

SIGMOD DaMoN 23.06.2014

# Main Memory Adaptive Indexing for Multi-core Systems

**Felix Martin Schuhknecht**

Victor Alvarez

Jens Dittrich

Stefan Richter

Information Systems Group

Saarland University

<https://infosys.uni-saarland.de/>

# Problem: Answer Range Queries


# Problem: Answer Range Queries

```
select A
from R
where R.A >= 10 and R.A < 20
```

# Problem: Answer Range Queries

```
select A
from R
where R.A >= 10 and R.A < 20
```


One extreme: Scan + Filter



# Problem: Answer Range Queries

```
select A
from R
where R.A >= 10 and R.A < 20
```


One extreme: Scan + Filter




# Problem: Answer Range Queries

```
select A
from R
where R.A >= 10 and R.A < 20
```

One extreme: Scan + Filter




Other extreme: Index




# Problem: Answer Range Queries

```
select A
from R
where R.A >= 10 and R.A < 20
```

One extreme: Scan + Filter




Other extreme: Index



# Index: When to build?


One extreme: **At once (Traditional Indexing)**






# Index: When to build?

One extreme: **At once (Traditional Indexing)**



# Index: When to build?

One extreme: *At once (Traditional Indexing)*




# Index: When to build?

One extreme: *At once* (Traditional Indexing)




# Index: When to build?

One extreme: *At once (Traditional Indexing)*




# Index: When to build?

One extreme: *At once (Traditional Indexing)*



# Index: When to build?


One extreme: *At once (Traditional Indexing)*



# Index: When to build?

One extreme: **At once (Traditional Indexing)**


Other extreme: **Incrementally at query time (Adaptive Indexing)**



# Index: When to build?

One extreme: **At once (Traditional Indexing)**

Other extreme: **Incrementally at query time (Adaptive Indexing)**






# Index: When to build?

One extreme: *At once (Traditional Indexing)*

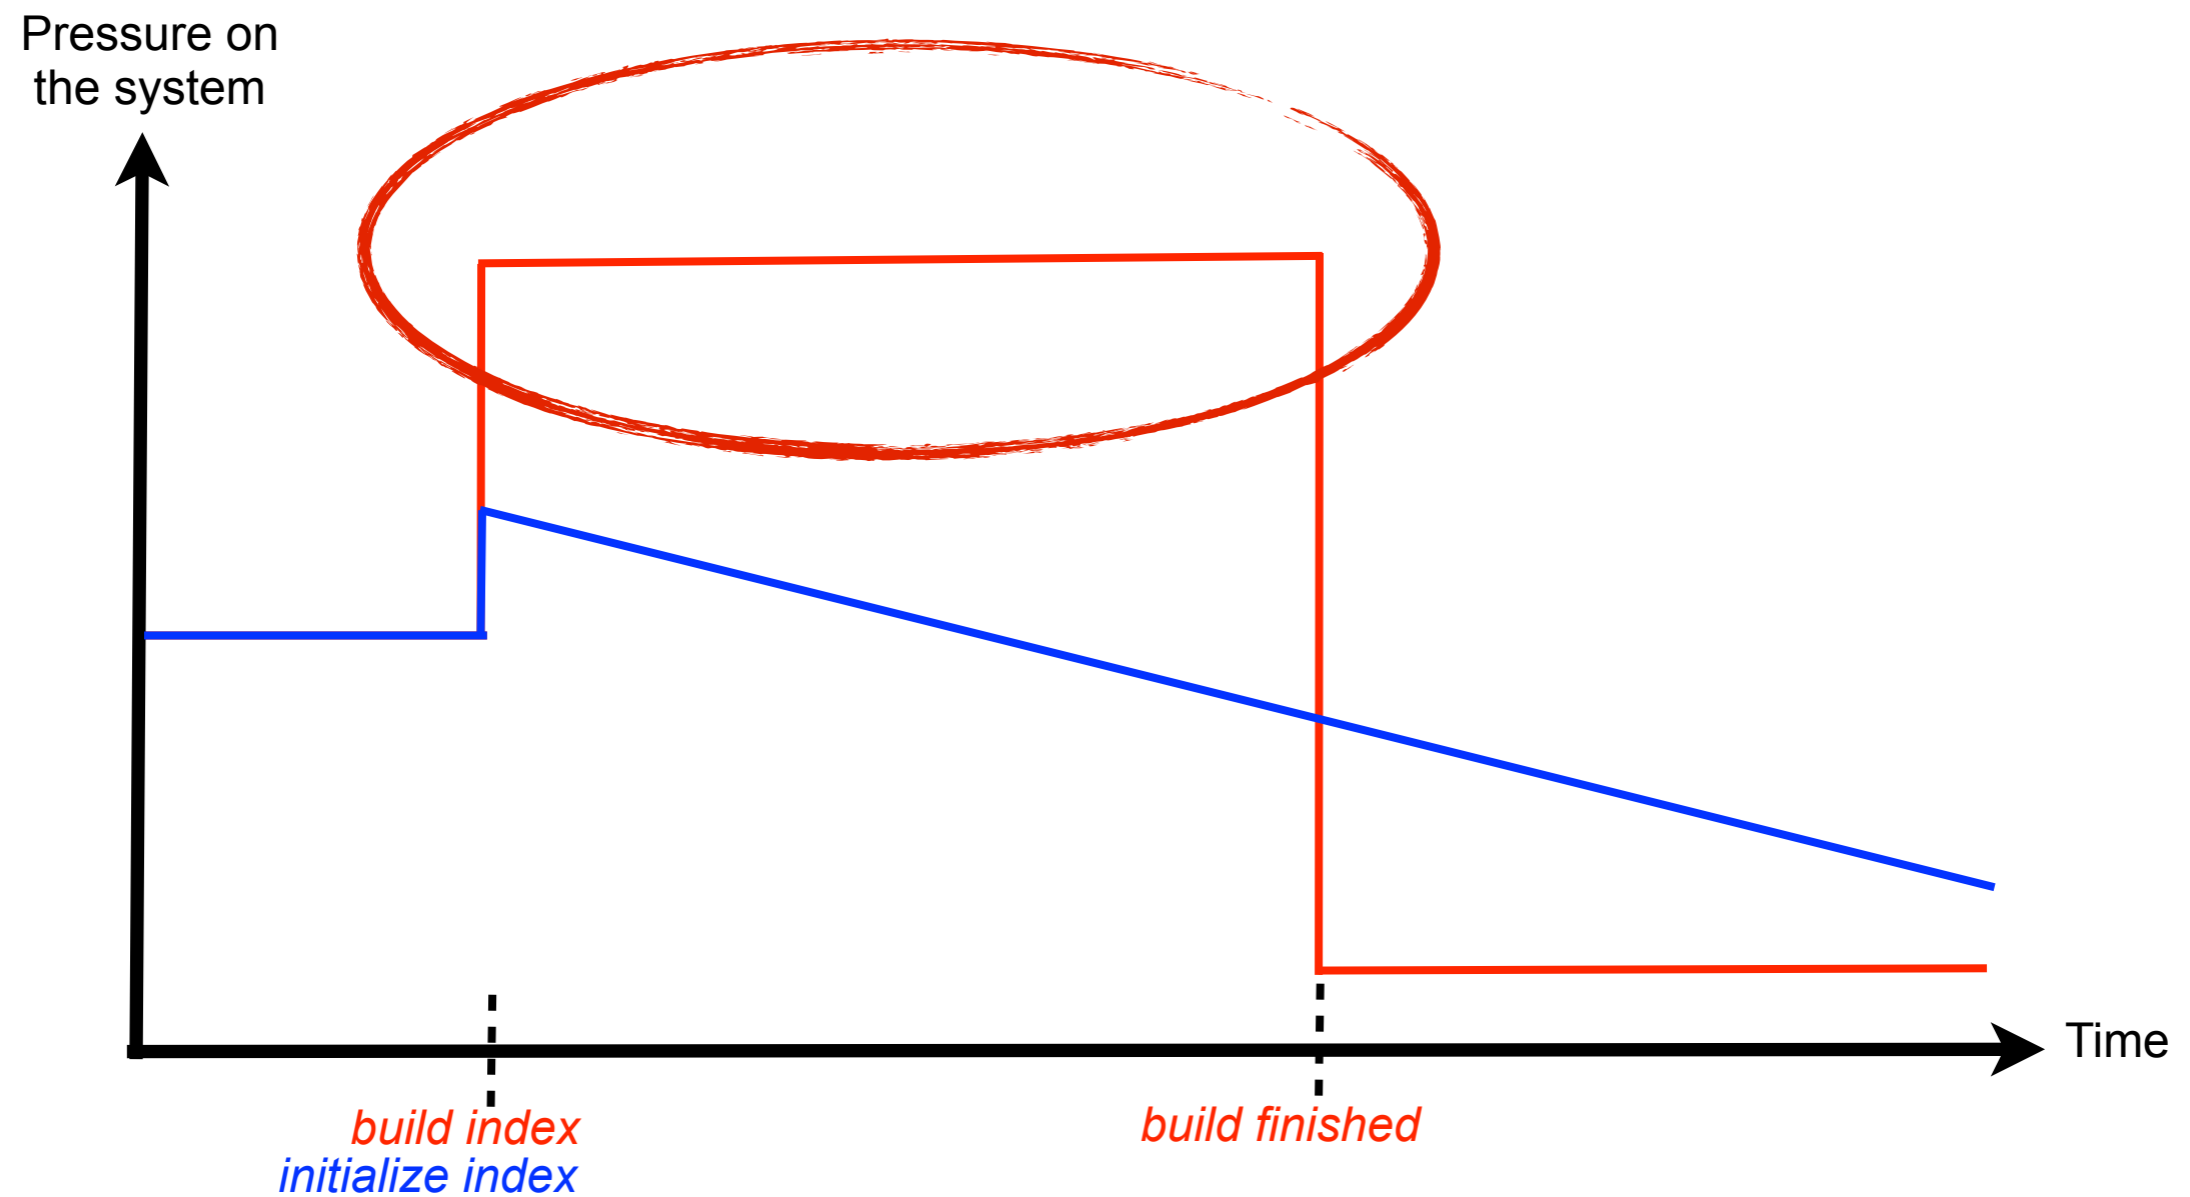
Other extreme: *Incrementally at query time (Adaptive Indexing)*



# Index: When to build?

One extreme: *At once (Traditional Indexing)*

Other extreme: *Incrementally at query time (Adaptive Indexing)*

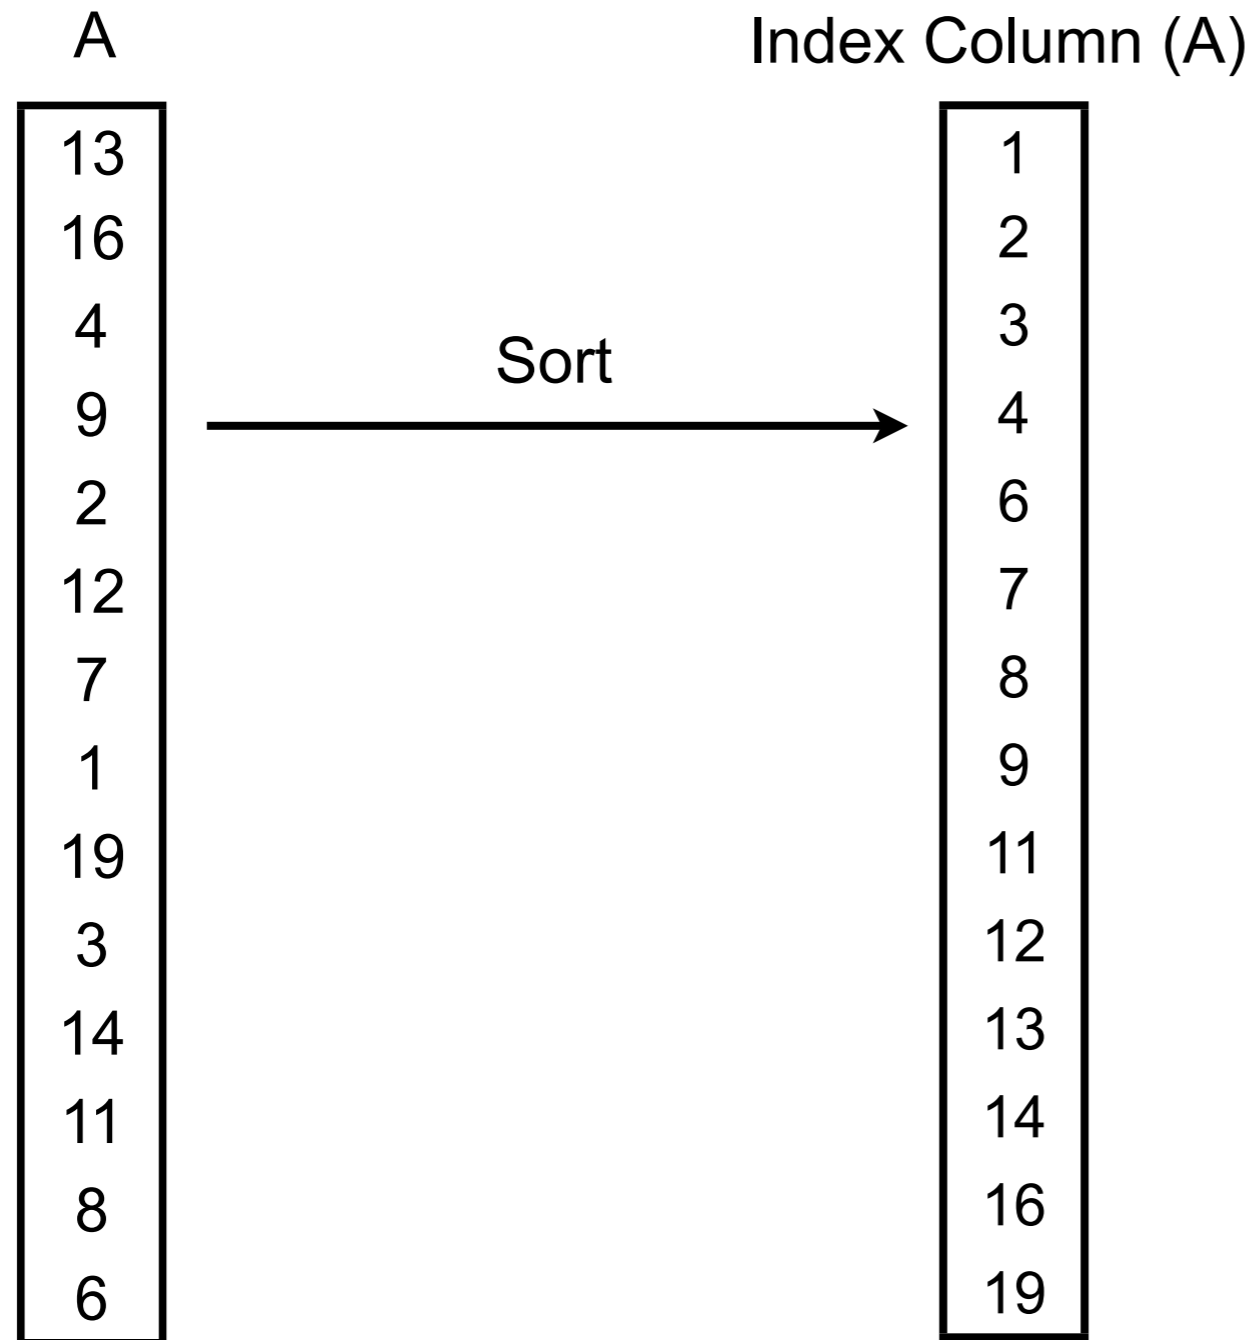


# Traditional Indexing: Sort + Binary Search

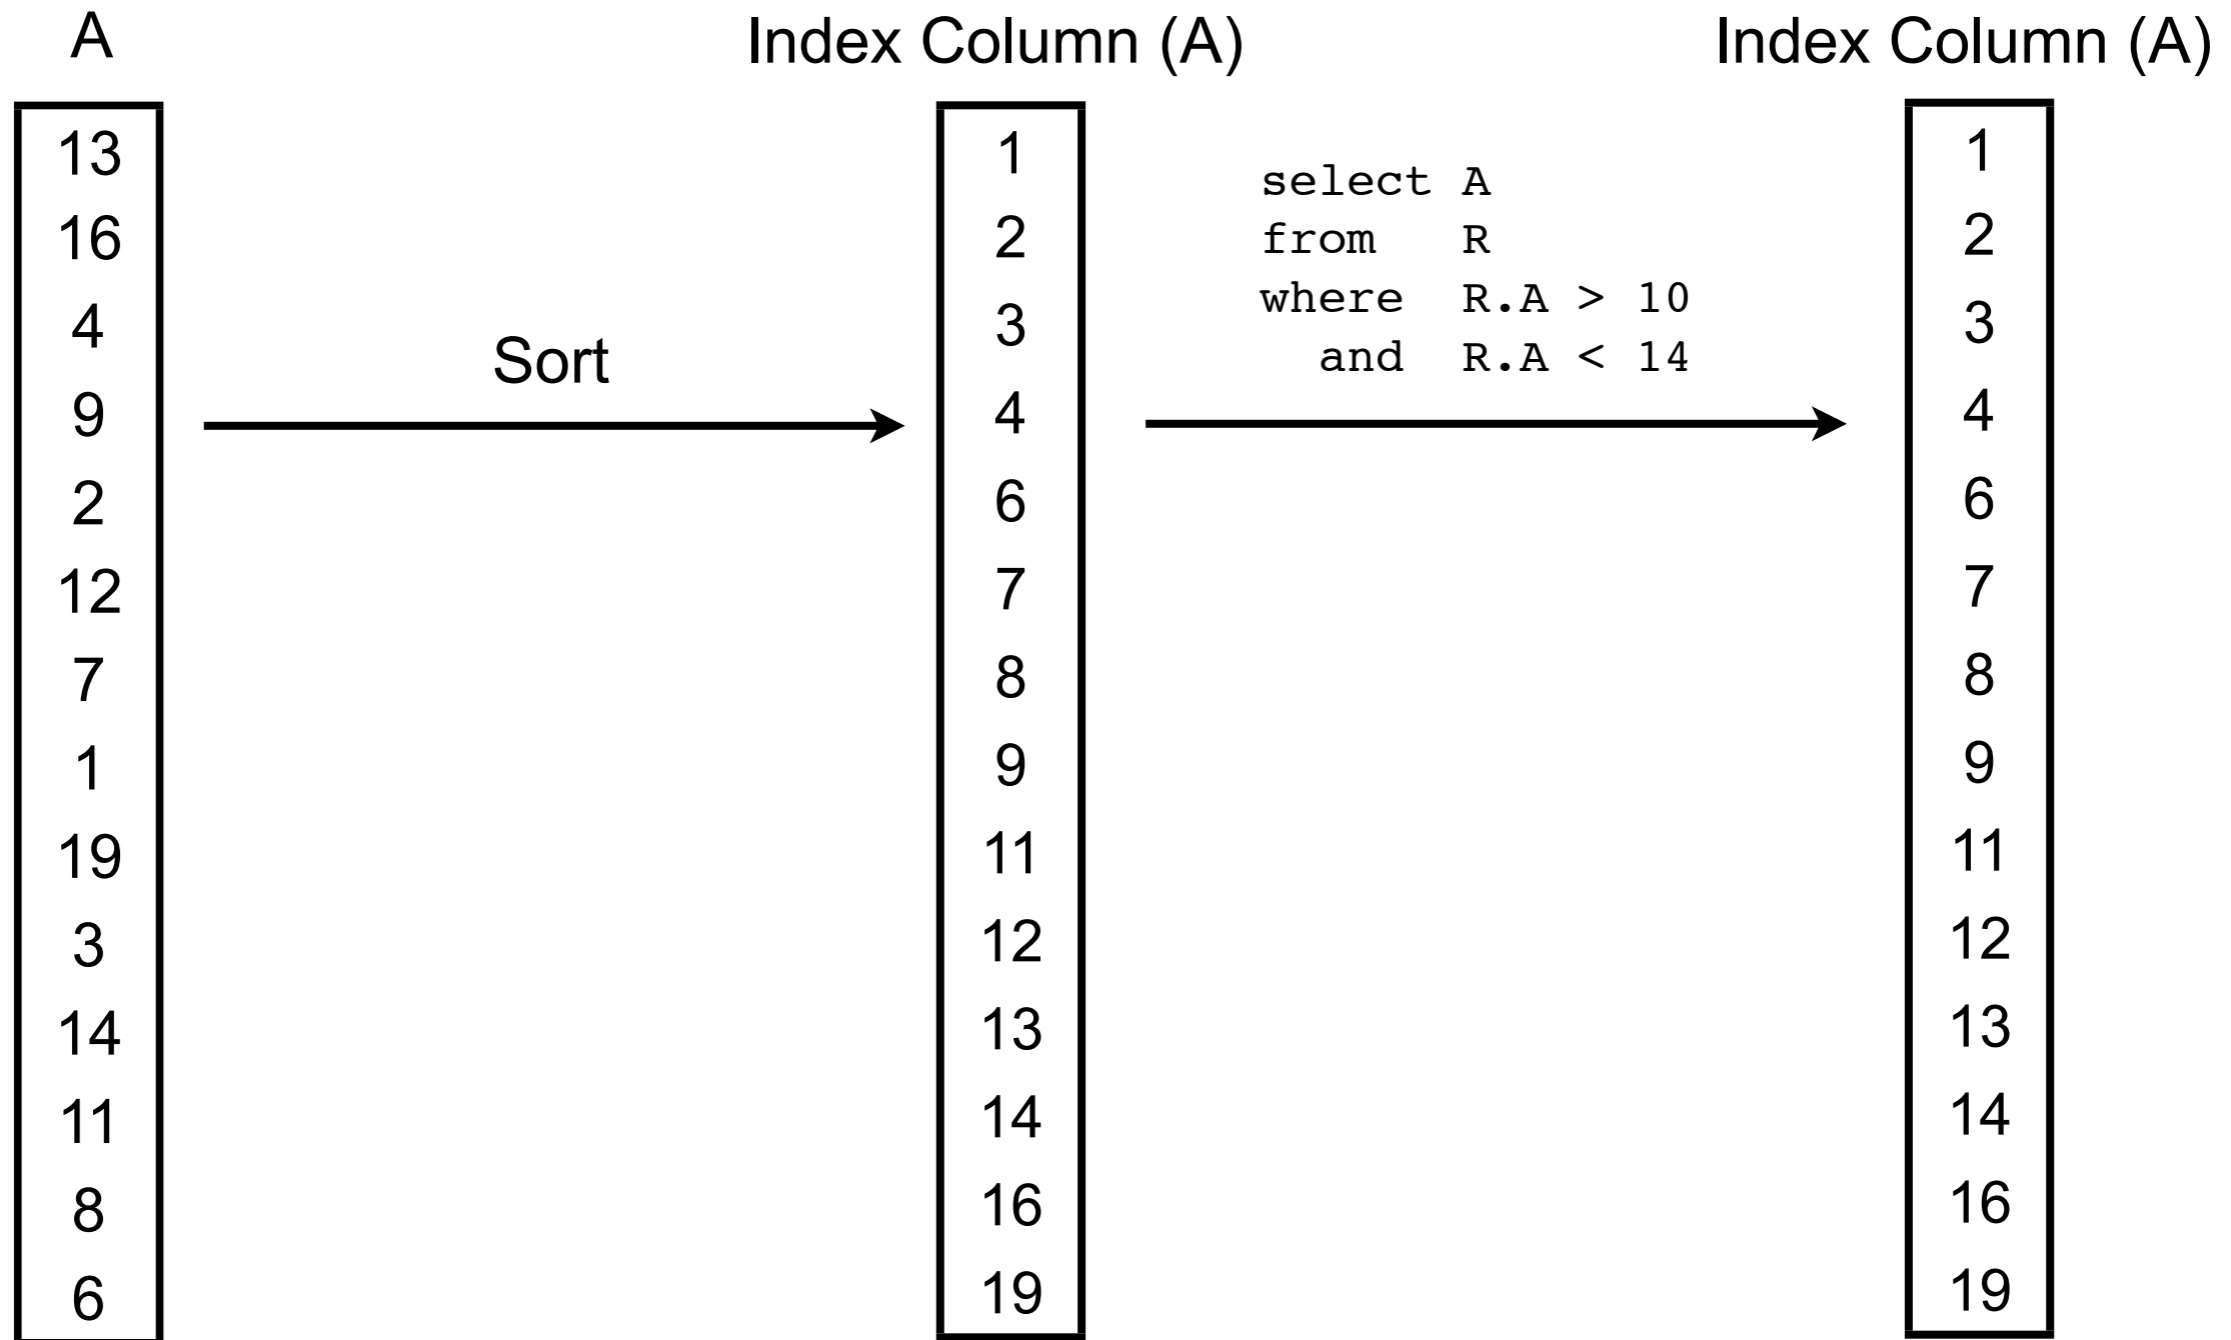
A

13
16
4
9
2
12
7
1
19
3
14
11
8
6

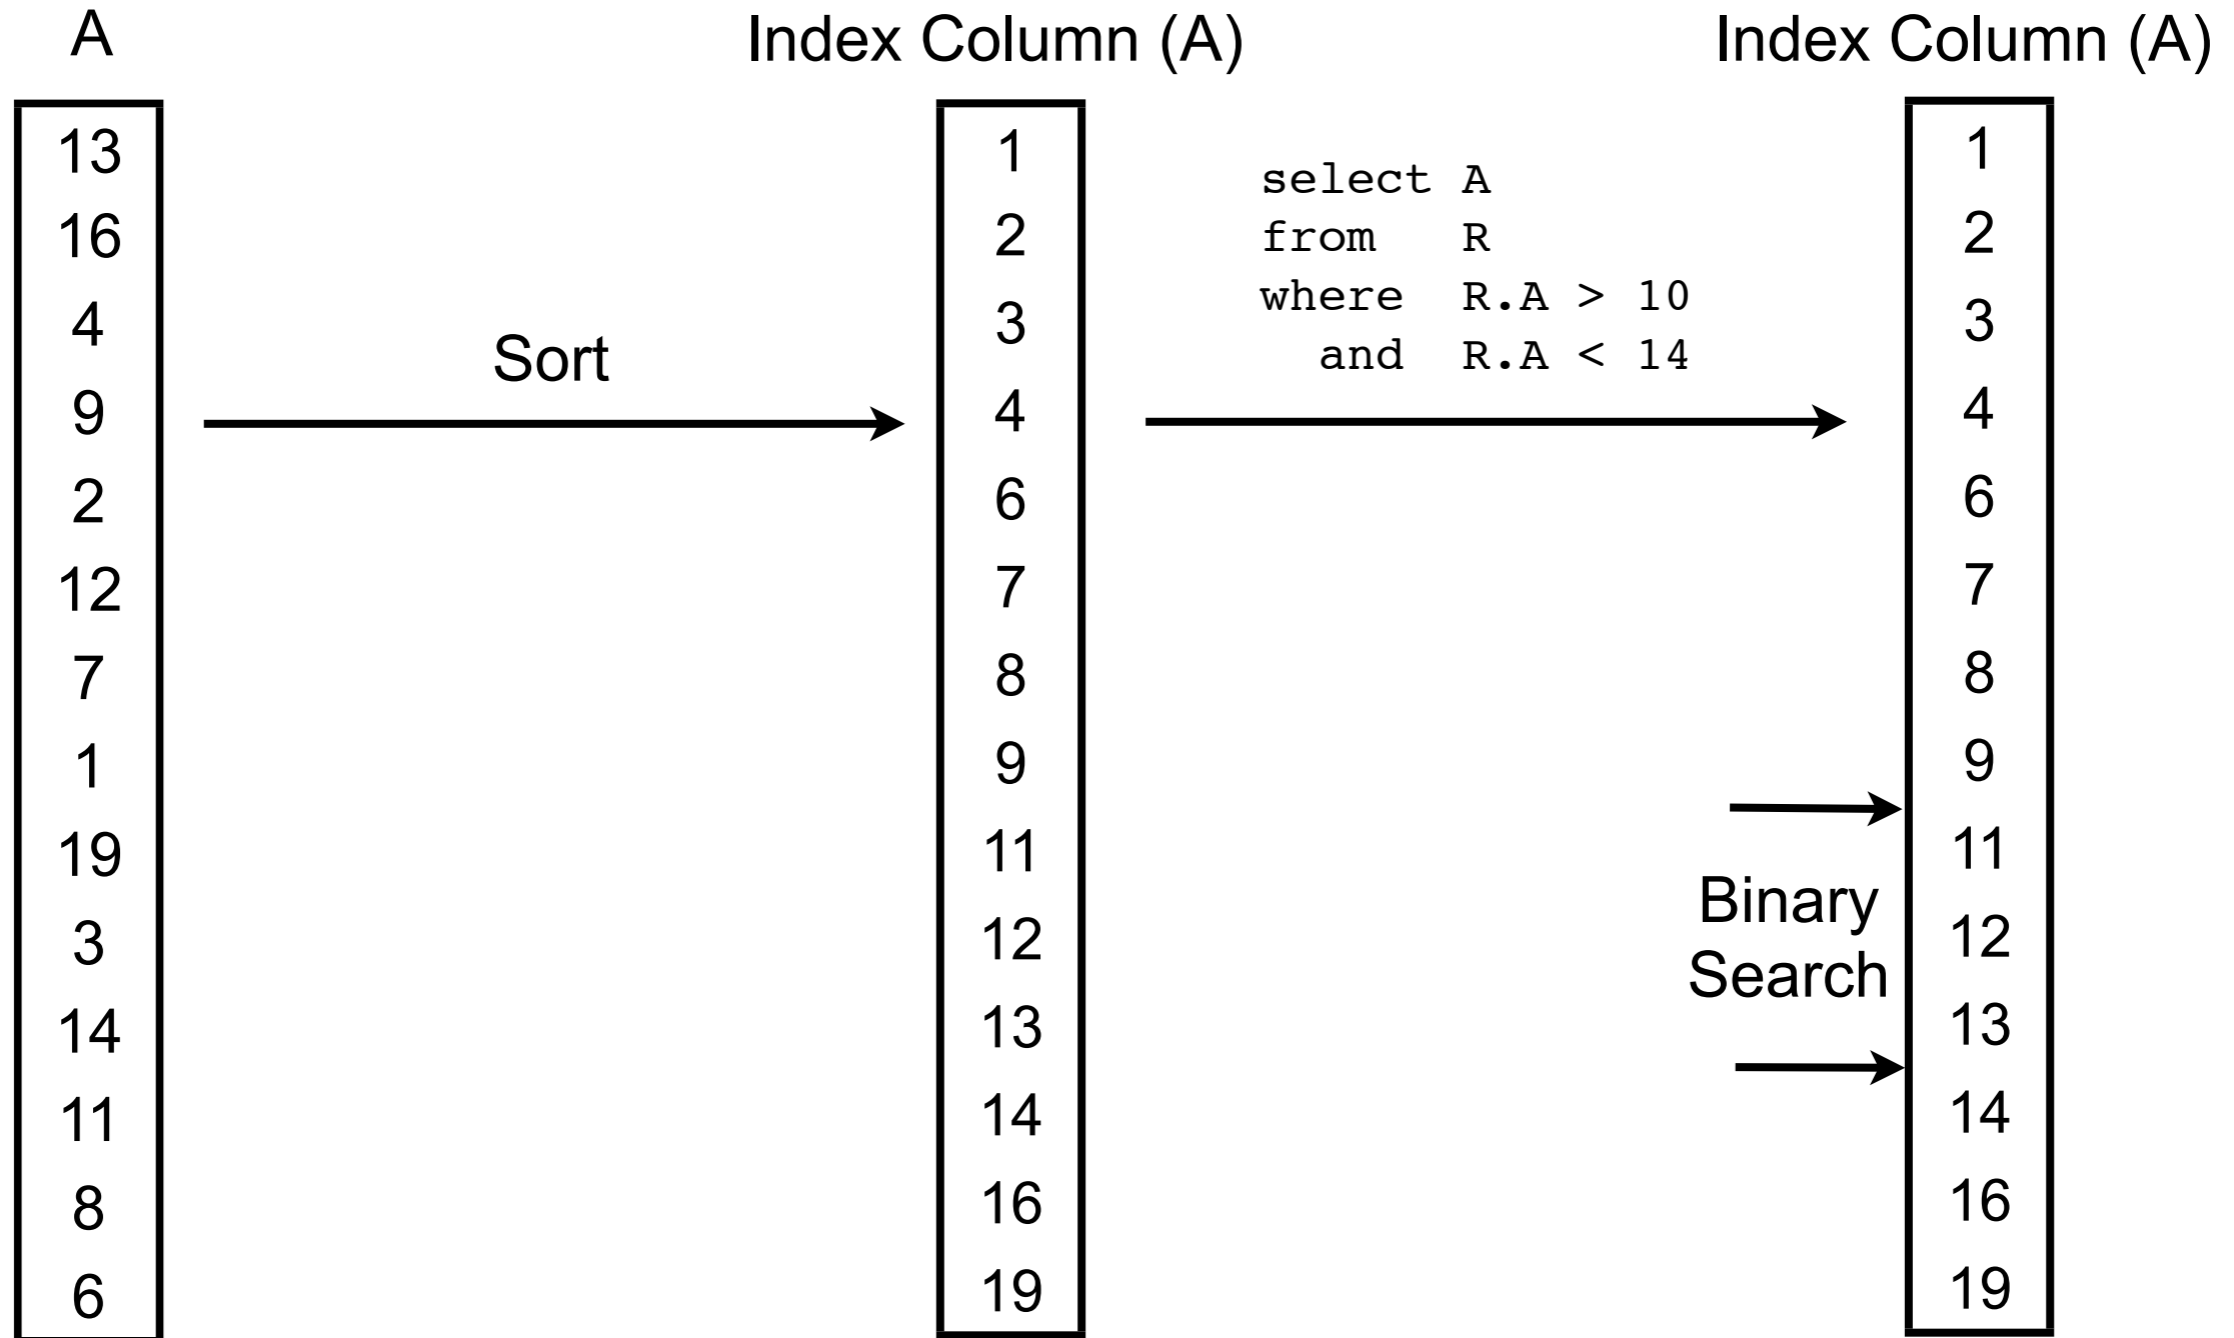
# Traditional Indexing: Sort + Binary Search



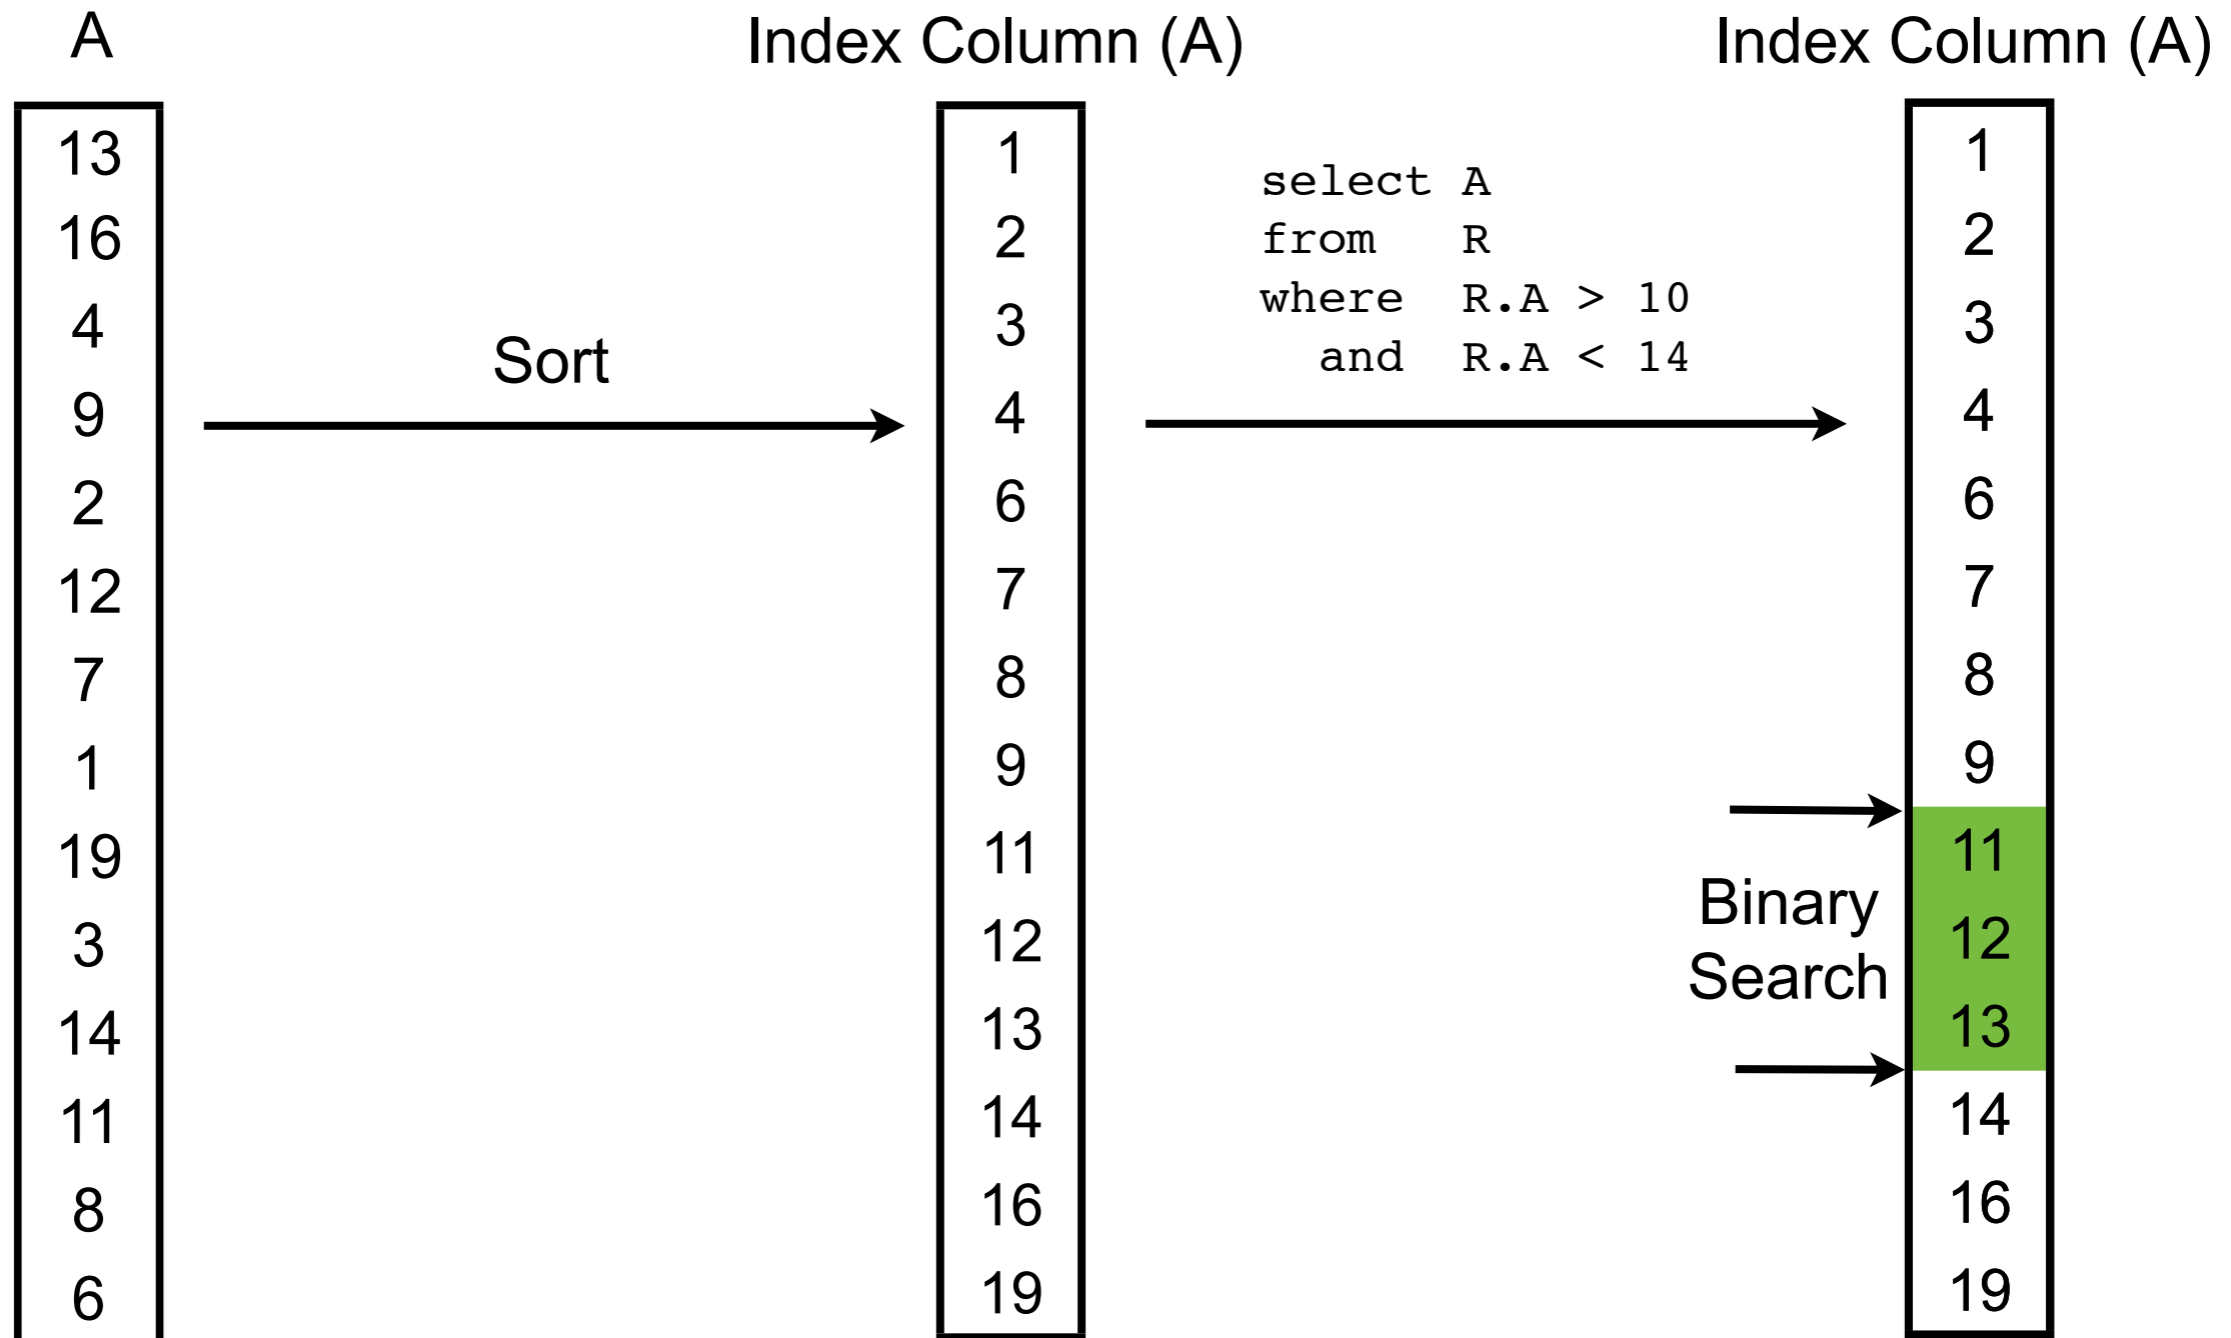
# Traditional Indexing: Sort + Binary Search



