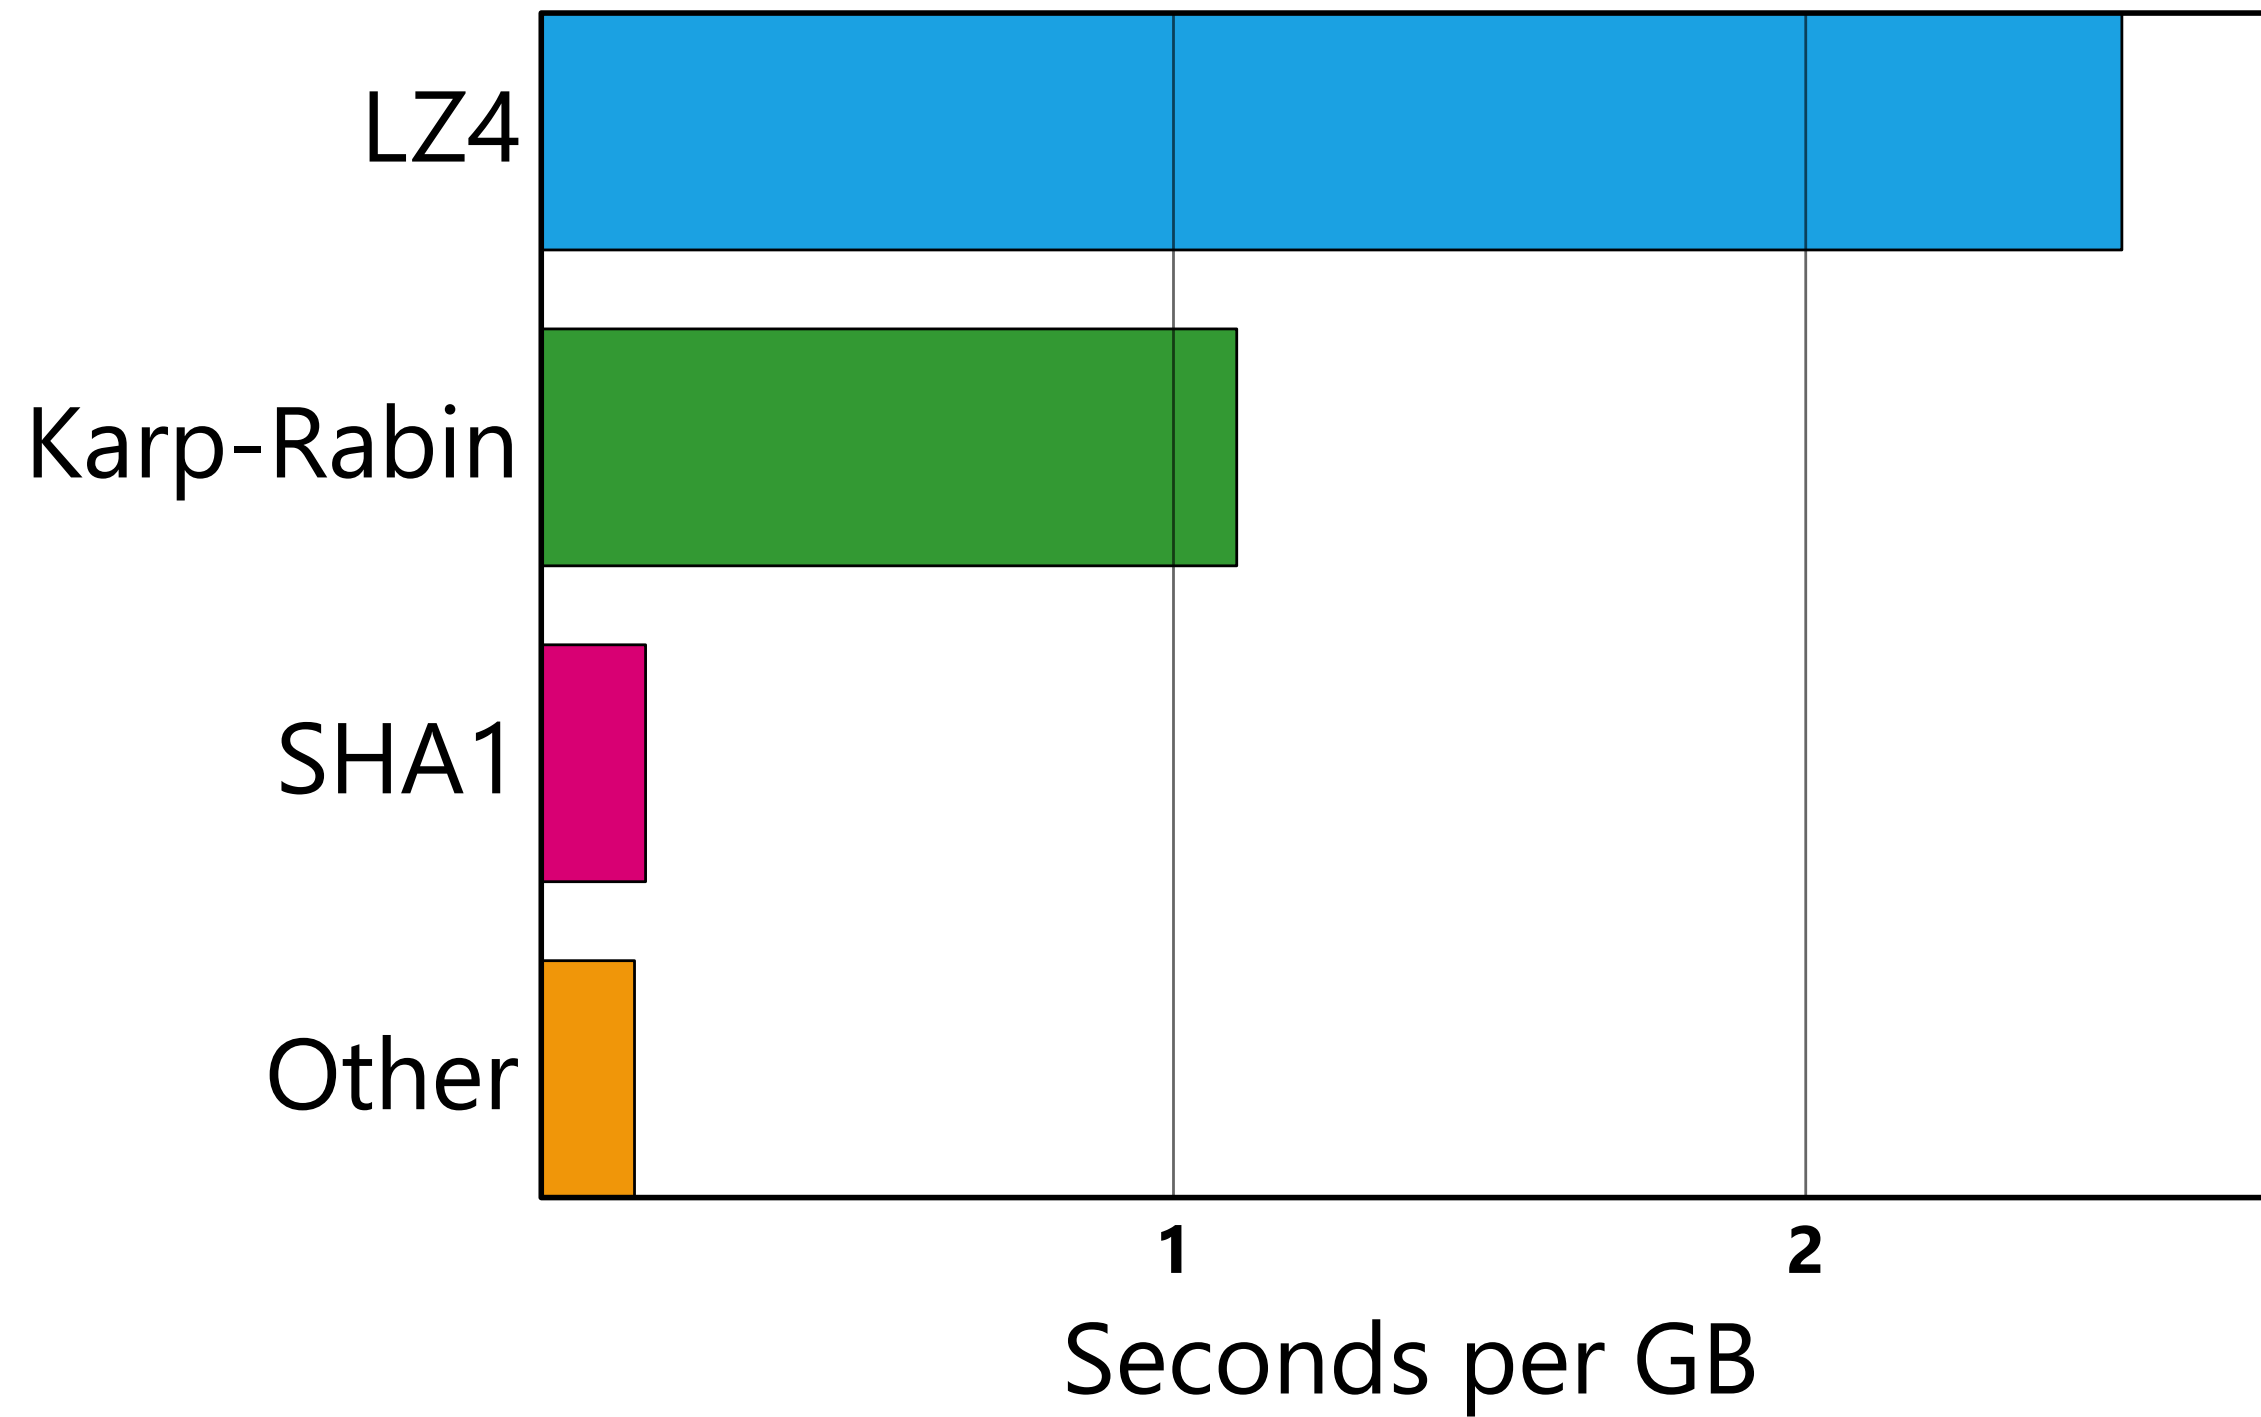
The Opera ghost really existed. He was **not**, **as** was long believed, a creature of the imagination of the artists, the superstition of the managers, or a **p**roduct of the *absurd* and impressionable brains of the young ladies of the ballet, their mothers, the box-keepers, the cloak-room attendants or the concierge. Yes, he existed in flesh and blood, although he assumed the complete appearance of a real phantom; that is to say, **of** a spectral shade.