# Traditional Indexing: Sort + Binary Search



# Traditional Indexing: Sort + Binary Search



# Adaptive Indexing: Standard Cracking




# Adaptive Indexing: Standard Cracking

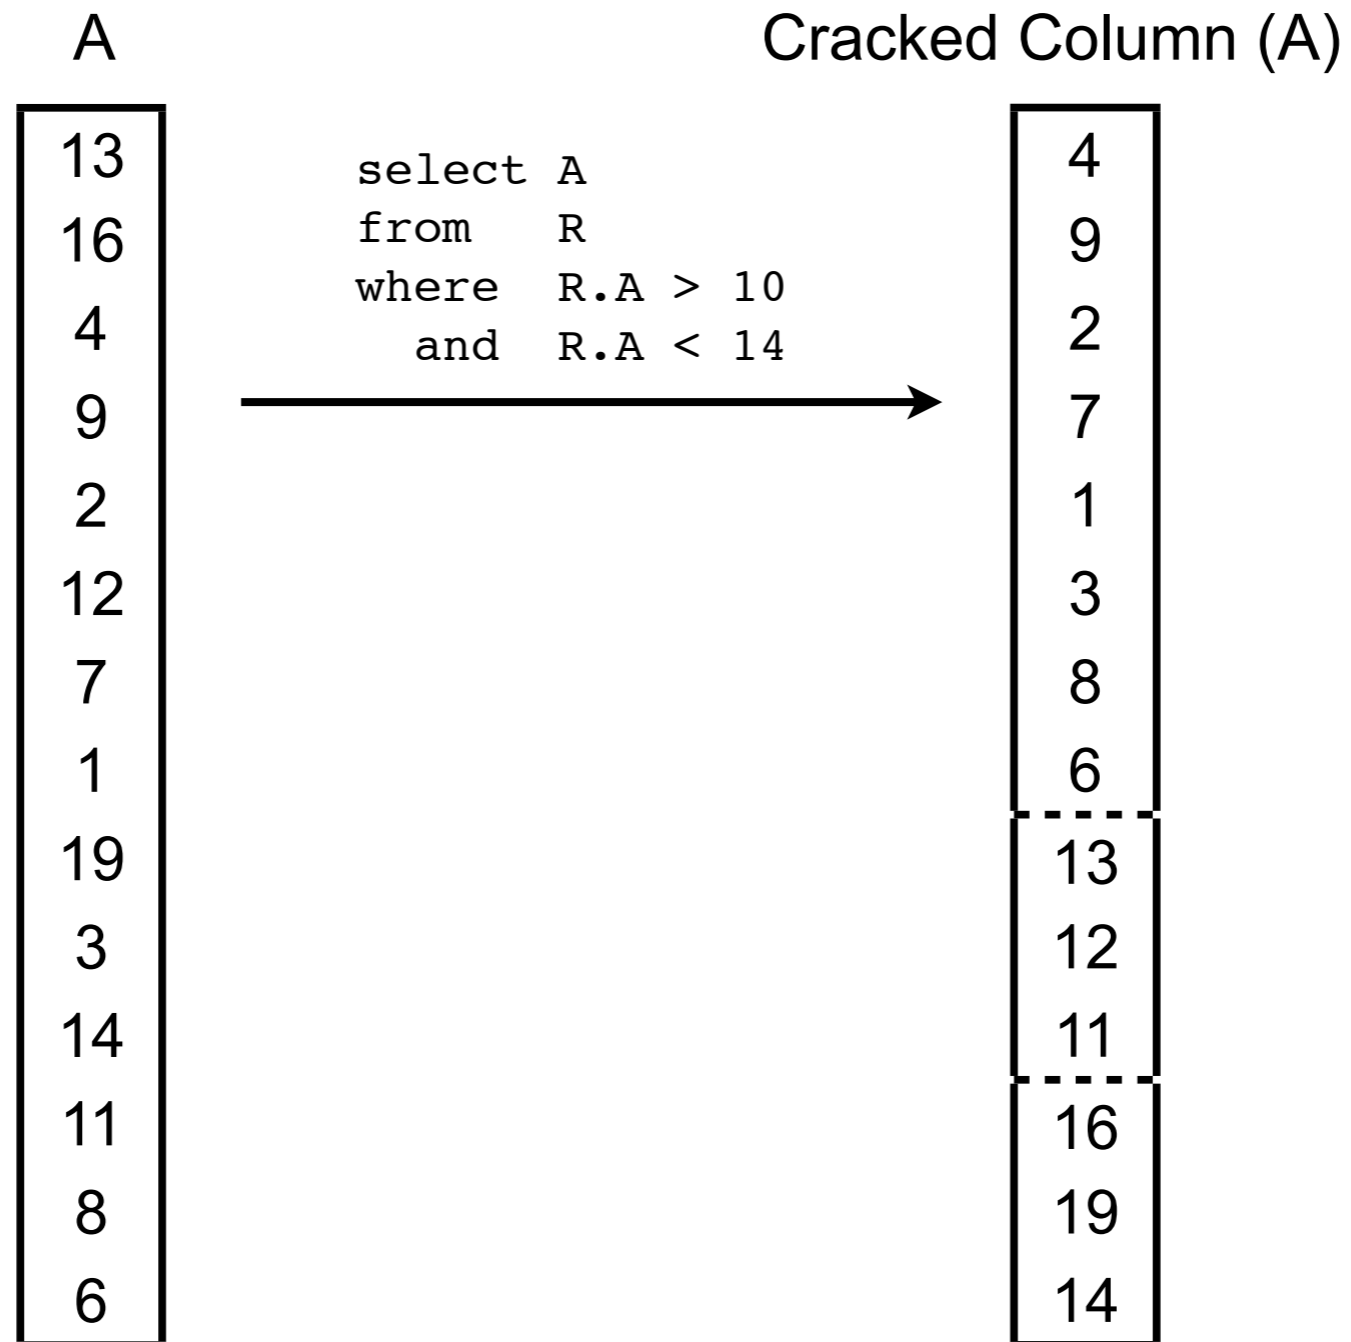
A

13
16
4
9
2
12
7
1
19
3
14
11
8
6

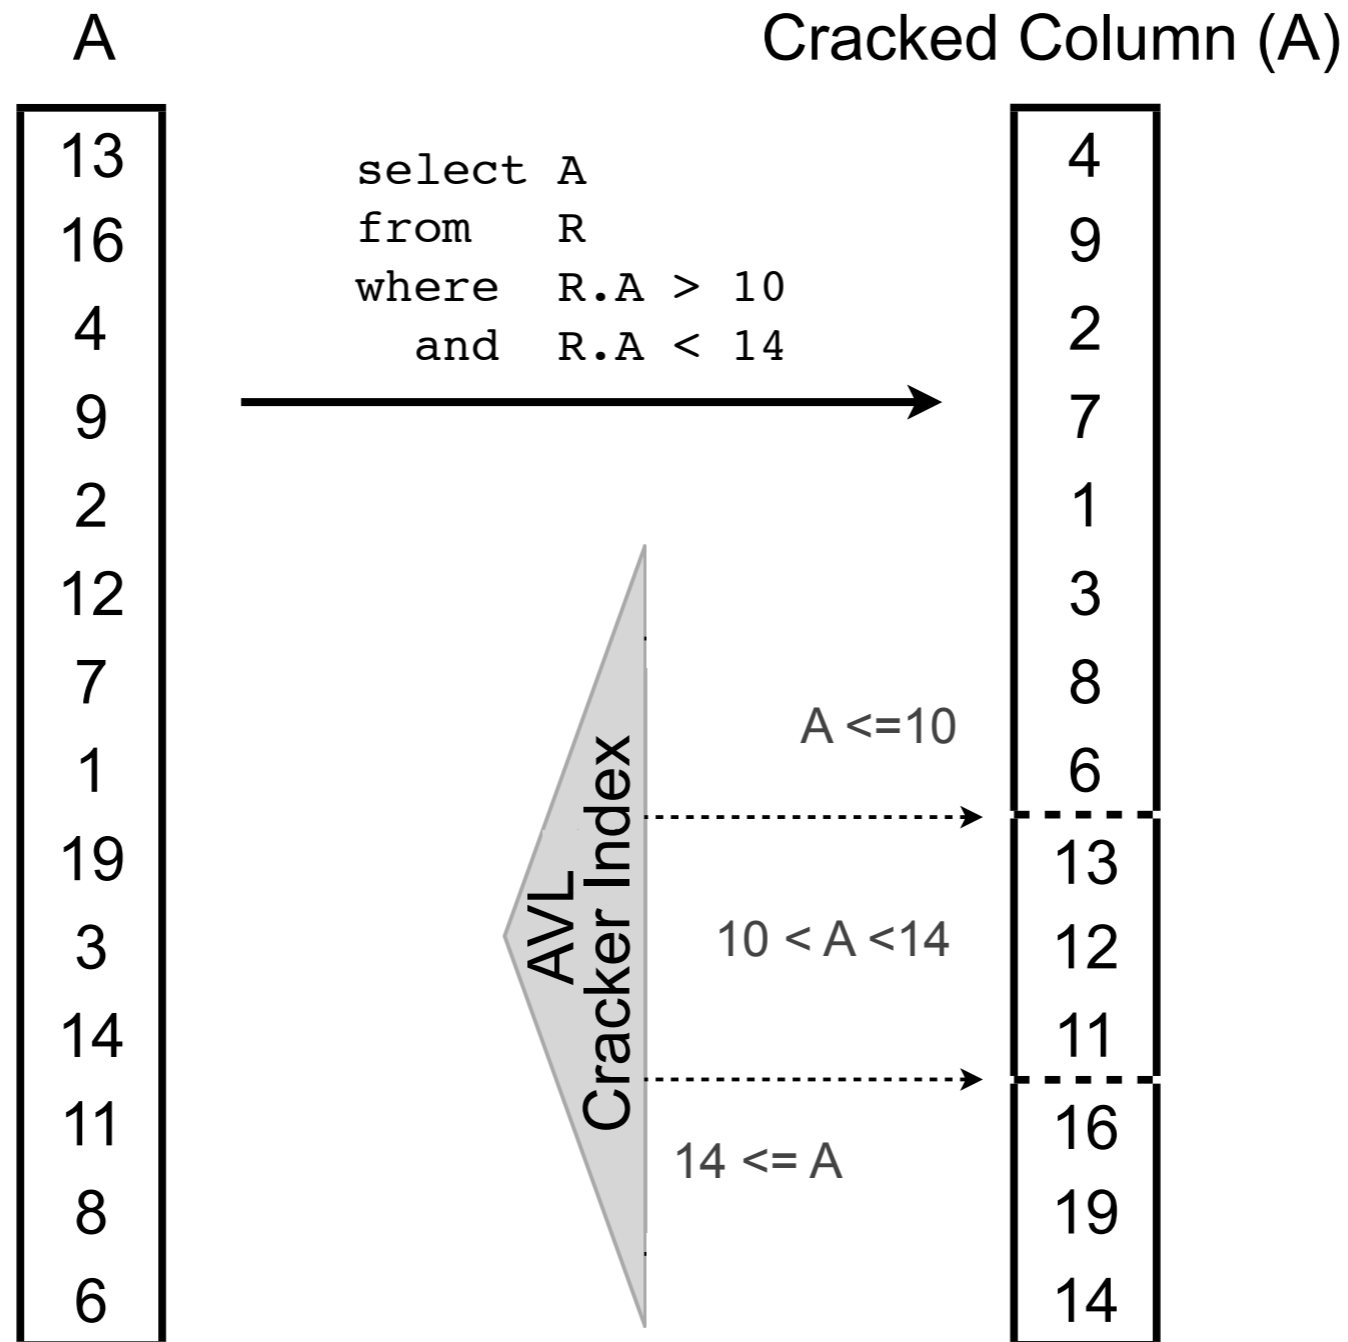
# Adaptive Indexing: Standard Cracking



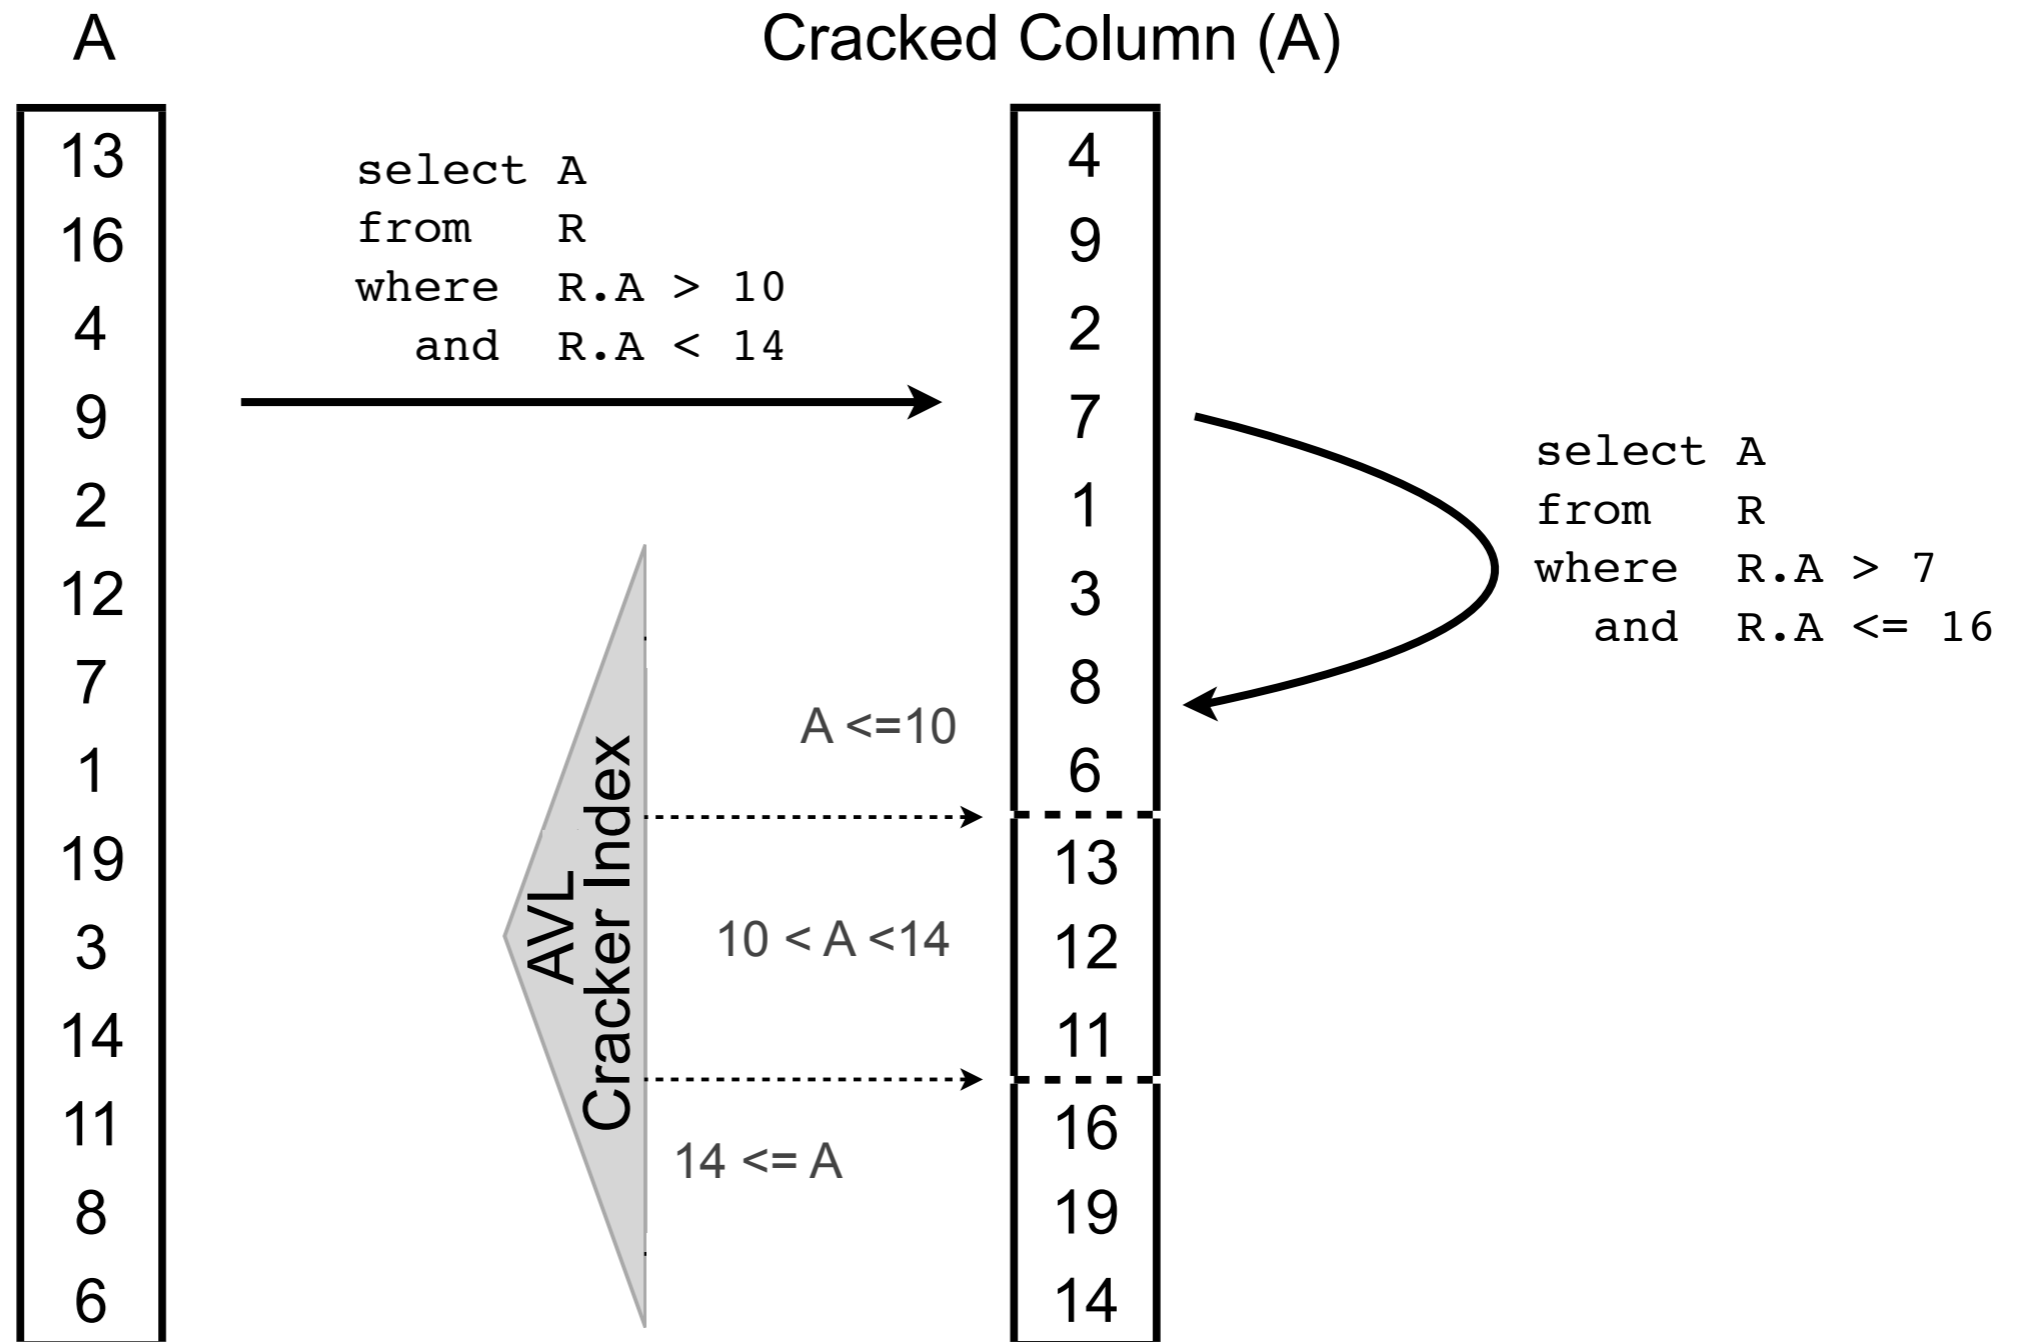
# Adaptive Indexing: Standard Cracking



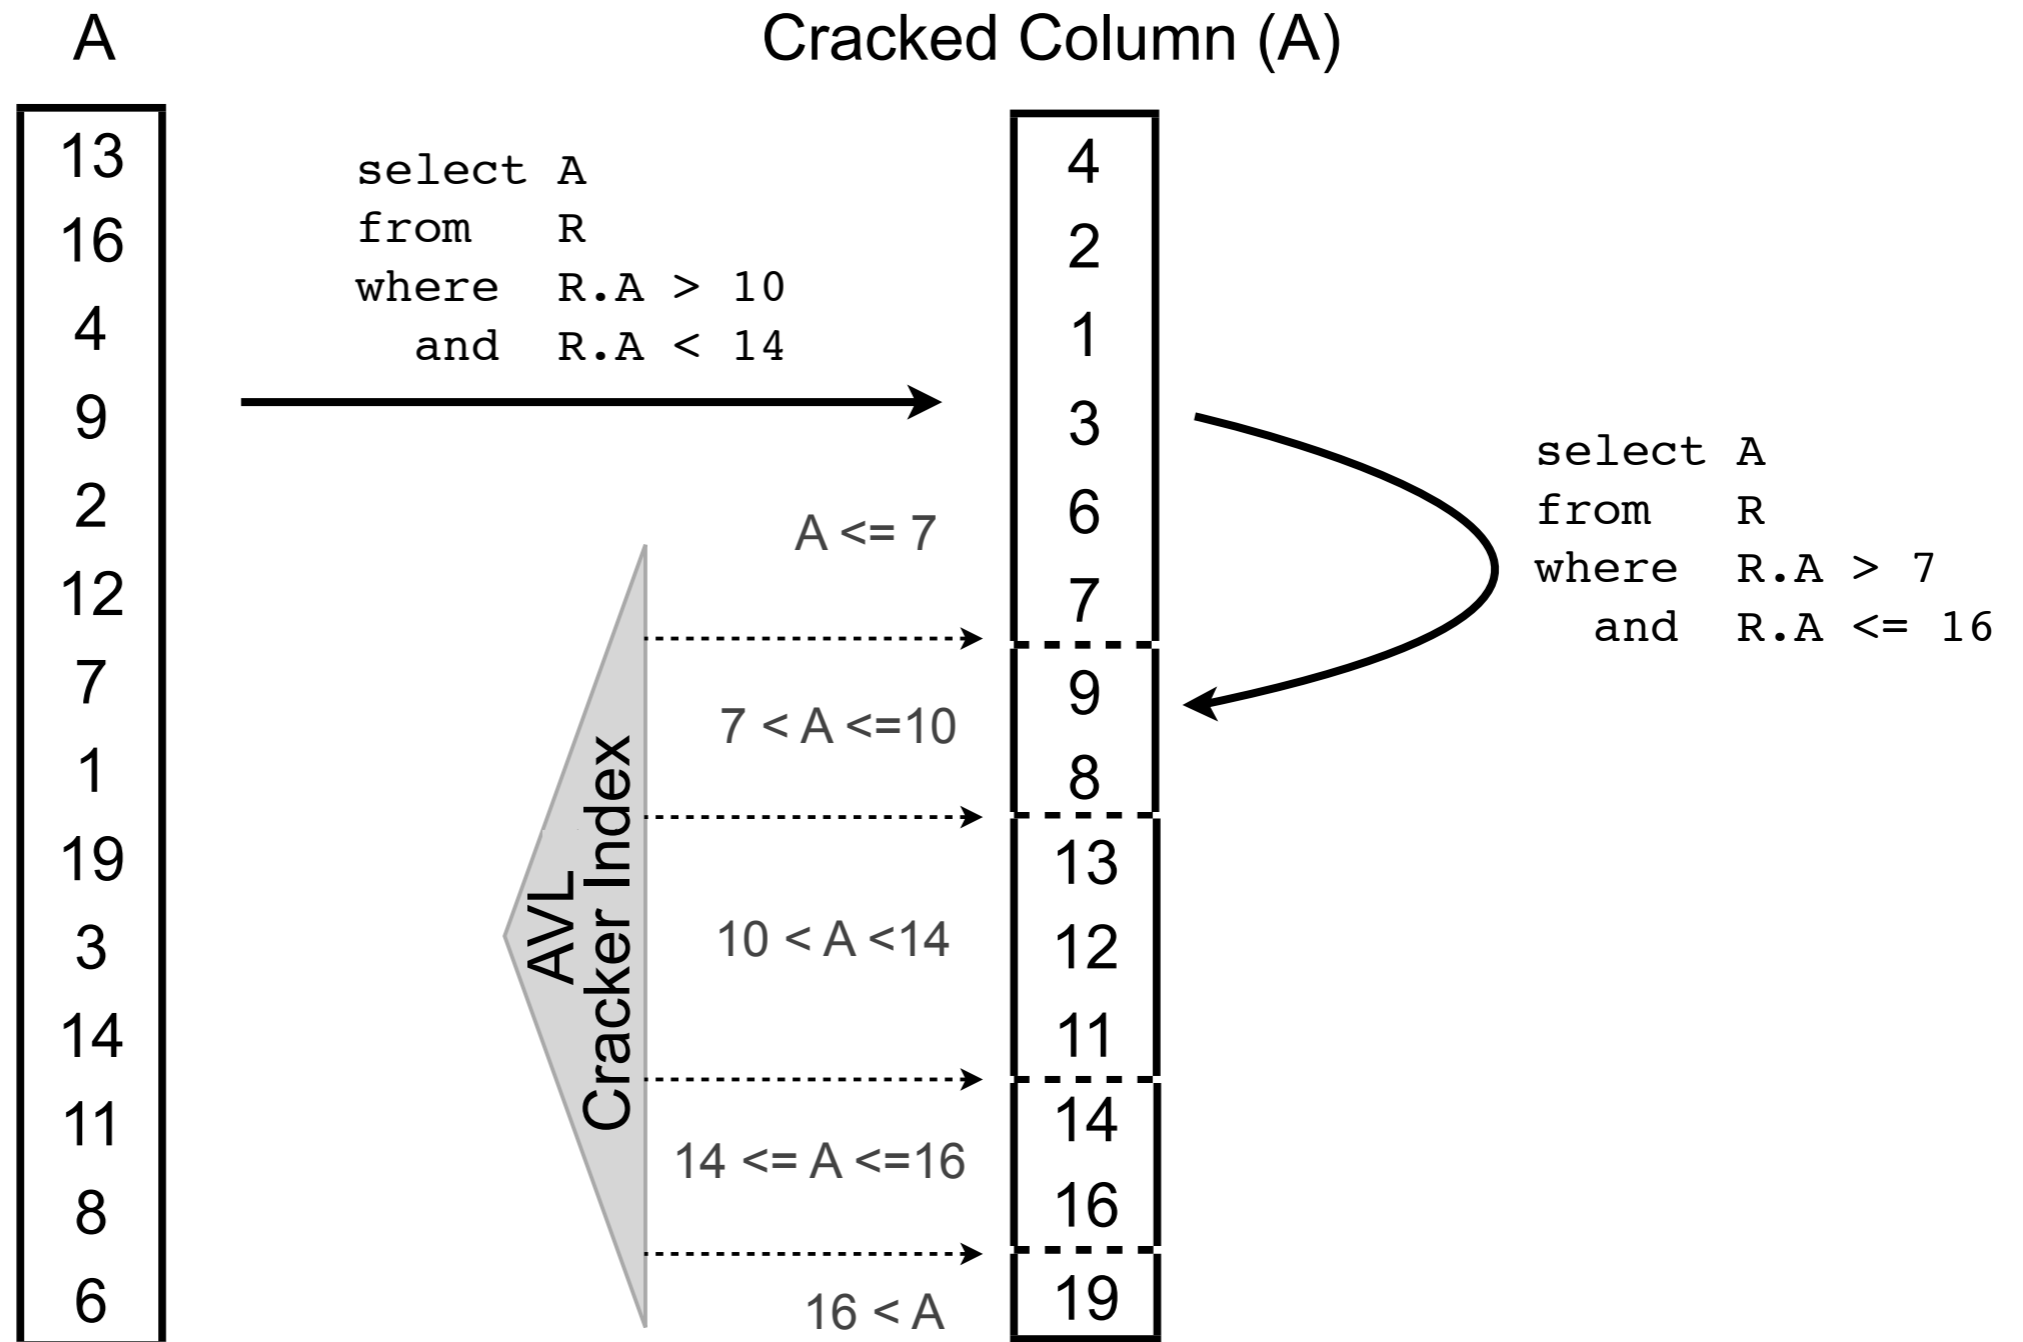
# Adaptive Indexing: Standard Cracking



# Adaptive Indexing: Standard Cracking



# Adaptive Indexing: Standard Cracking



# Motivation

Standard Cracking

# Motivation

Standard Cracking

Stochastic Cracking

Hybrid Cracking

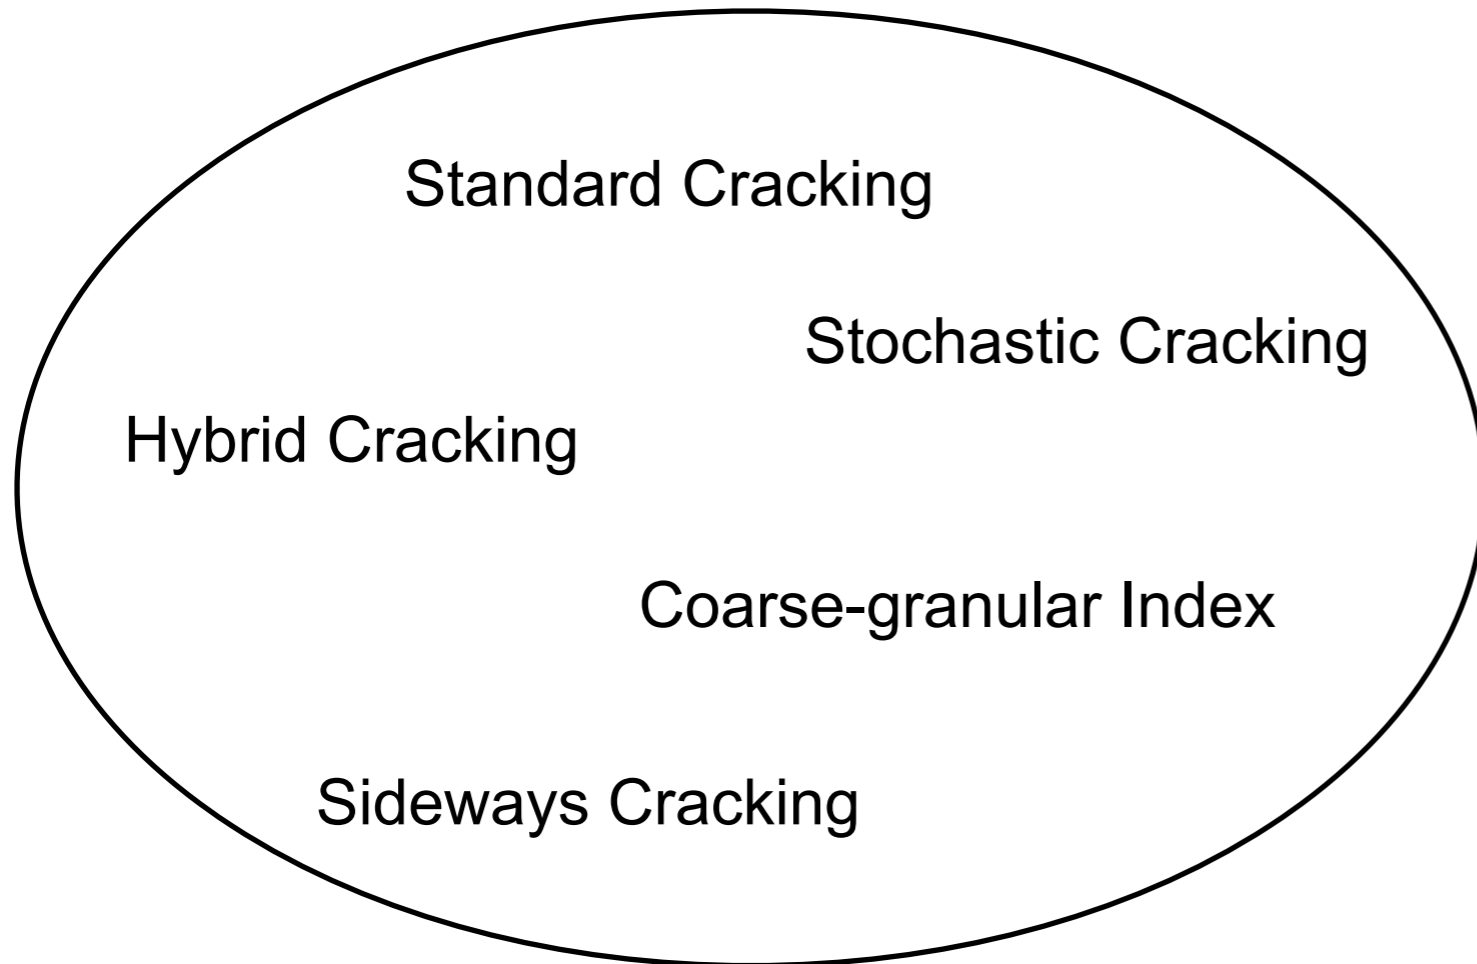
Coarse-granular Index

Sideways Cracking




# Motivation

Single-threaded adaptive algorithms



# Motivation

Single-threaded adaptive algorithms




Quicksort

Radixsort

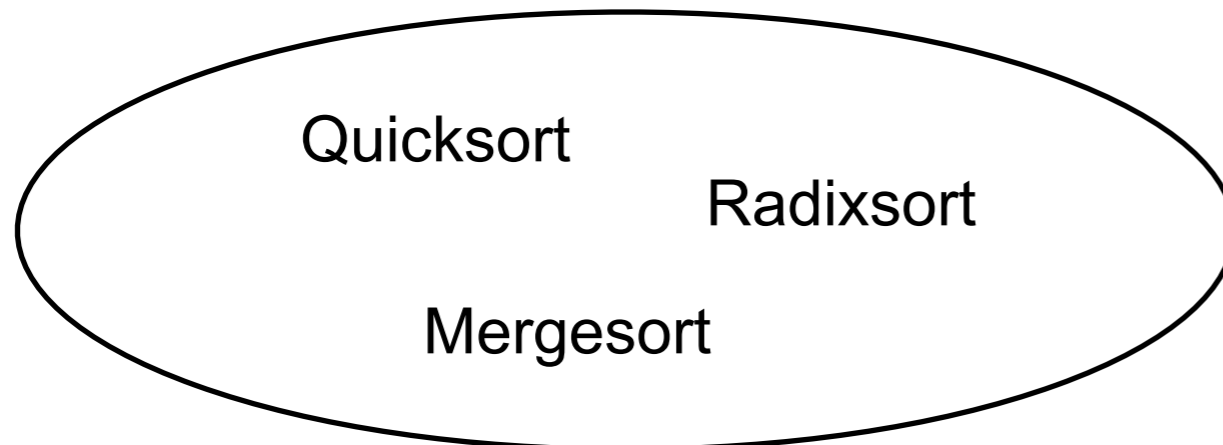
Mergesort

# Motivation

Single-threaded adaptive algorithms

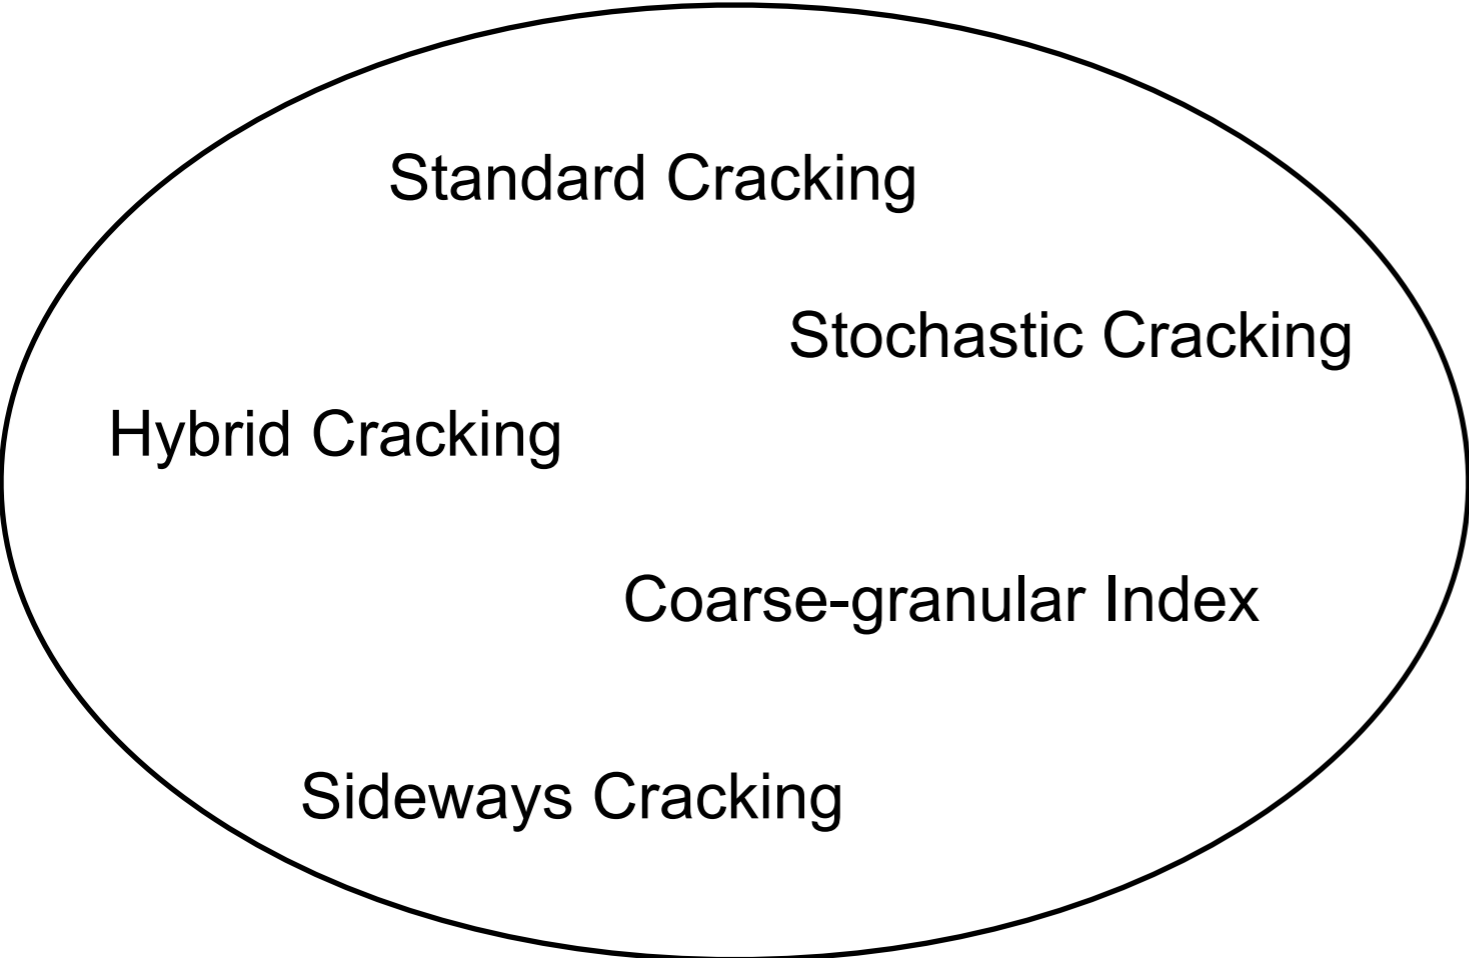


Single-threaded sorting algorithms




# Motivation

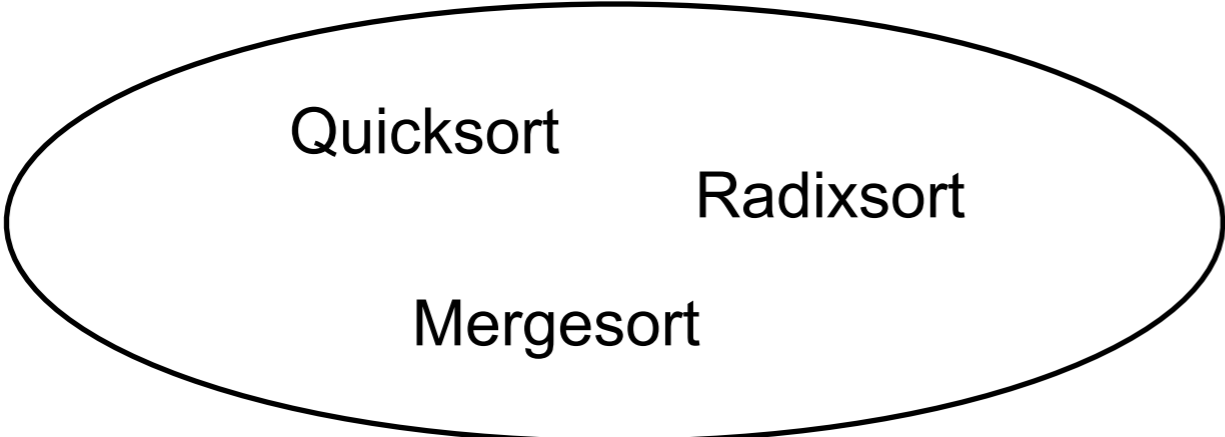
Single-threaded adaptive algorithms



Multi-threaded adaptive algorithms




Single-threaded sorting algorithms

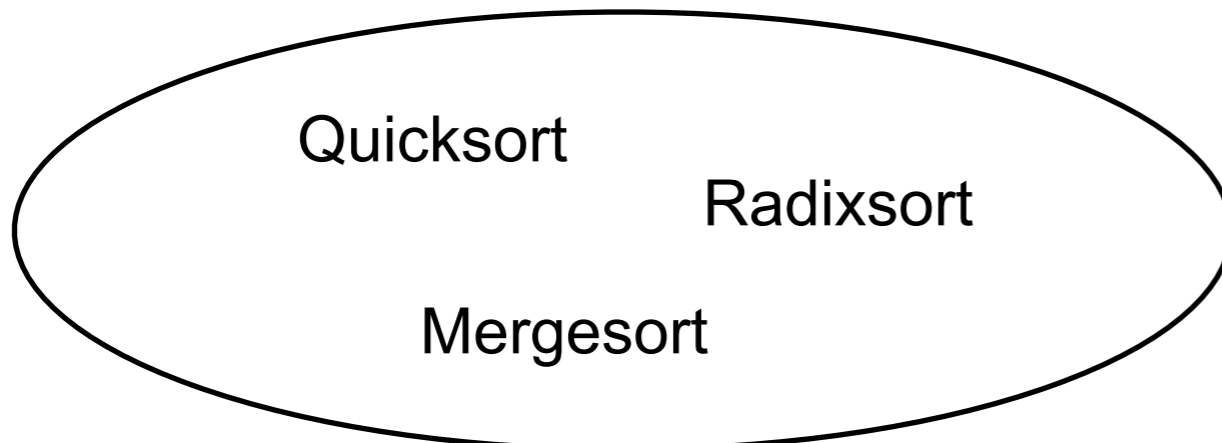


# Motivation



Single-threaded adaptive algorithms



Single-threaded sorting algorithms




Multi-threaded adaptive algorithms



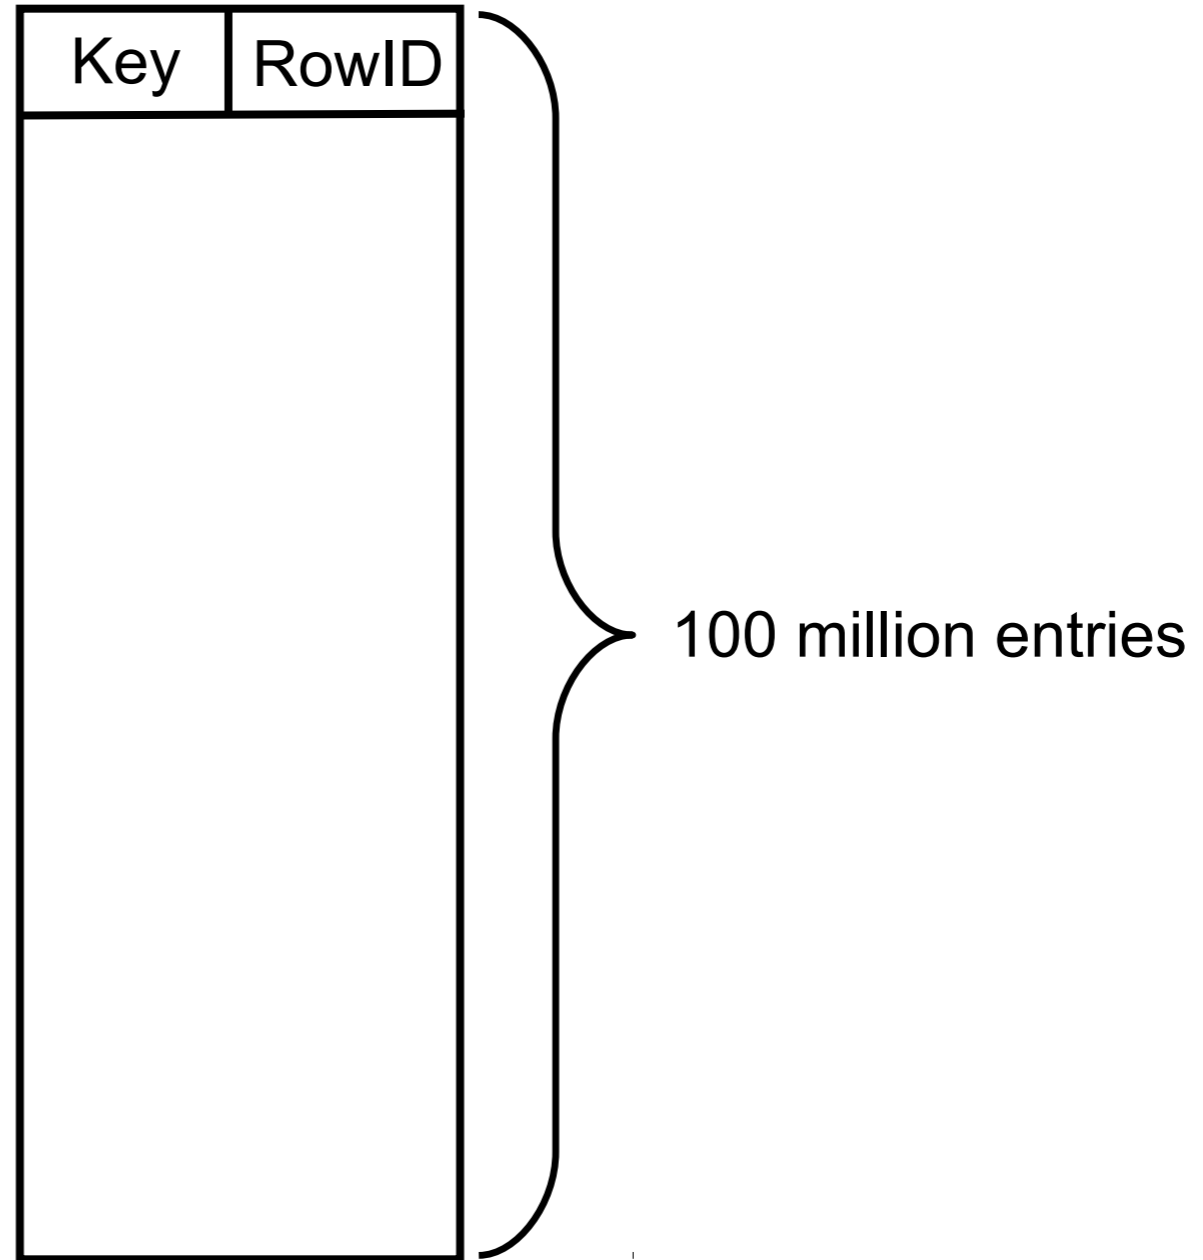
# Setup



# Setup



# Setup

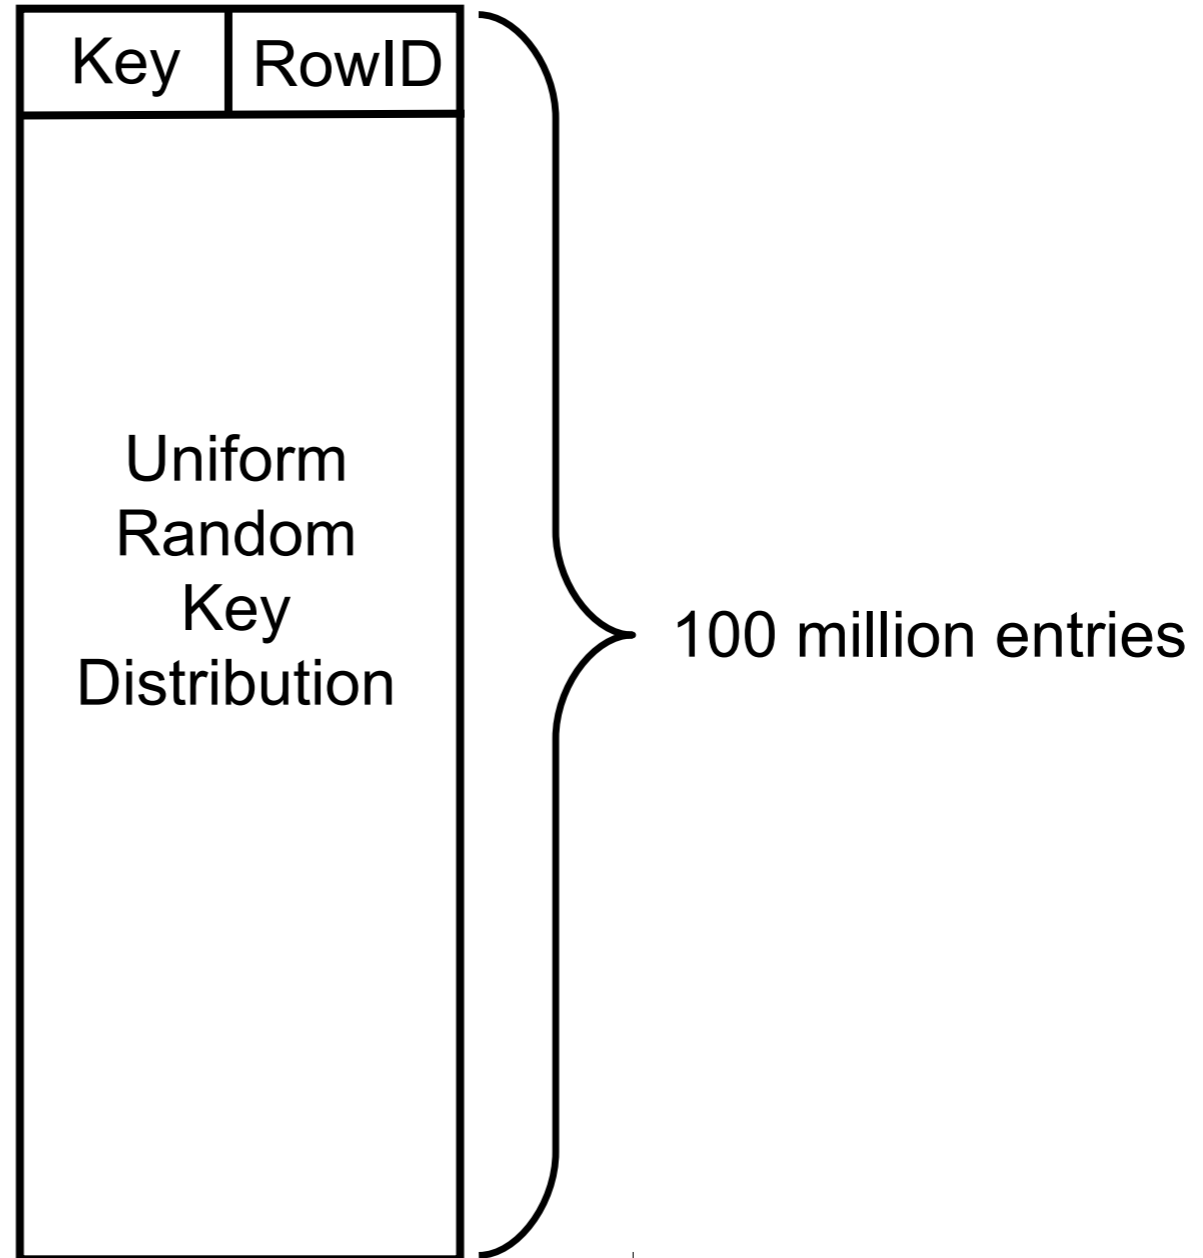




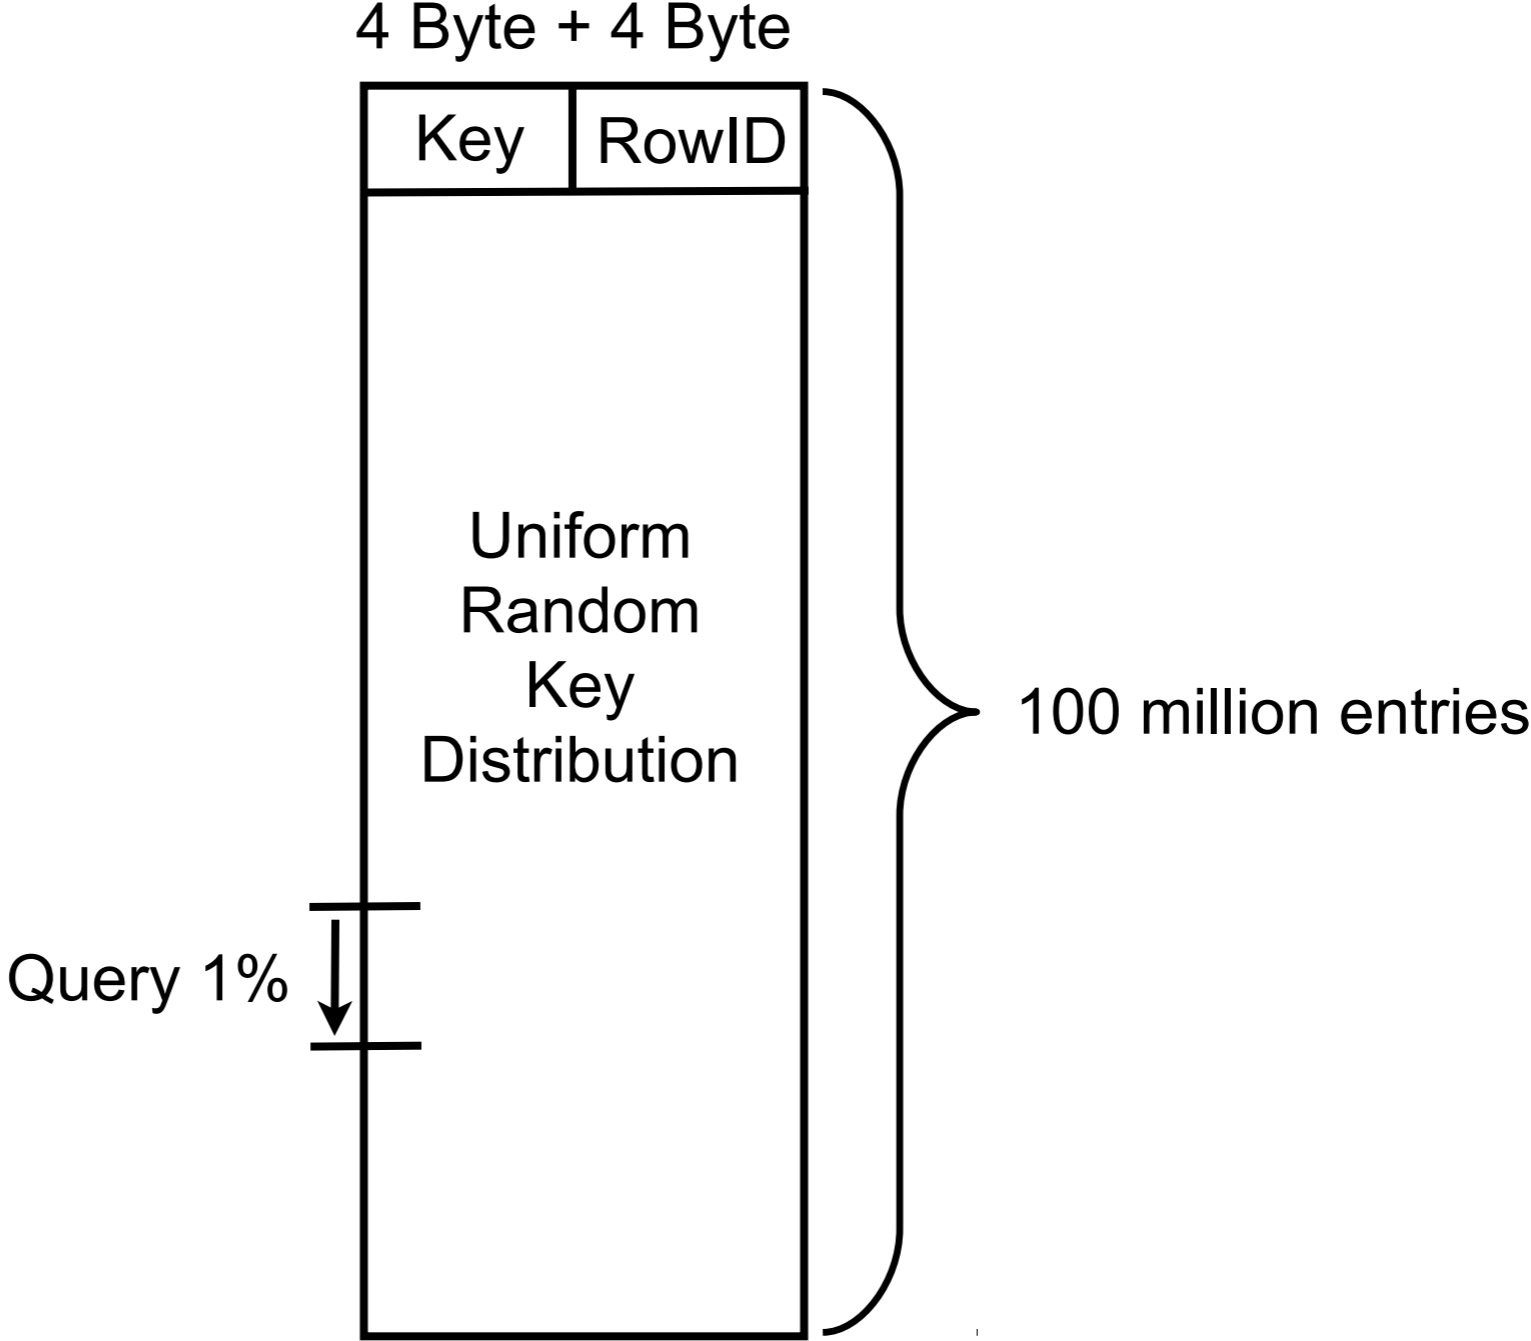


# Setup

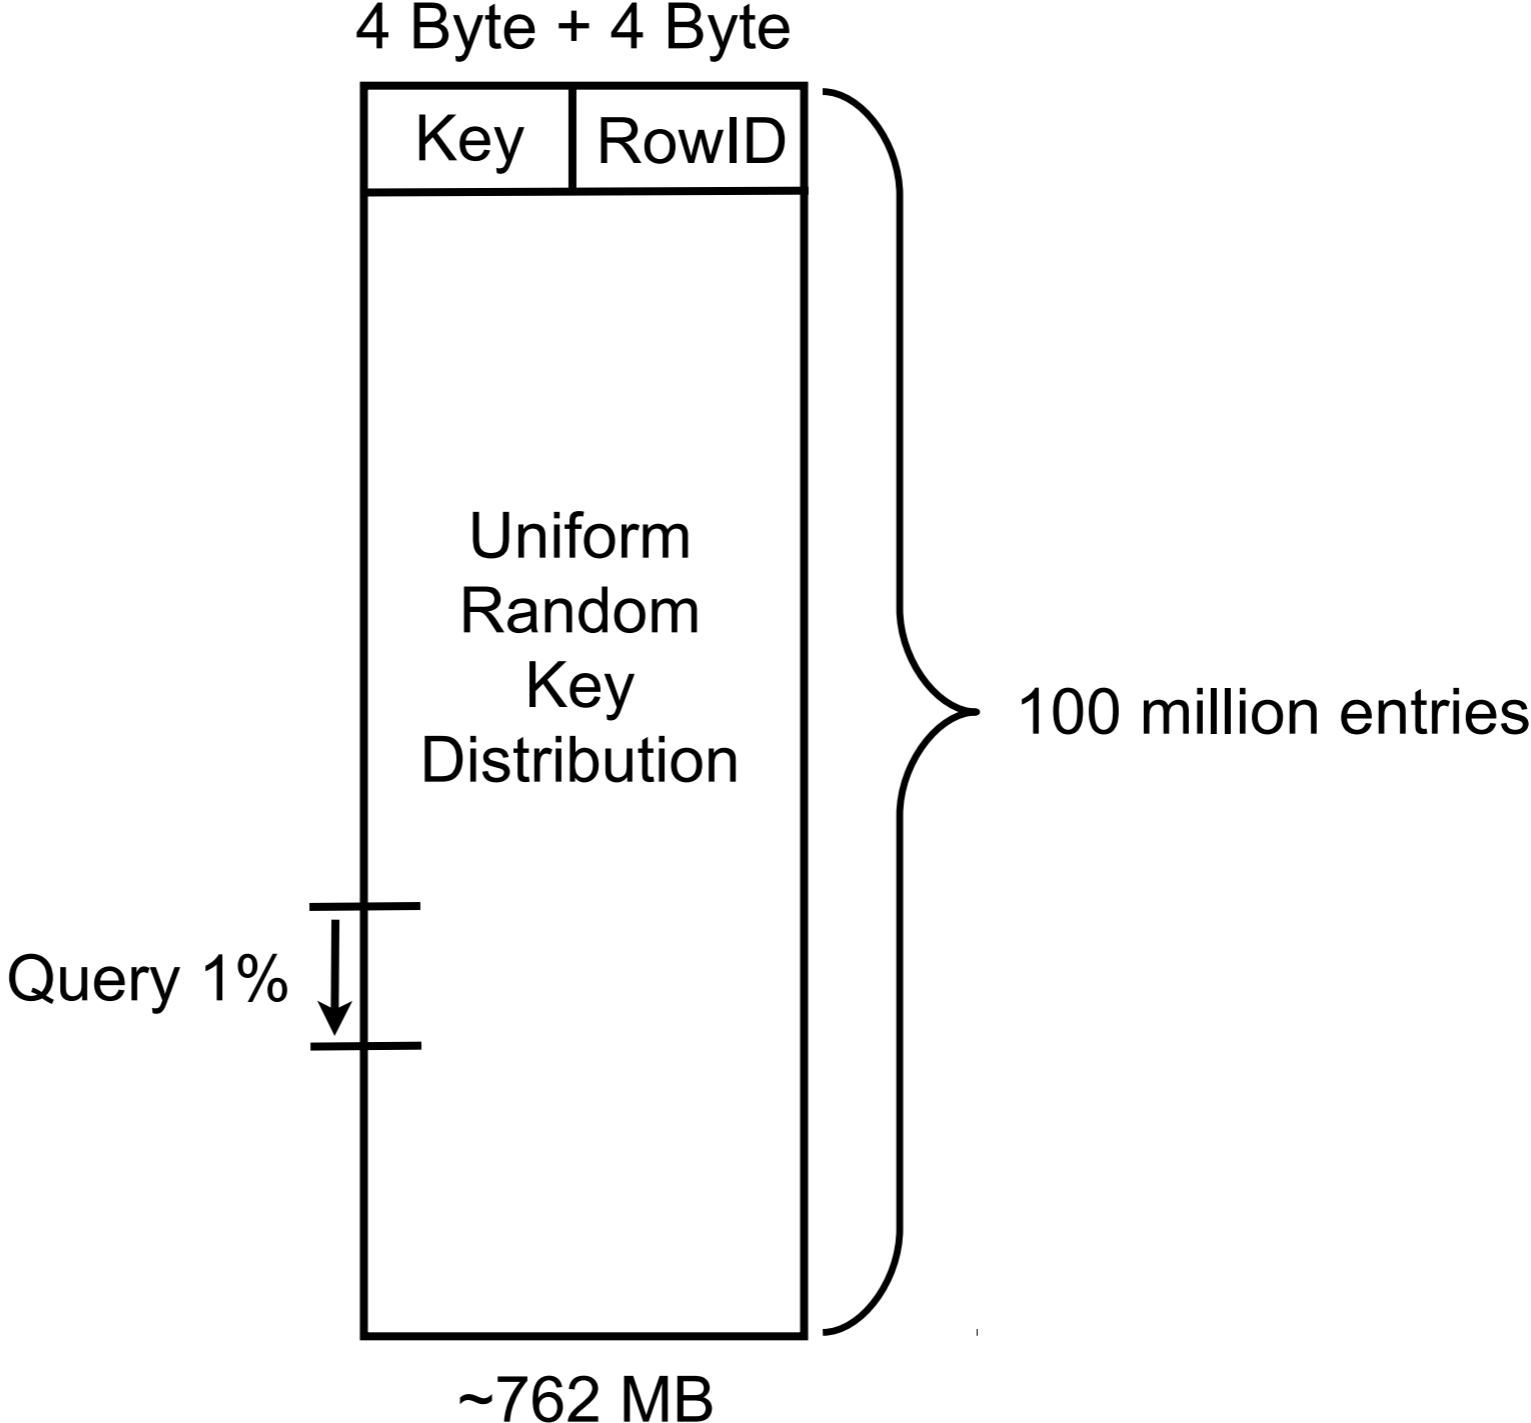
4 Byte + 4 Byte




# Setup



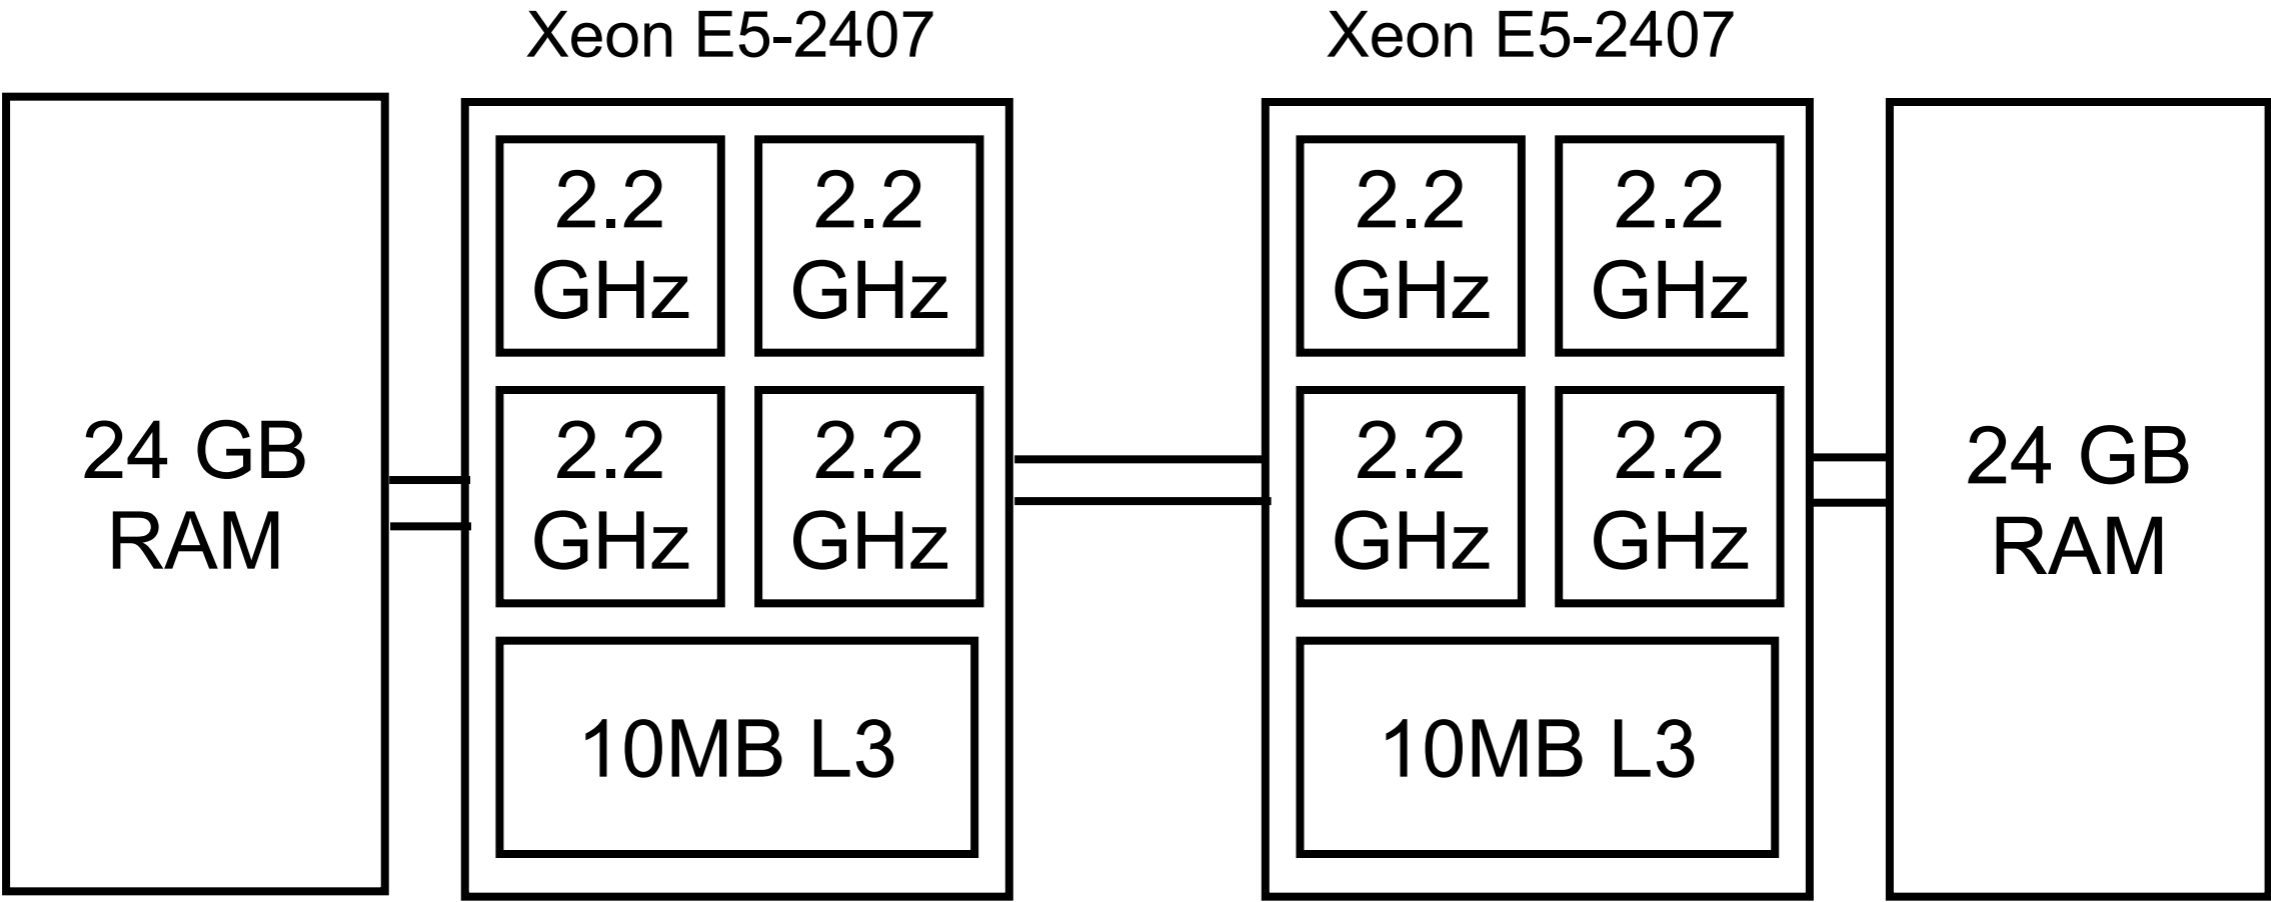
# Setup



# Setup




# Setup




↖  
No Turbo, no HyperThreading

# Single-threaded algorithms




# Single-threaded algorithms





# Single-threaded algorithms




Multi-threaded environments?




Multi-threaded algorithms!

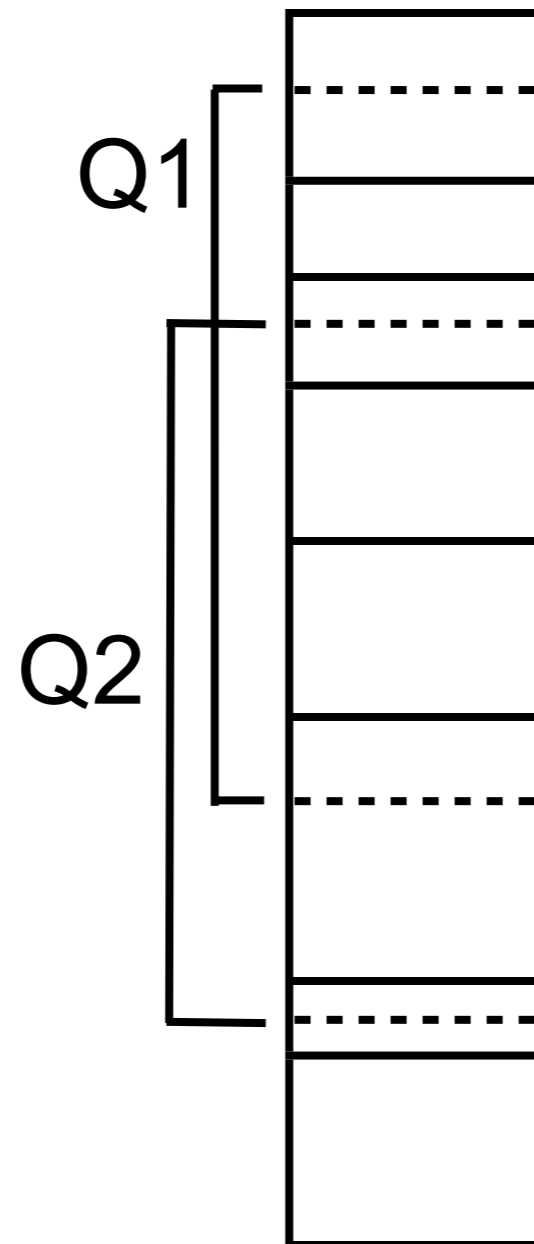
# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)




# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)




# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)



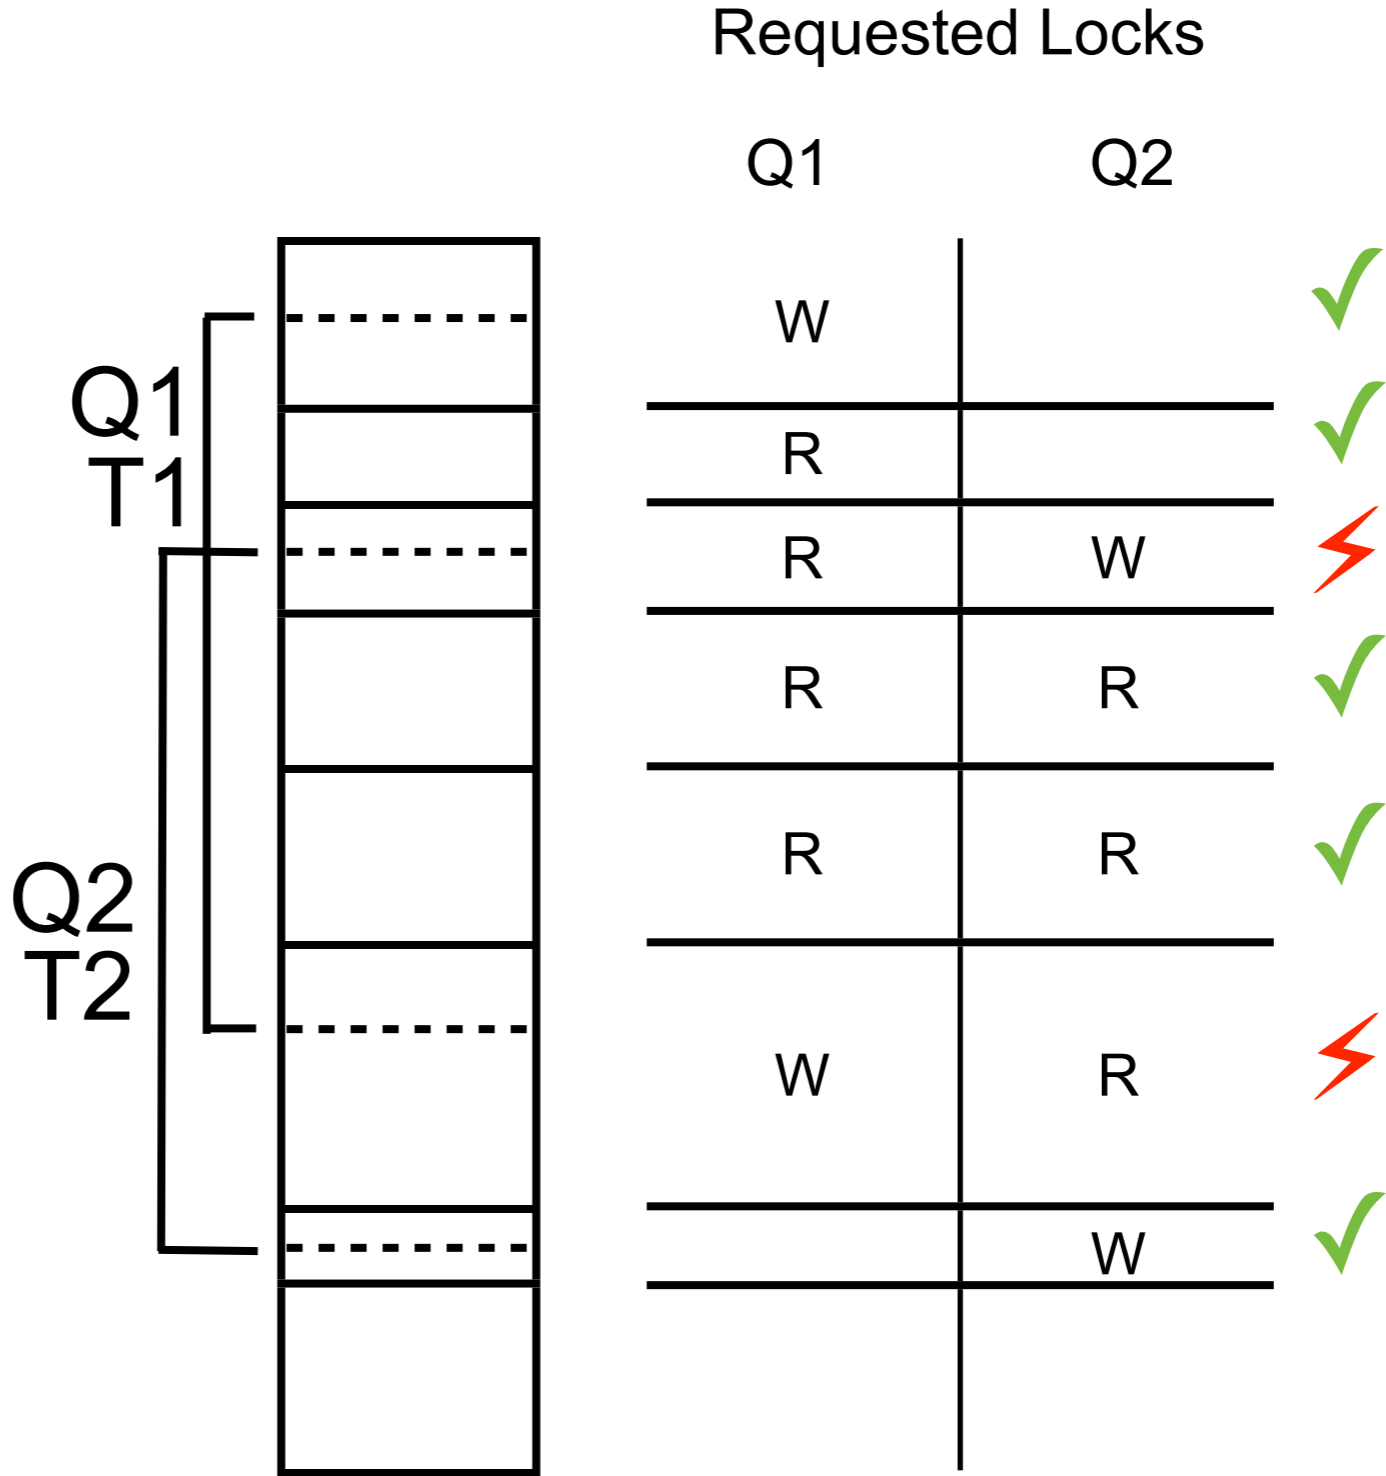
# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)



# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)

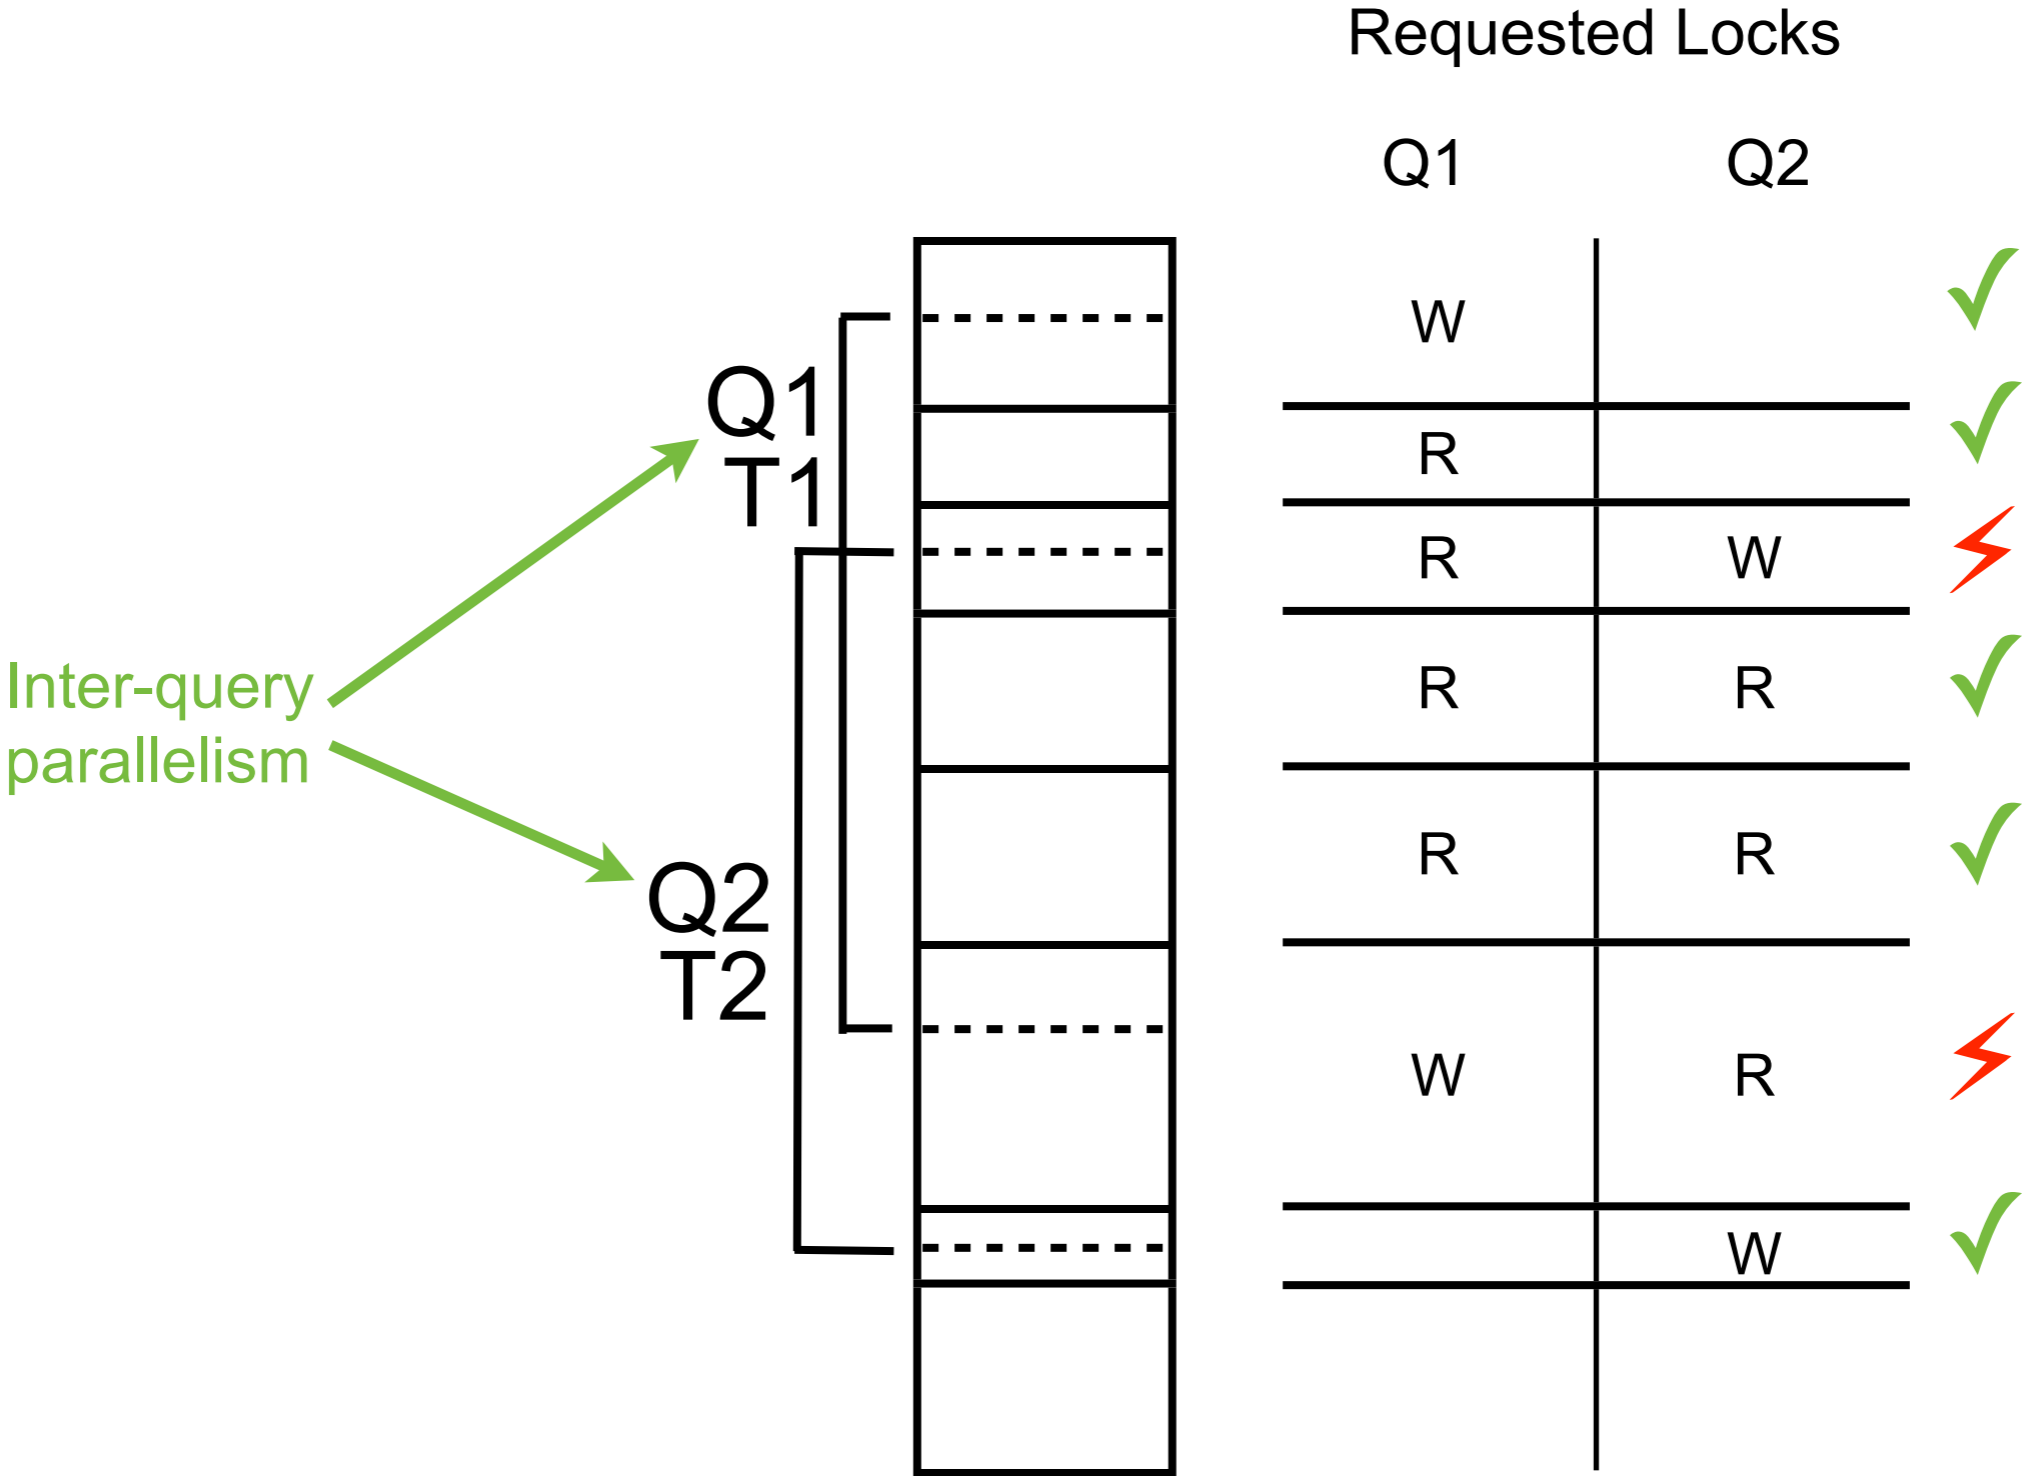


# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)

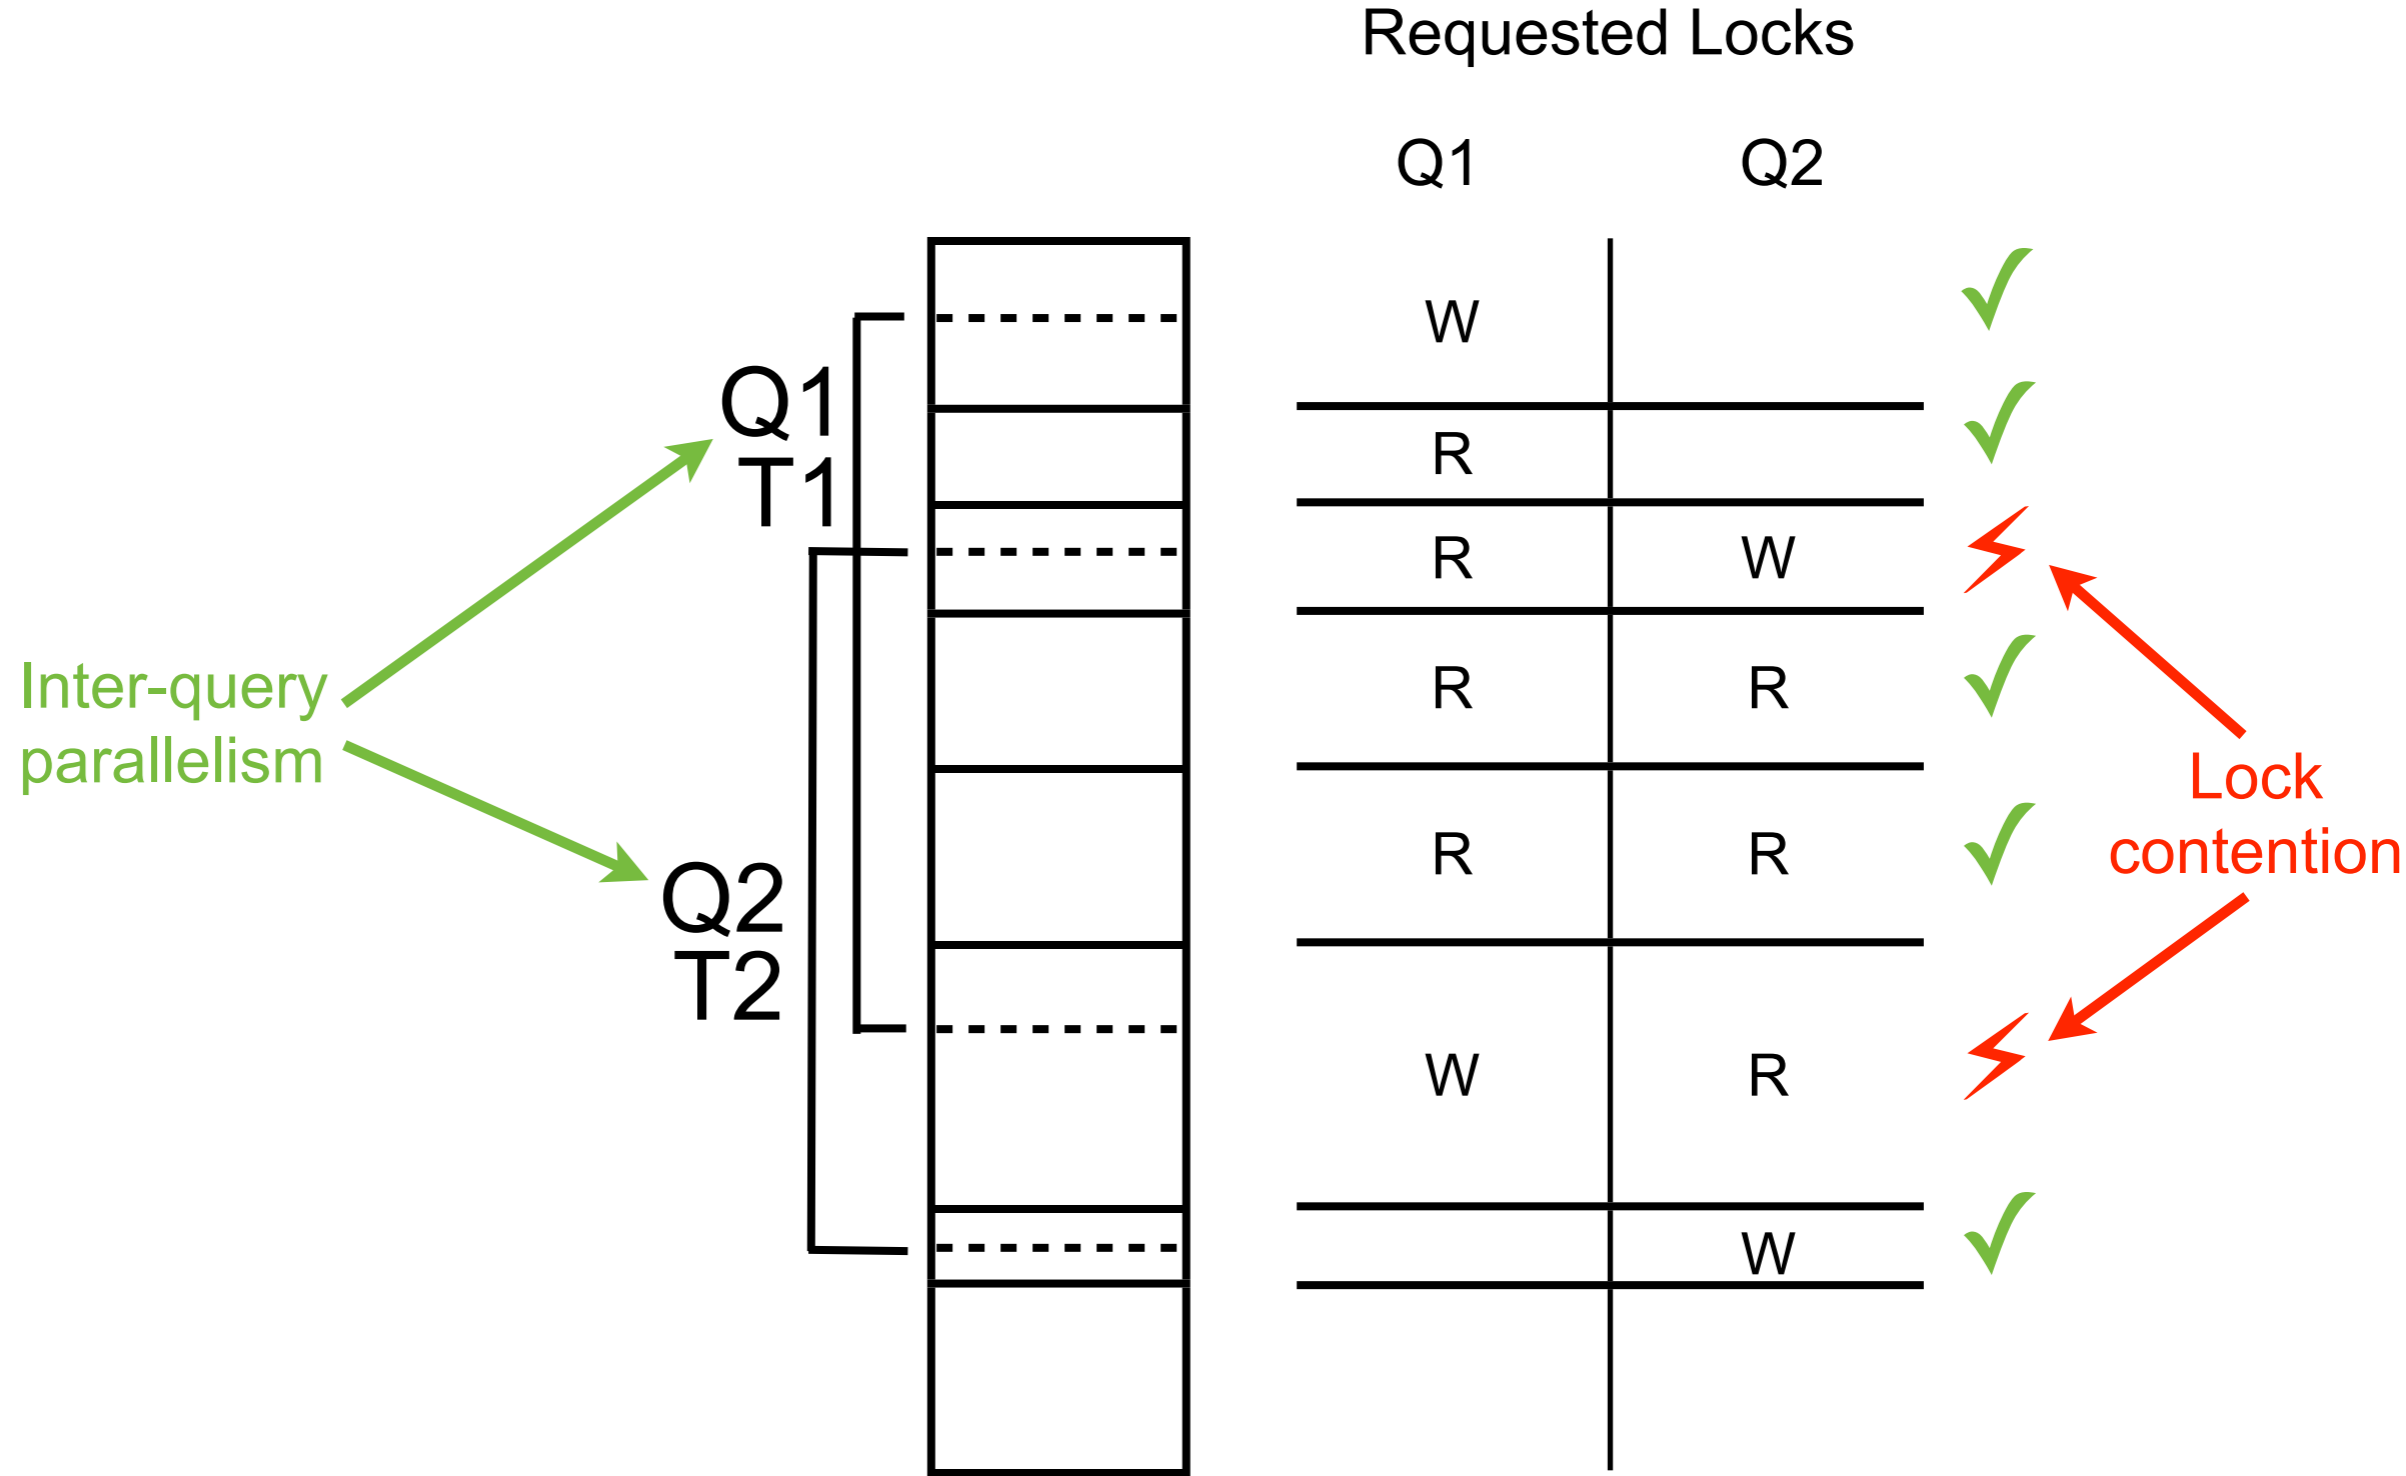




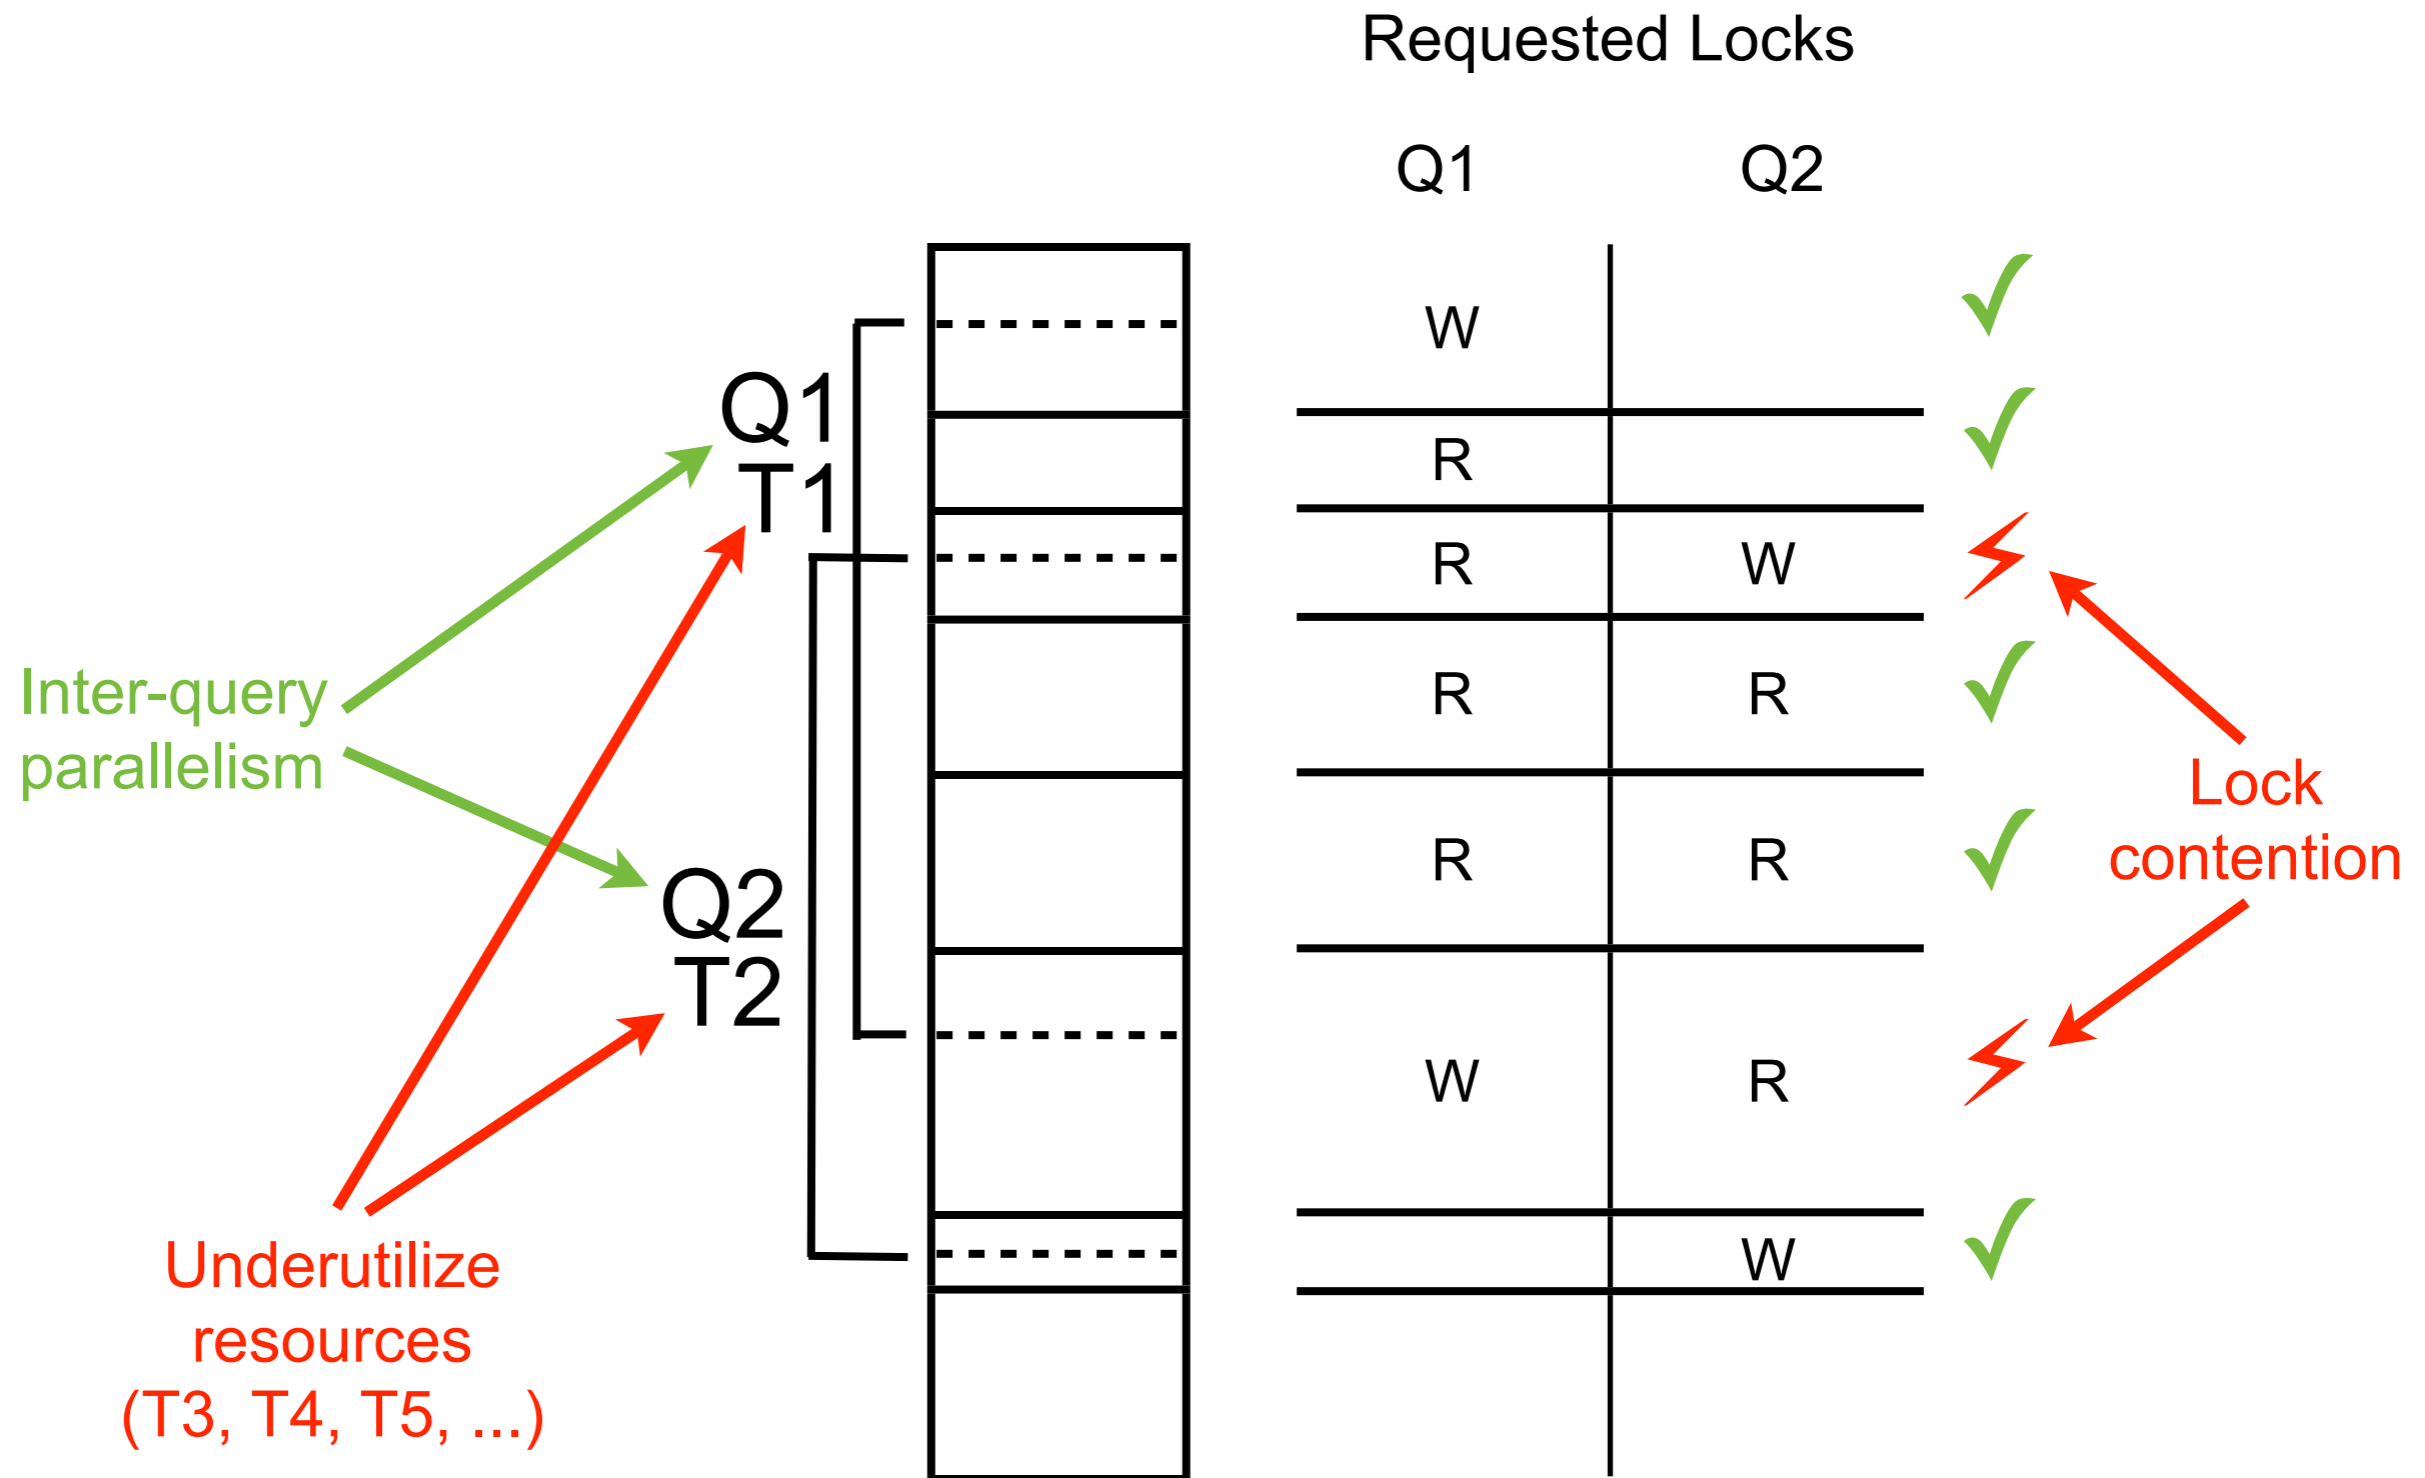
# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)



# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)




# Multi-threaded algorithms: Parallel Standard Cracking (P-SC)




# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)

Query




# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)

Query




k Chunks

# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)



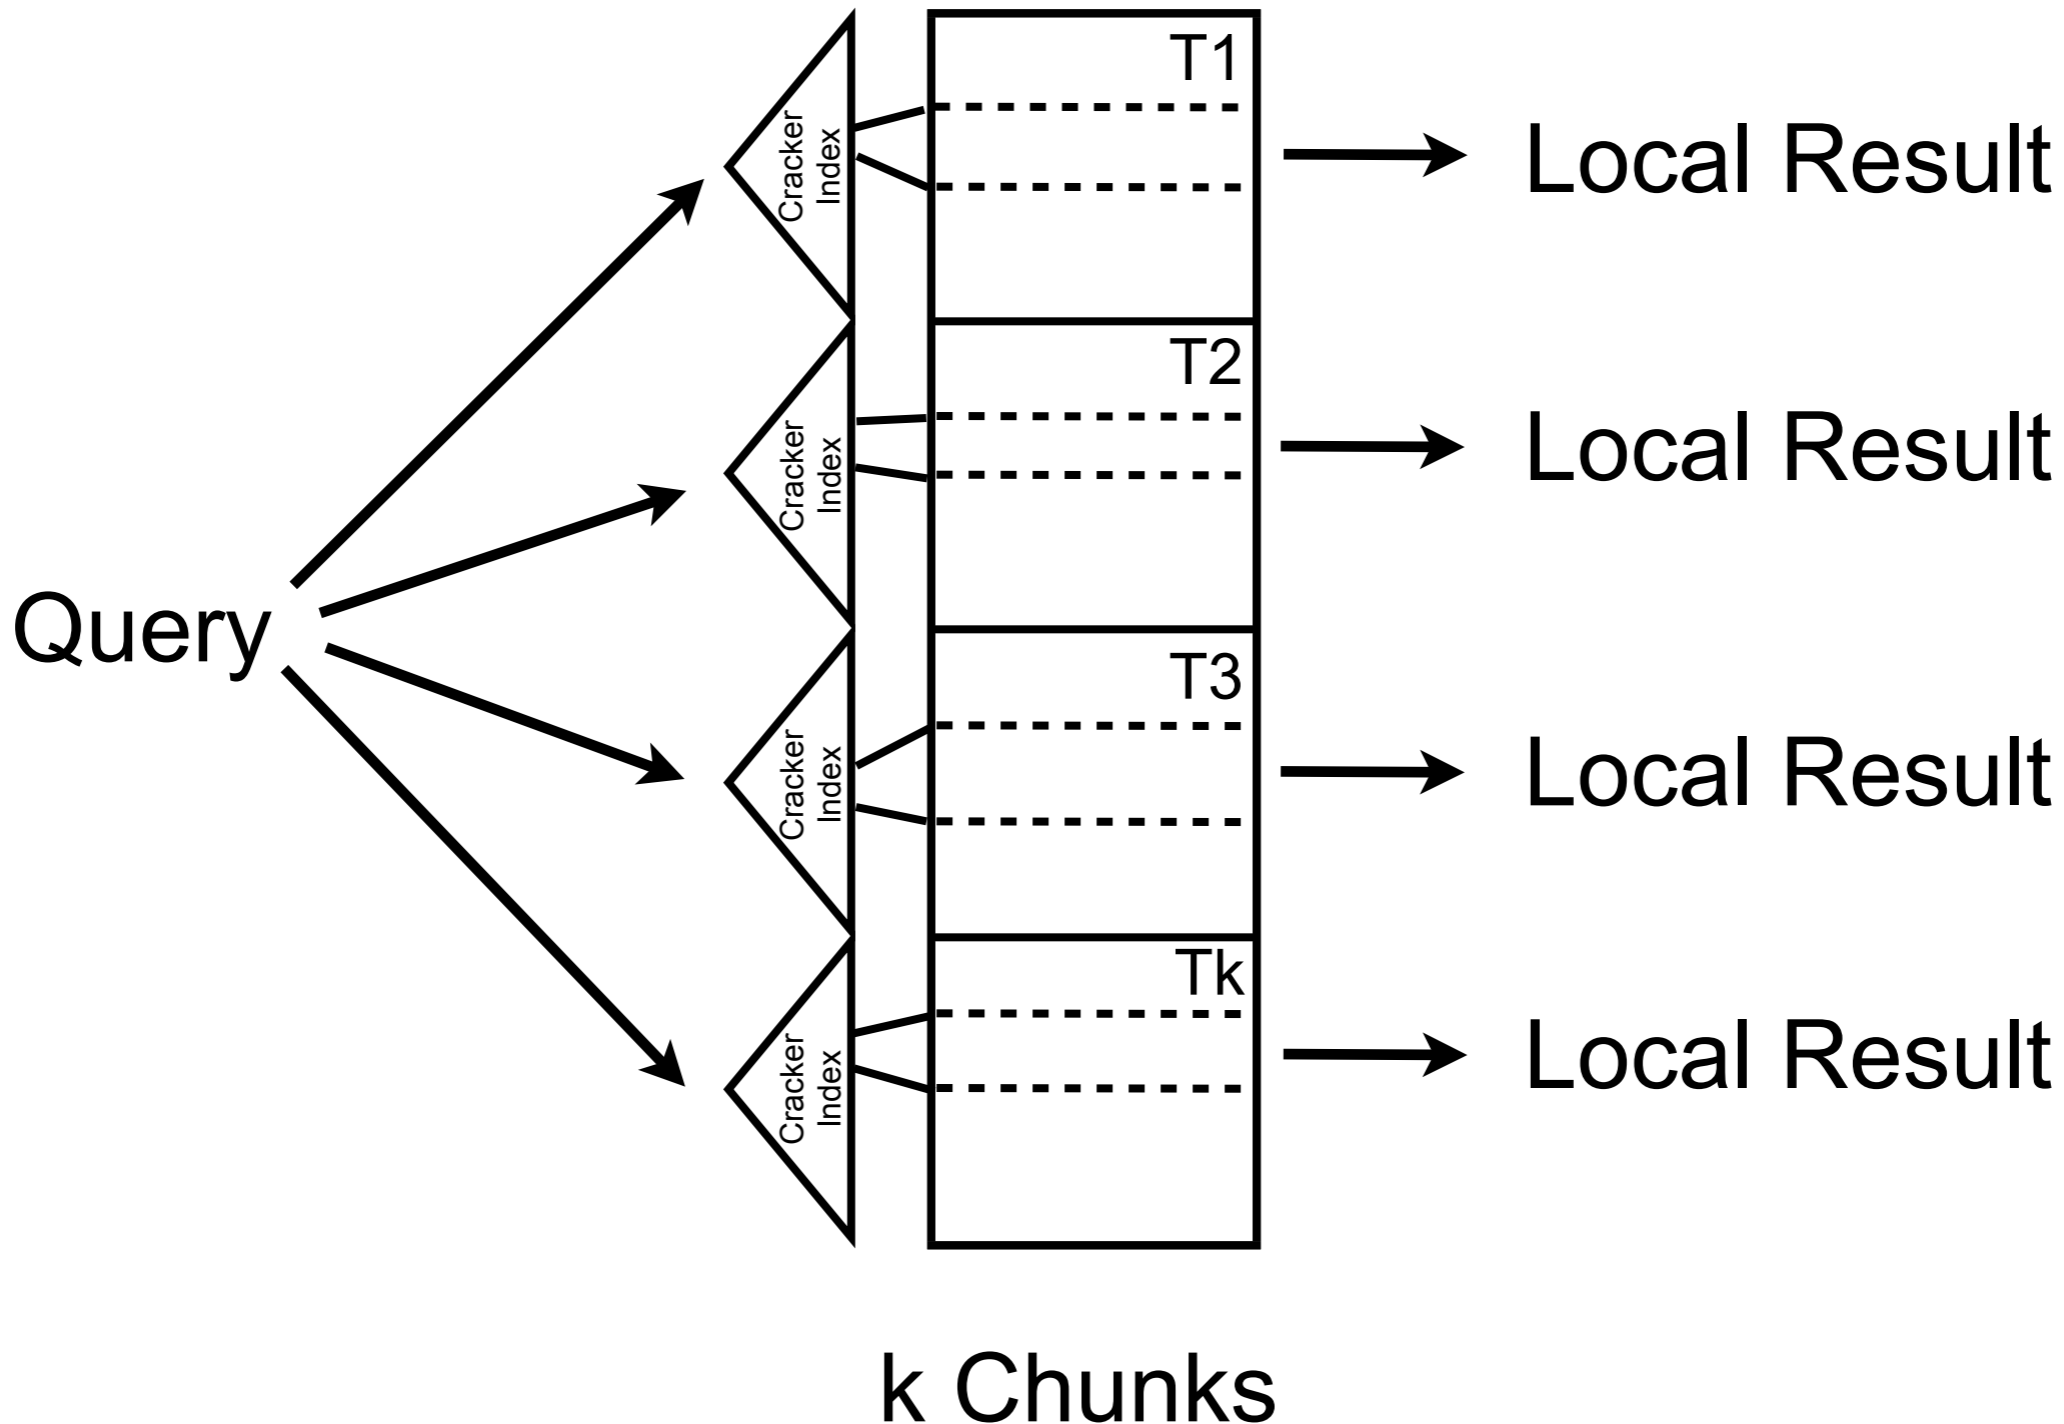
k Chunks

# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)




k Chunks

# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)






# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)




k Chunks

# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)

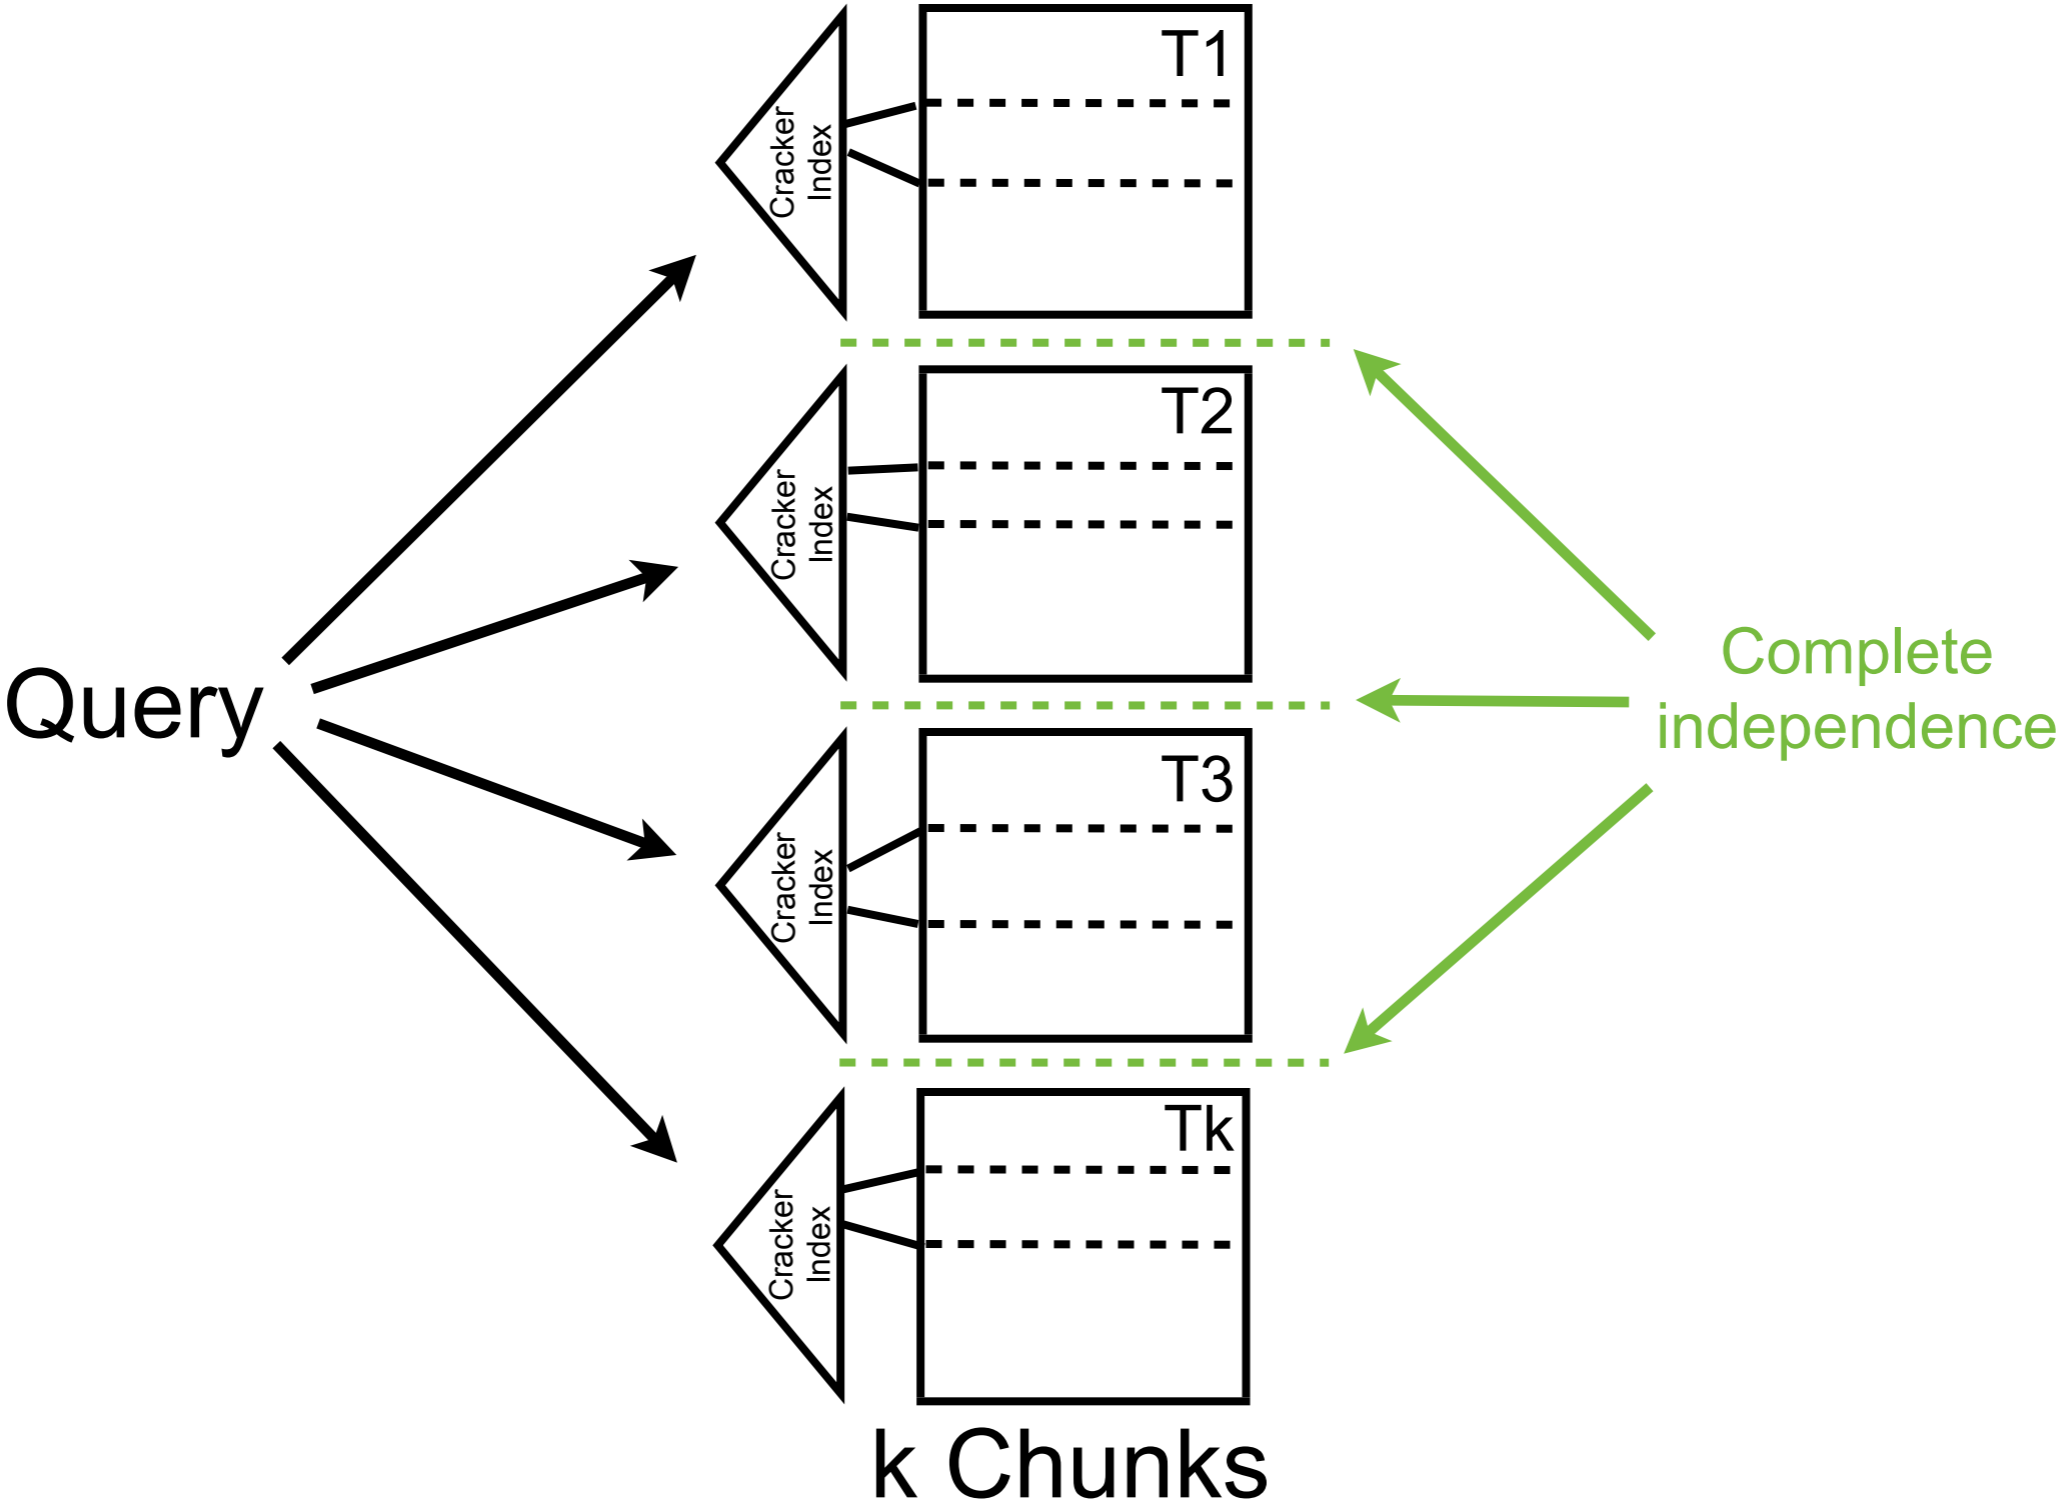


k Chunks


# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)



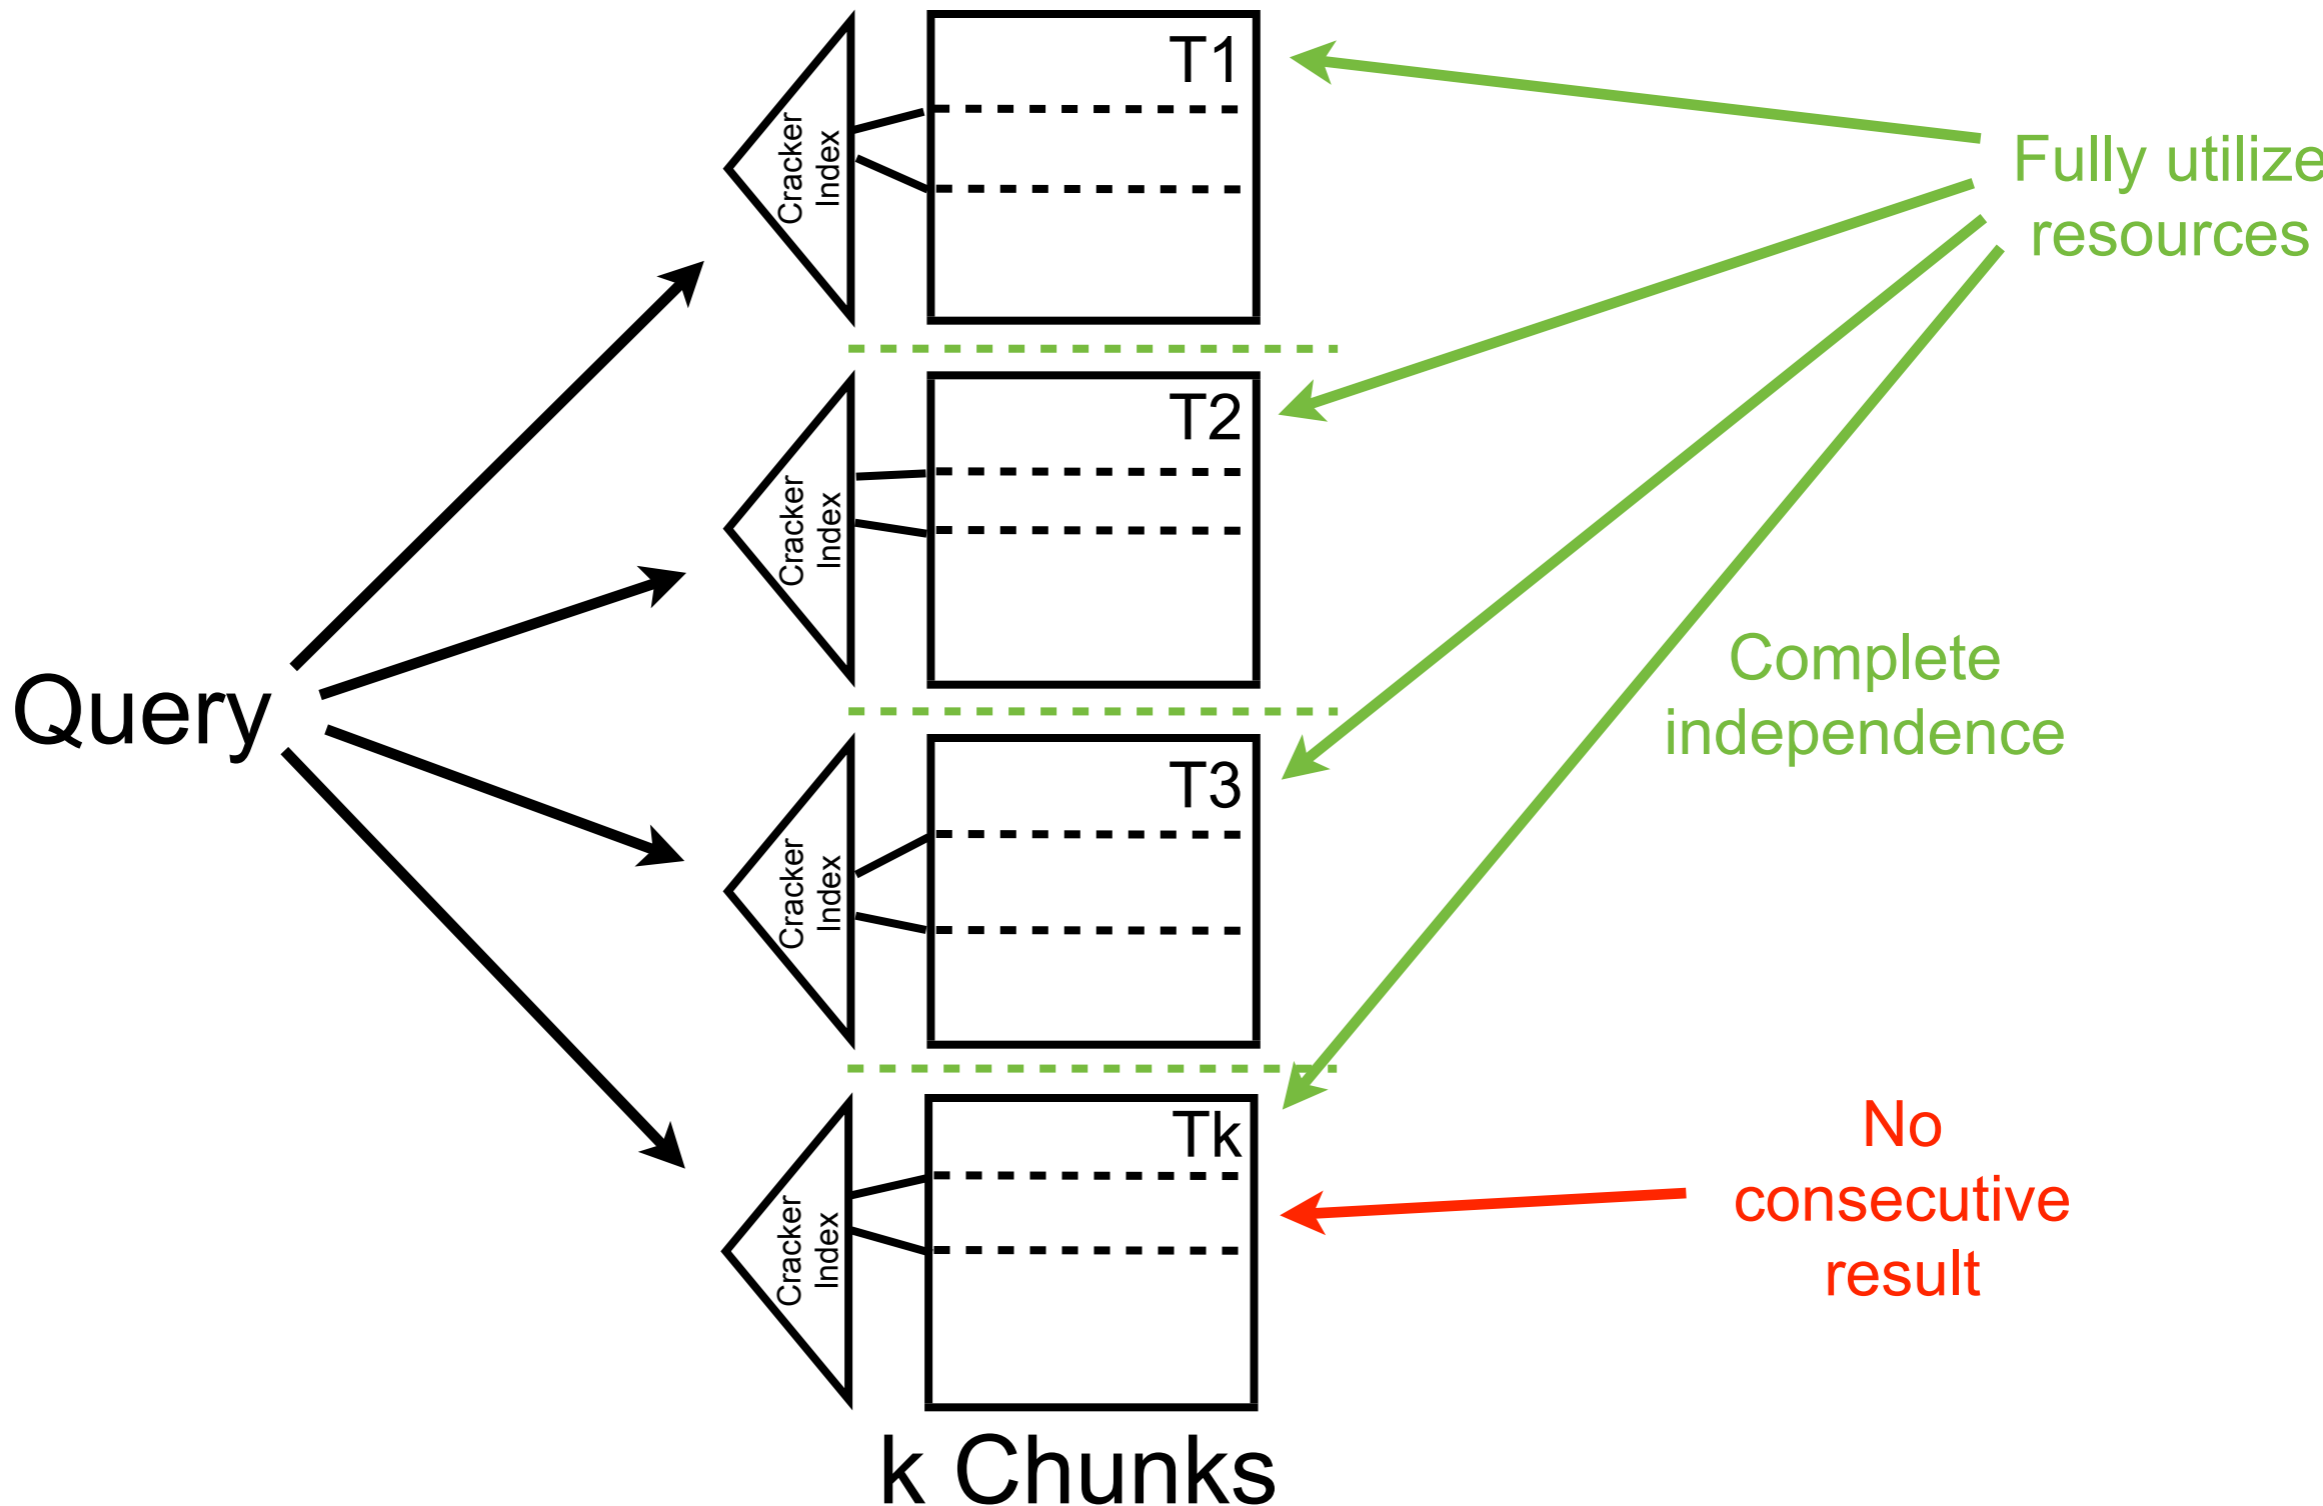
# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)



# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)

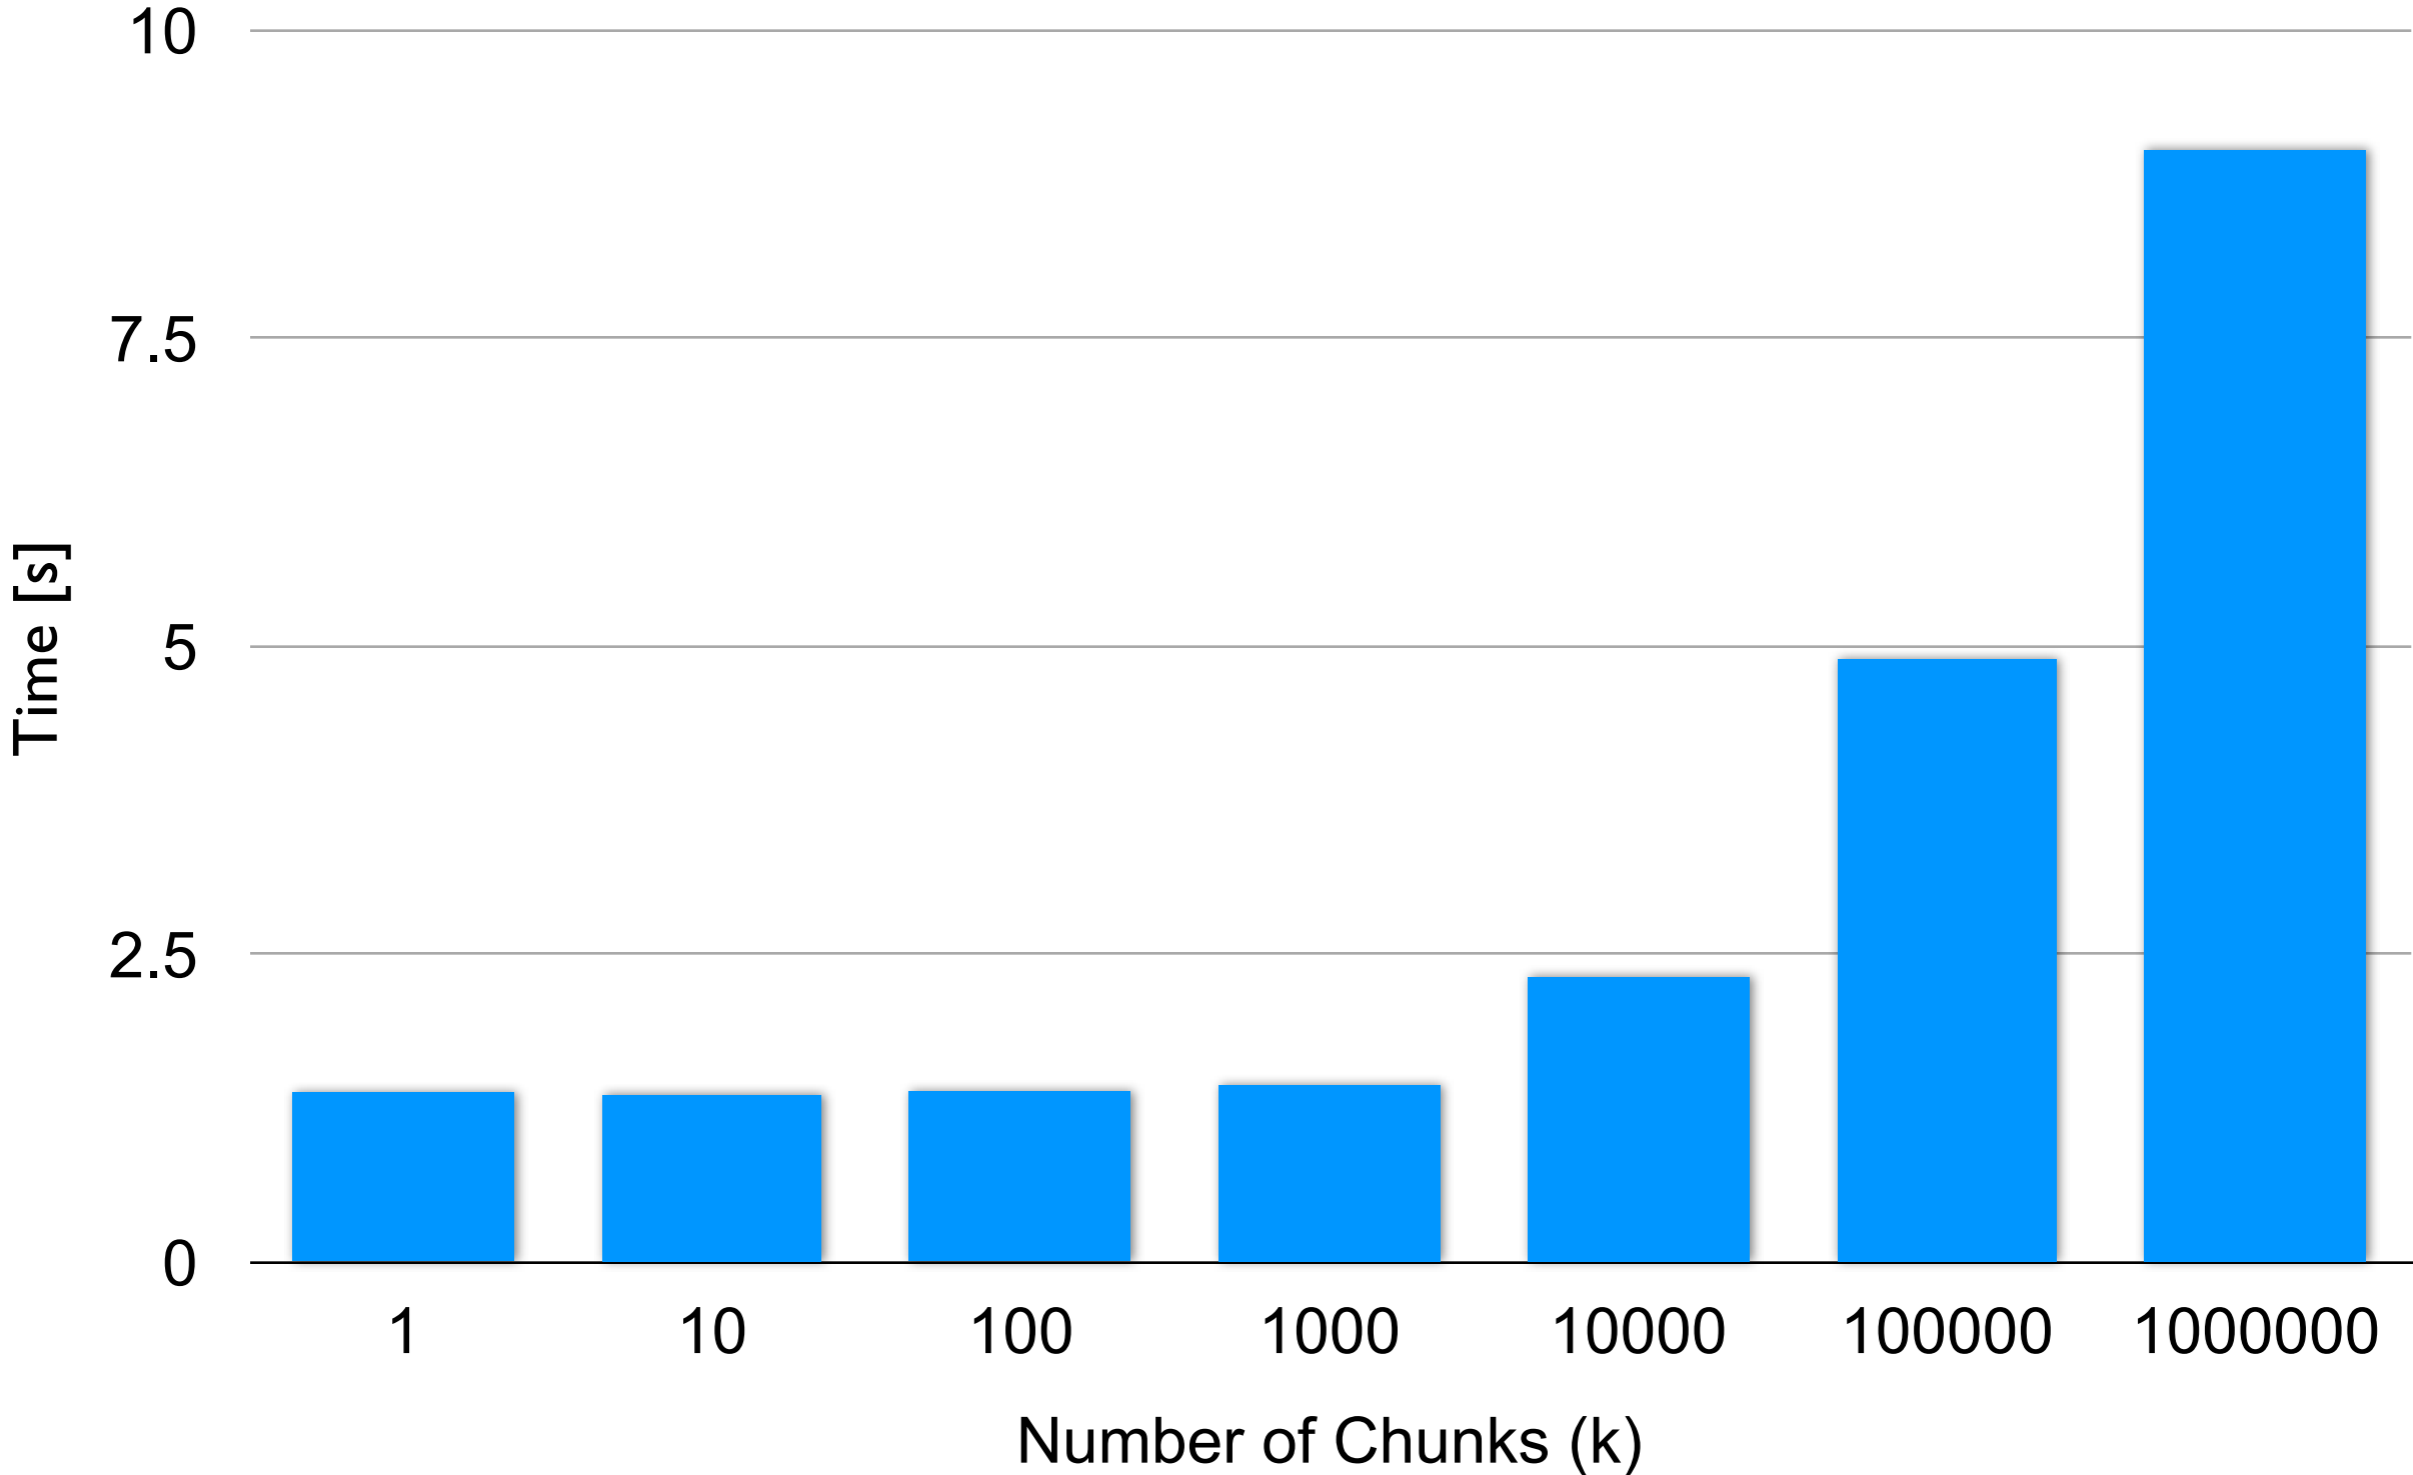


# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)




# Micro Benchmark

Reading 1% from k locations using one thread




# Micro Benchmark

Reading 1% from k locations using one thread






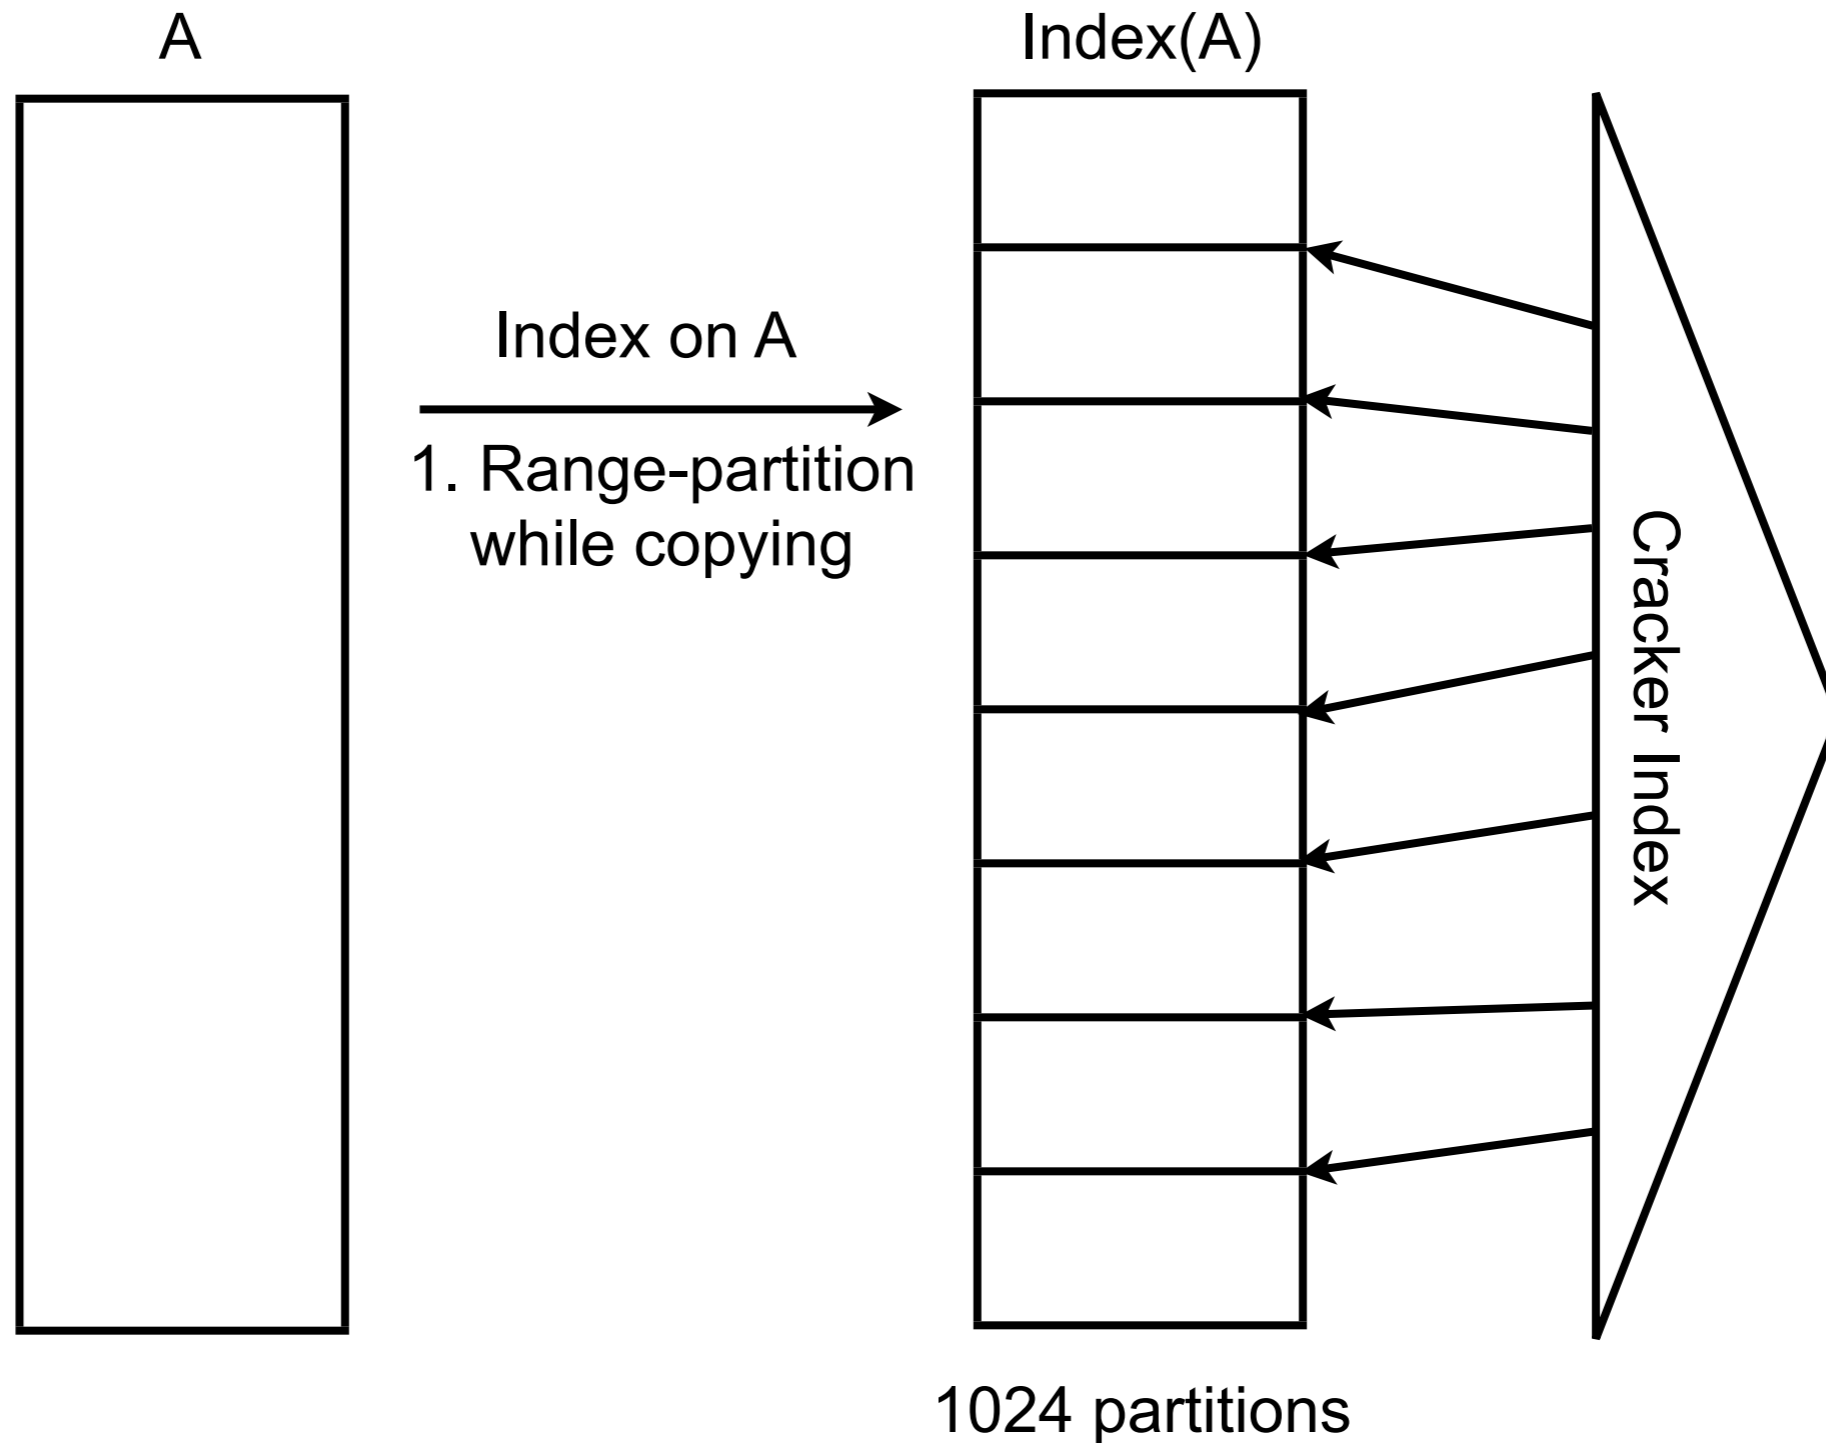
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



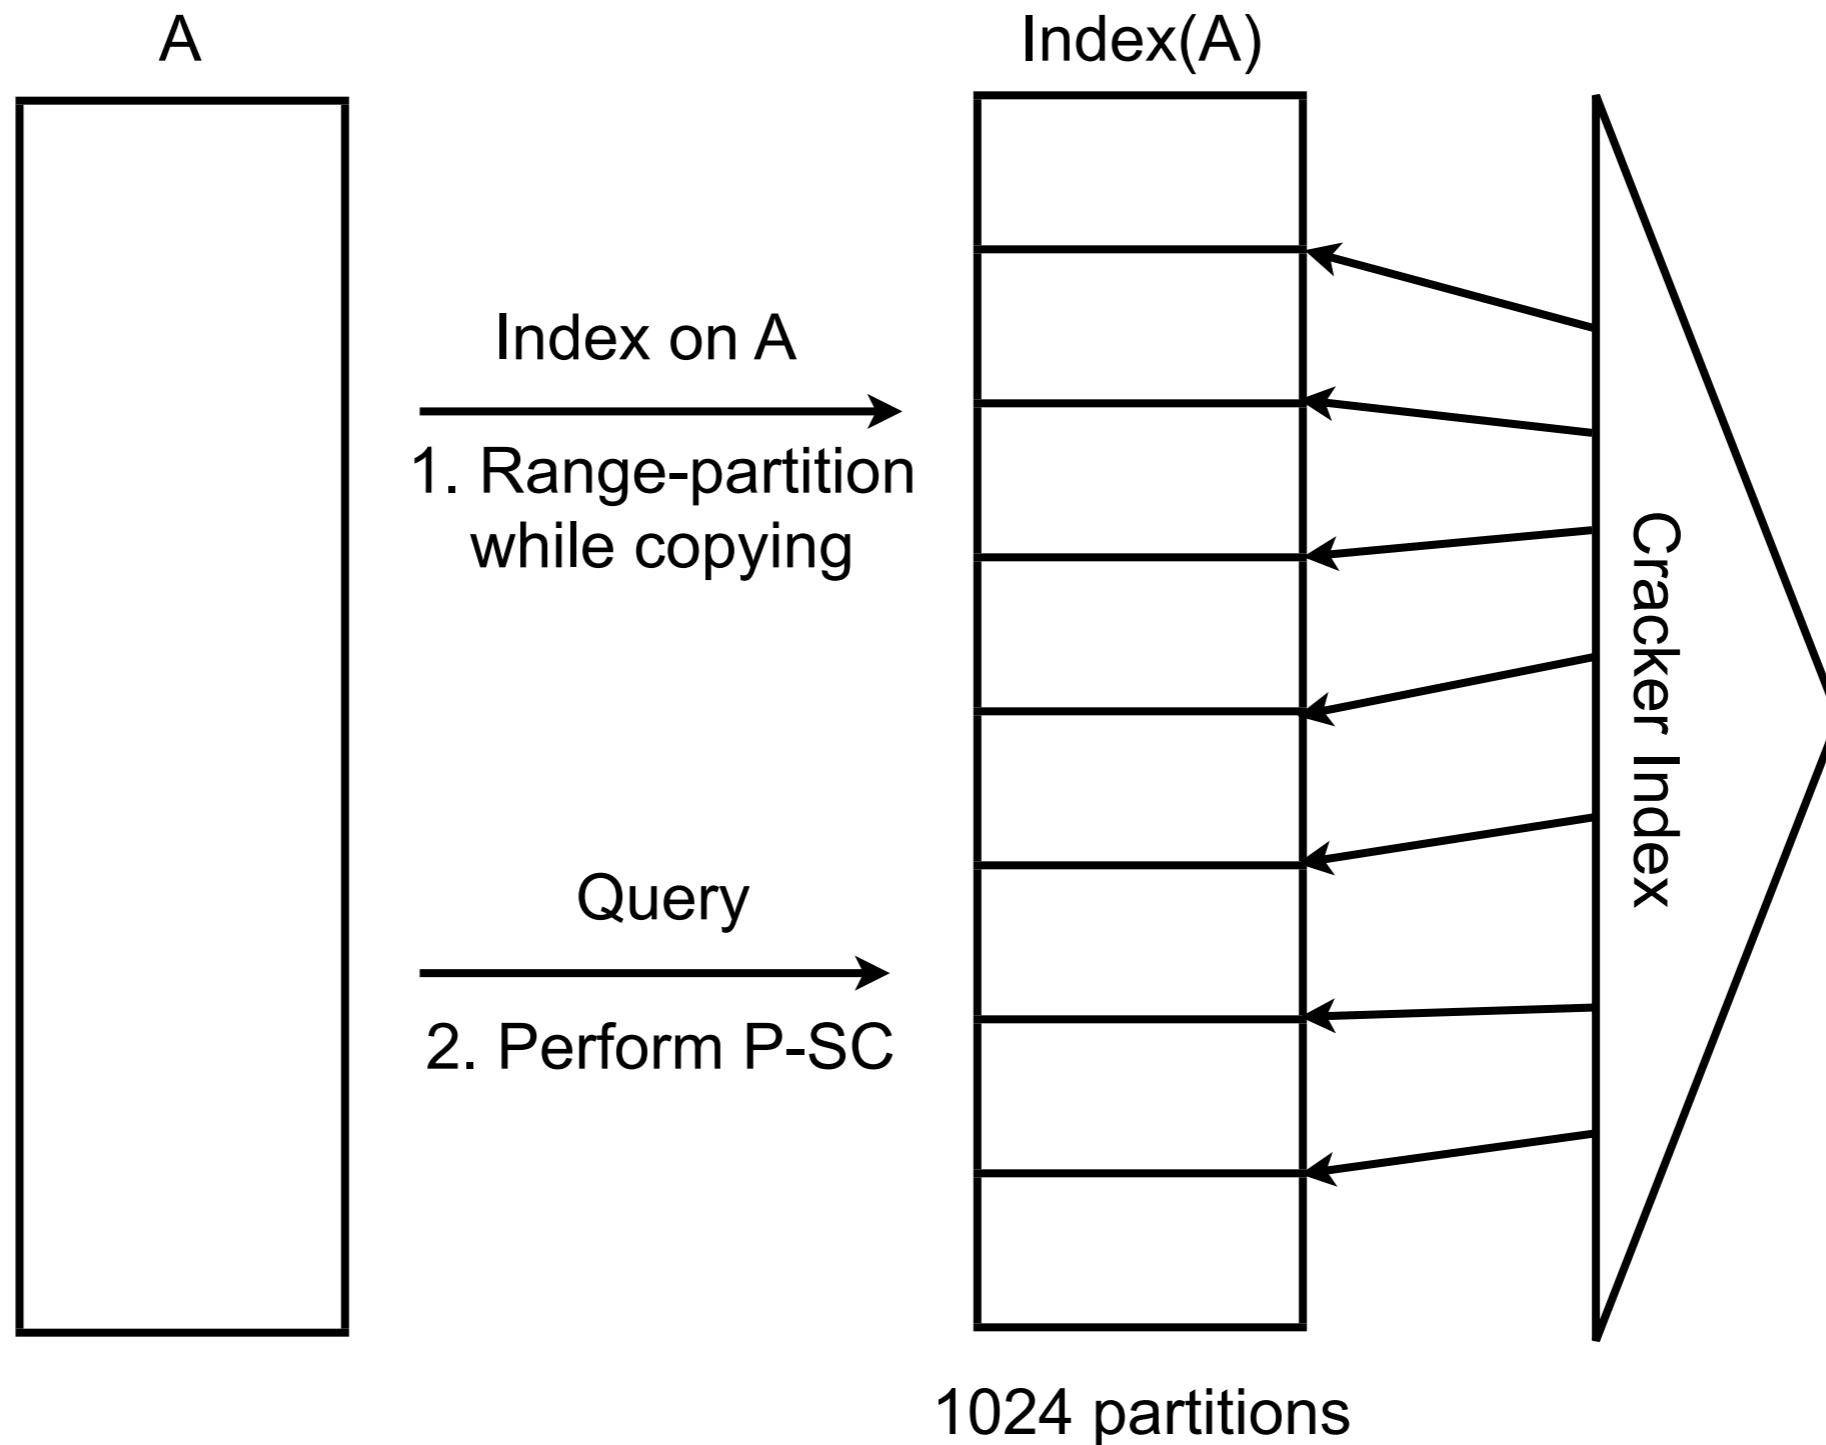
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



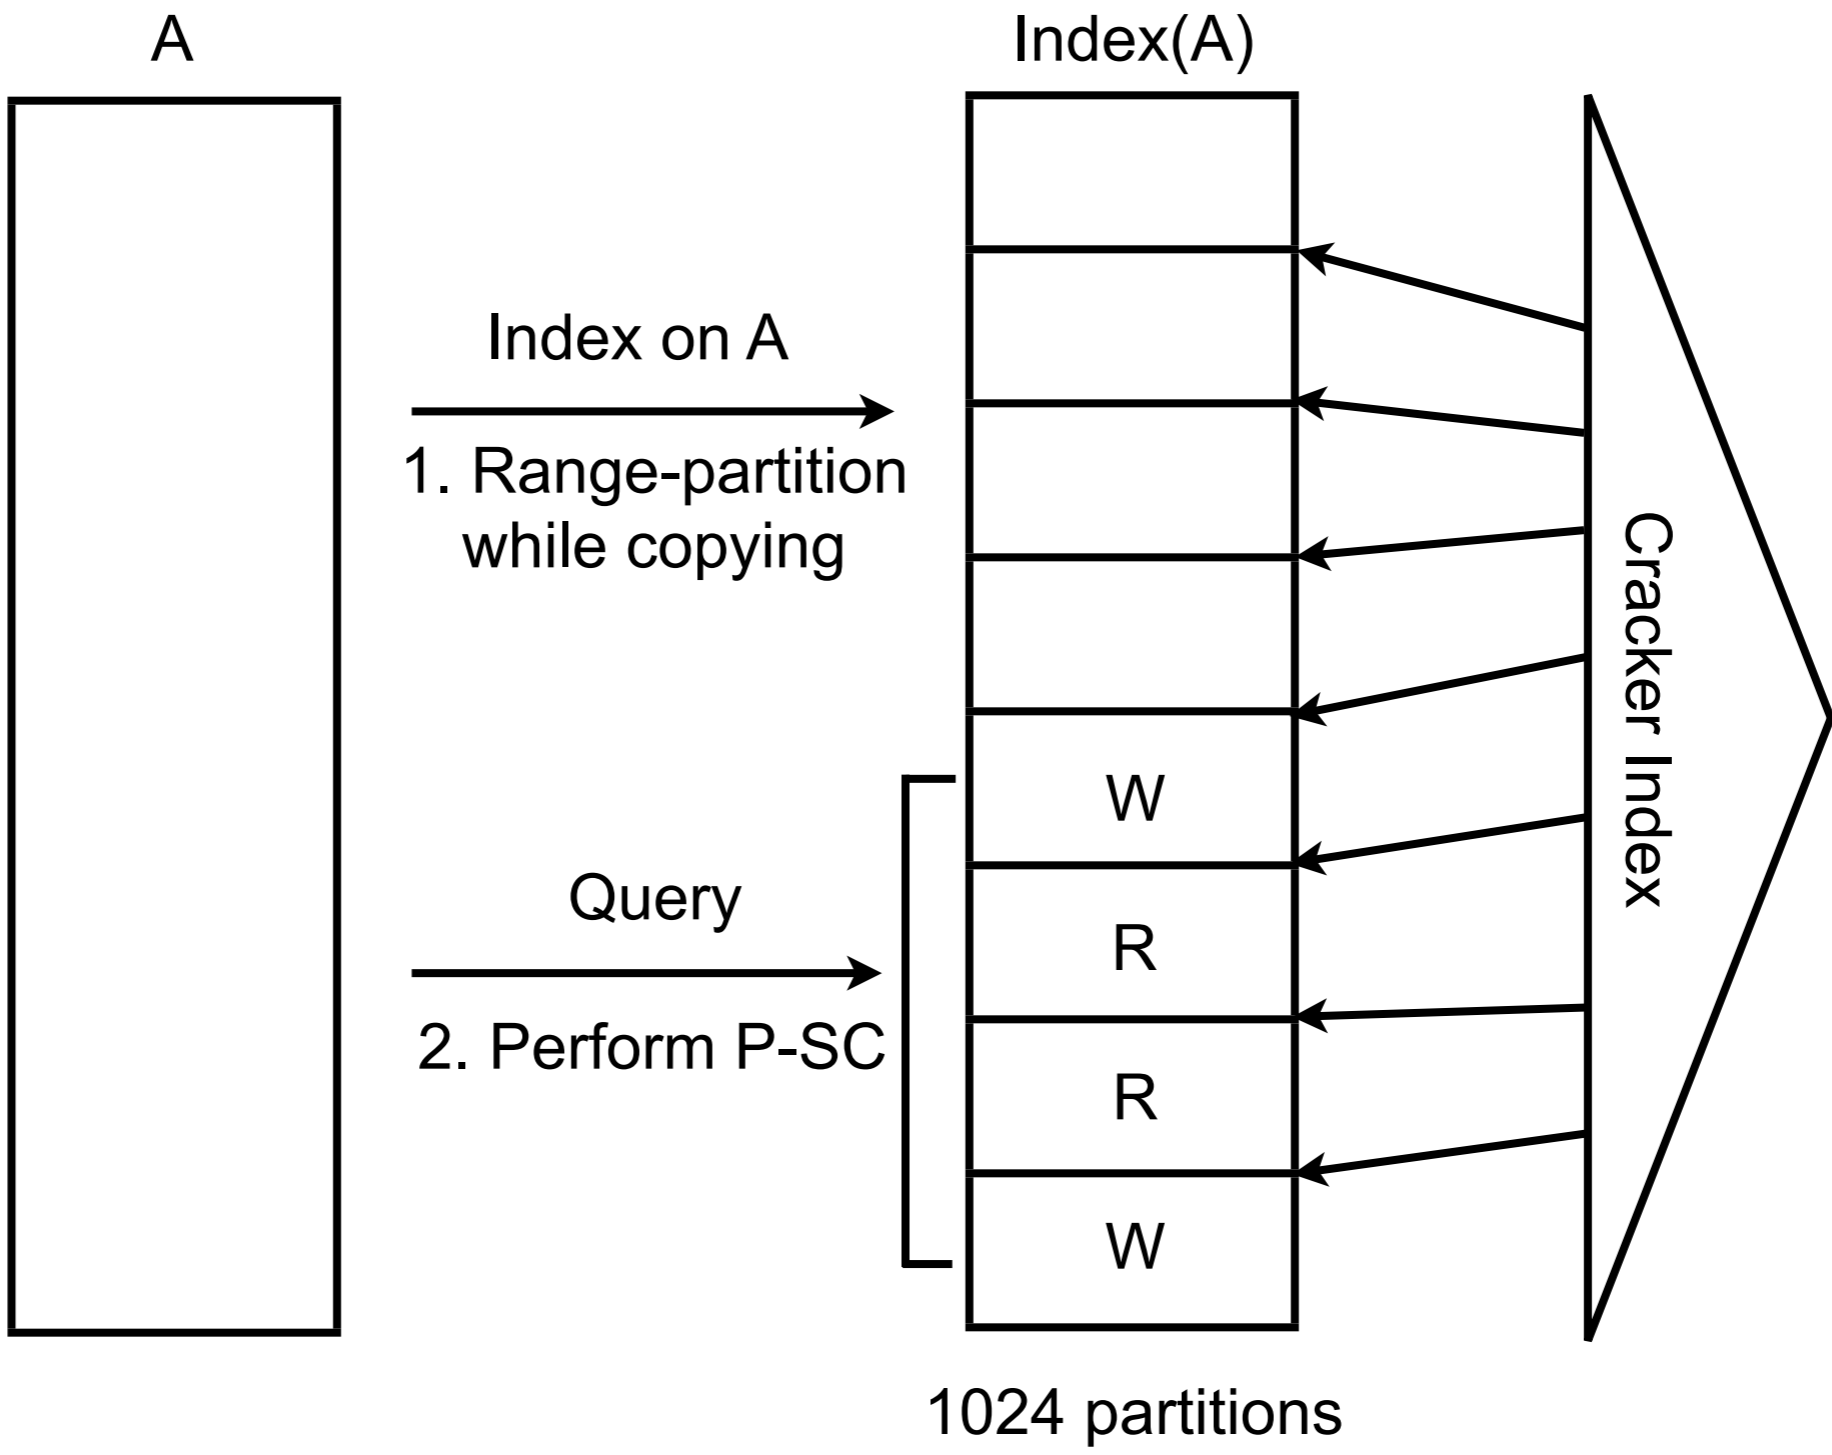
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