# *Background - Deduplication Performance*

In 2017 we worked on a deduplication engine,  
and we tried to improve its performance.







# *Chunking Problem*

Given a stream of bytes, divide it into chunks for deduplication.

1. Output identical chunks for identical data
2. Good chunk size distribution.
3. Good performance.
4. Works for any input (photo, DB, text, random, etc...)



# *Traditional Solutions*

Karp-Rabin  
Cyclic Polynomial



$$H_k = \text{Hash}(X_{k-63}, X_{k-62}, X_{k-61}, \dots, X_{k-2}, X_{k-1}, X_k)$$

If **Criteria**( $H_k$ ) holds then mark a boundary after  $X_k$

$$H_k = \text{Hash}(X_{k-63}, X_{k-62}, X_{k-61}, \dots, X_{k-2}, X_{k-1}, X_k)$$

$$H_{K+1} = \text{Hash}(X_{k-62}, X_{k-61}, X_{k-60}, \dots, X_{k-1}, X_k, X_{k+1})$$

$$\begin{aligned}
 H_k &= \text{Hash}( X_{k-63} \quad X_{k-62} \quad X_{k-61} \quad \dots \quad X_{k-2} \quad X_{k-1} \quad X_k ) \\
 H_{k+1} &= \text{Hash}( X_{k-62} \quad X_{k-61} \quad X_{k-60} \quad \dots \quad X_{k-1} \quad X_k \quad X_{k+1} ) \\
 &= \text{RollHash}( X_{k-63}, H_k, X_{k+1} )
 \end{aligned}$$