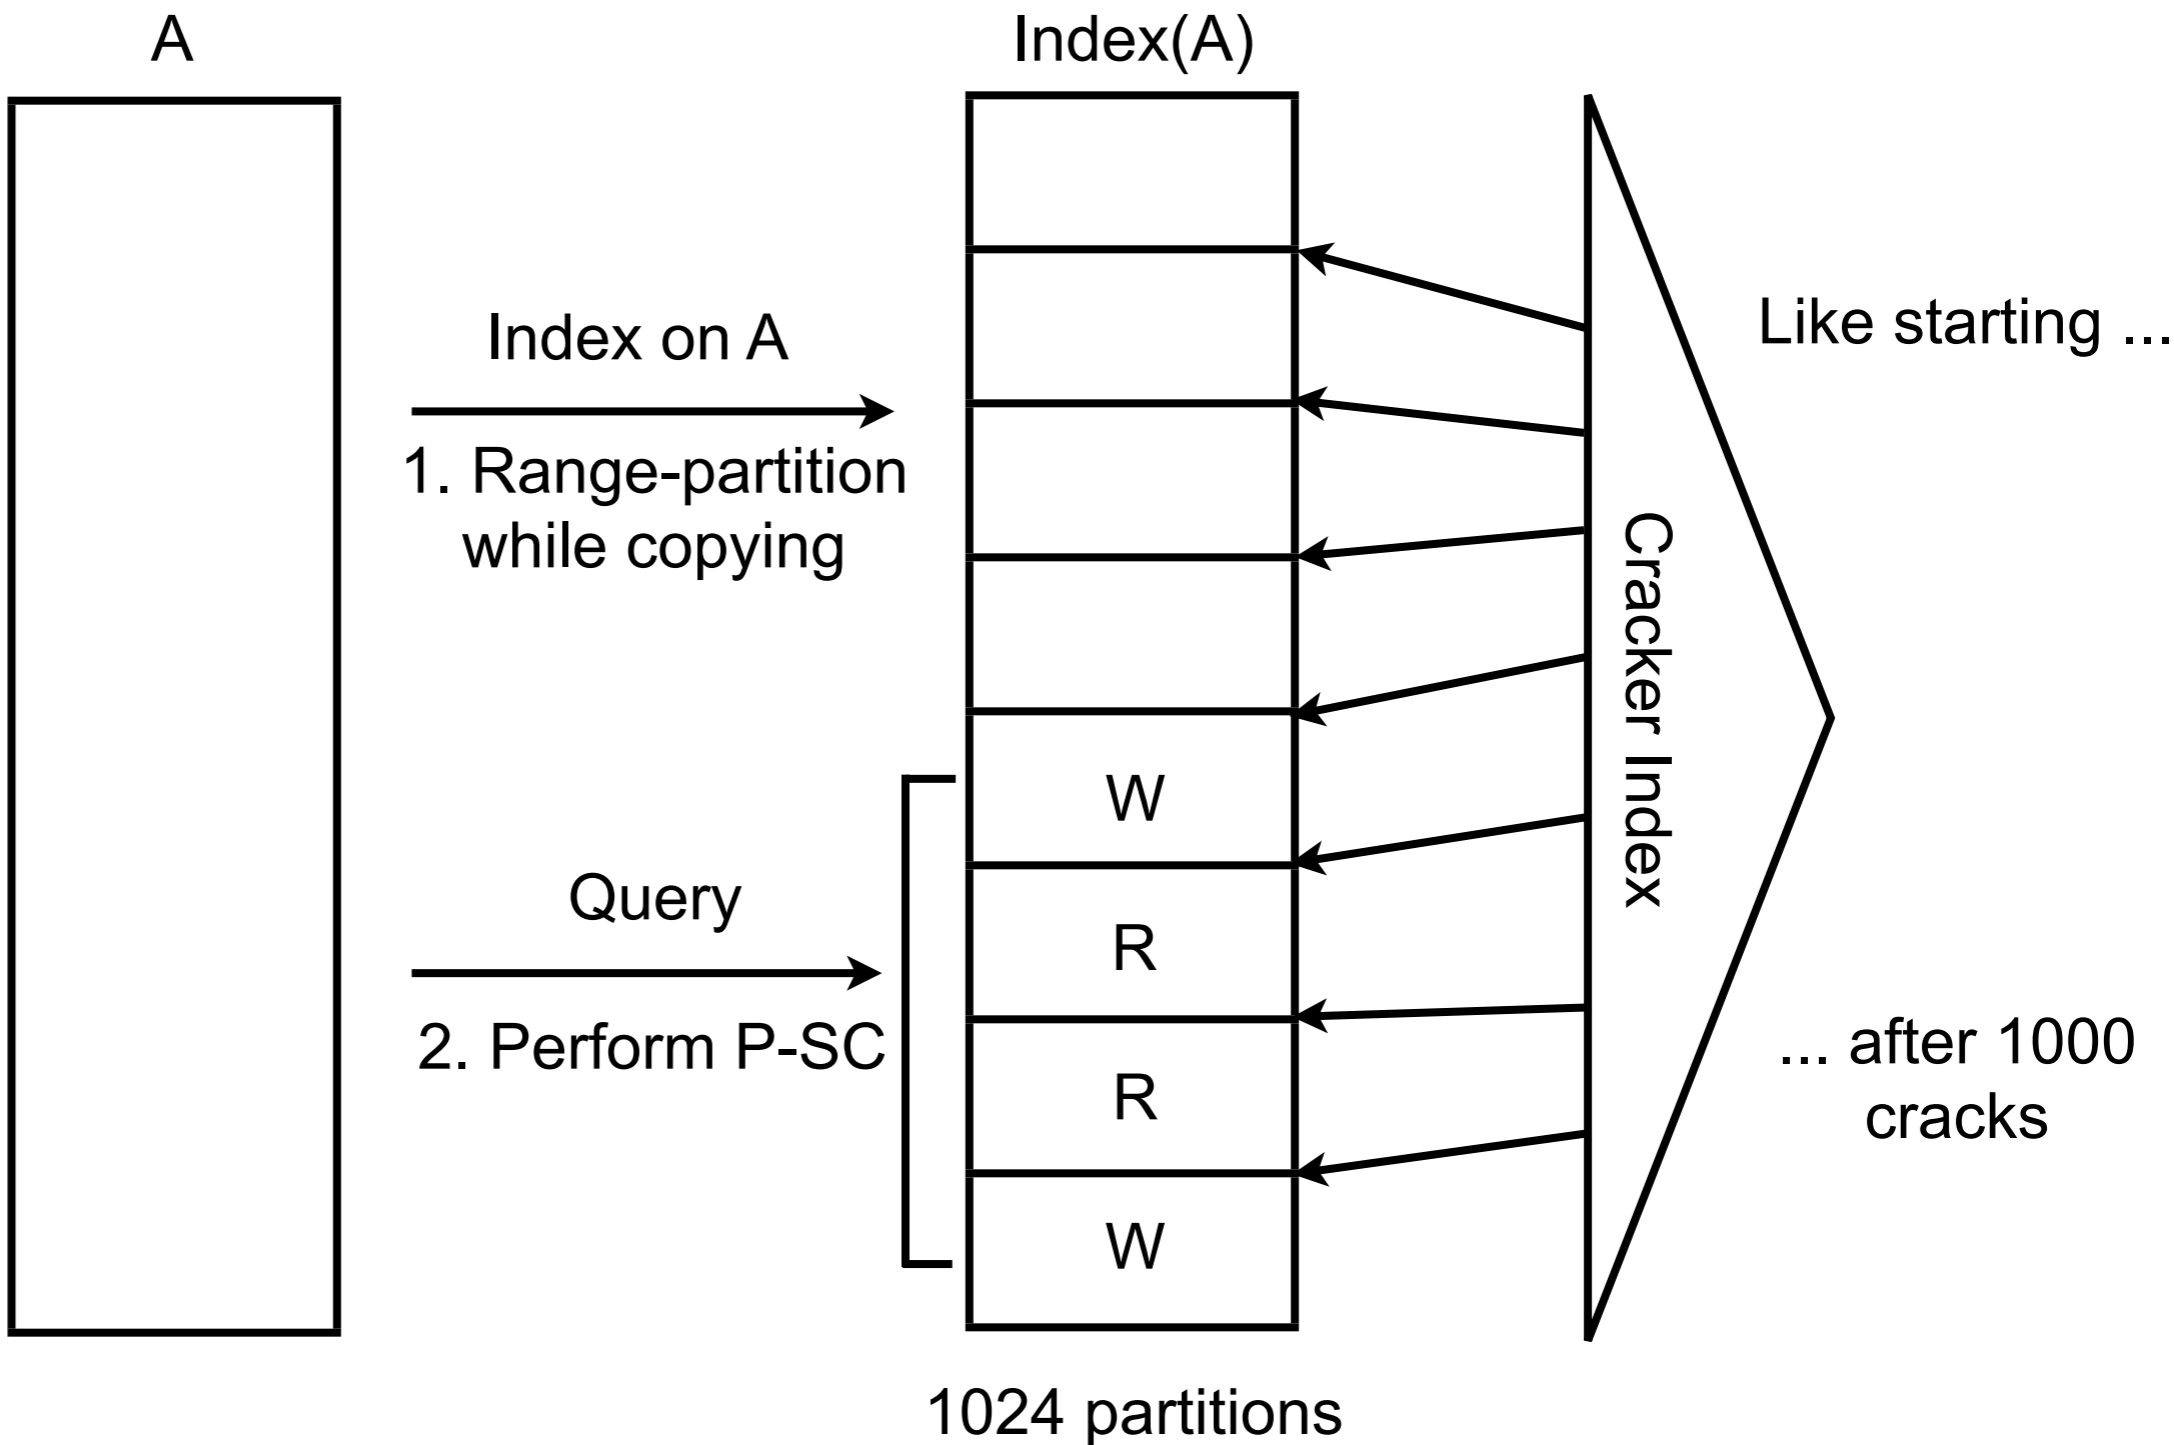
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



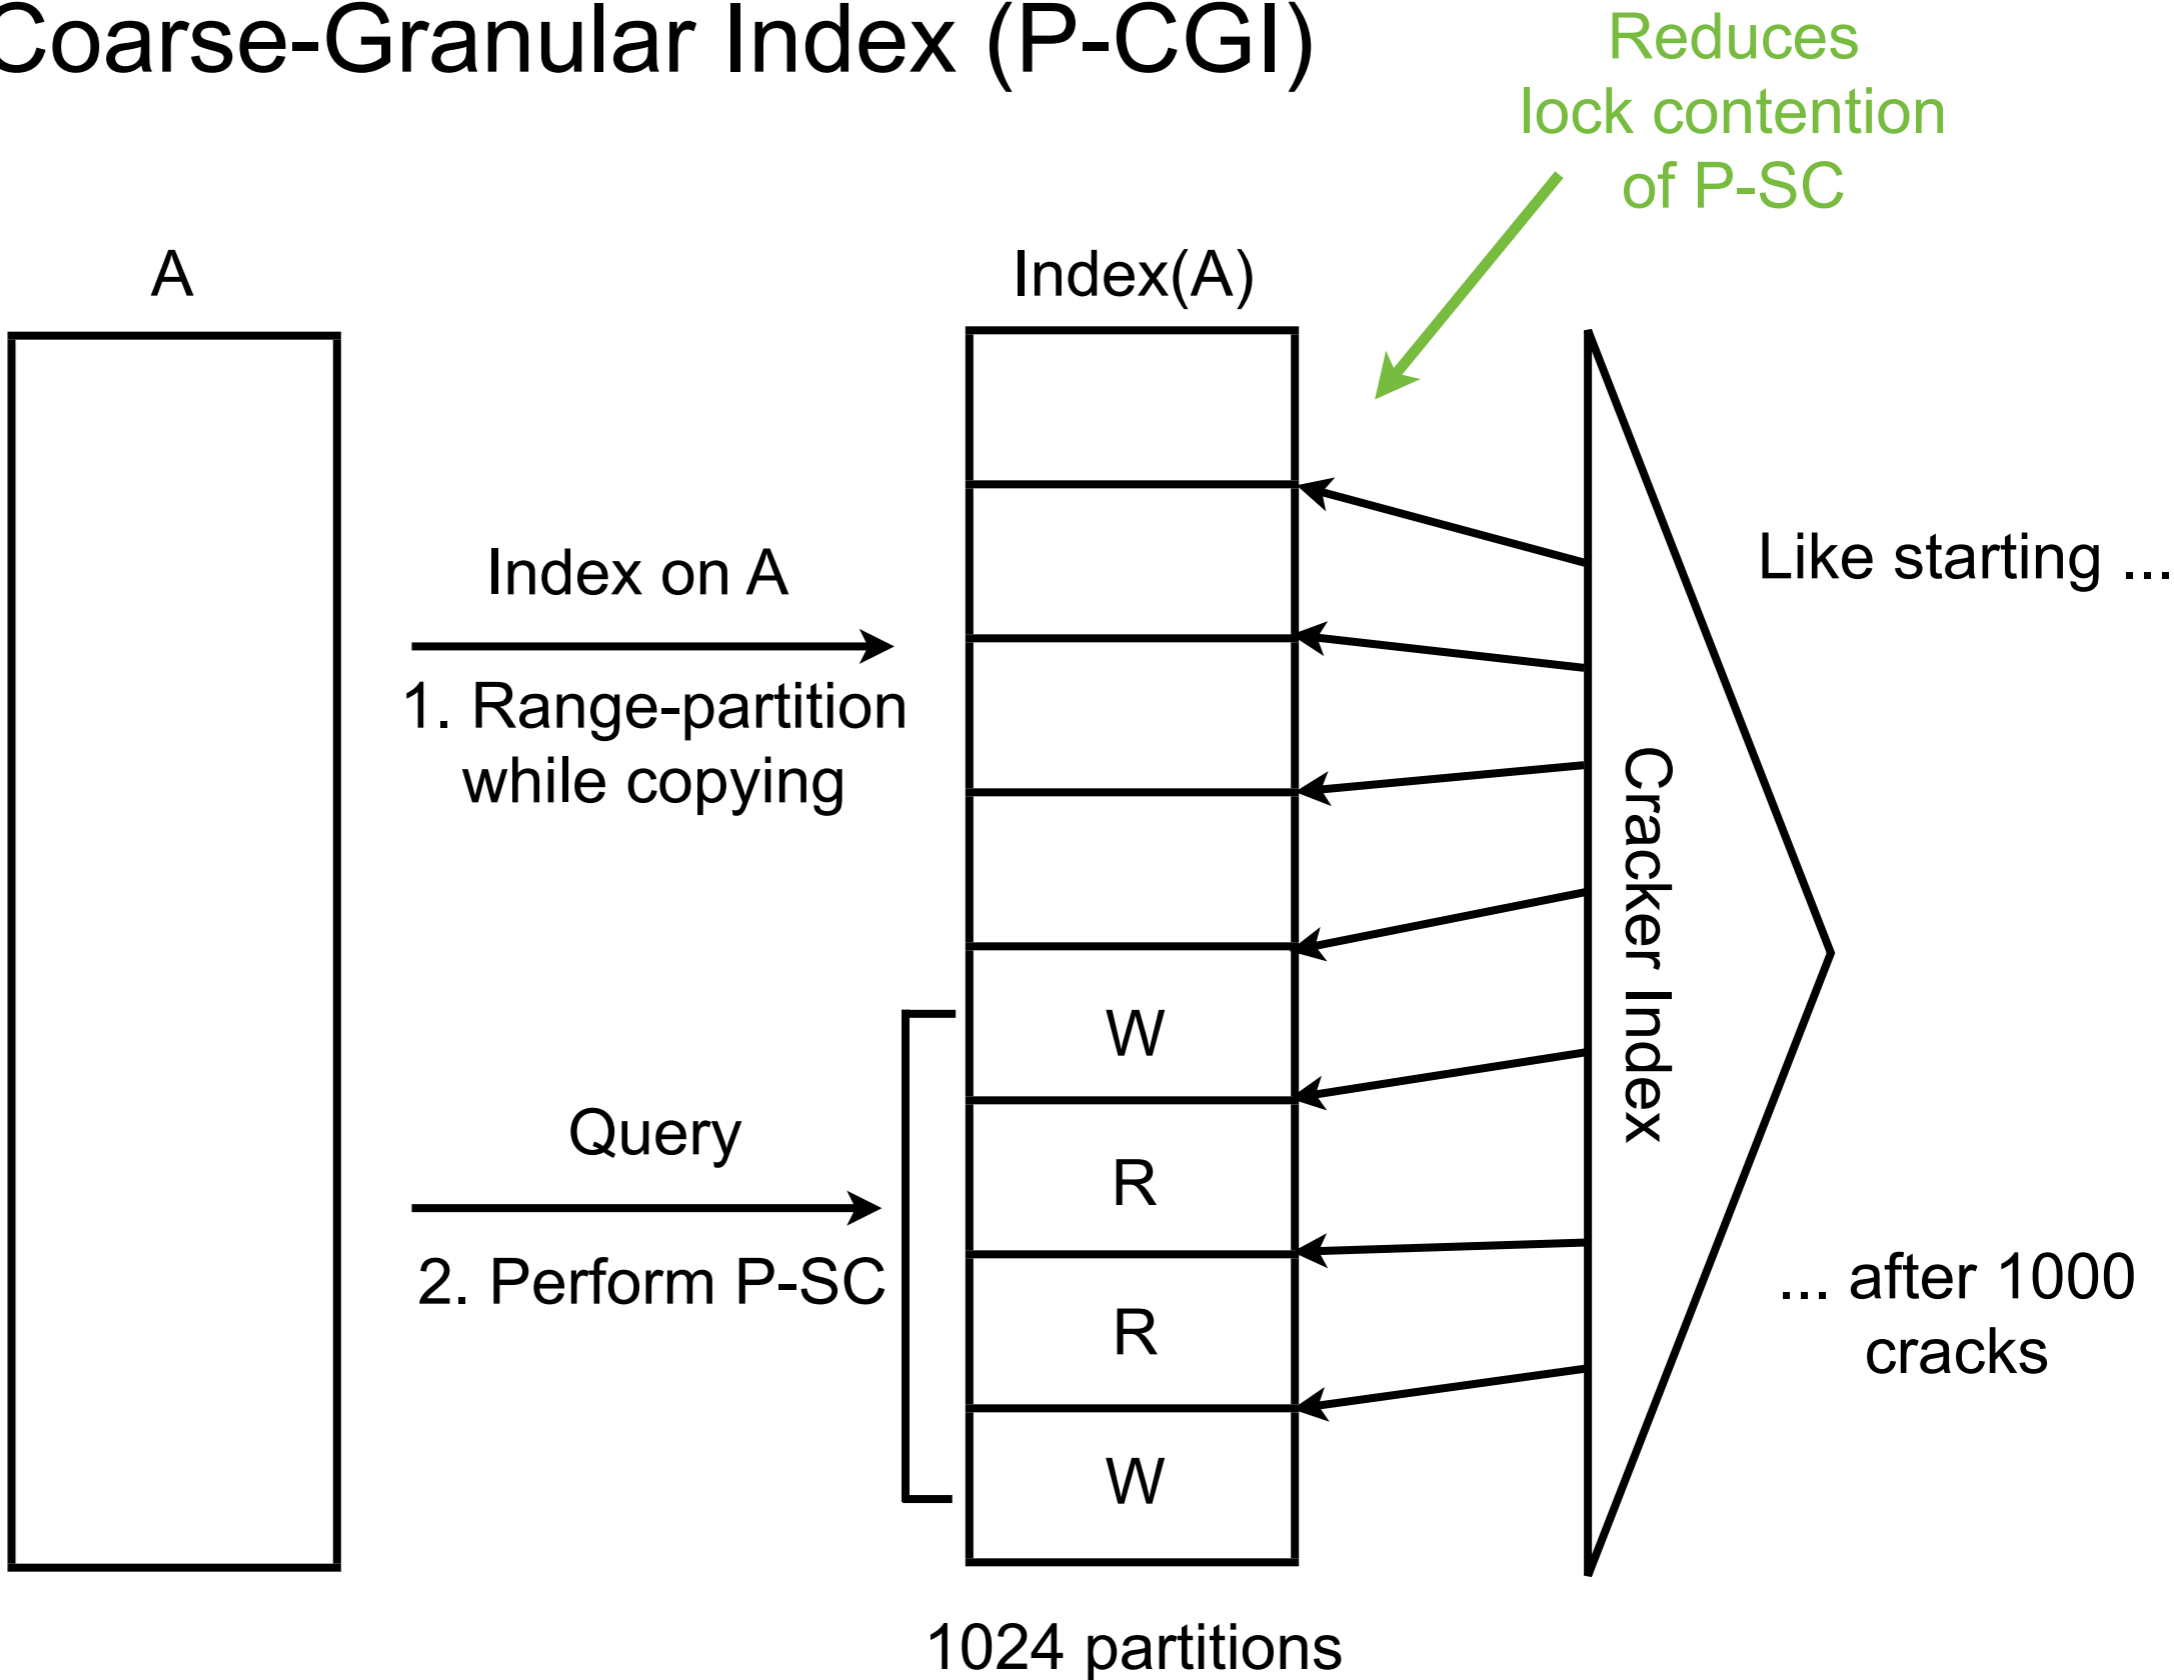
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



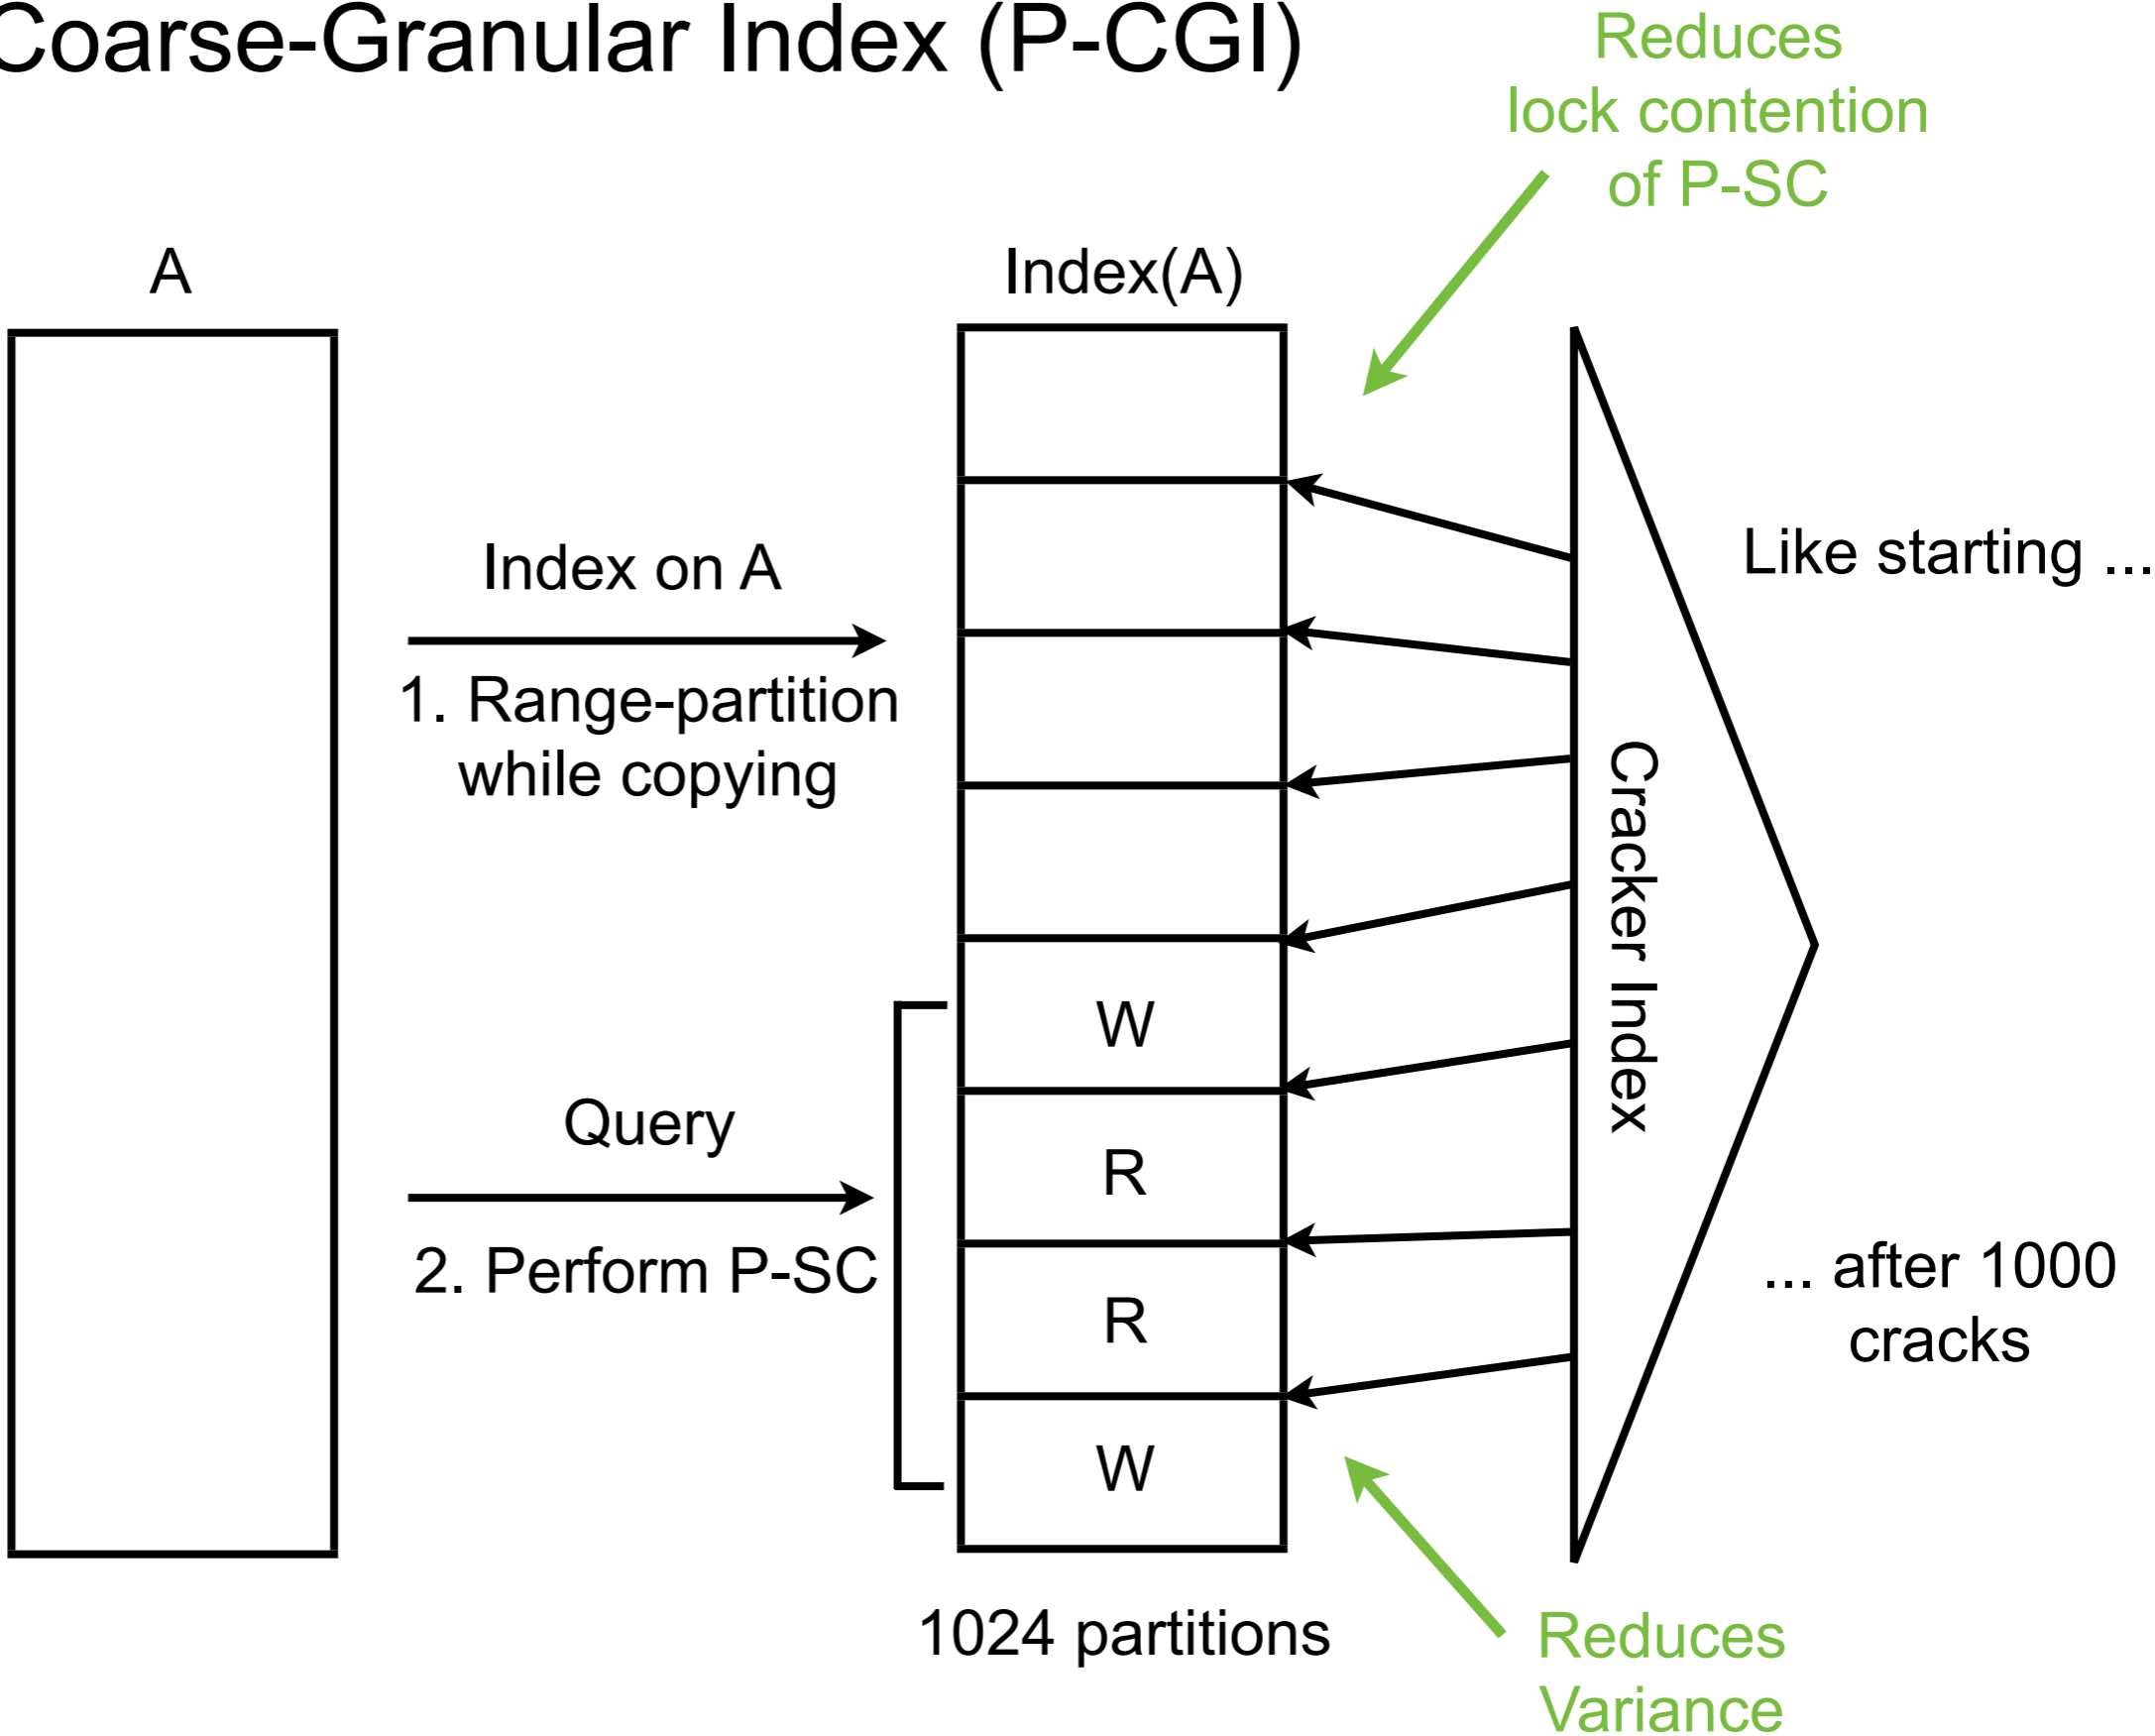
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)

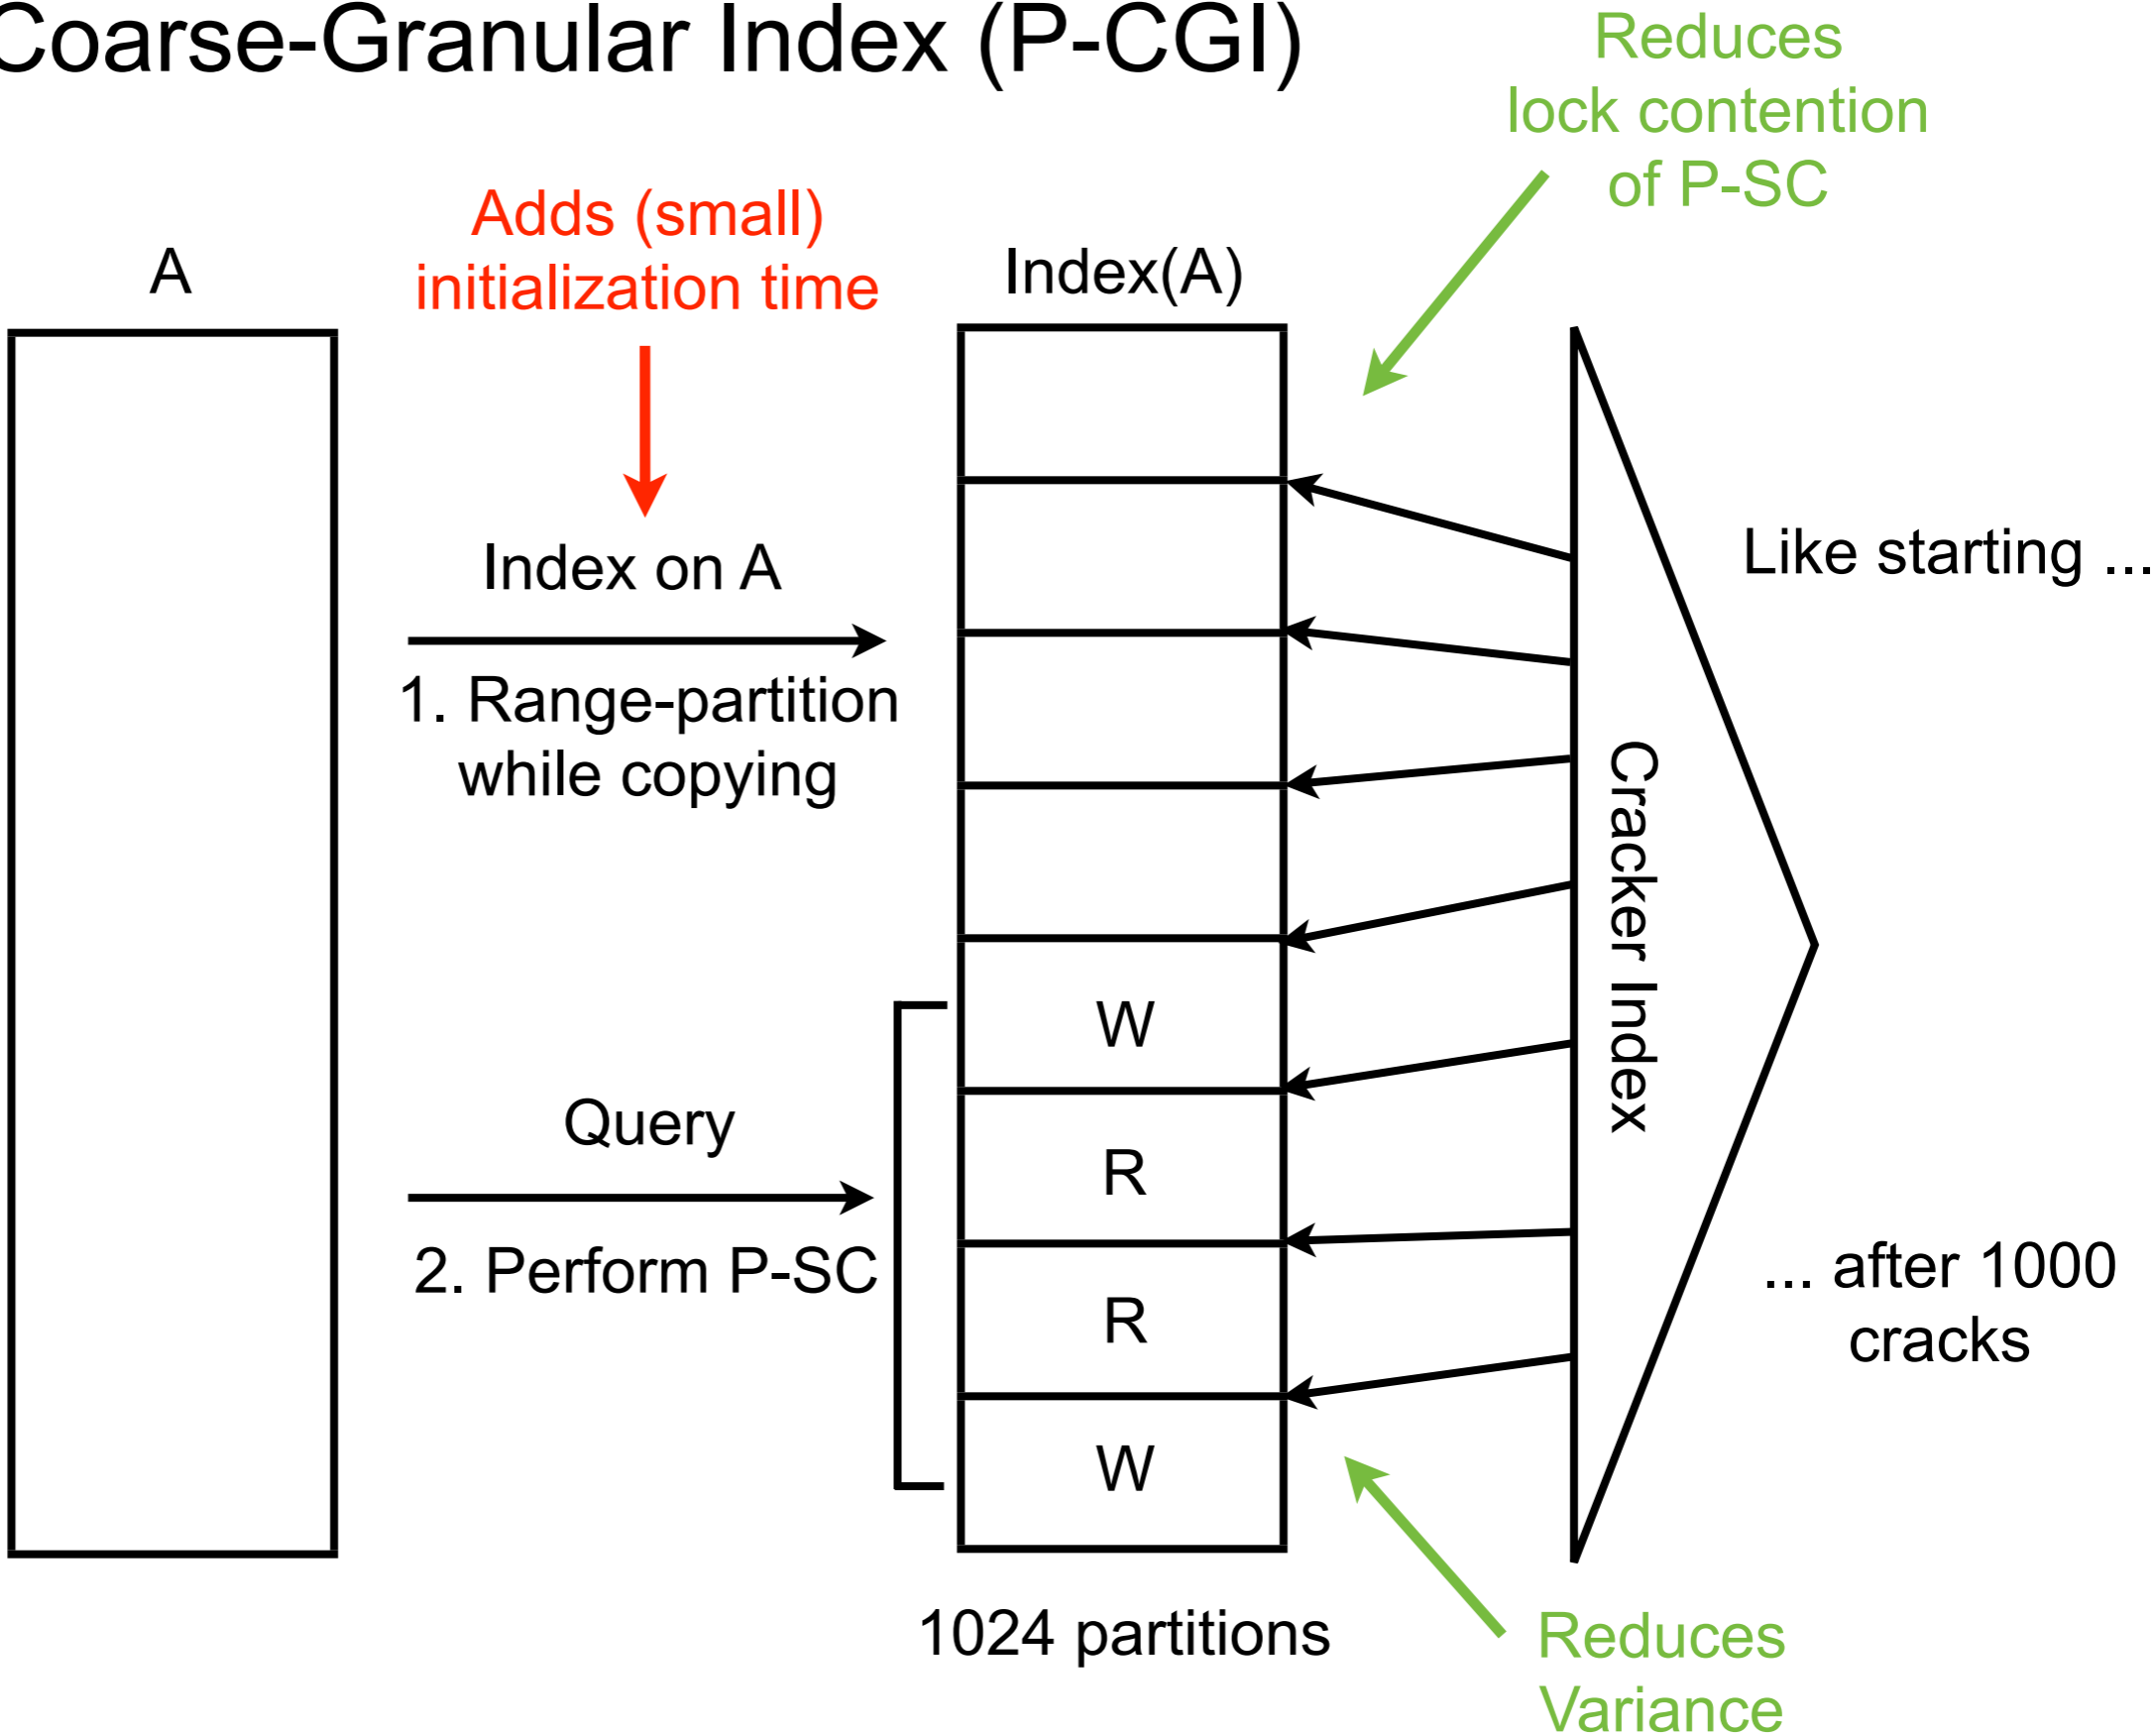


# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)

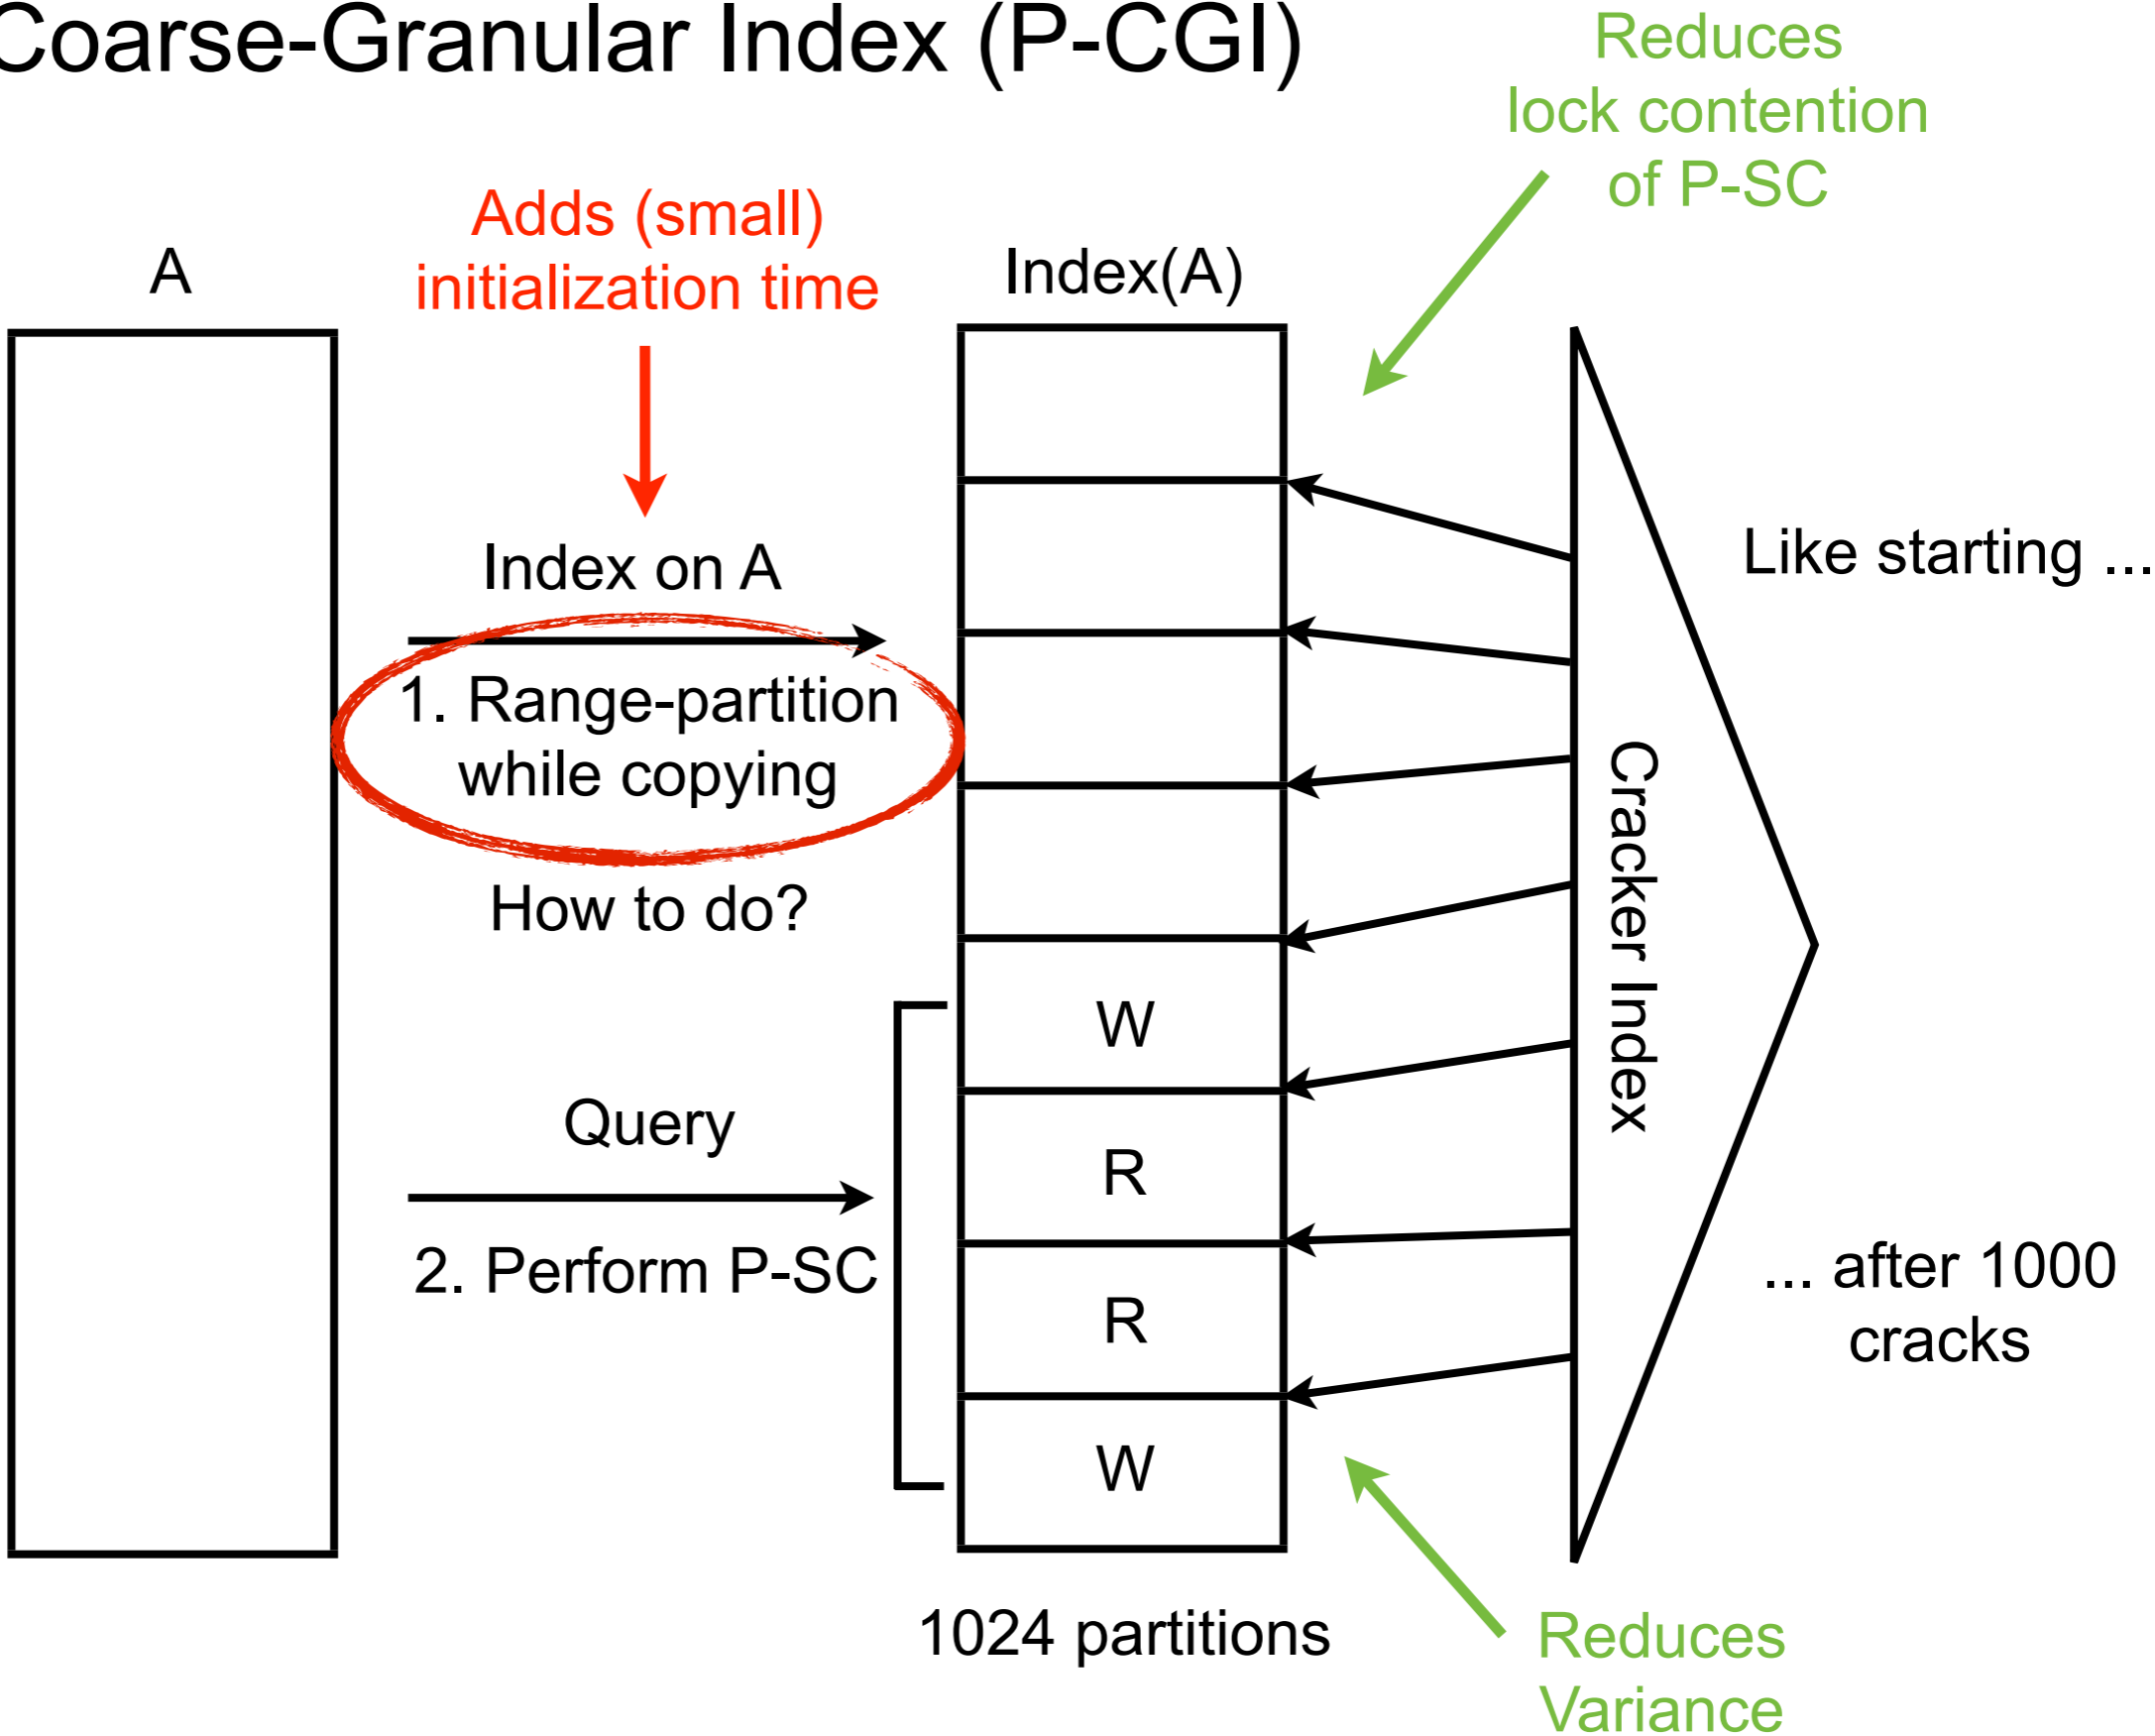




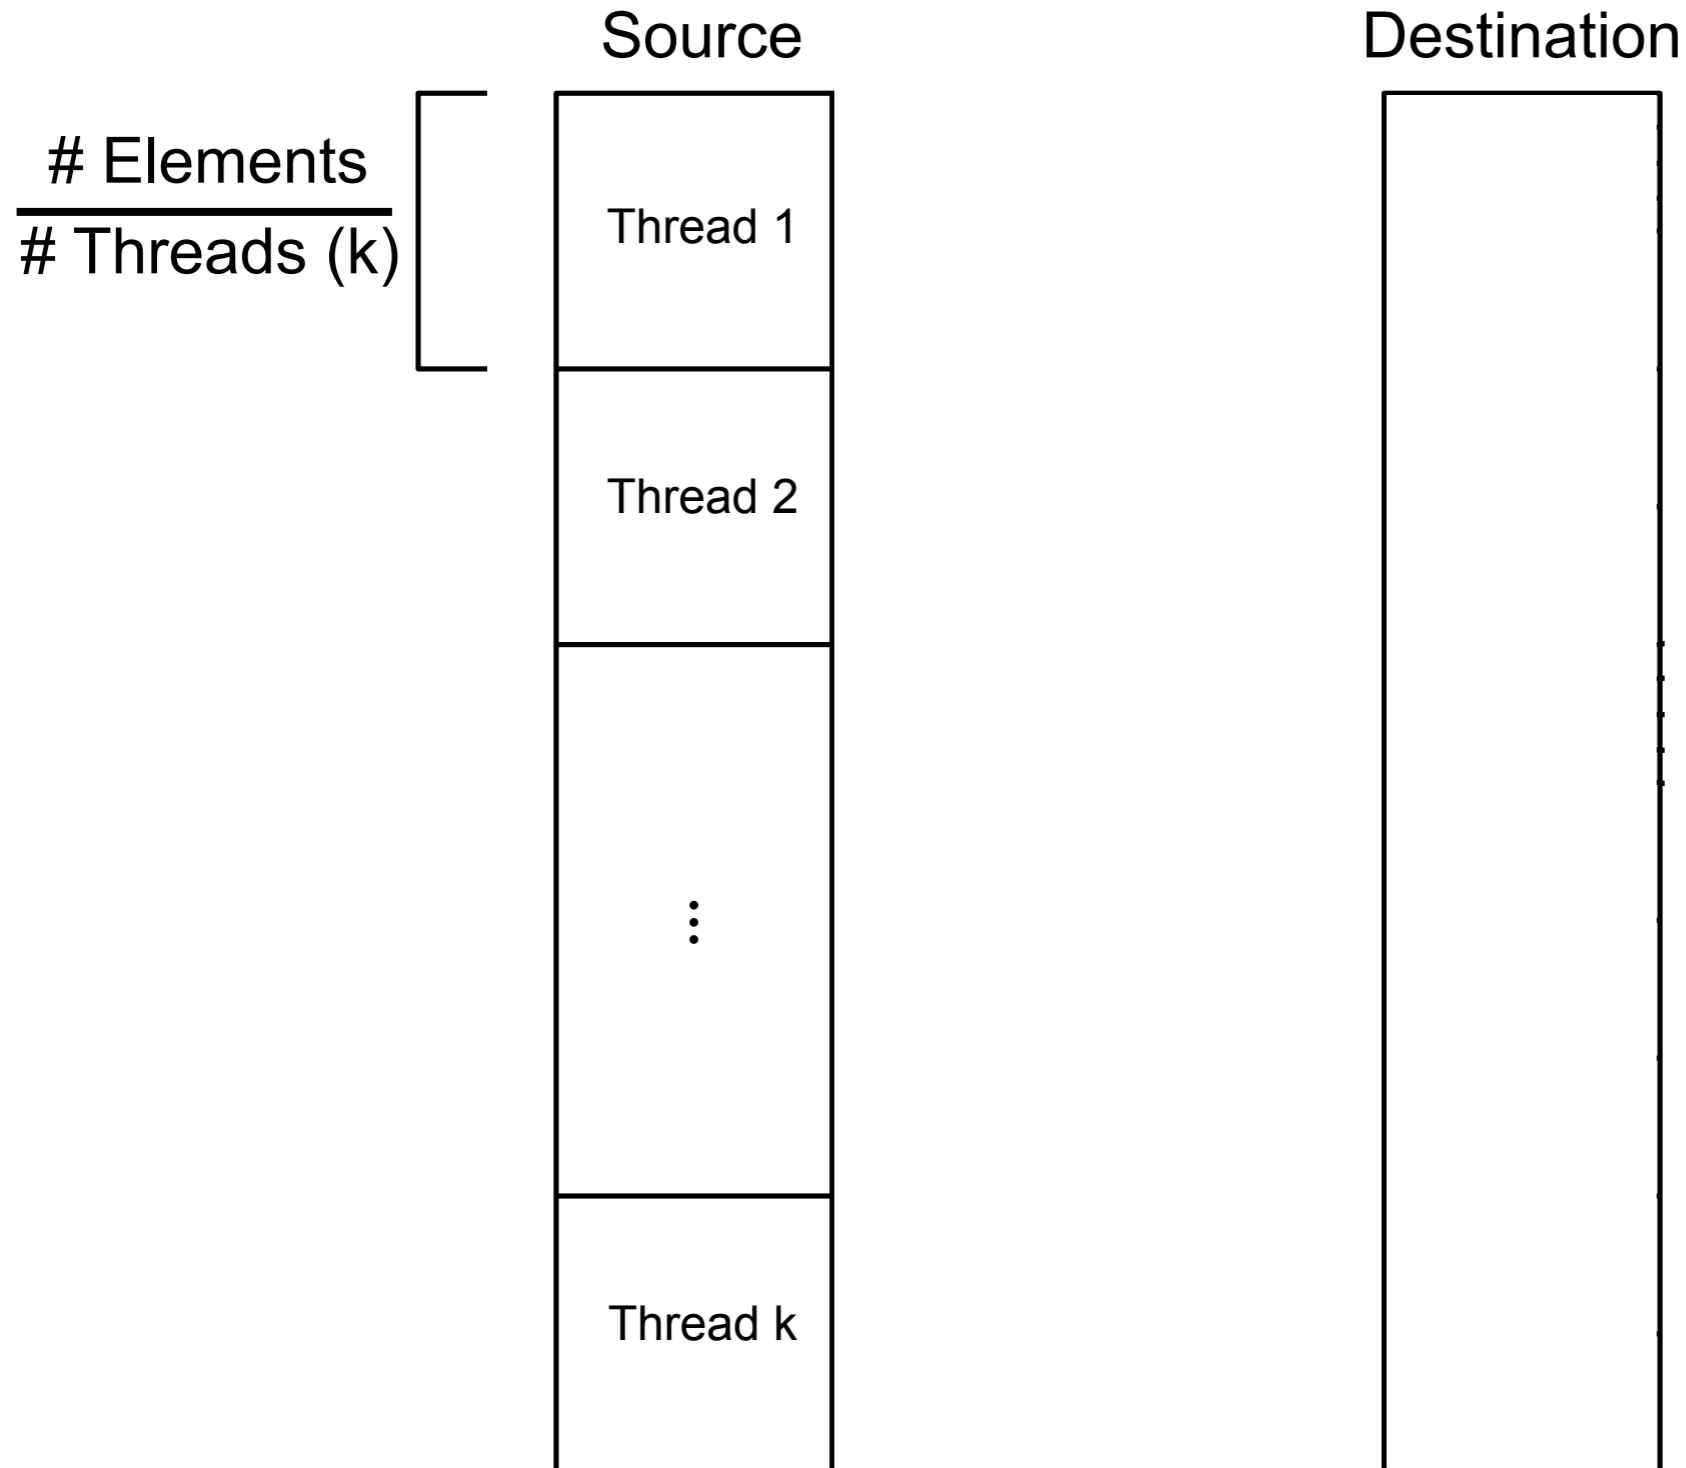
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



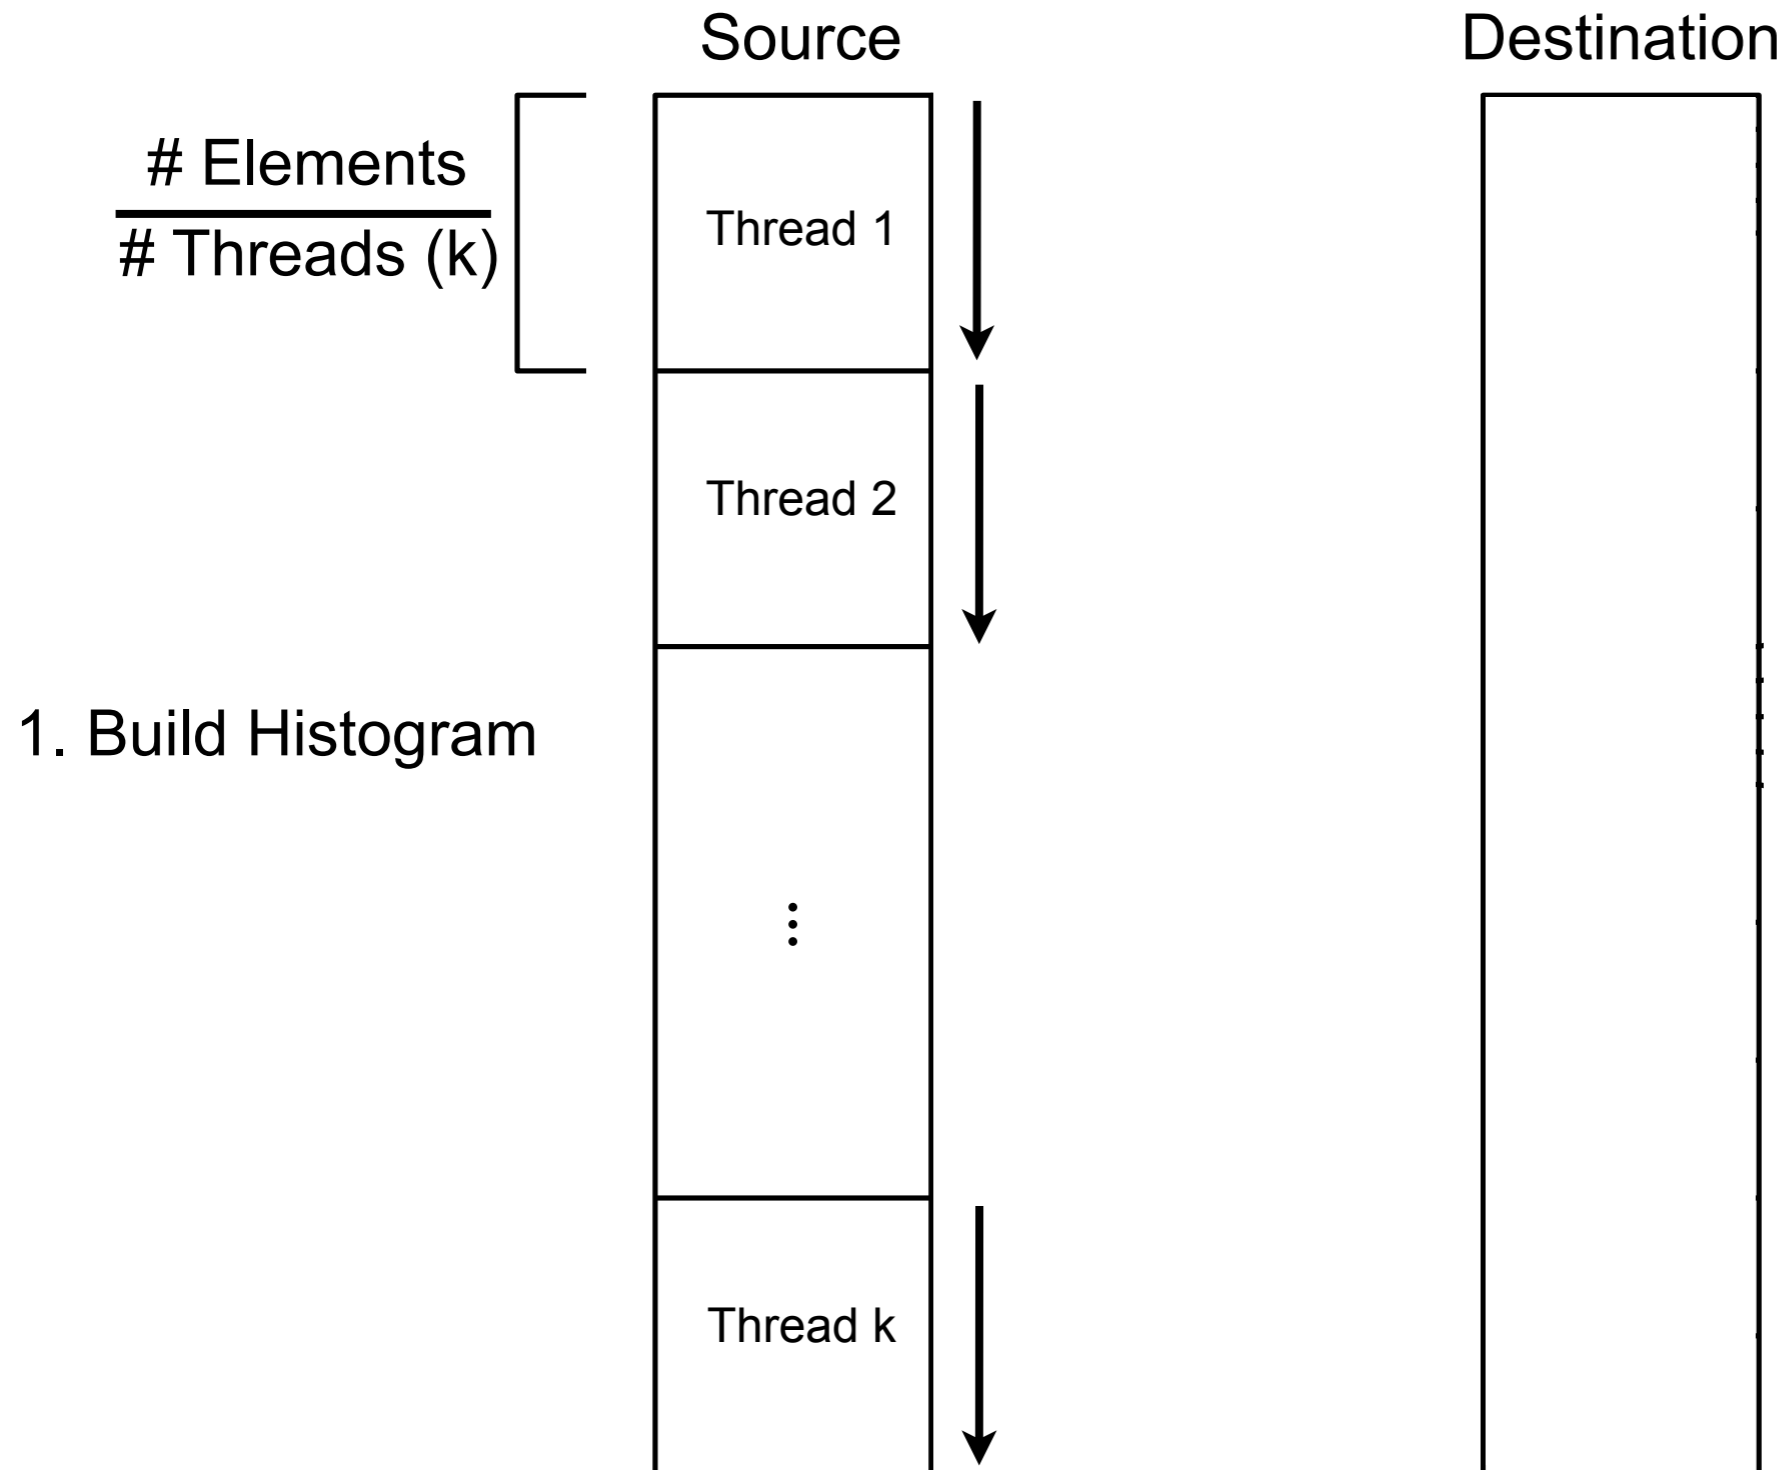
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI)



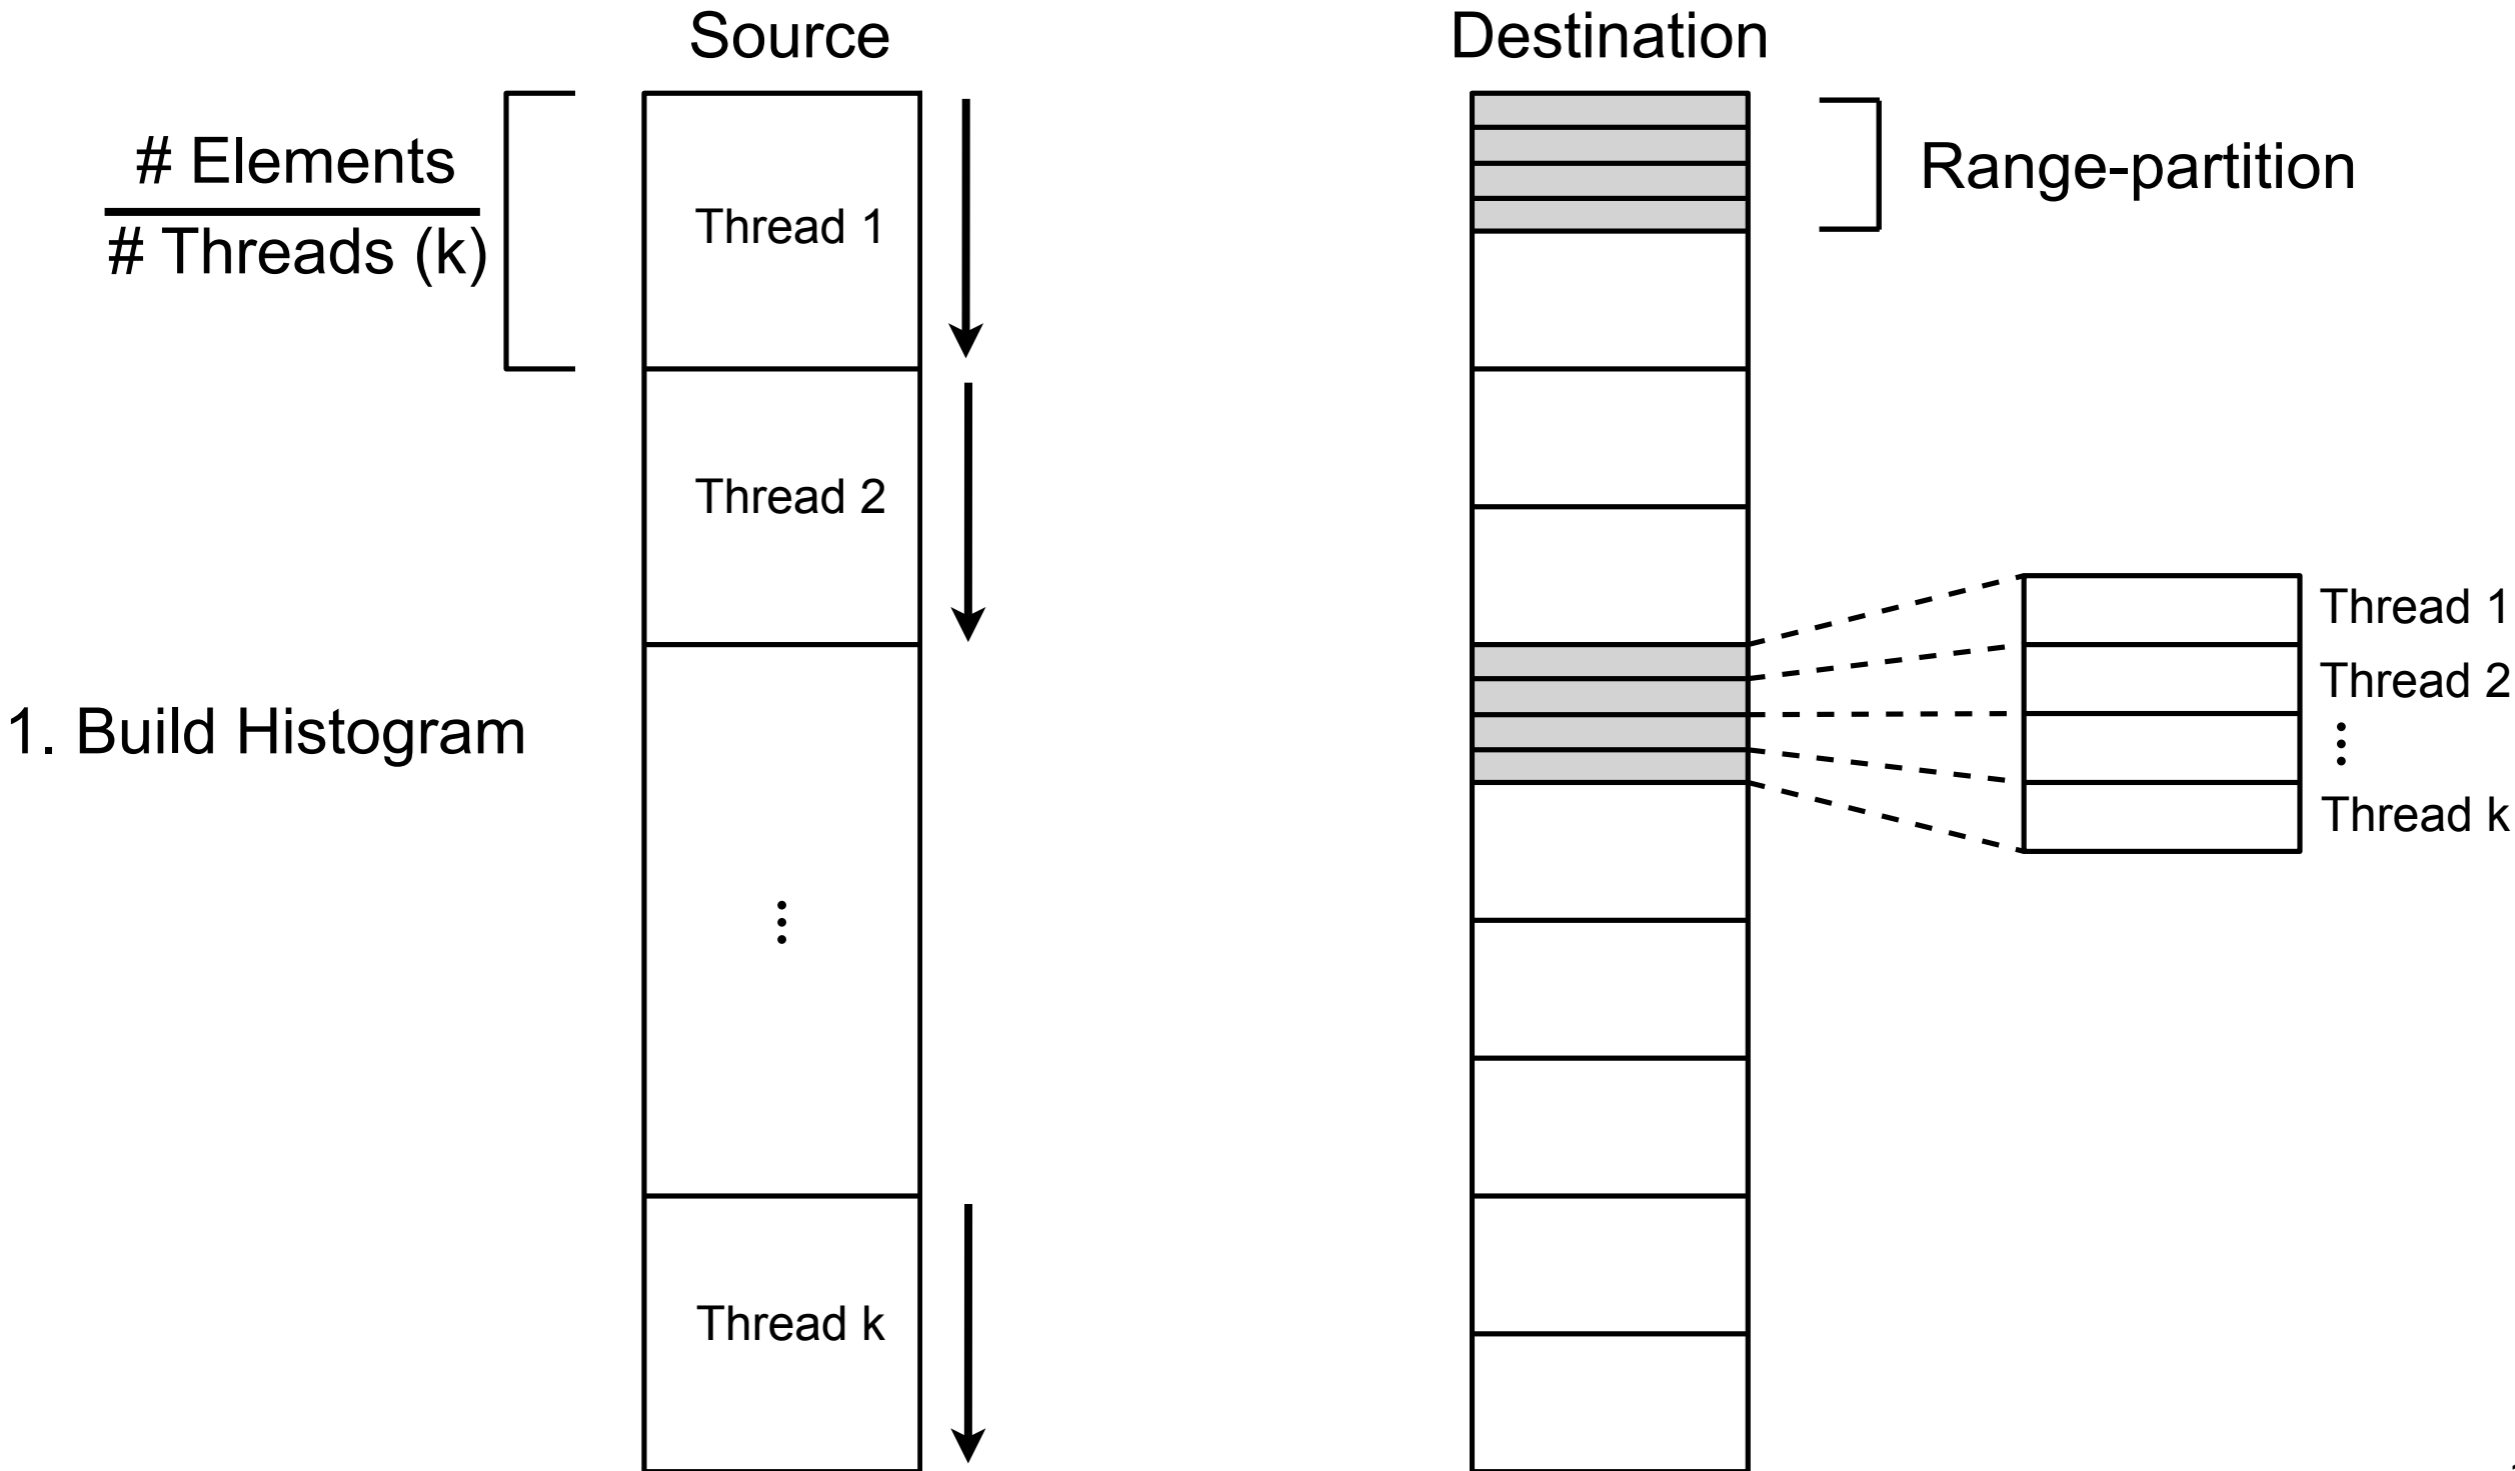
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning



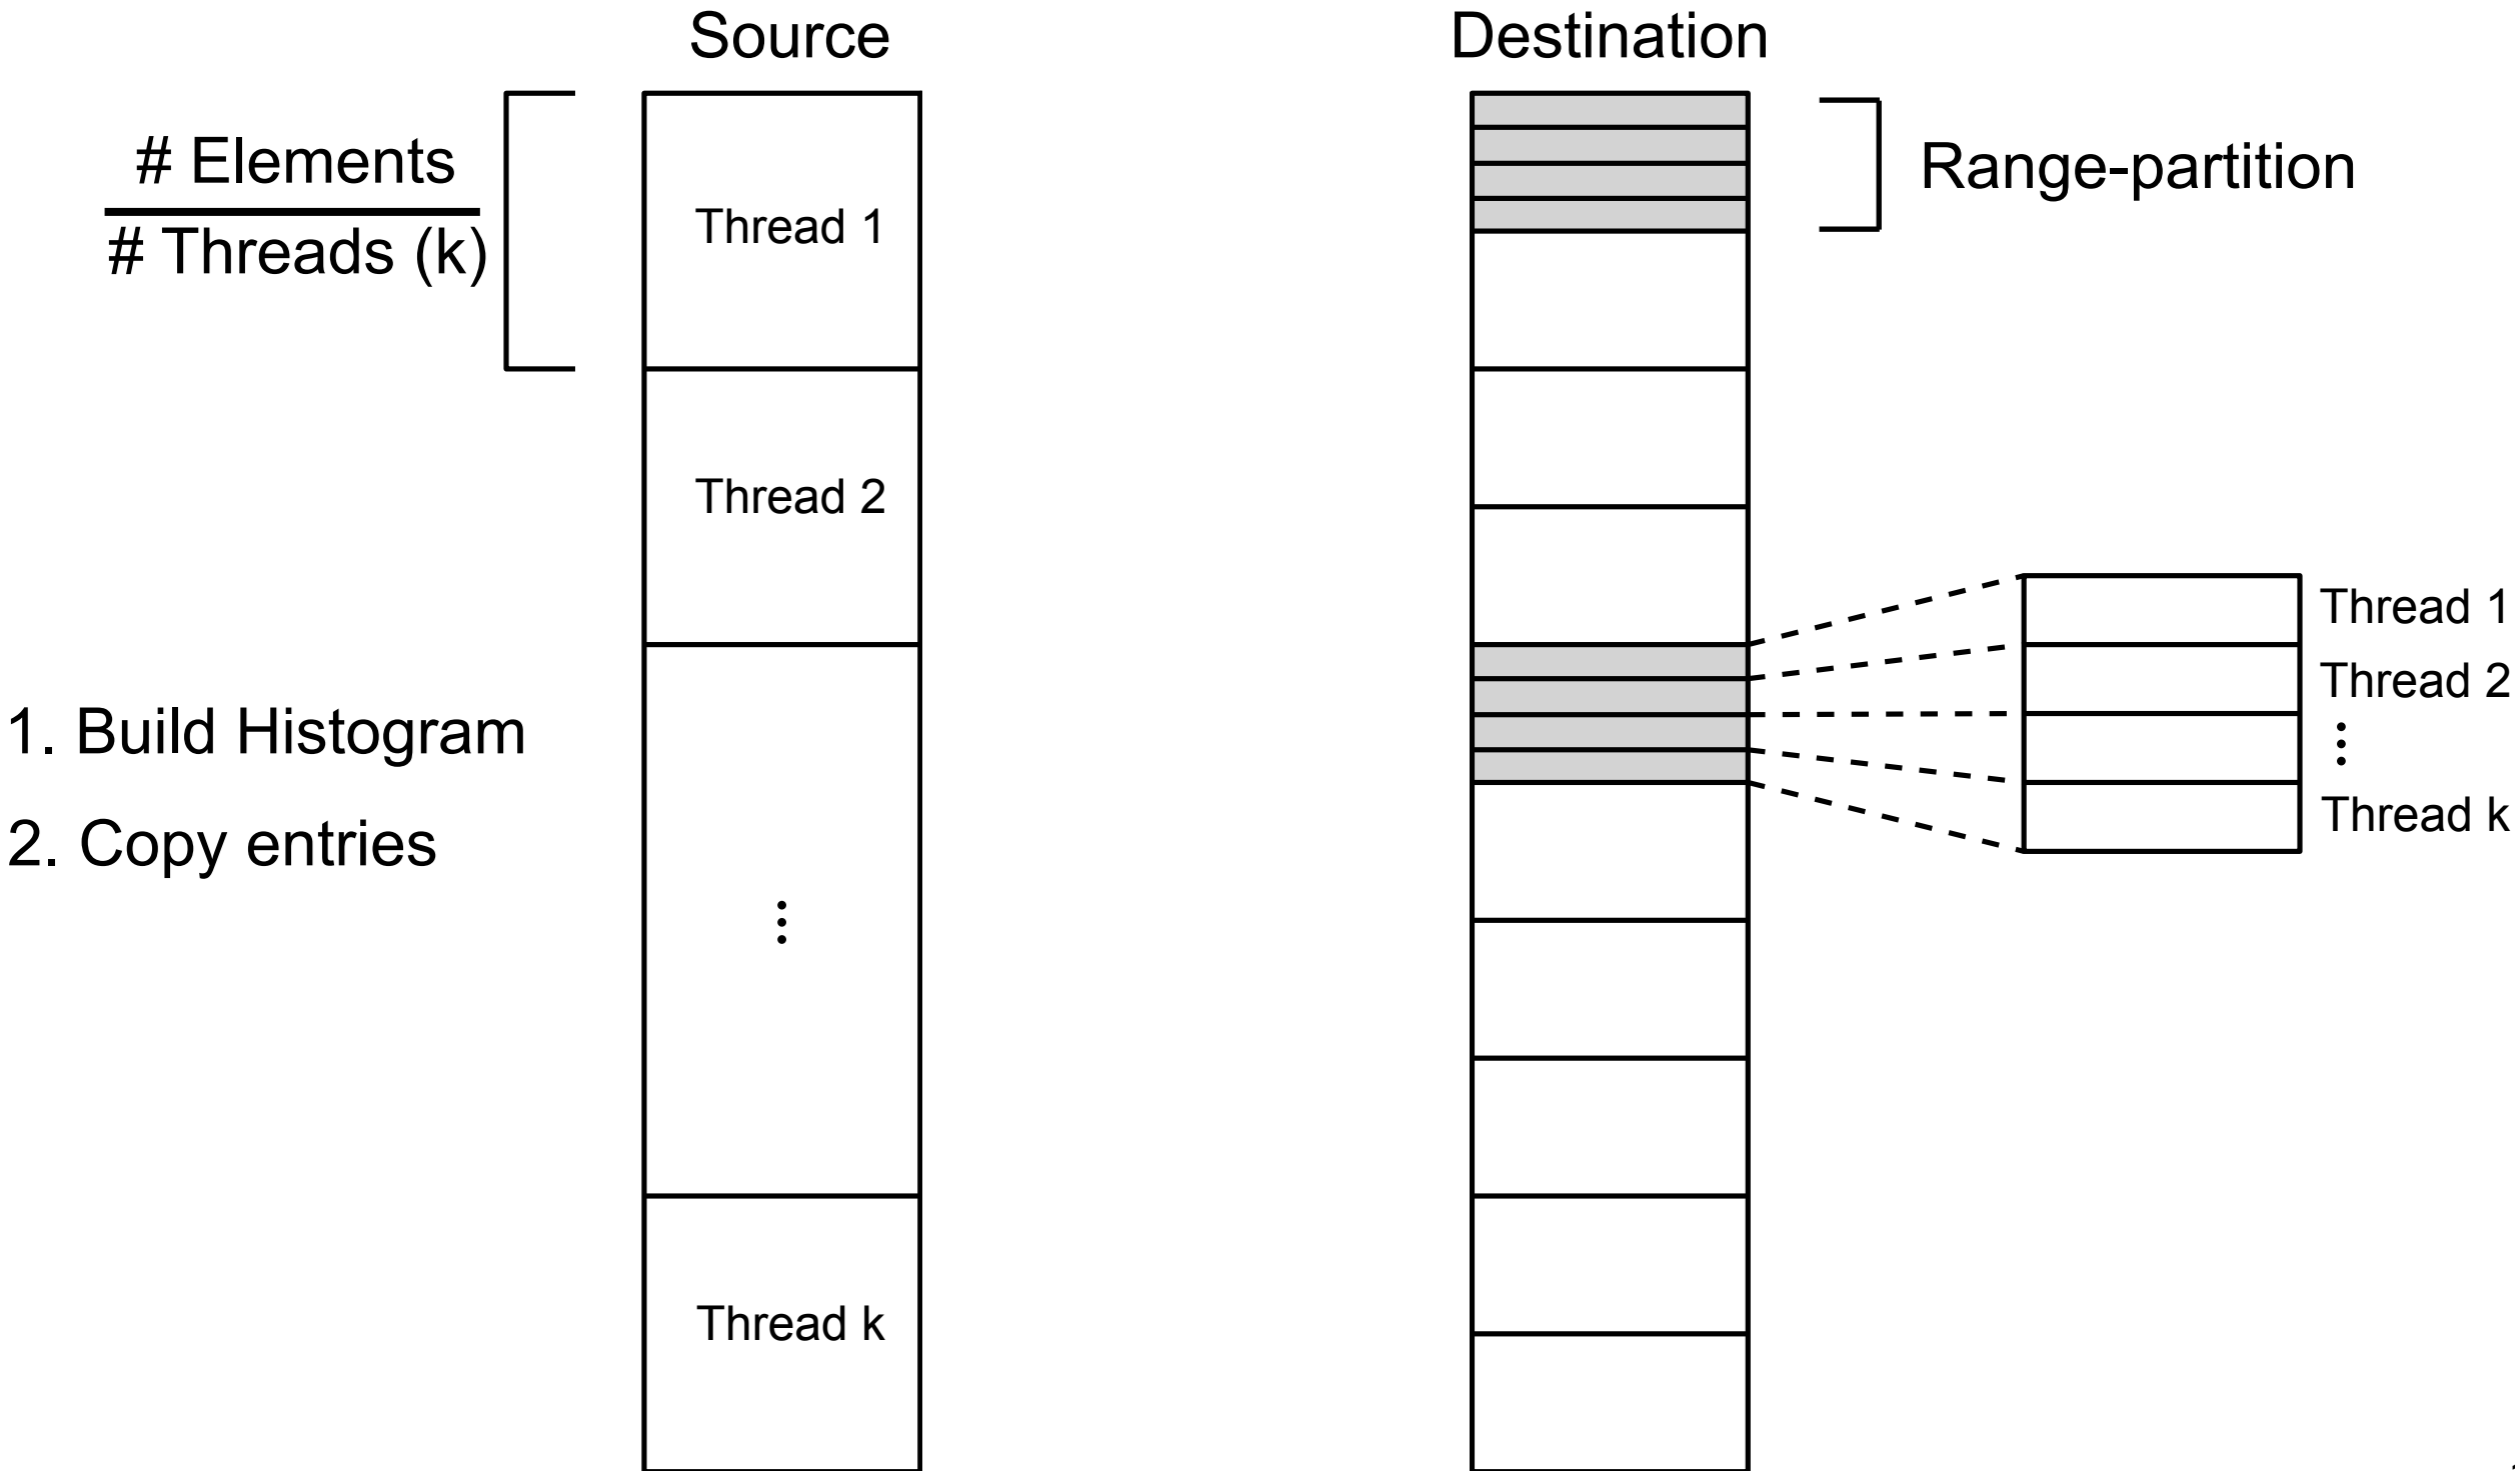
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning



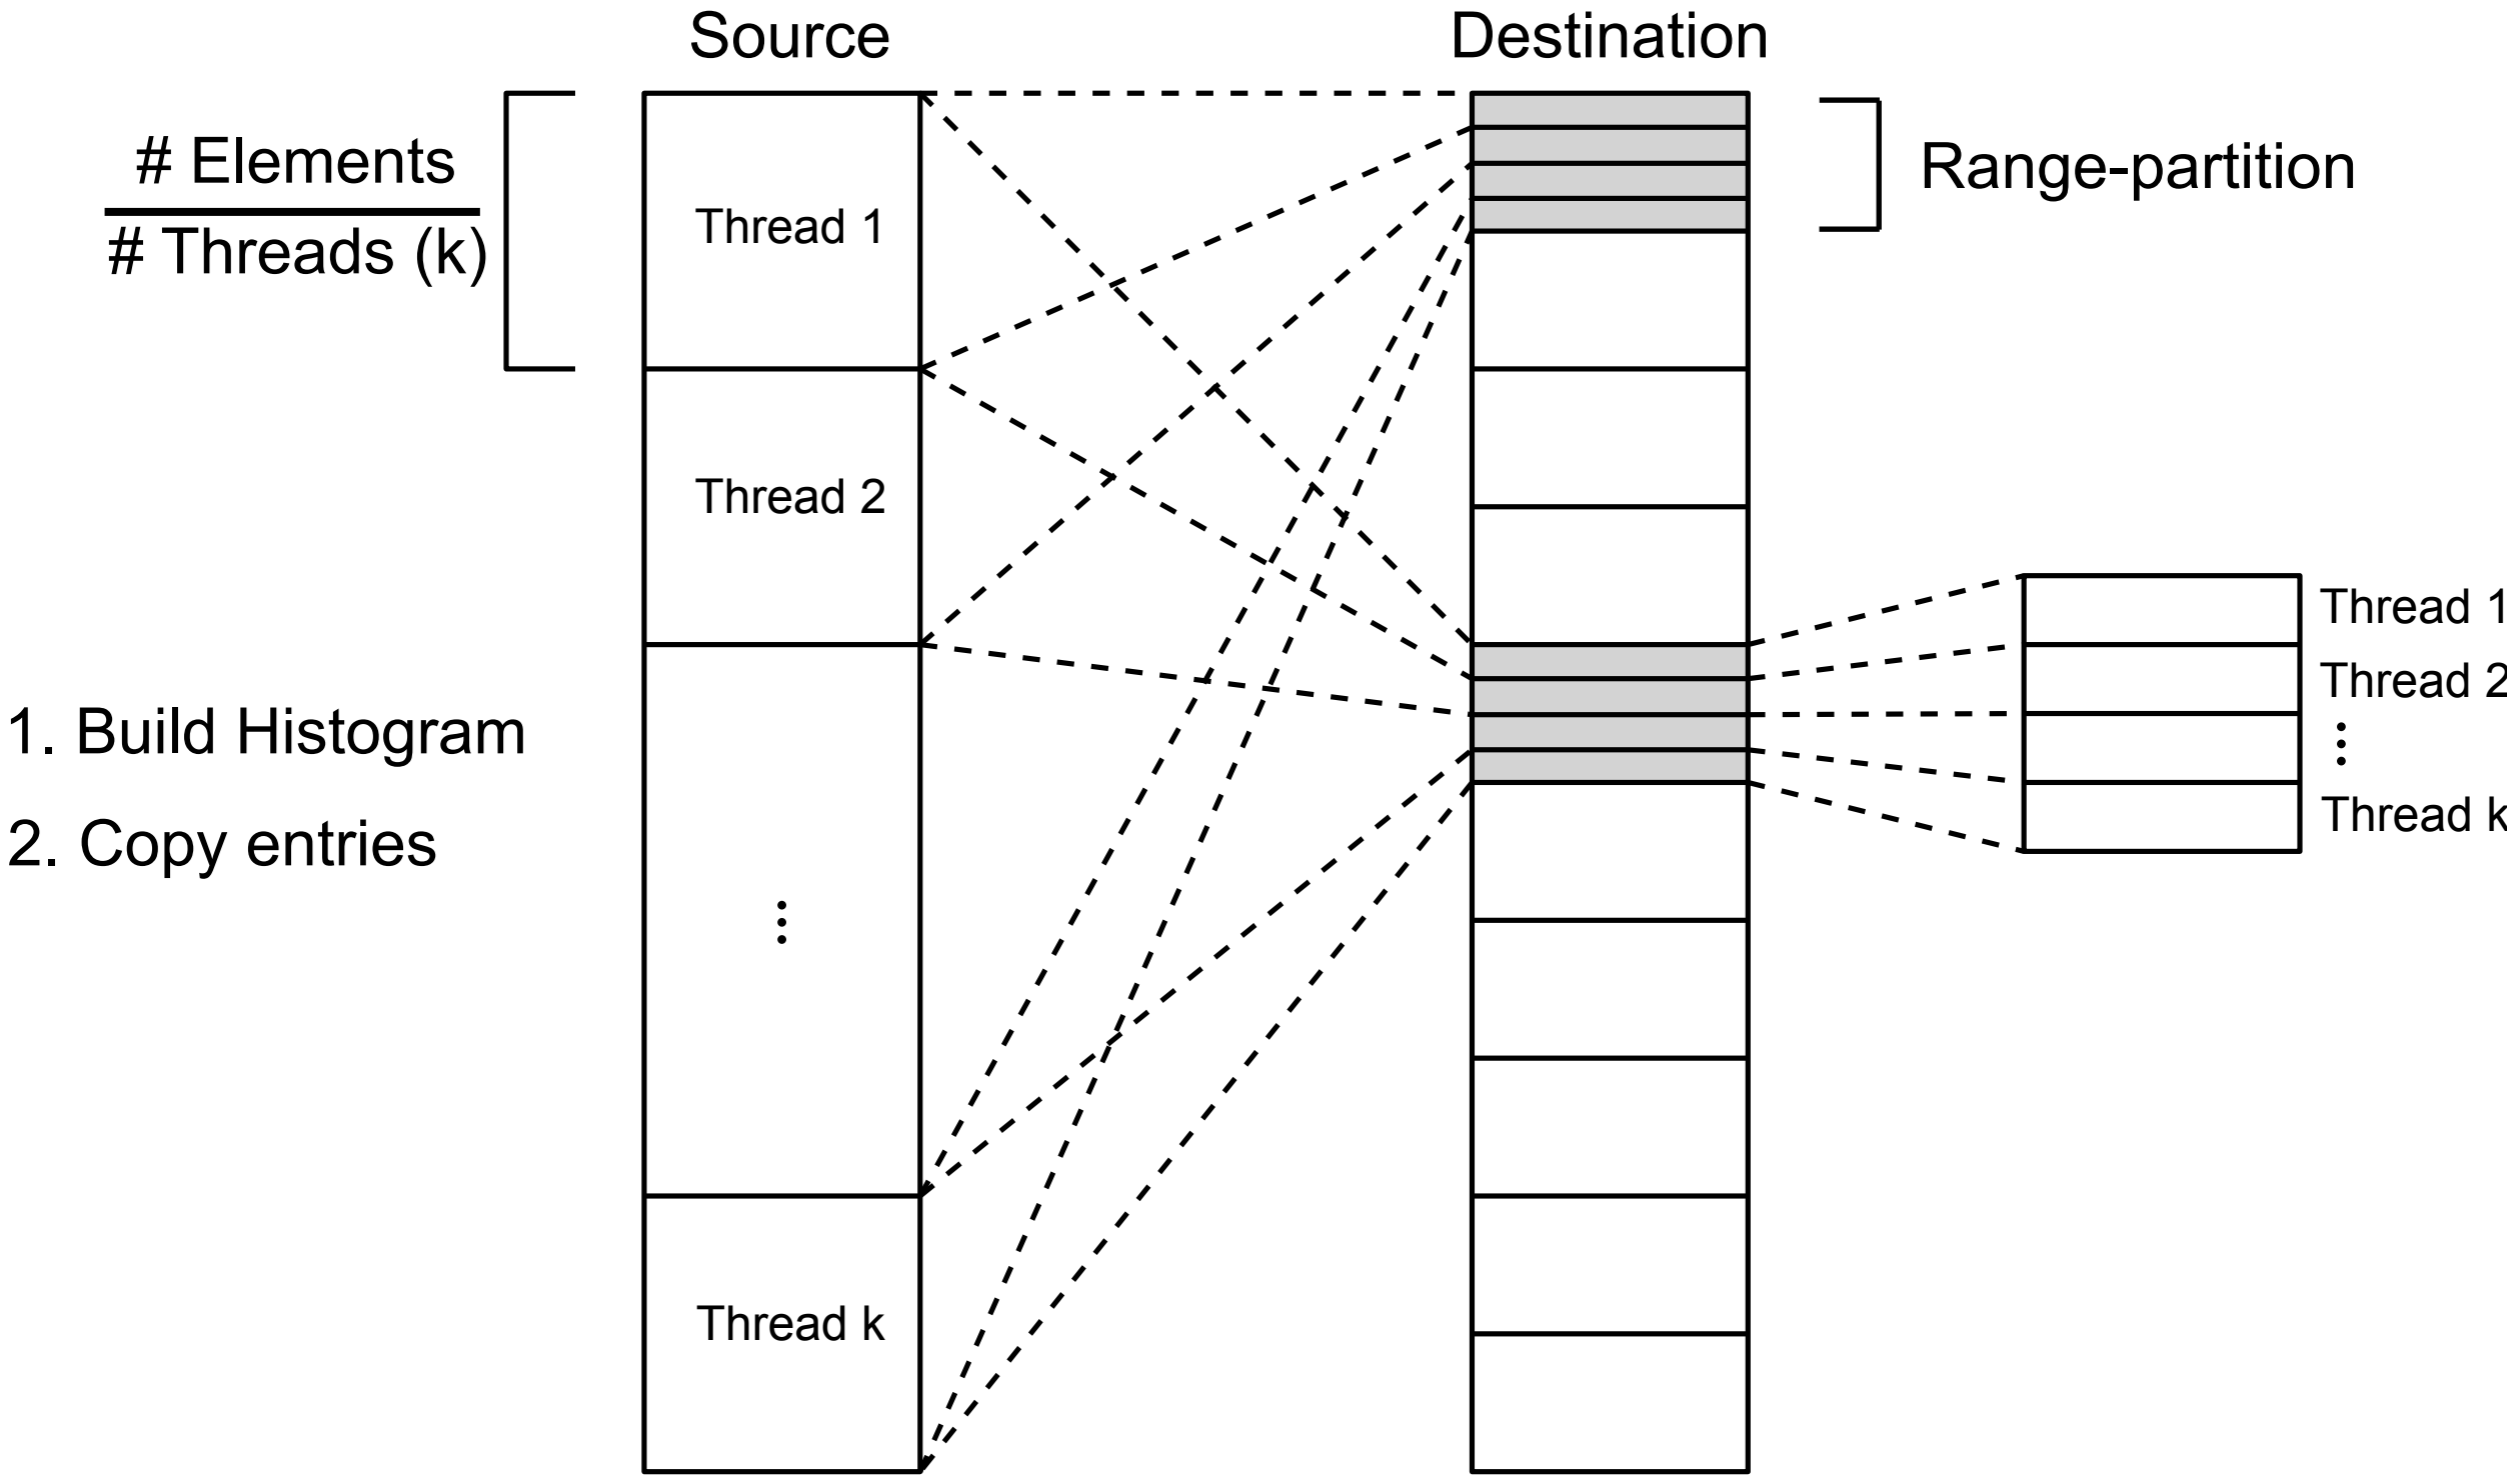
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning



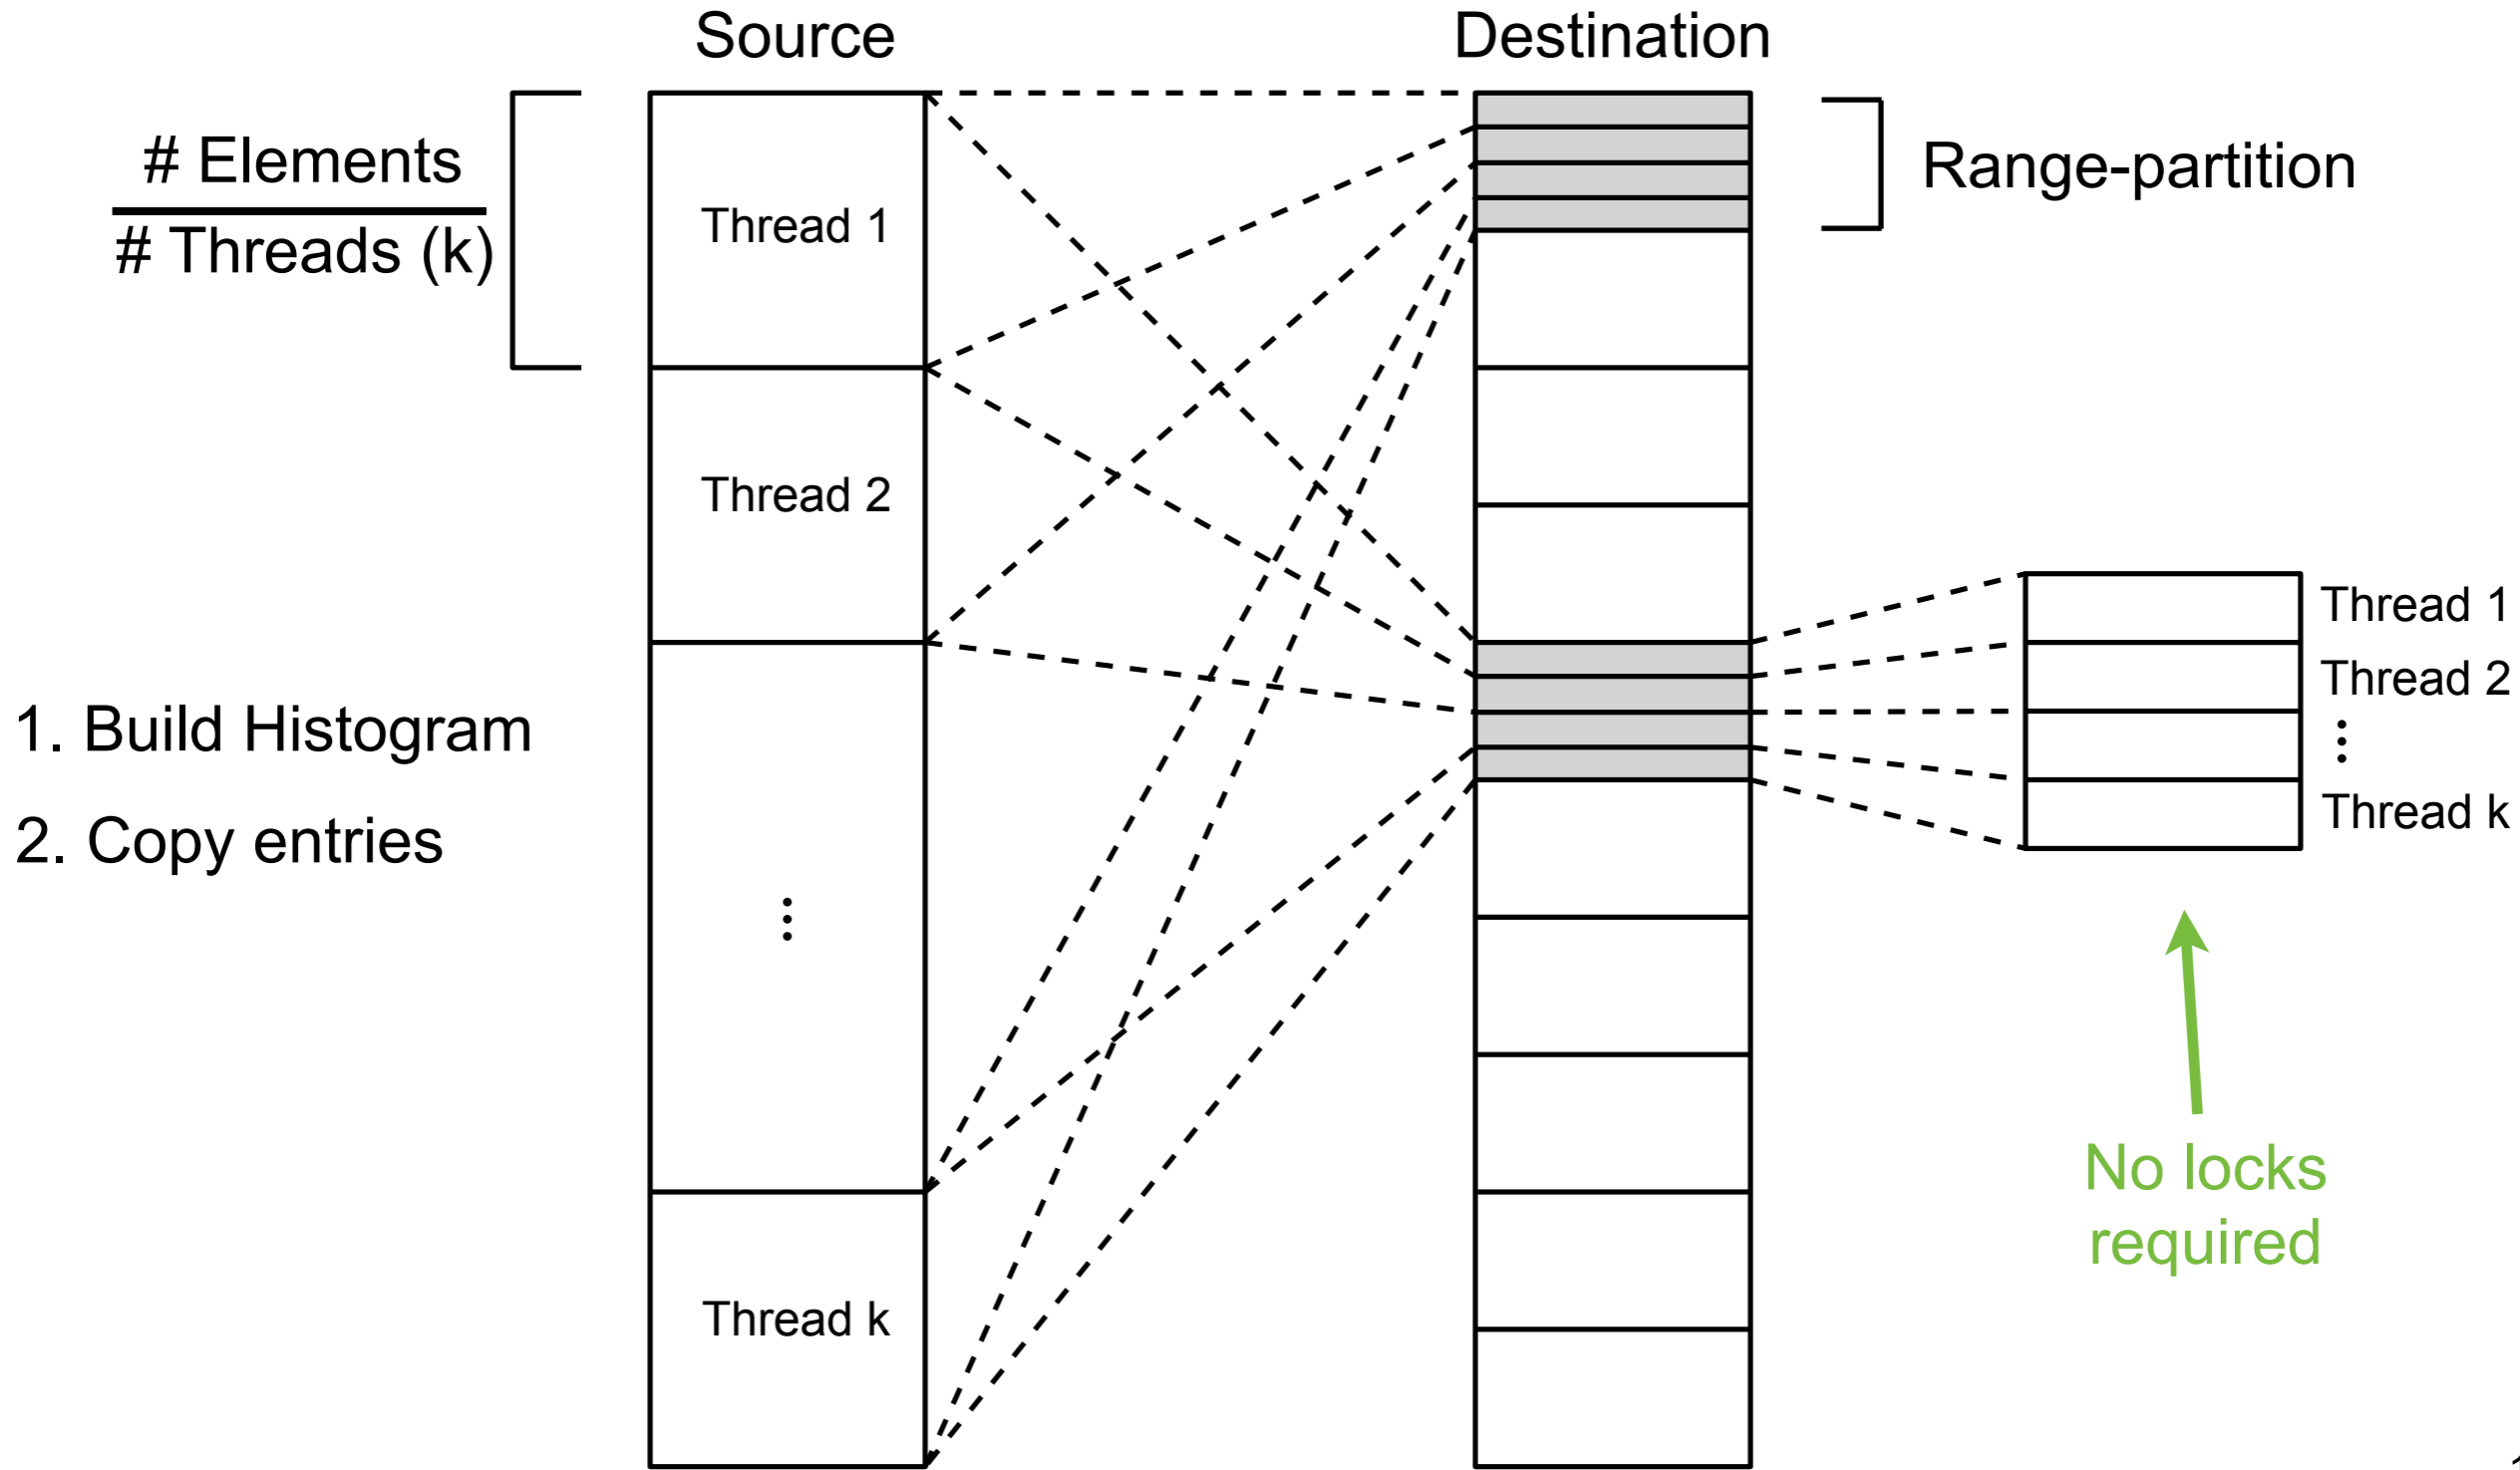
# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning



# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning

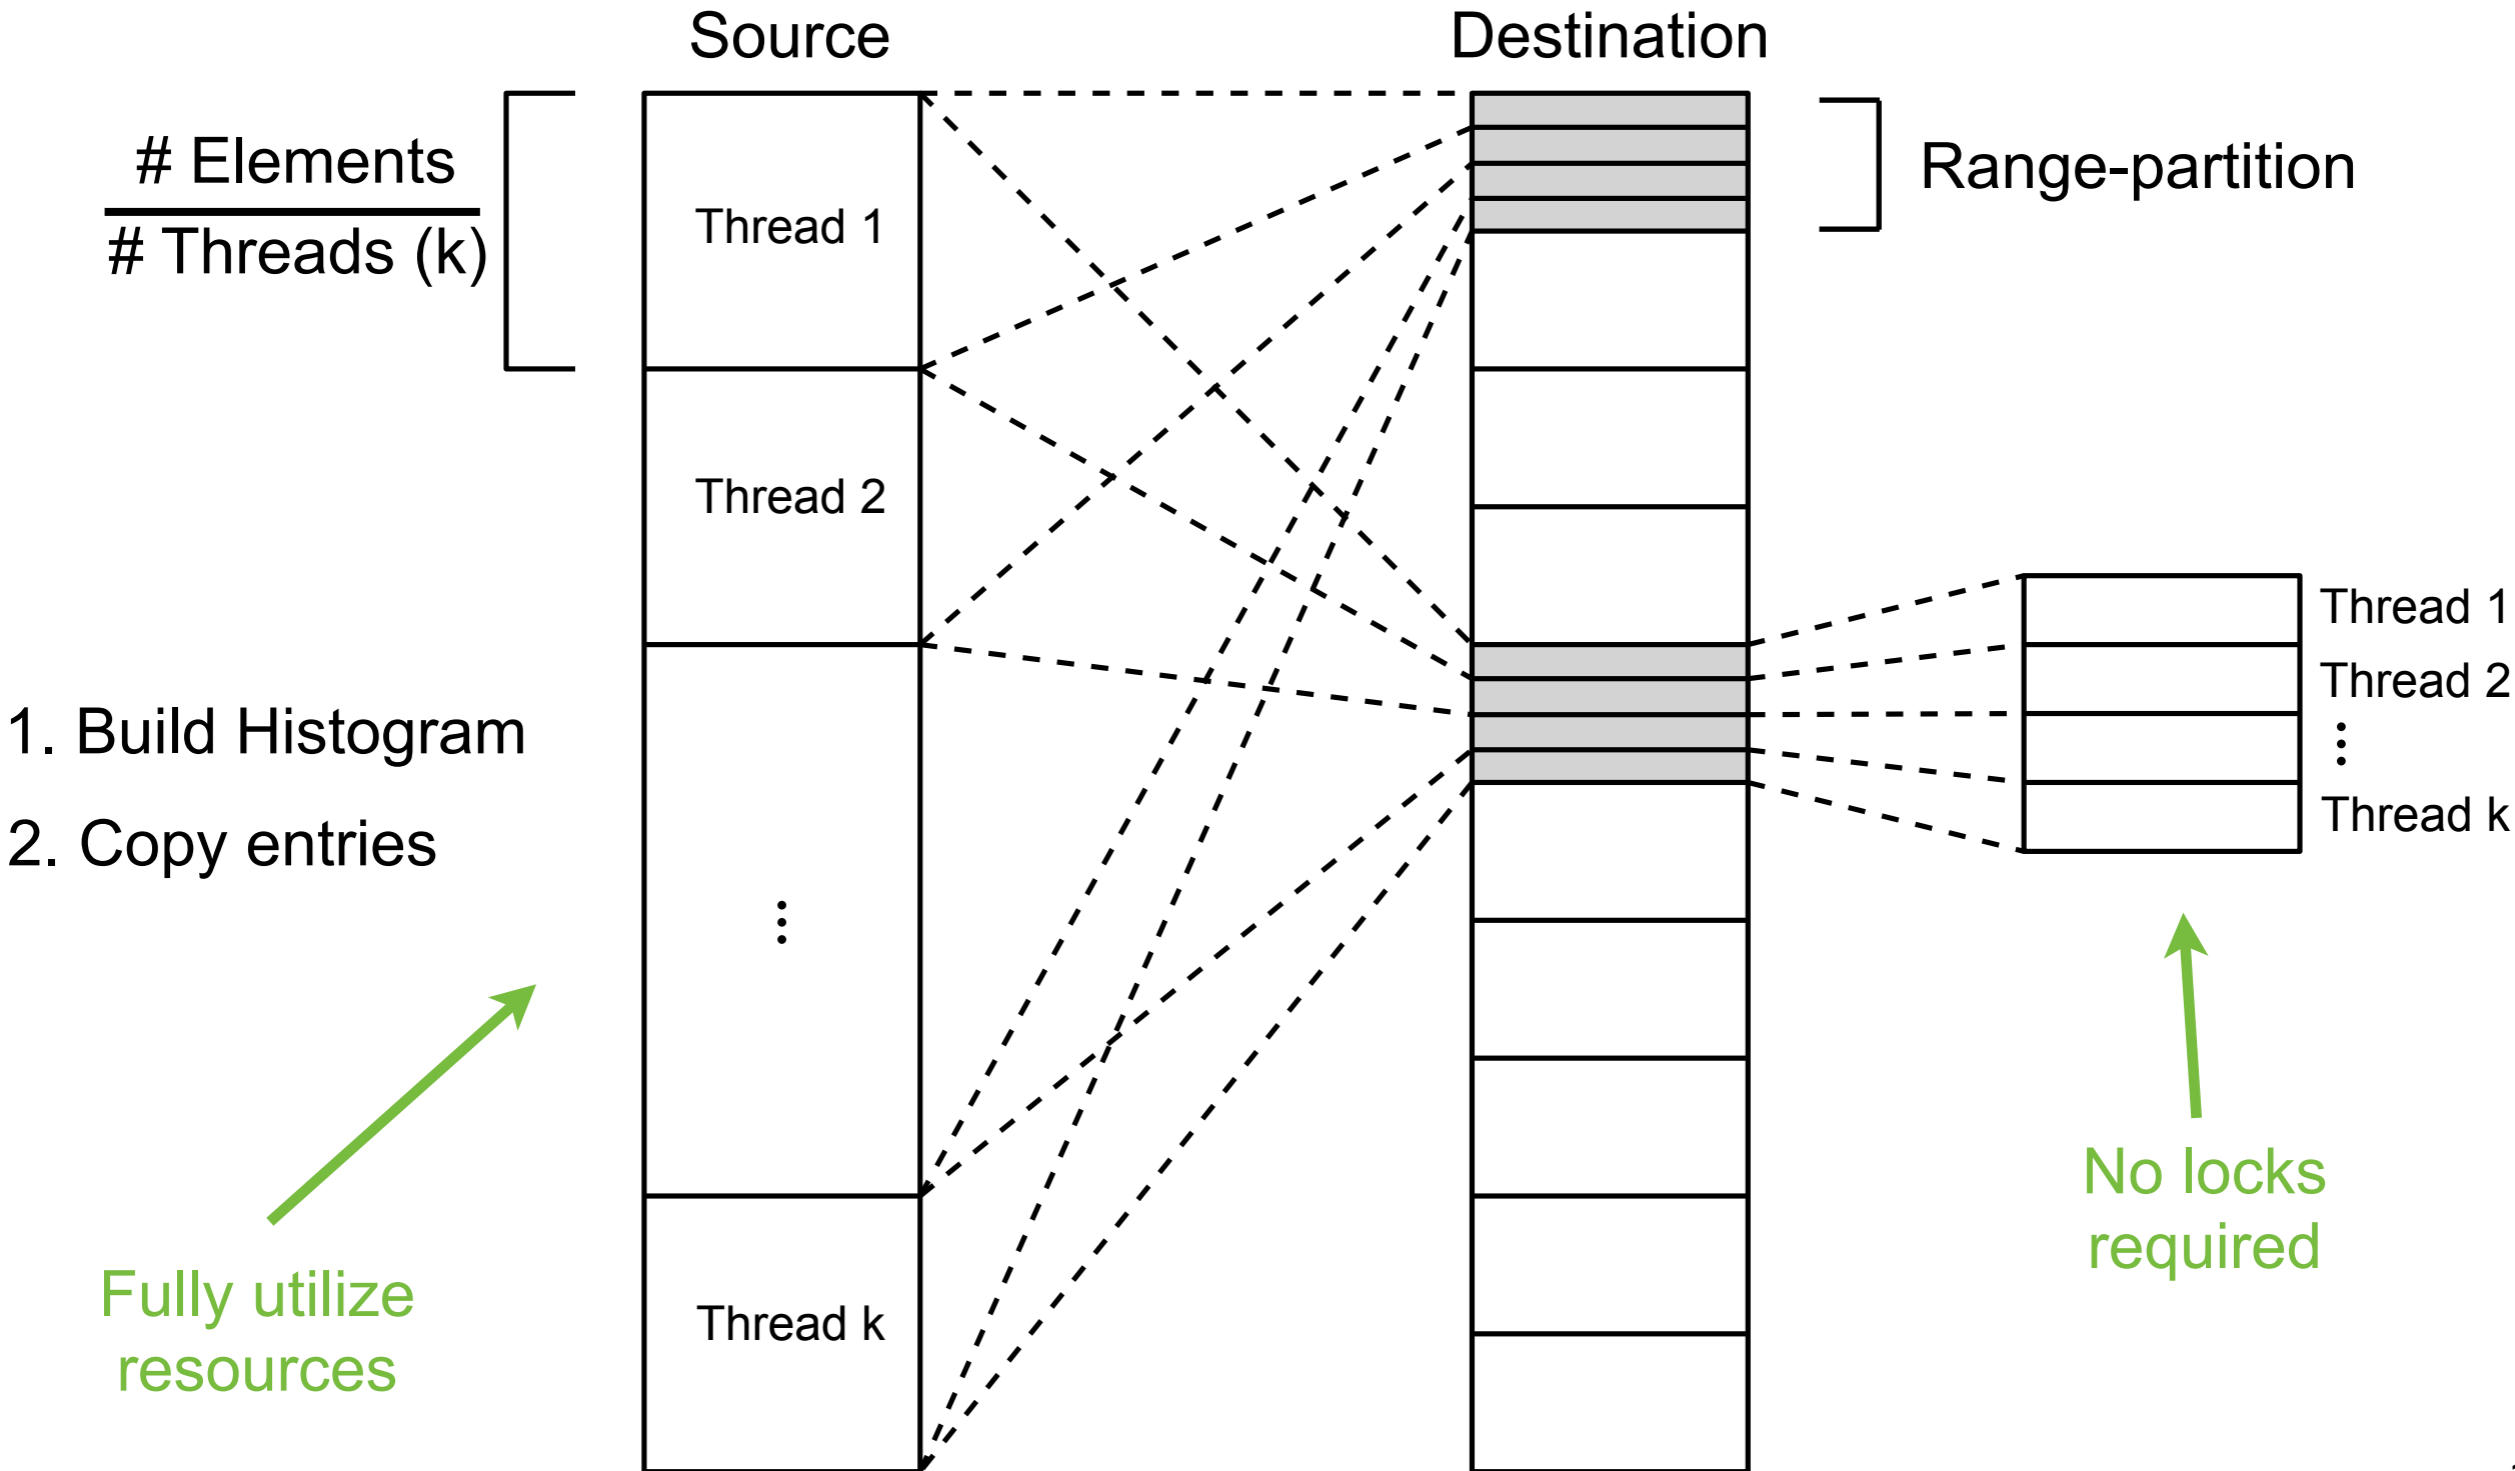


# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning

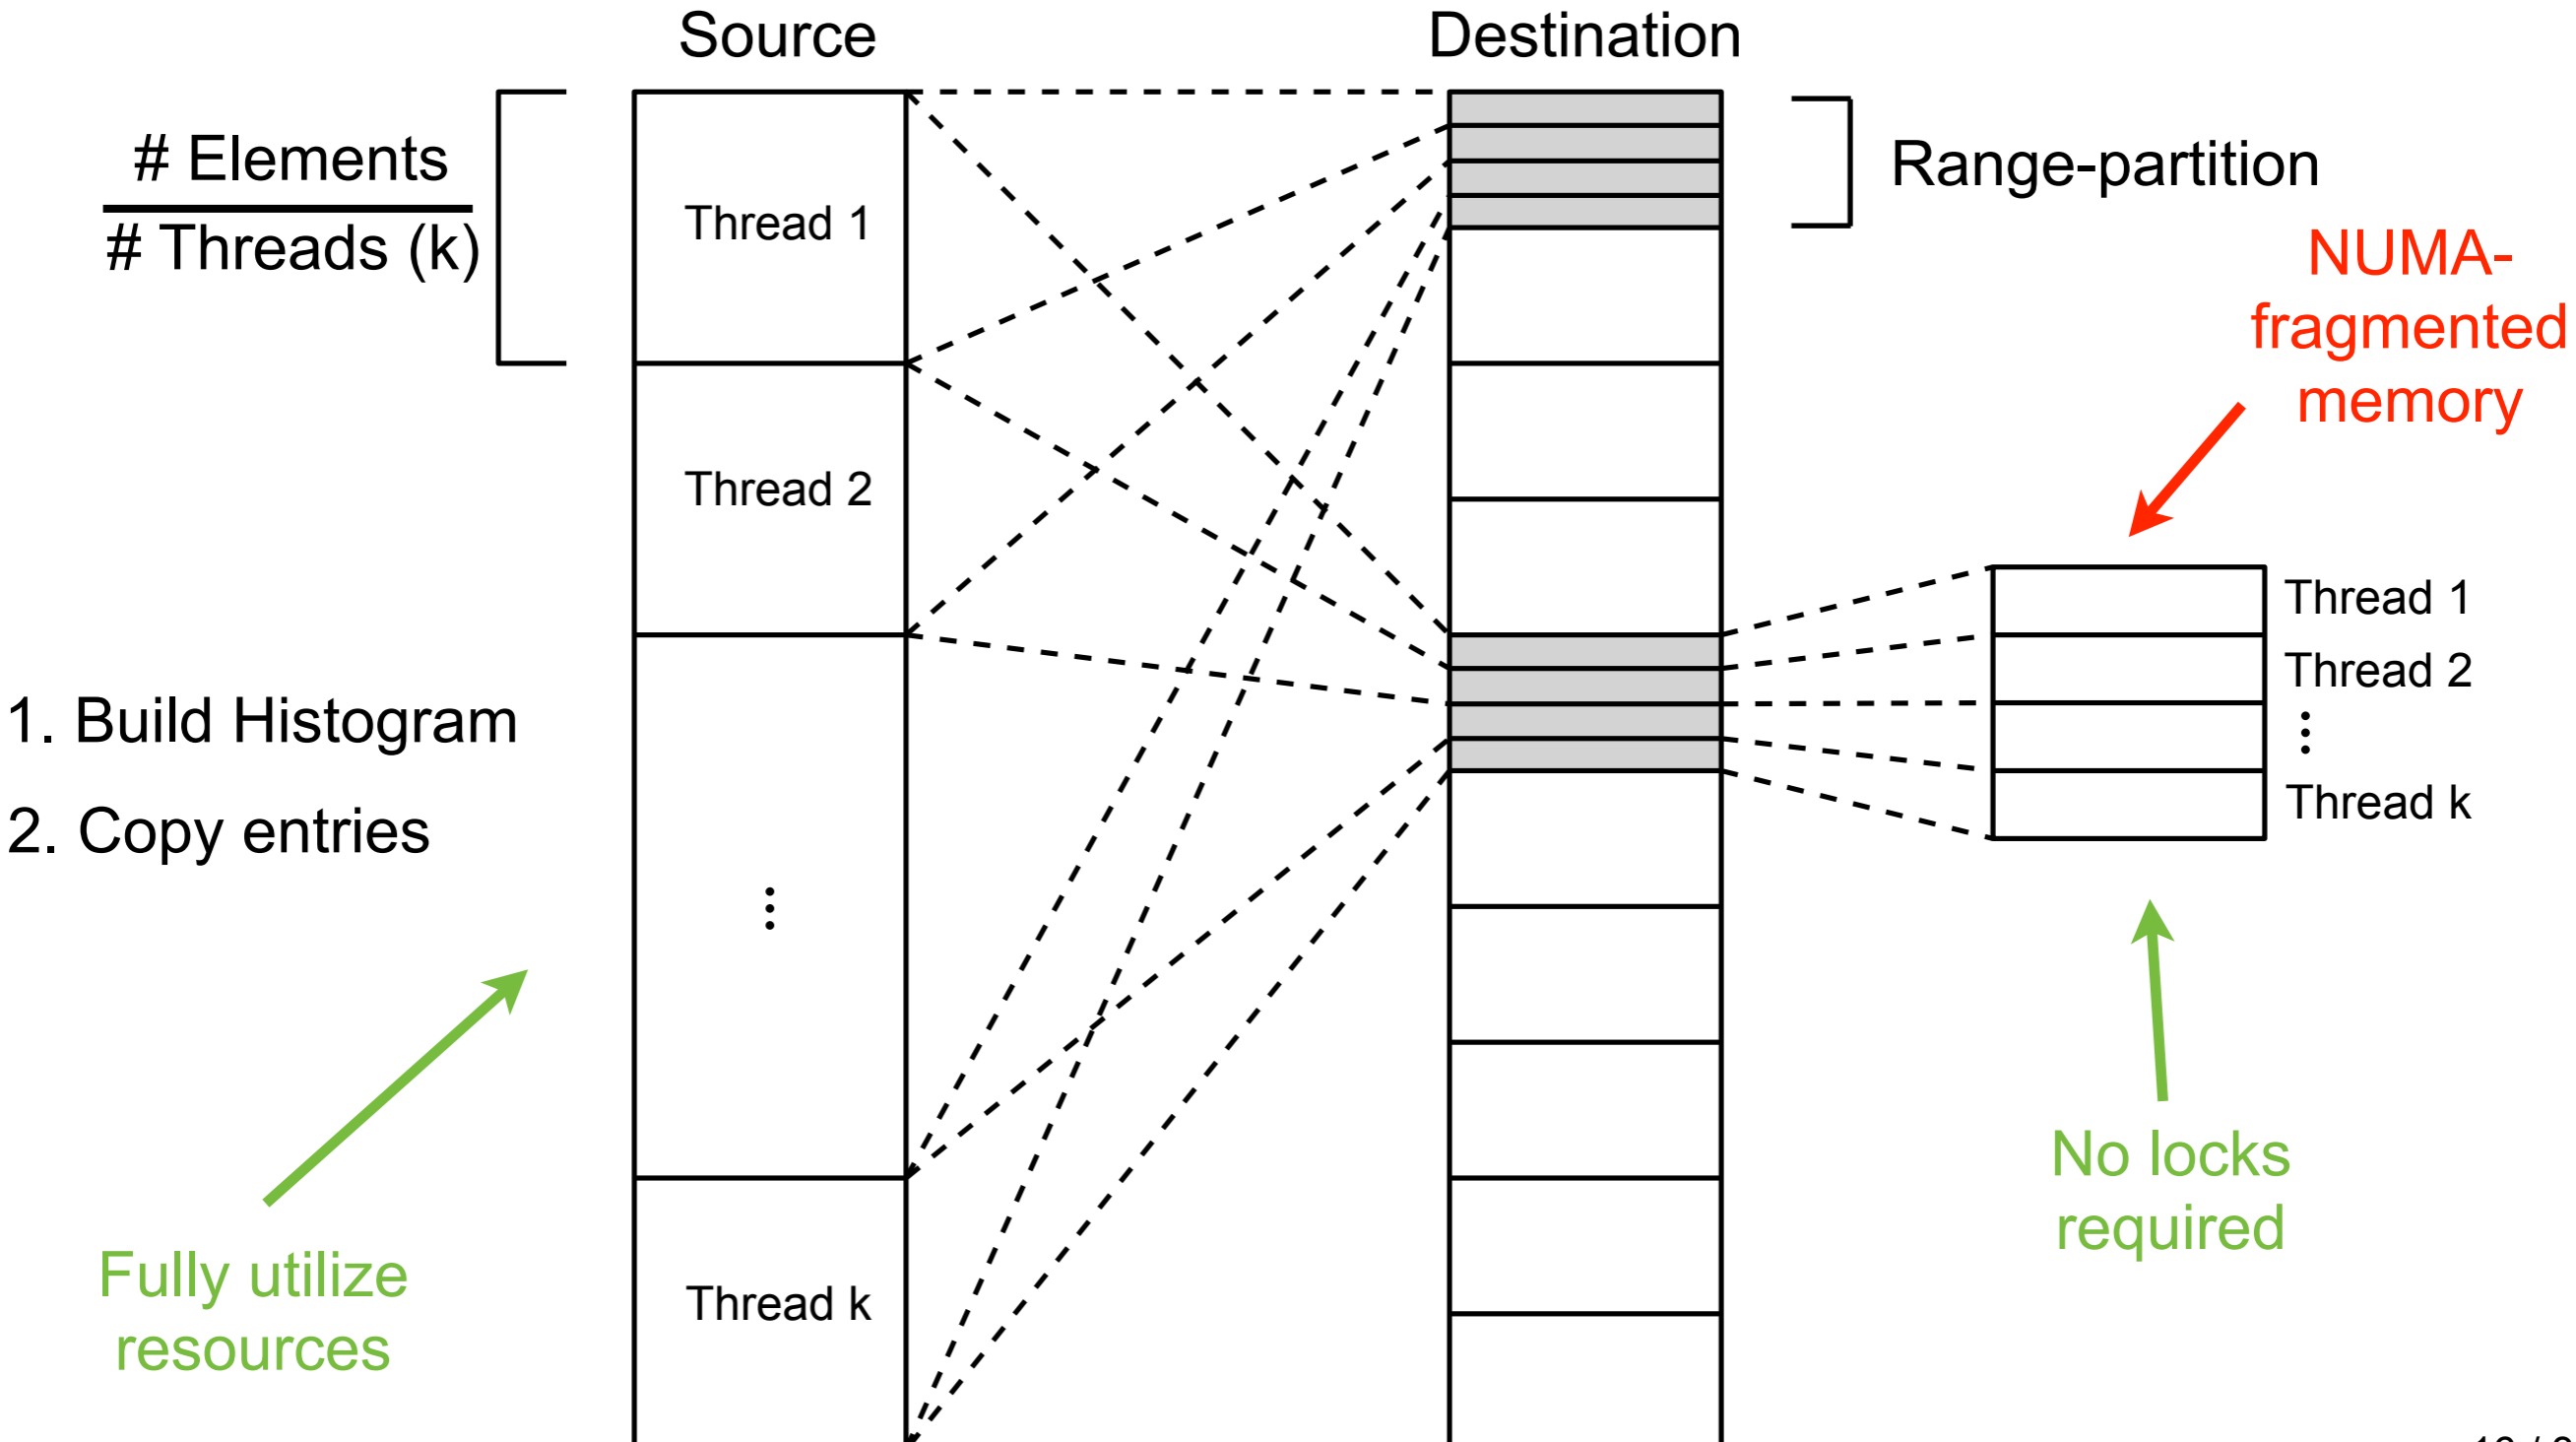





# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning



# Multi-threaded algorithms: Parallel Coarse-Granular Index (P-CGI): Parallel Range Partitioning

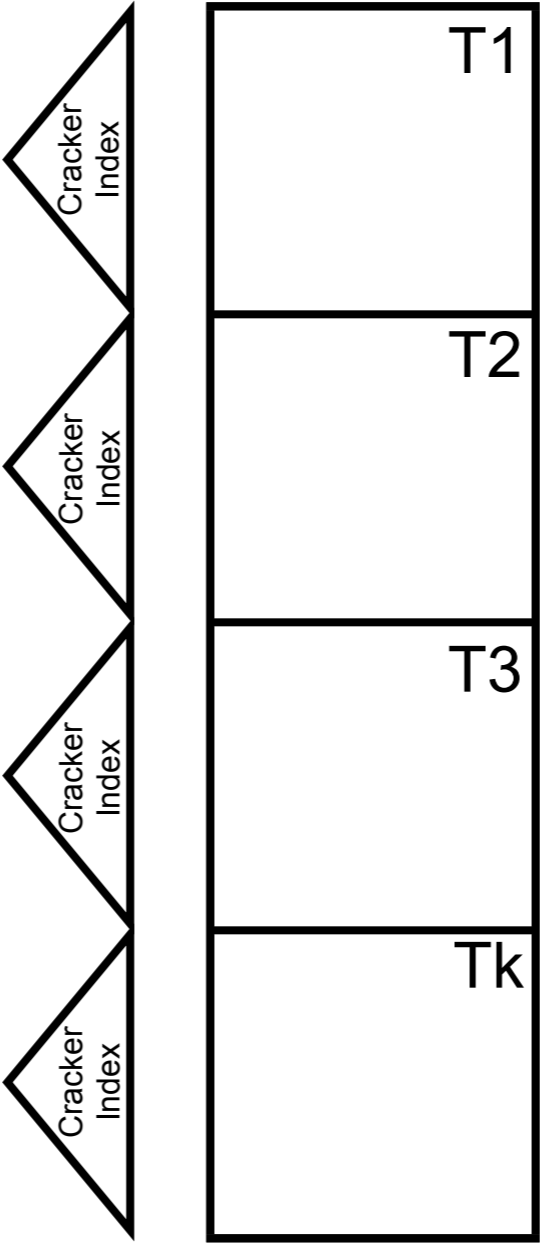


# Multi-threaded algorithms: Parallel-chunked Standard Cracking (P-CSC)



k Chunks


# Multi-threaded algorithms: Parallel-chunked Coarse-Granular Index (P-CCGI)



k Chunks


# Multi-threaded algorithms: Parallel-chunked Coarse-Granular Index (P-CCGI)

Range-partitioning

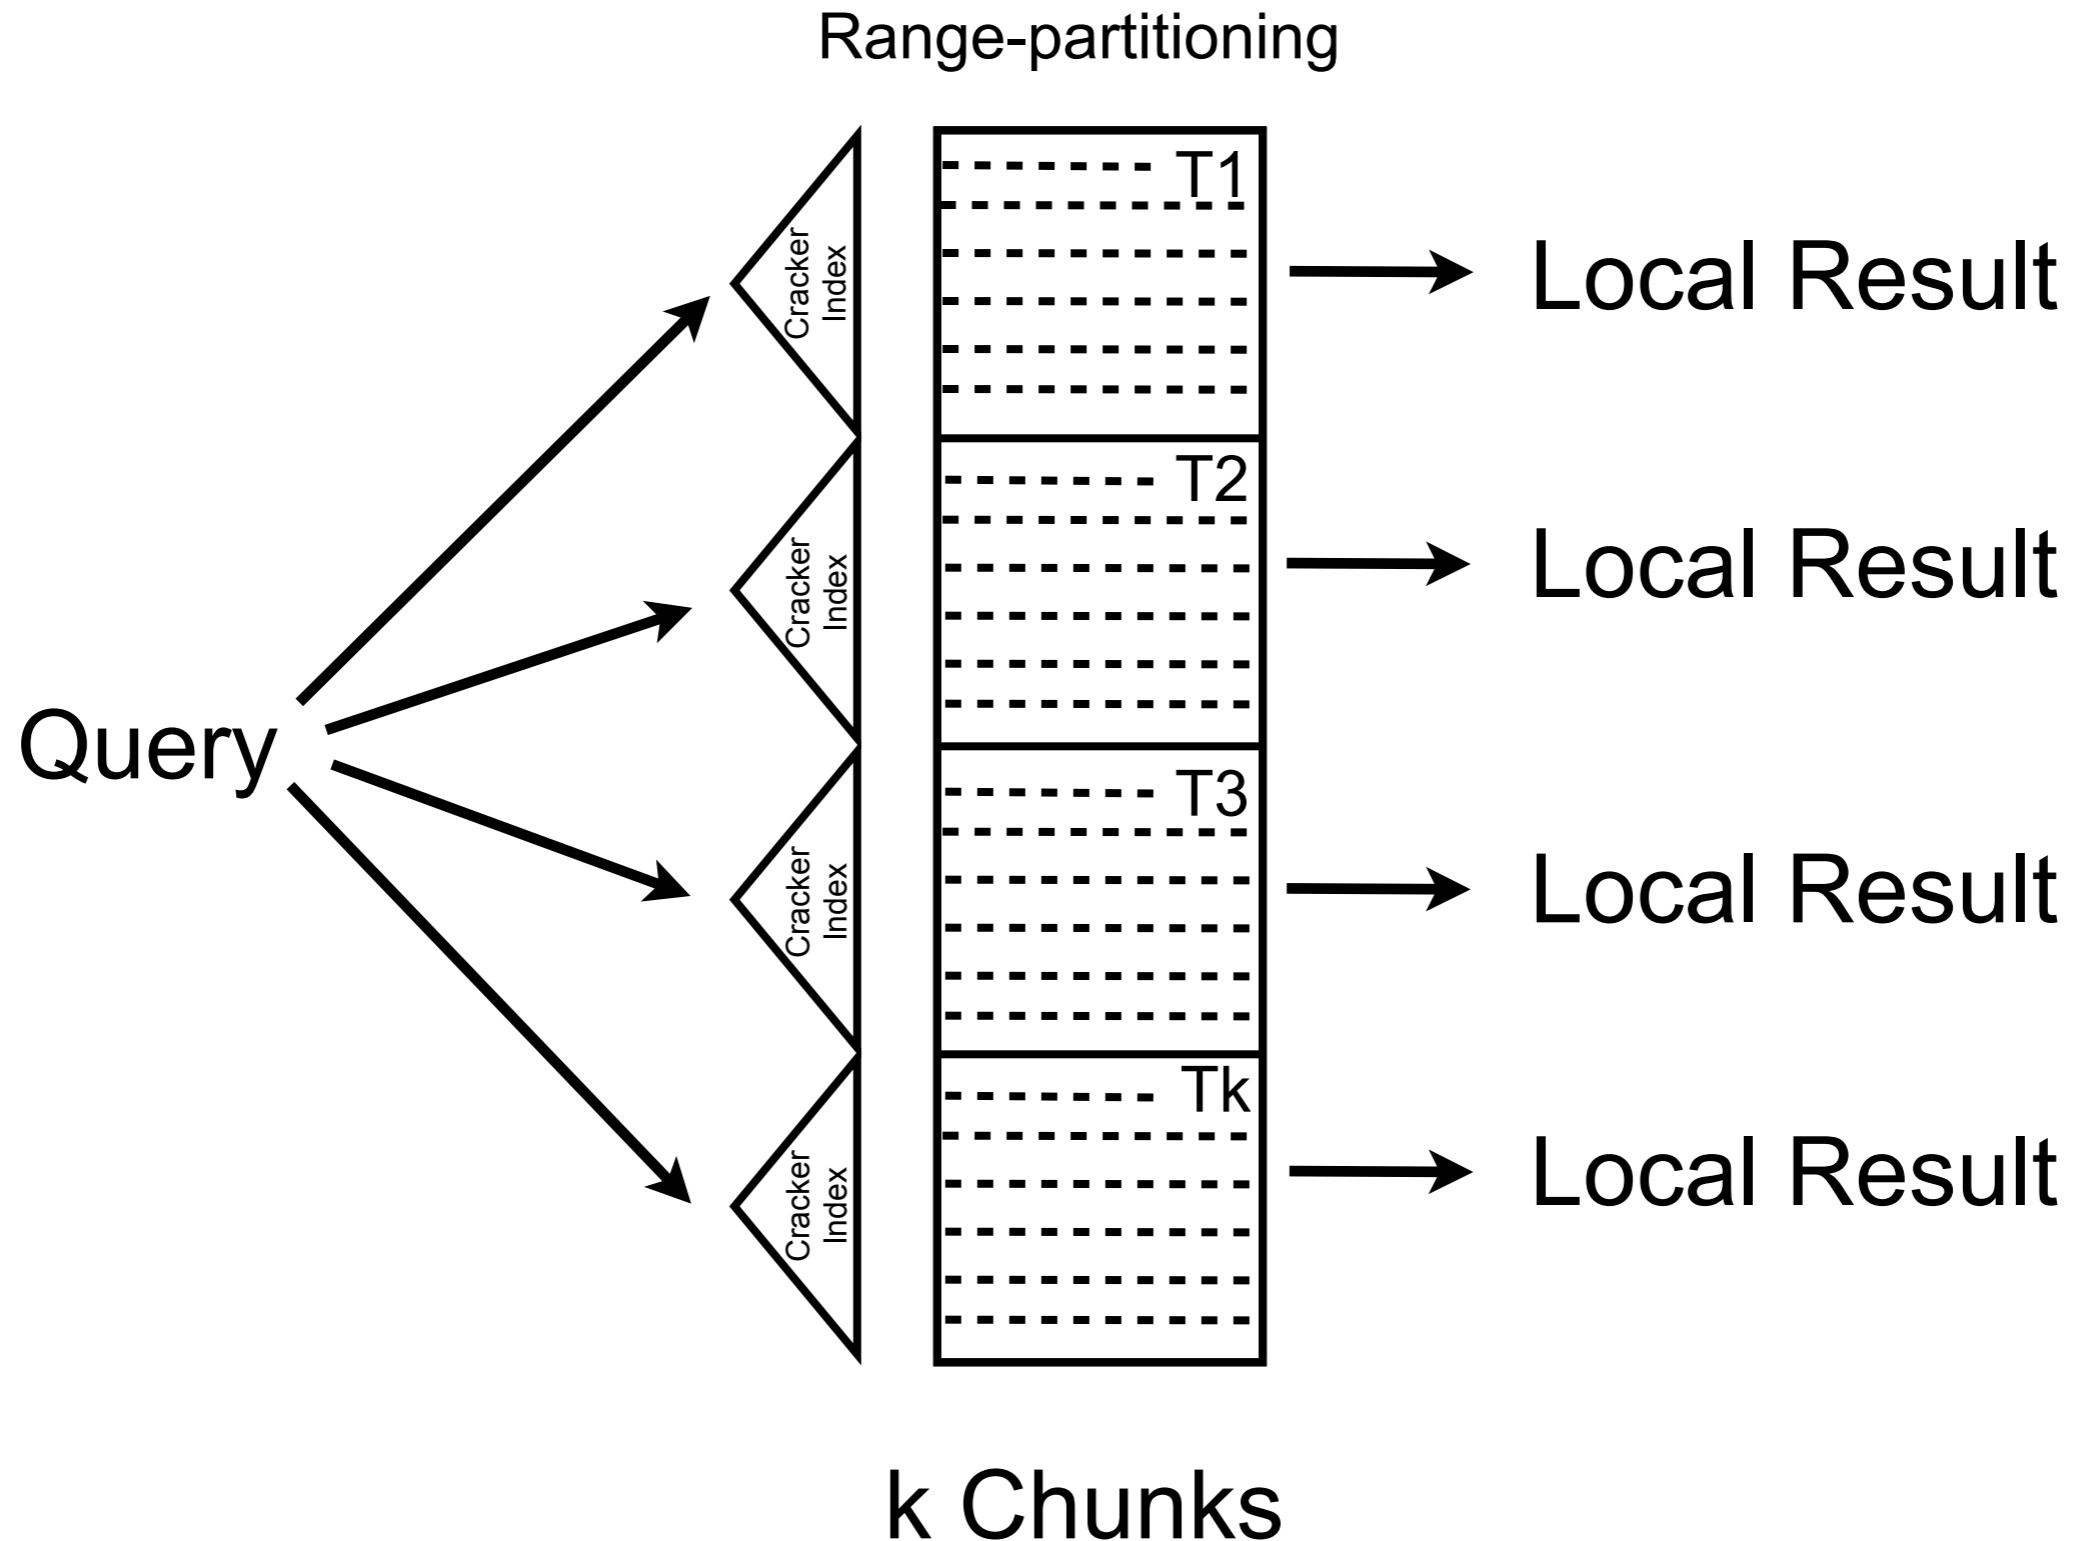


k Chunks

# Multi-threaded algorithms: Parallel-chunked Coarse-Granular Index (P-CCGI)




# Multi-threaded algorithms: Parallel-chunked Coarse-Granular Index (P-CCGI)



Multi-threaded algorithms:

Parallel-chunked Coarse-Granular Index (P-CCGI)


P-CSC + Range Partitioning



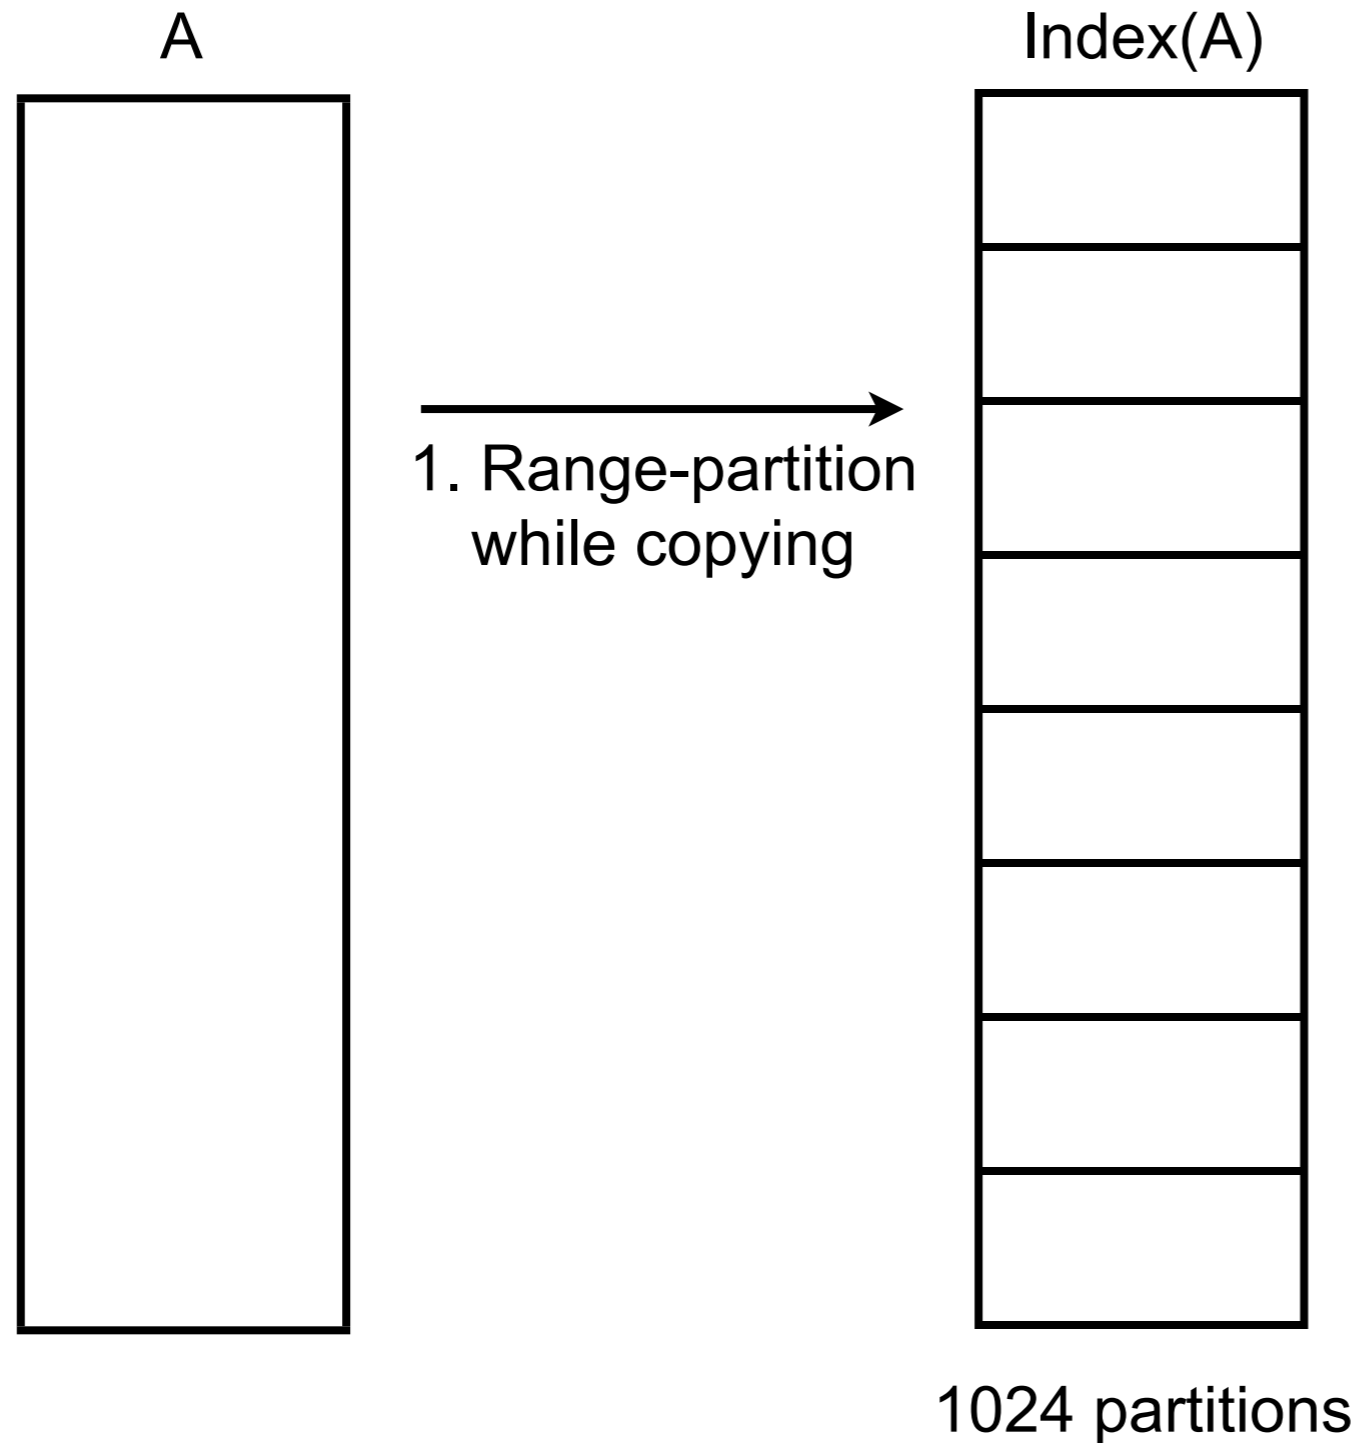


# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)