## Karp-Rabin

$$h_i = \sum_{j=0}^{j=63} p^j x_{i-j} \quad \text{mod } N$$

$$h_{i+1} = p^{64} x_{i-64} + p h_i + x_i \quad \text{mod } N$$

## Cyclic Polynomial

$$h_i = \bigoplus_{j=0}^{j=63} \text{rotate}(x_{i-j}, j)$$

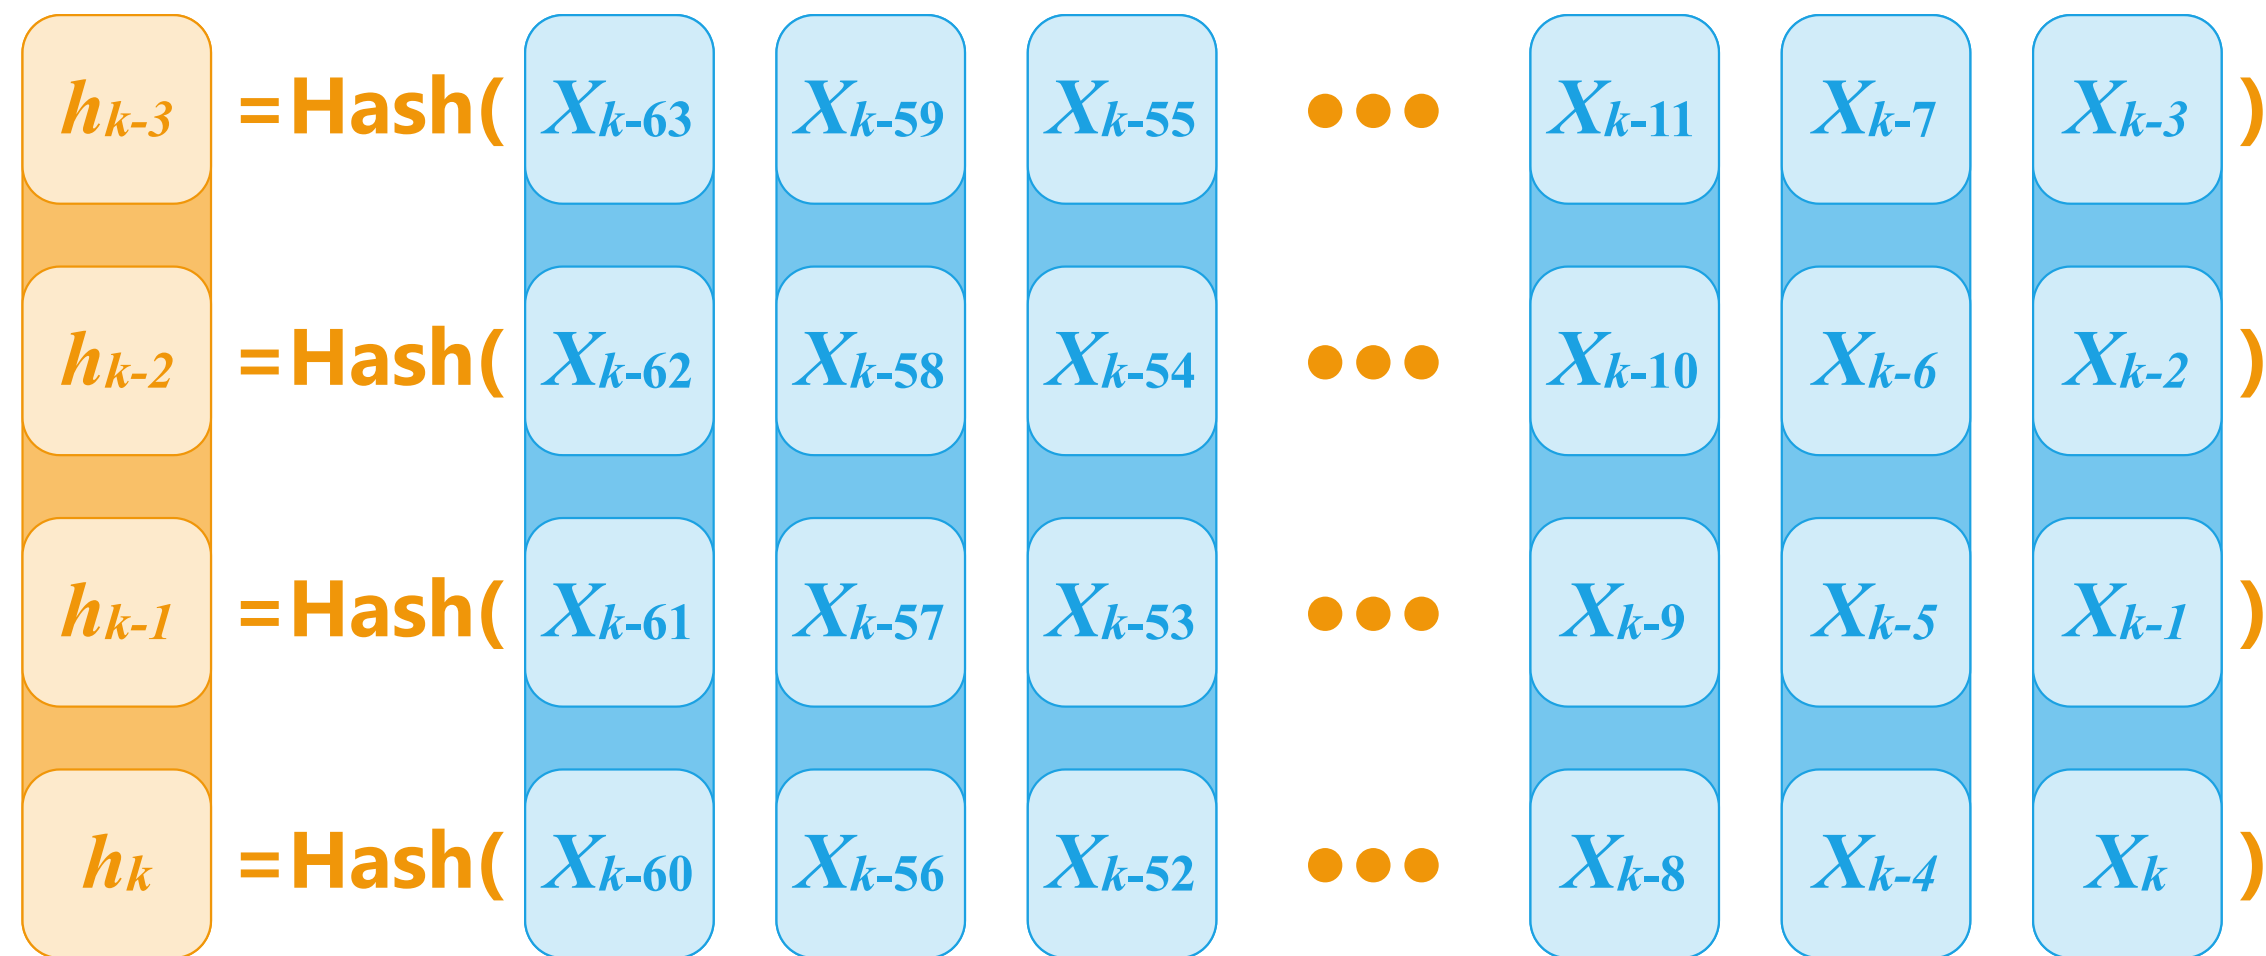
$$h_{i+1} = x_{i-64} + \text{rotate}(h_i, 1) + x_i$$


# *Proposed Solution*

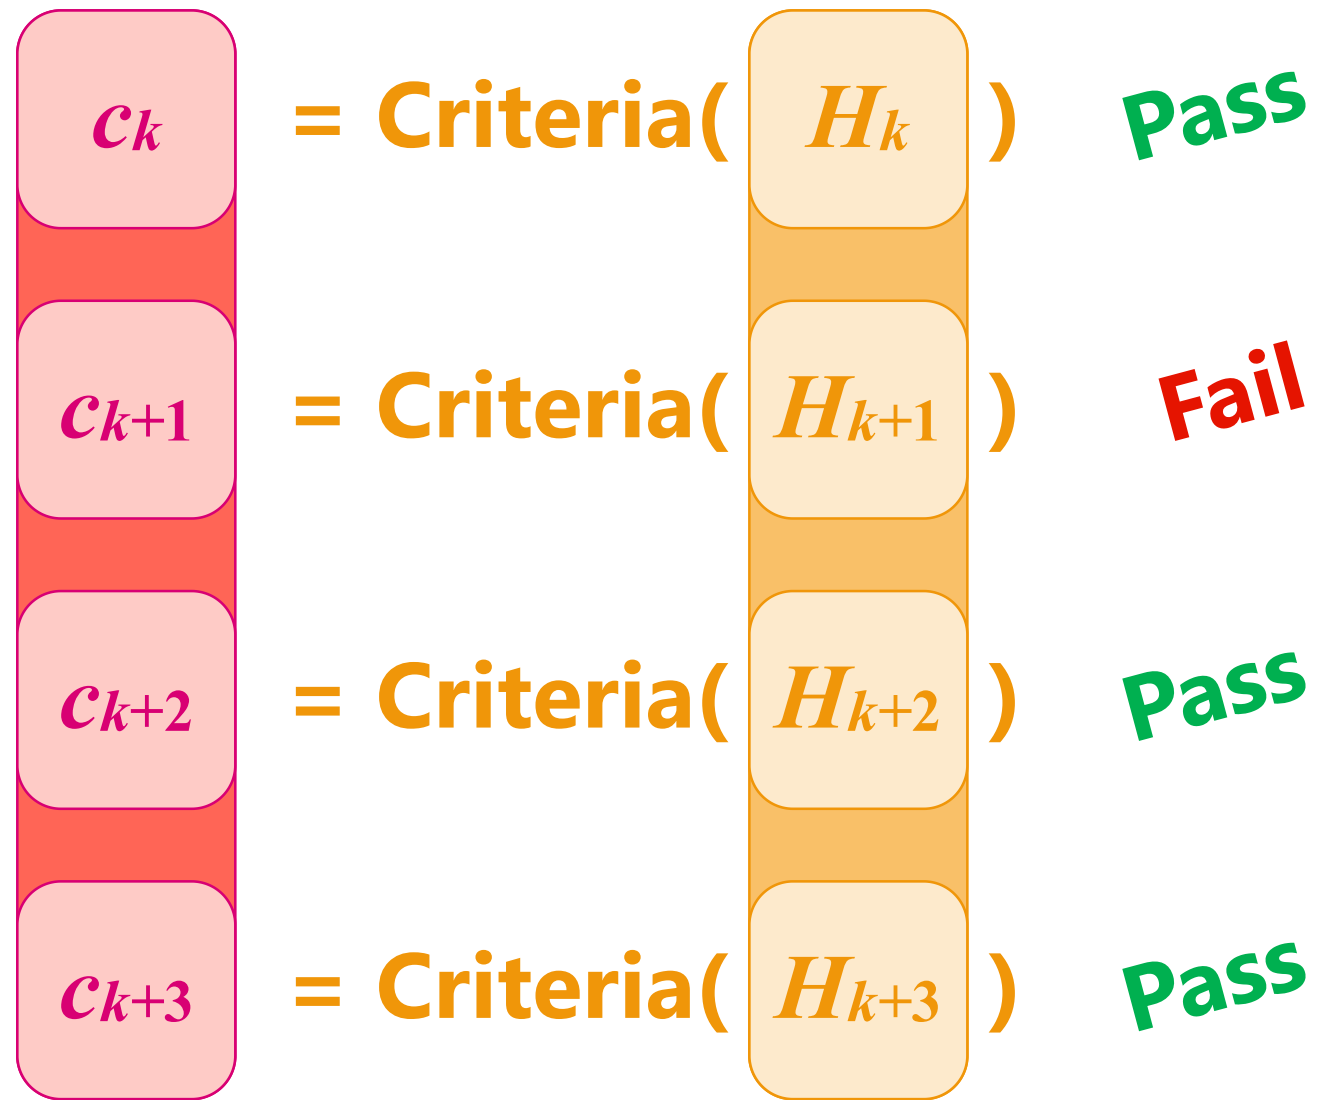
## **How does it work:**


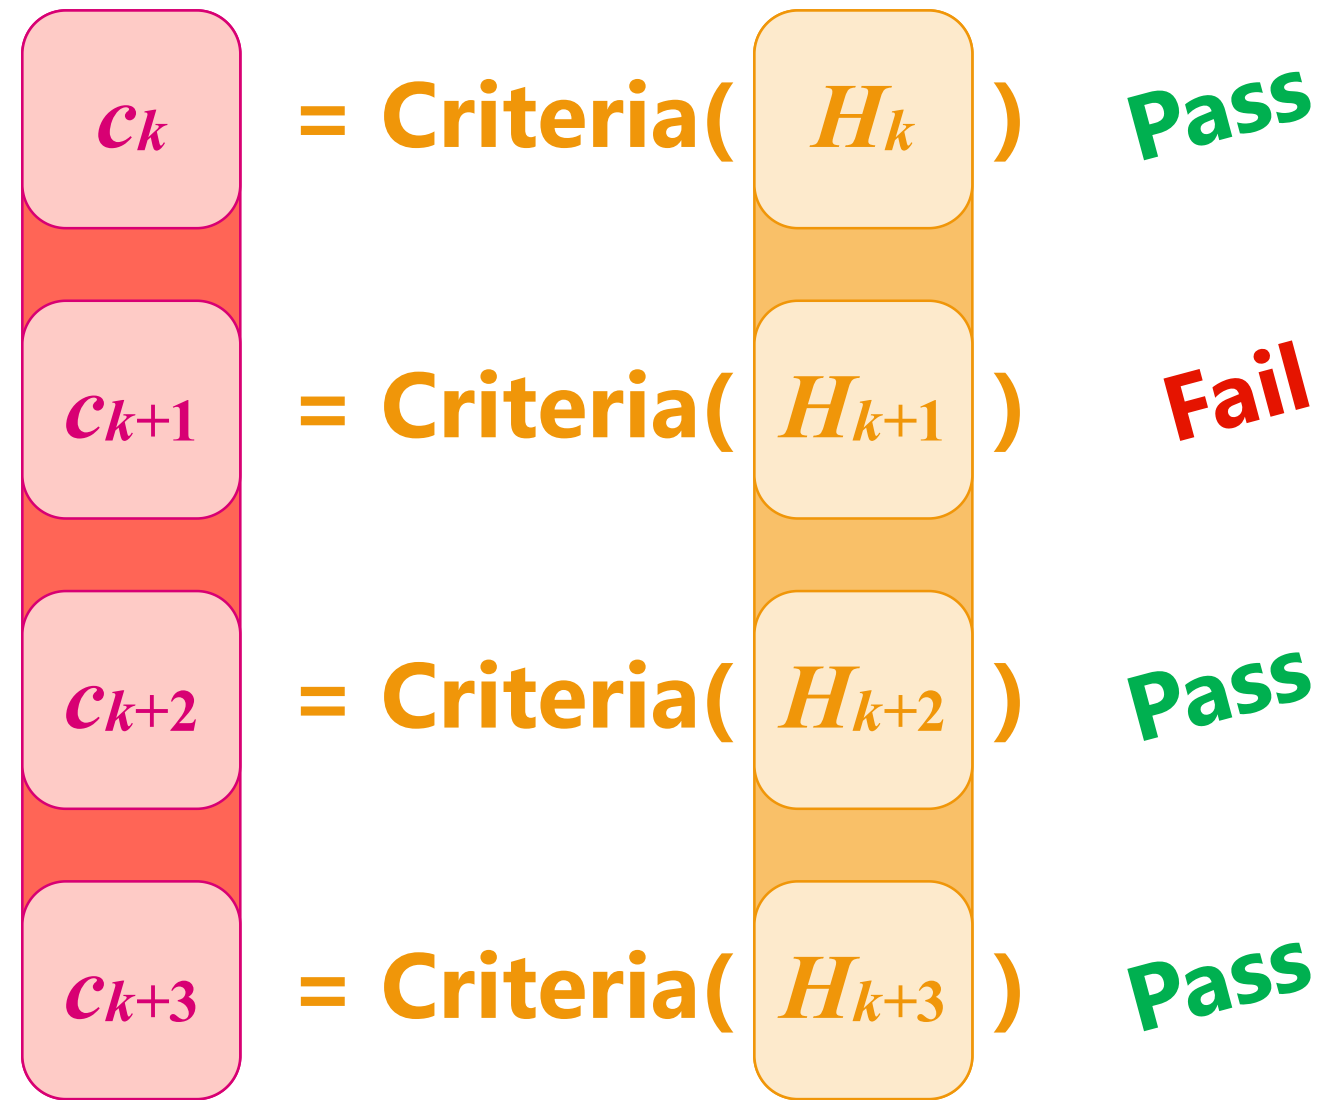
1. Work with rolling vectors
2. Calculate a hash of byte size
3. Calculate the criteria, in a way that:
  - Number of calculations are constant unrelated to the vector size
  - Can find a cutting point at a byte offset











**$c_{k-4} \dots c_{k-1}$**

<b>1</b>	<b>1</b>	<b>0</b>	<b>0</b>
----------	----------	----------	----------

**$c_k \dots c_{k+3}$**


<b>0</b>	<b>1</b>	<b>0</b>	<b>0</b>
----------	----------	----------	----------

**$c_{k+4} \dots c_{k+7}$**

<b>0</b>	<b>0</b>	<b>0</b>	<b>1</b>
----------	----------	----------	----------




$c_{k-4} \dots c_{k-1}$




**Trailing  
Zeroes**

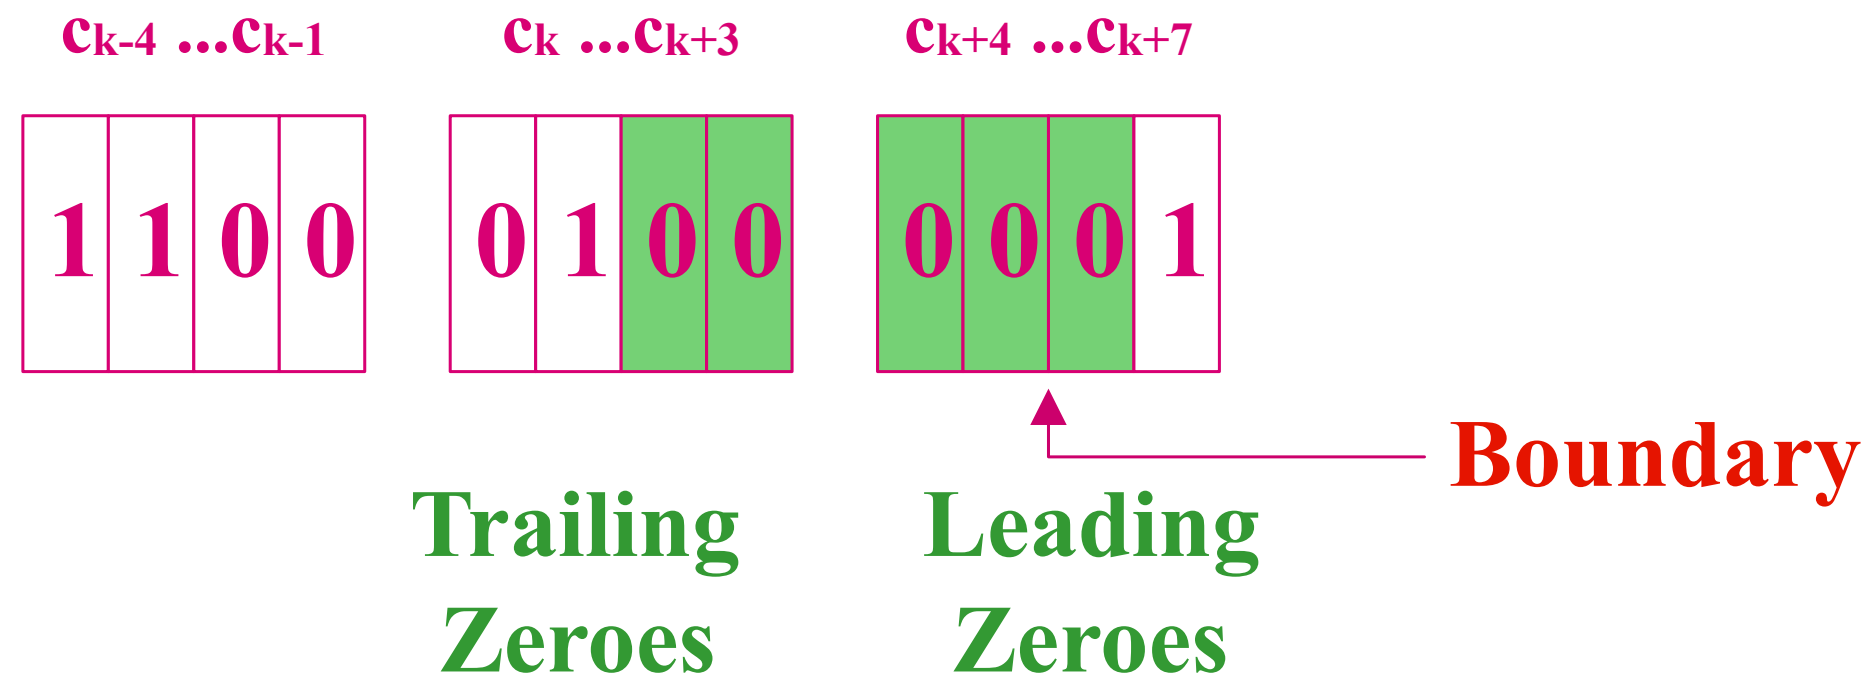
$c_k \dots c_{k+3}$




**Leading  
Zeroes**

$c_{k+4} \dots c_{k+7}$





# Measured Results





<b>Algorithm</b>	<b>Random Data</b>	<b>Corpus Data</b>
Karp-Rabin	975 MB/s	927 MB/s
Cyclic-Polynomial	1675 MB/s	1676 MB/s
Ours	6715 MB/s	7136 MB/s

<b>Chunking Alg.</b>	<i>Dedup Perf.</i>	<b>LZ4</b>	<b>SHA1</b>	<b>Other</b>	<b>Chunking</b>
Karp-Rabin	<i>262 MB/s</i>	63.9%	4.1%	3.8%	28.1%
Ours	<i>345 MB/s</i>	84.5%	5.4%	5.1%	4.7%

- Same Distribution
- Faster Chunking Performance
- Faster Overall Performance

# Future Work

*A chunking algorithm that is*

\*

**Past**

**Presented**

**Future**

---

**Backward Compatible**

N/A

$f(k, x) \neq f(l, x)$

$f(k, x) = f(l, x)$

---

**Work**

cn

cn

cn

---

**Speed**

cn

$c(n \log k) / k$

$c(n \log k) / k$





# Thanks

<https://github.com/dudejohnny/PSC2019>