A



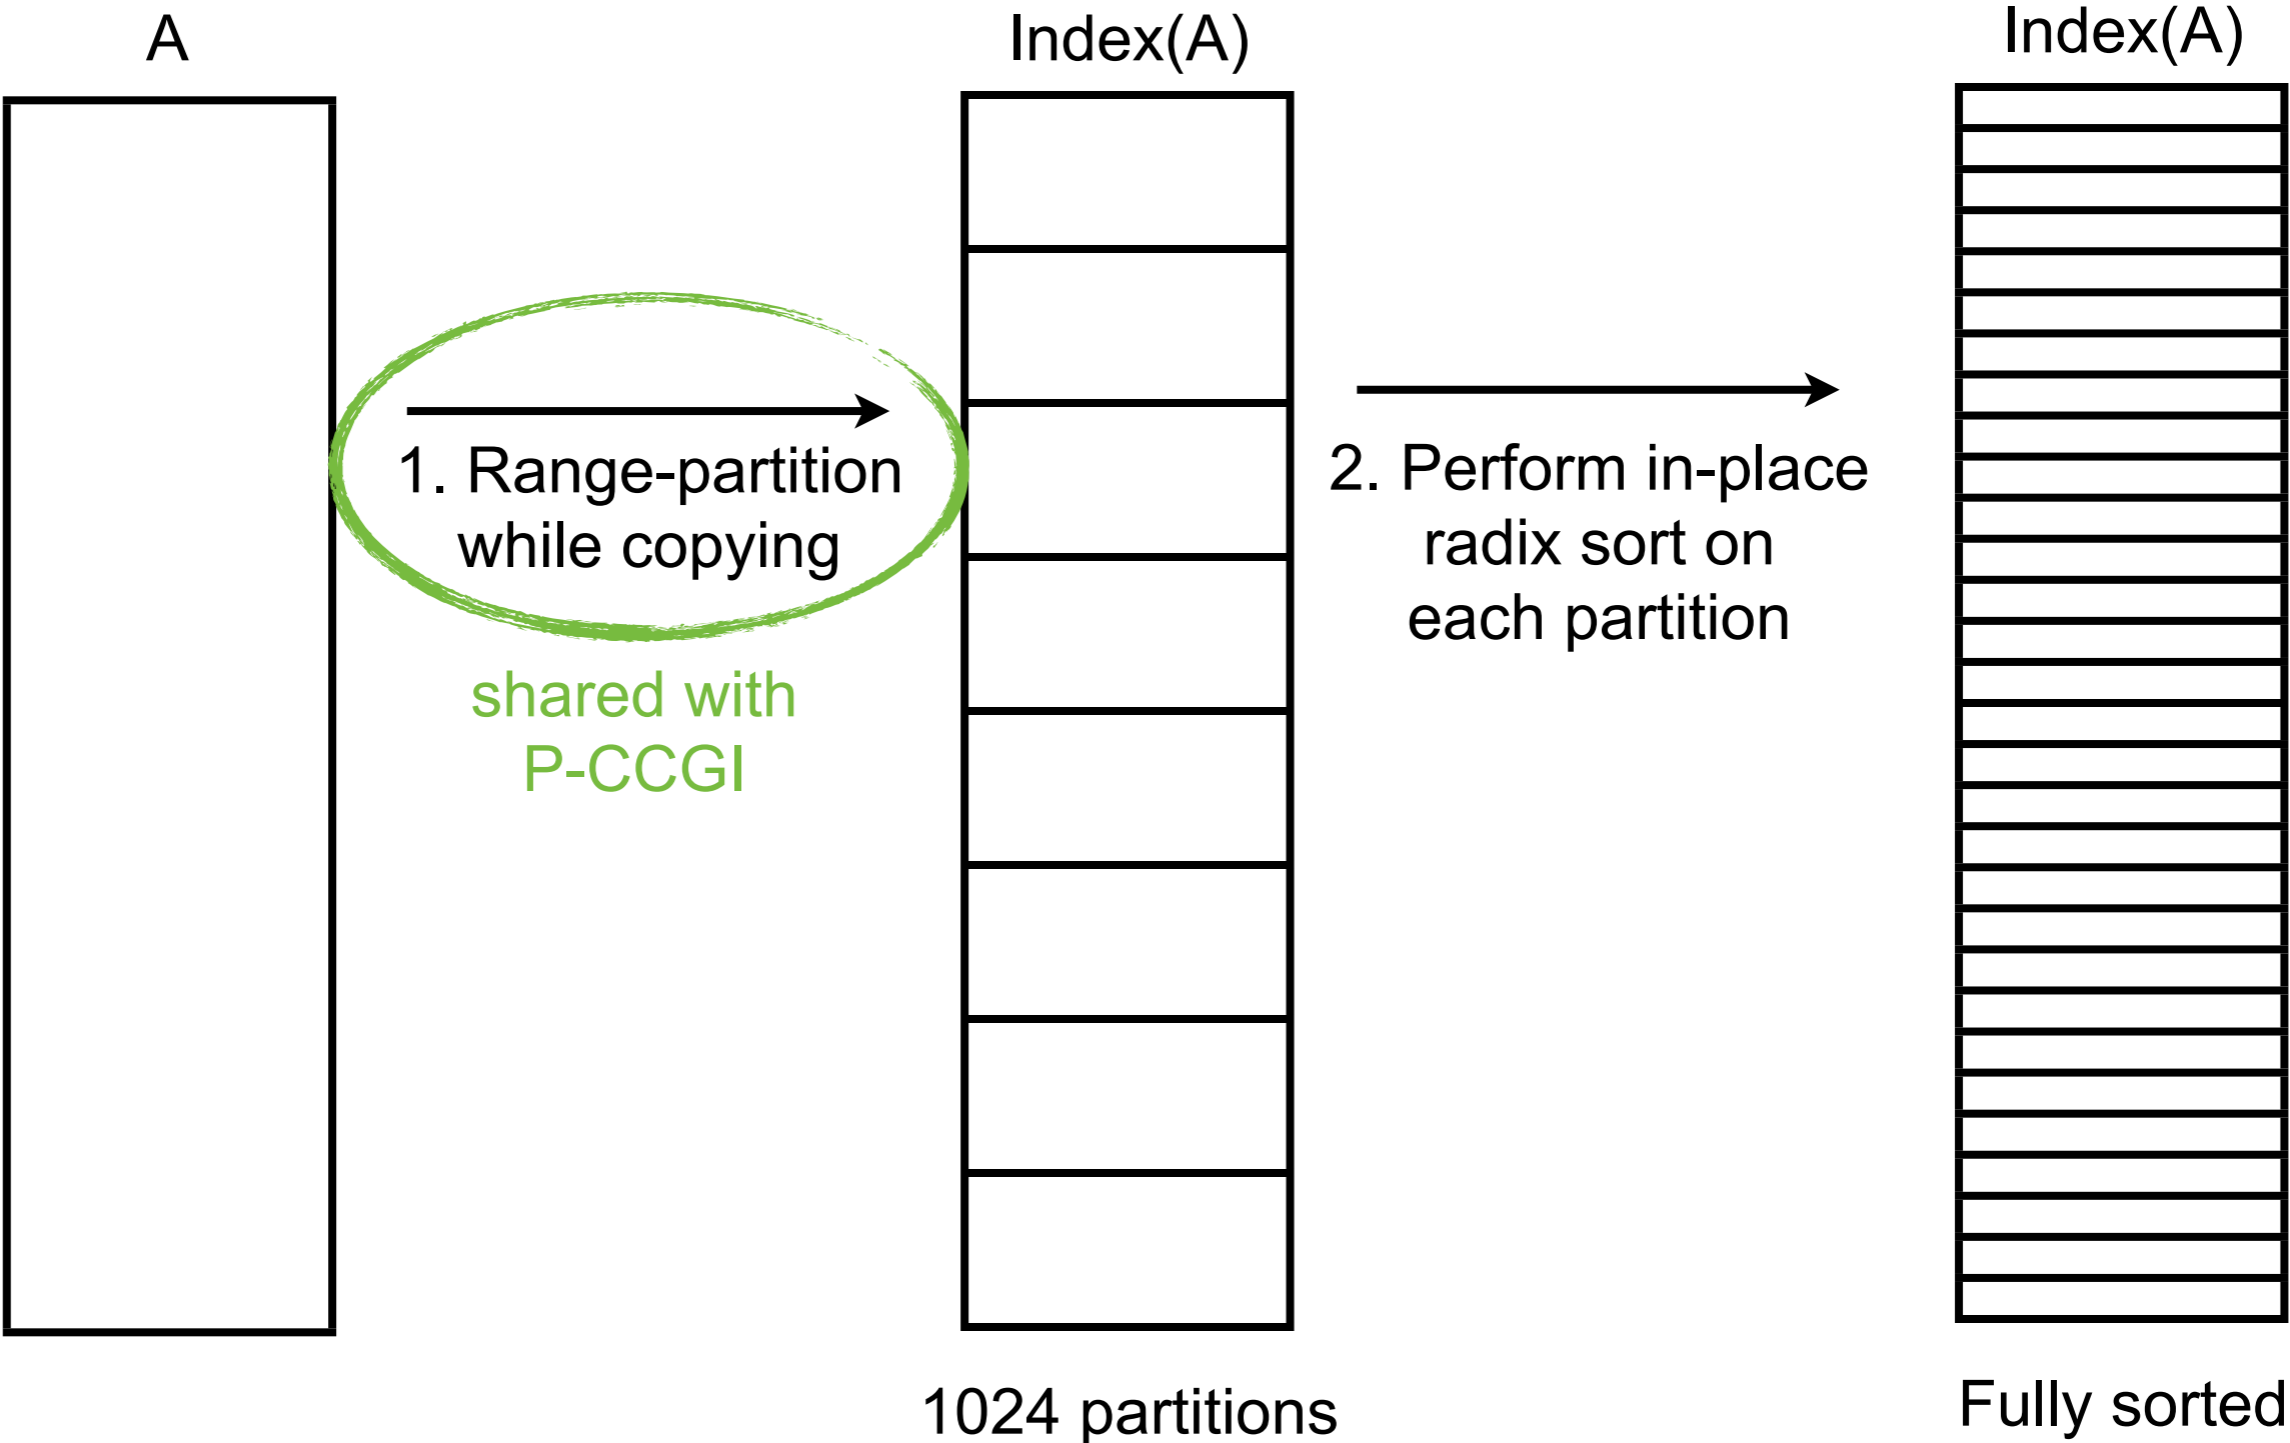
# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)



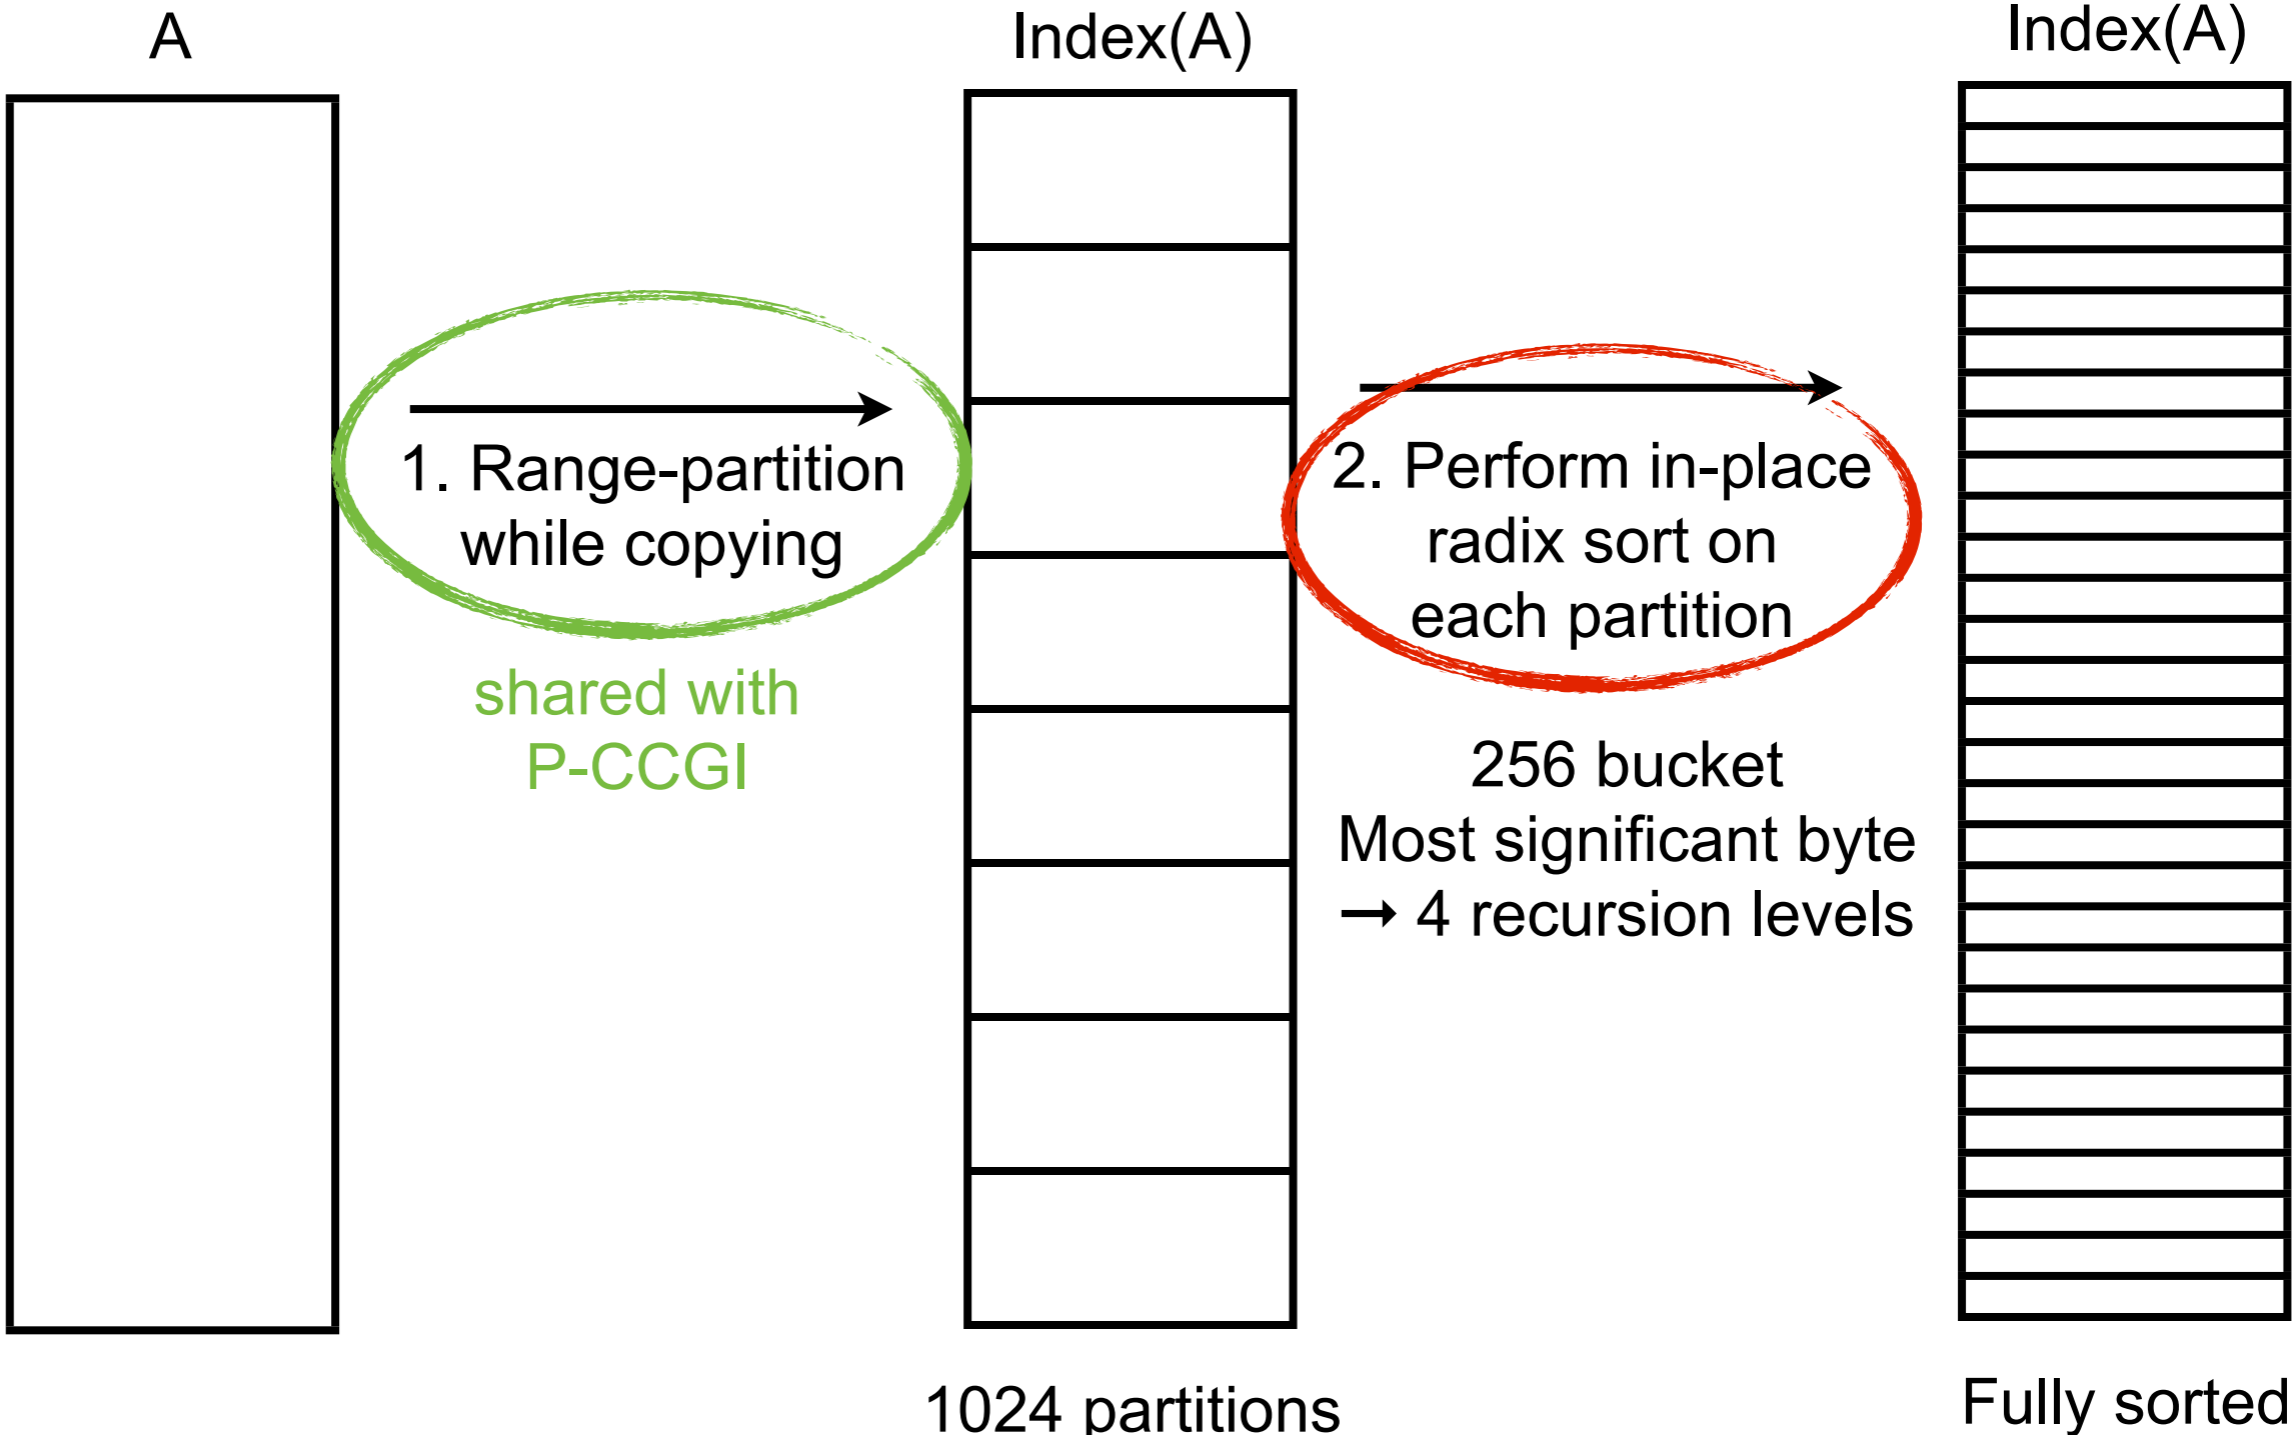
# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)




# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)



# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)

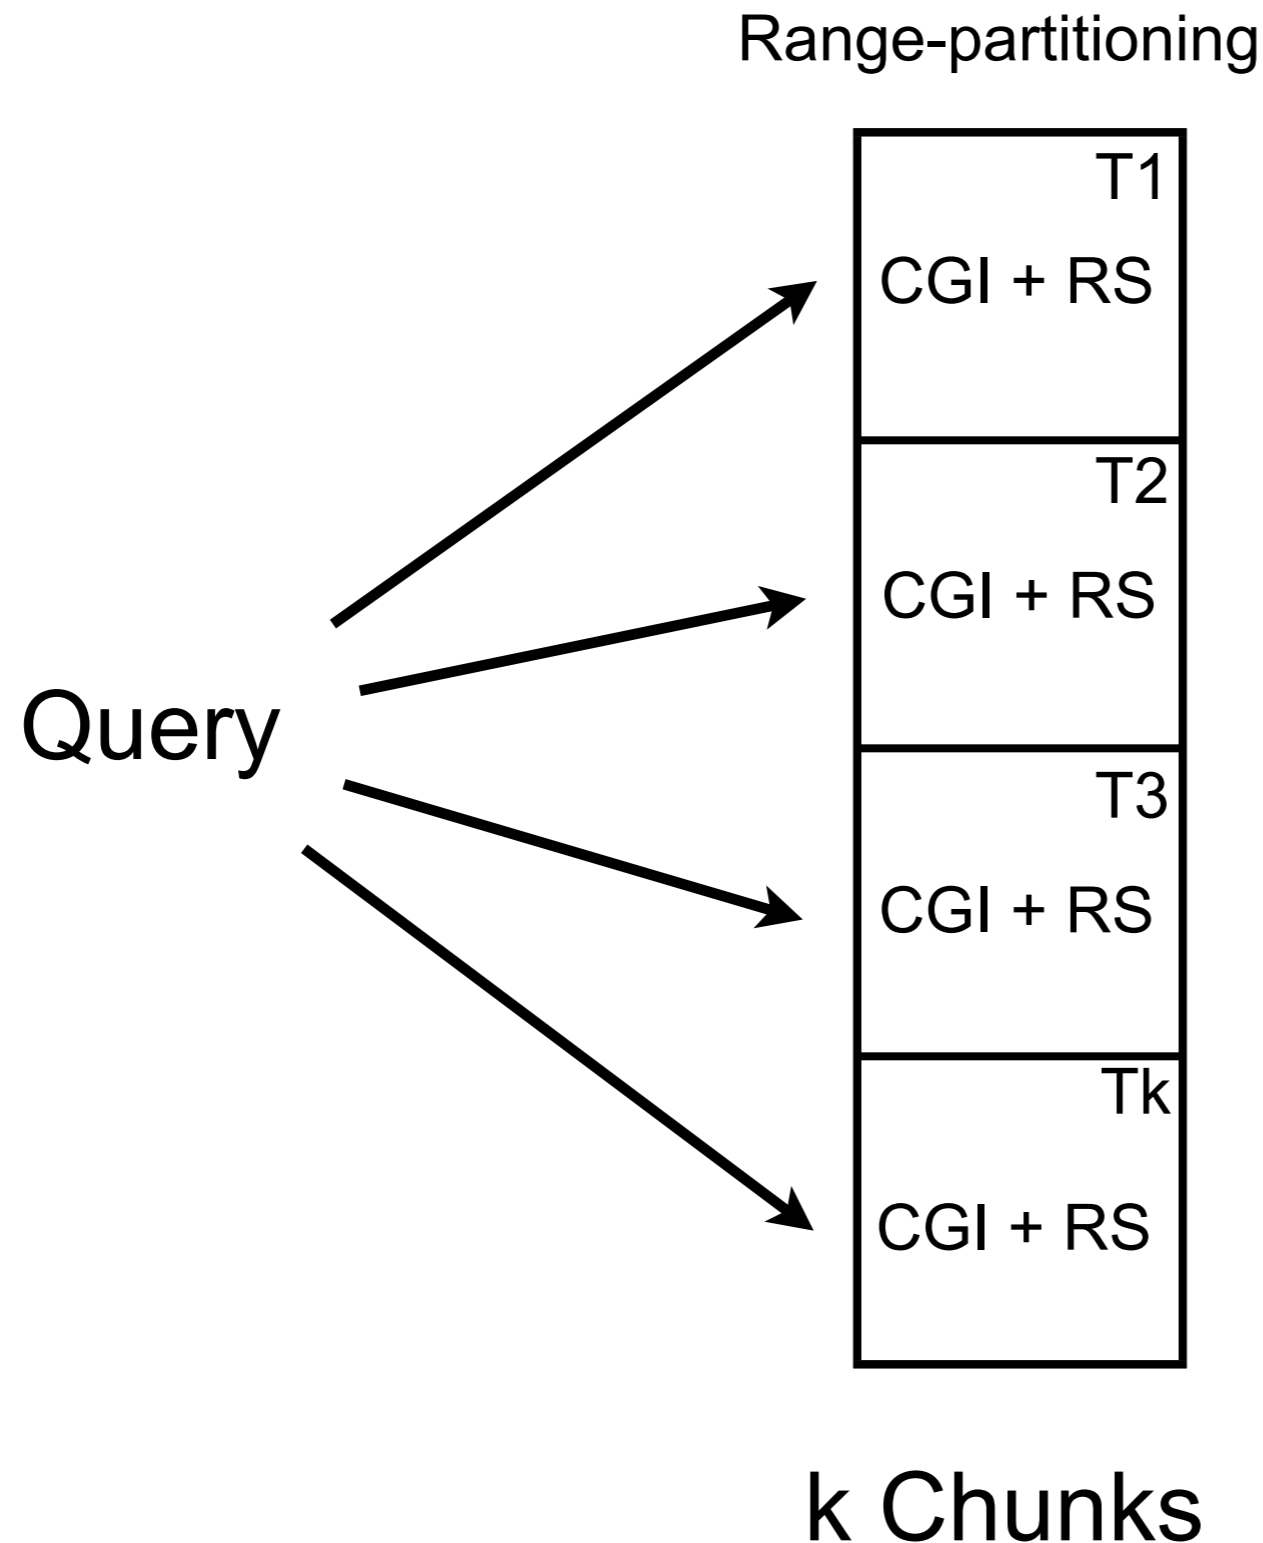


# Multi-threaded algorithms: Parallel Range-Partitioned Radix Sort (P-RPRS)



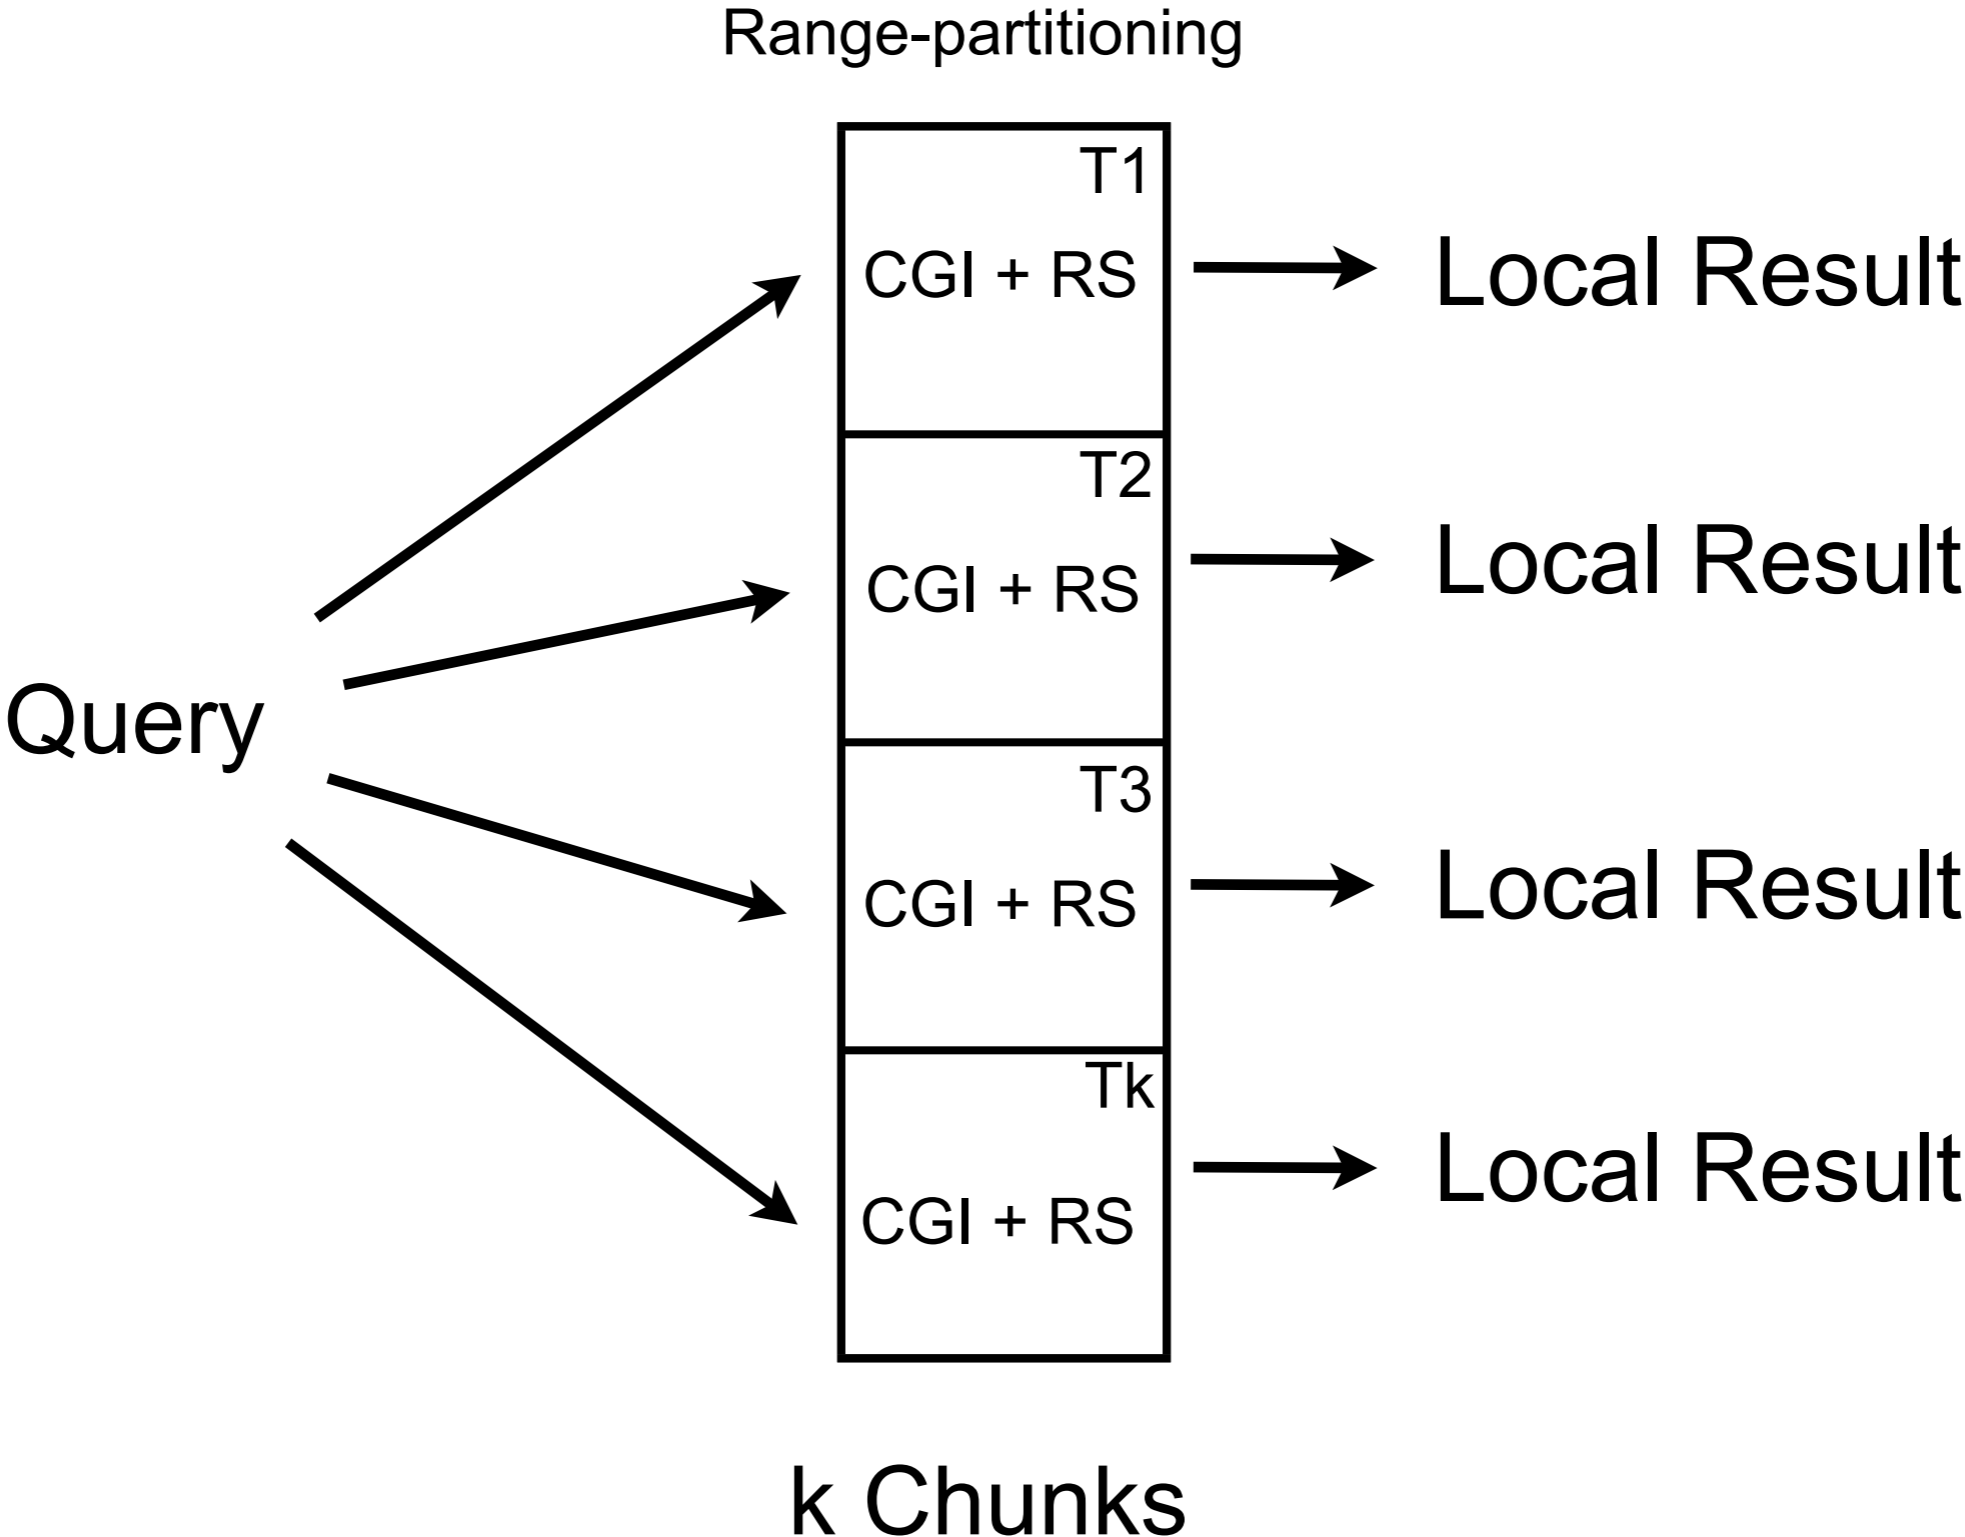
4 Cores / 8 Threads  
512 million 4 byte integers  
Uniform random distribution

# Multi-threaded algorithms: Parallel-chunked Range-Partitioned Radix Sort (P-CRS)



# Multi-threaded algorithms:

## Parallel-chunked Range-Partitioned Radix Sort (P-CRS)

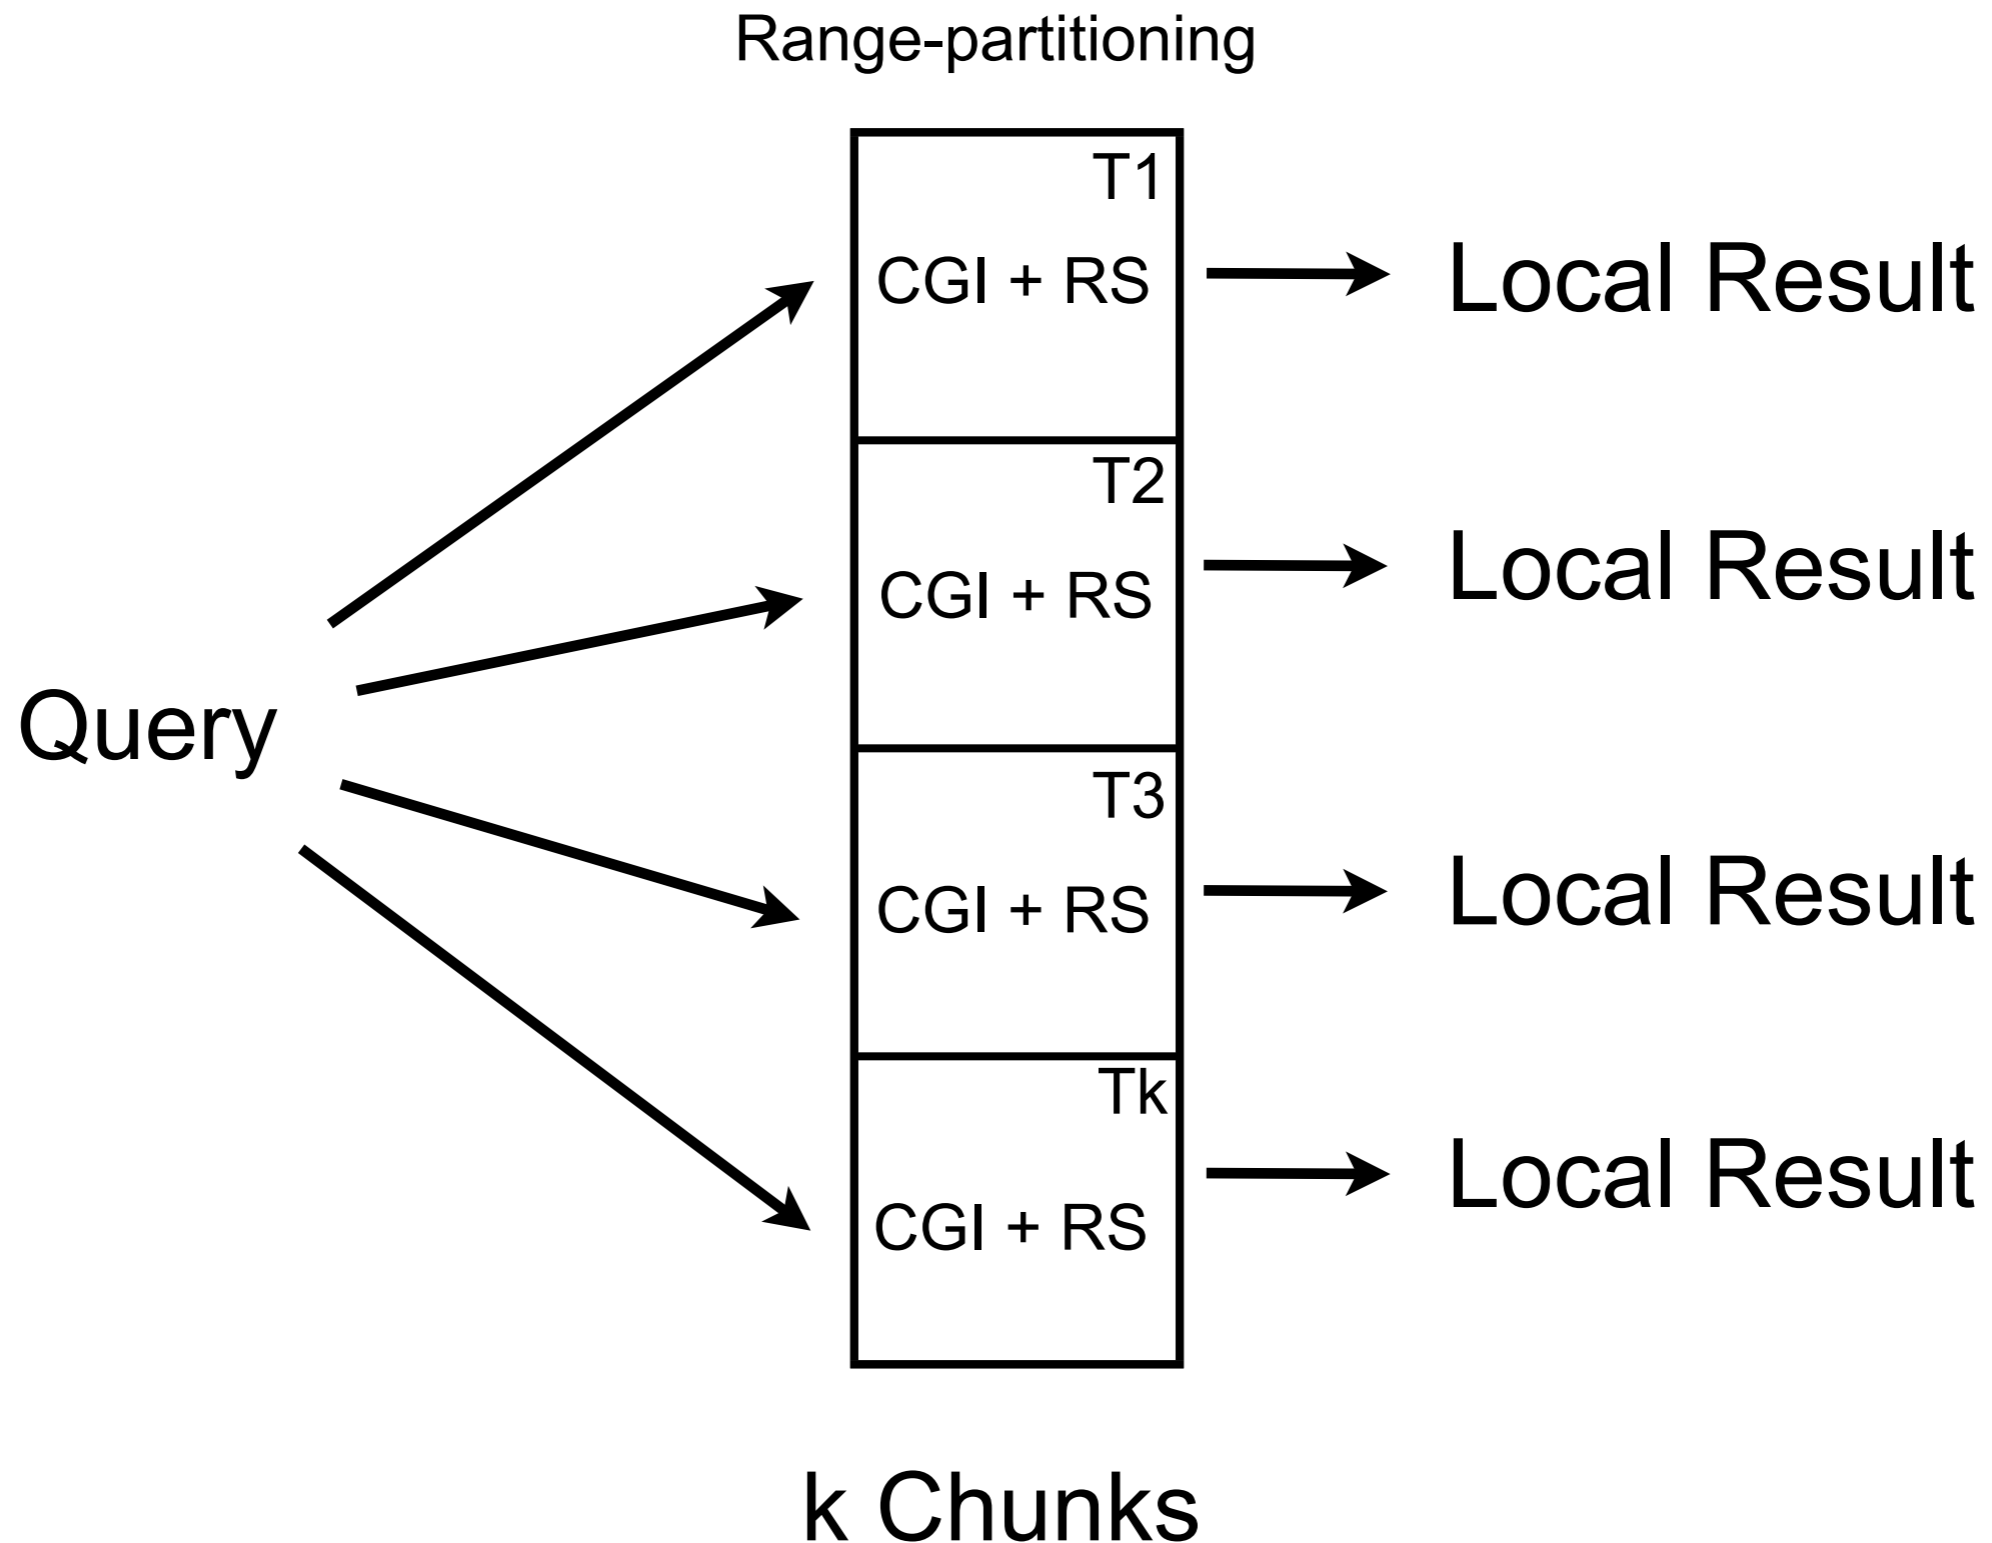




Multi-threaded algorithms:

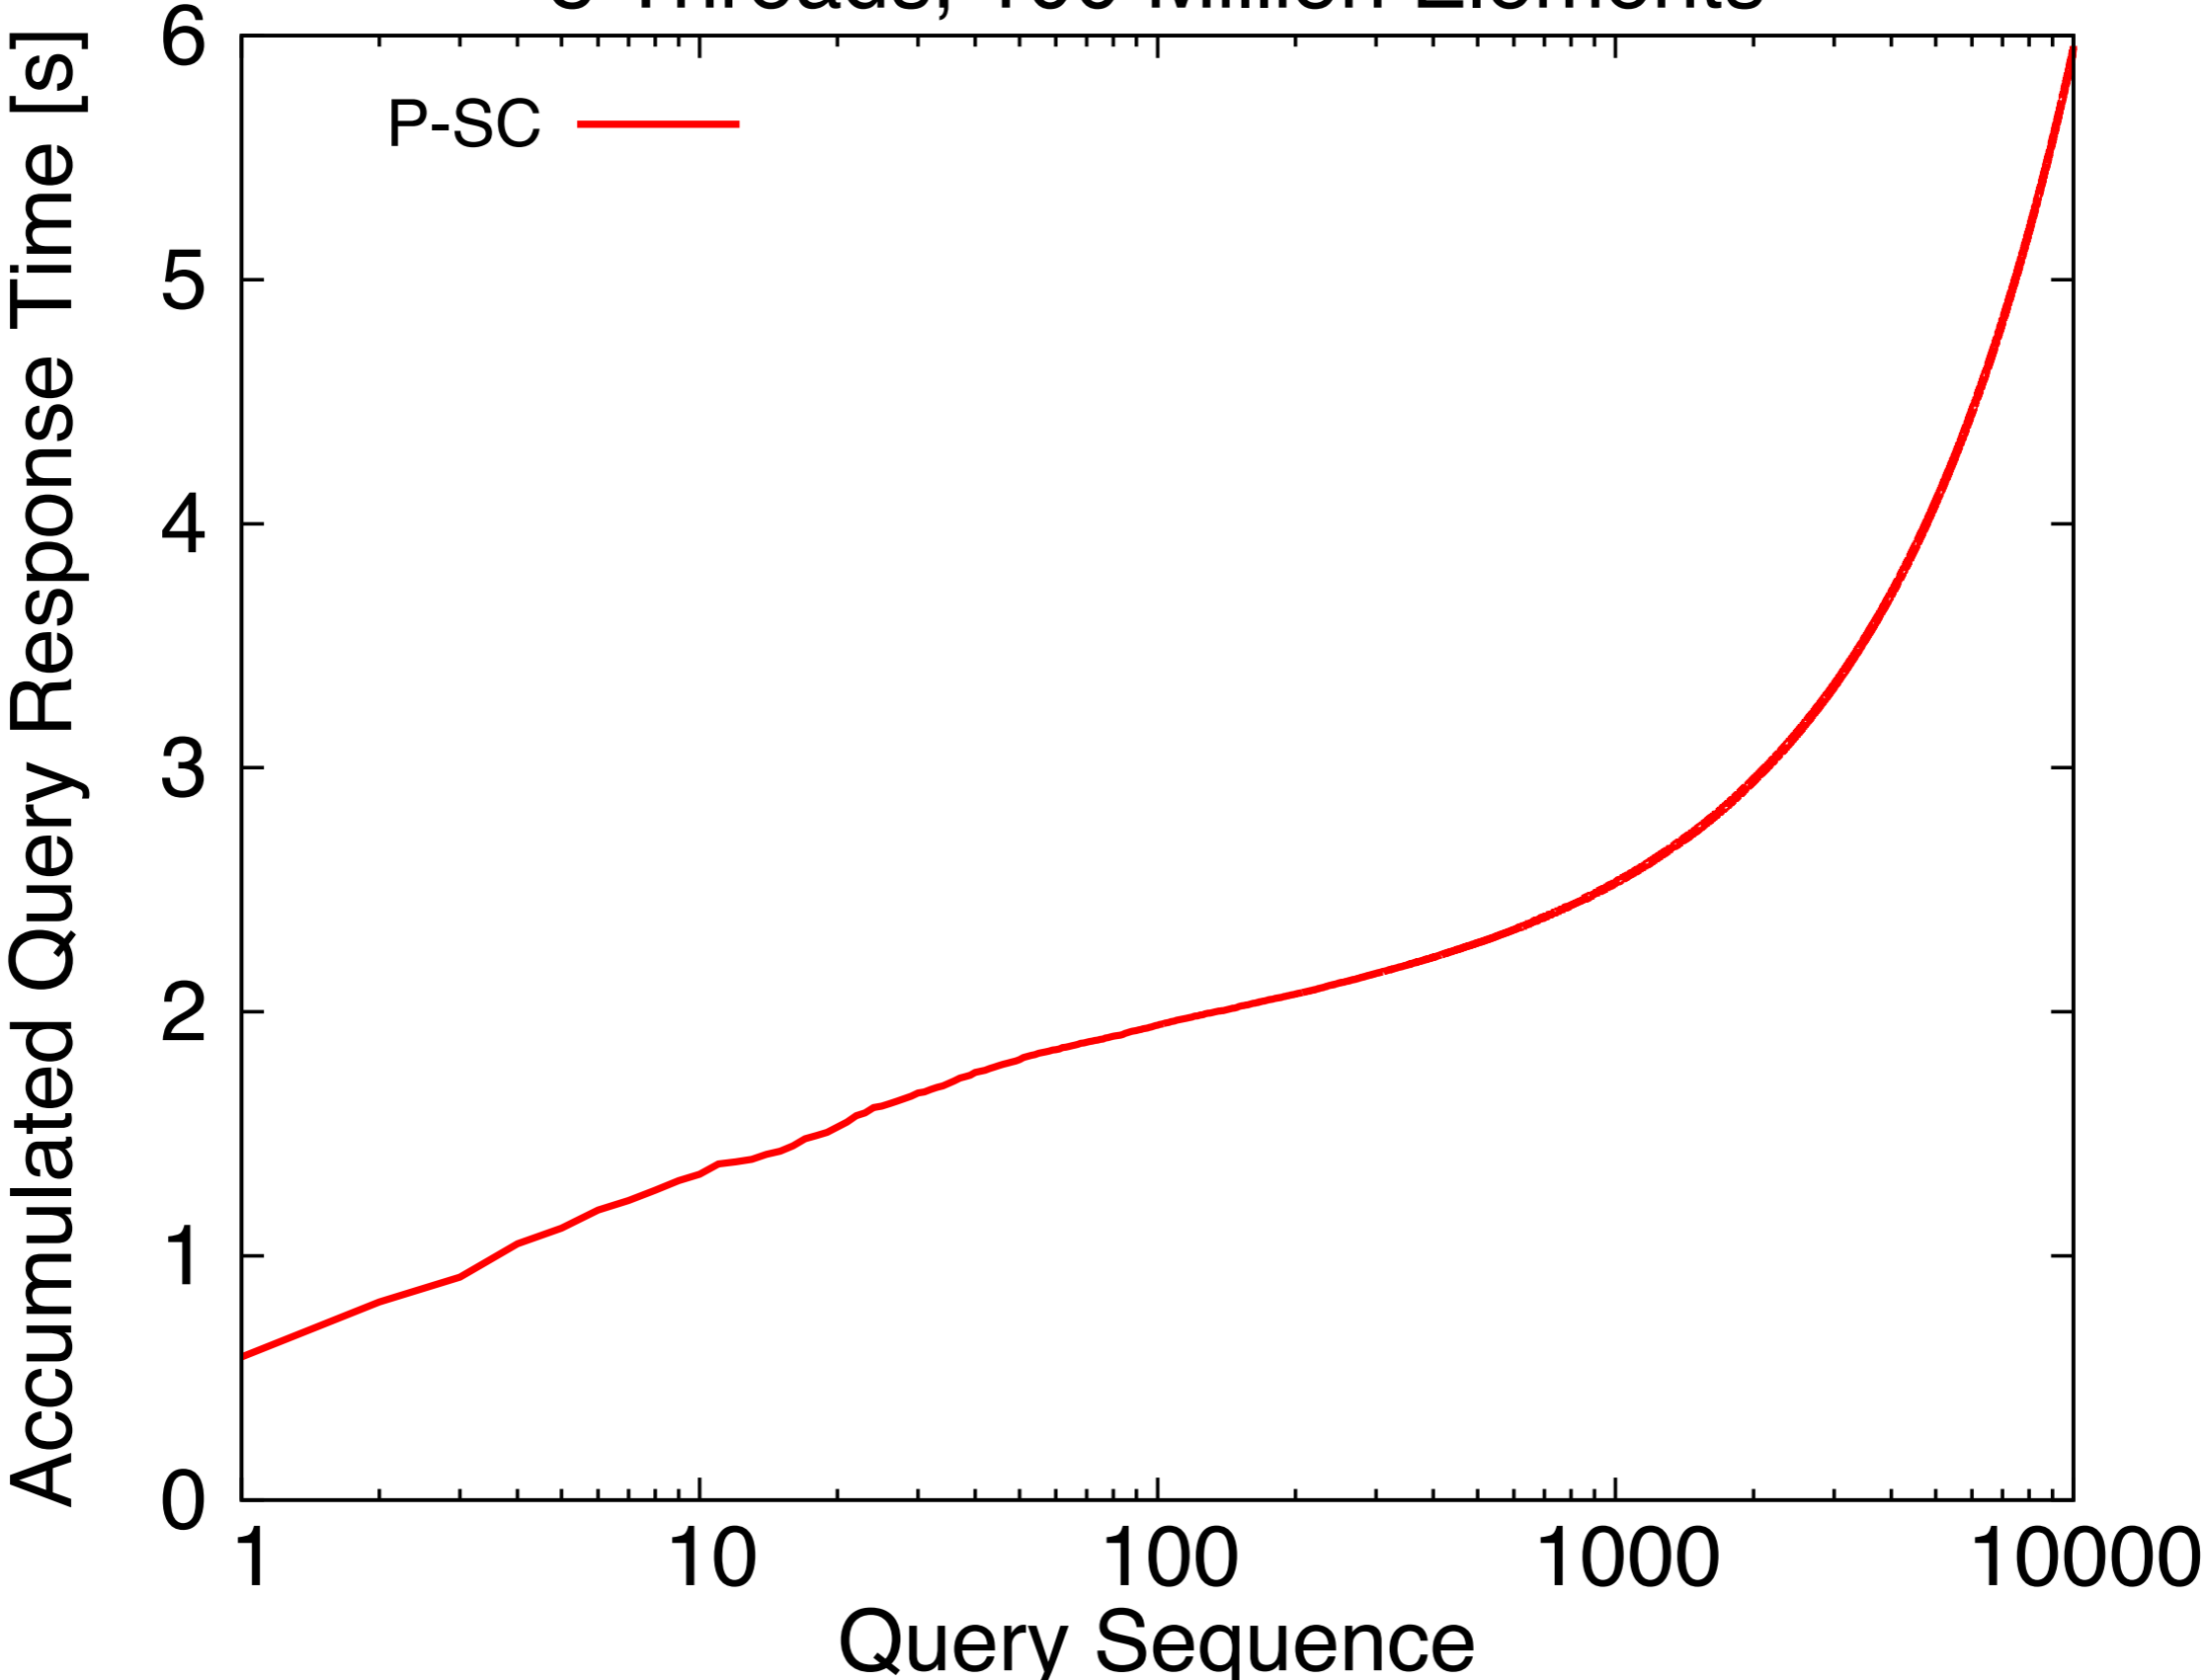
# Parallel-chunked Range-Partitioned Radix Sort (P-CRS)

P-RPRS + Chunking




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements



# Multi-threaded Results


8 Threads, 100 Million Elements






# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

8 Threads, 100 Million Elements



# Multi-threaded Results


8 Threads, 100 Million Elements






# Multi-threaded Results

8 Threads, 100 Million Elements




# Multi-threaded Results

## Factor Speedup from 1 to 8 Threads



# P-SC: Analysis

## Bandwidth




## Lock Time

Mutex	Wait Time (sec)
Piece lock	11.671
Cracker index lock	5.169
Total	16.84
Average (Total by 8)	2.105

# P-SC: Analysis

## Bandwidth




## Lock Time

Mutex	Wait Time (sec)
Piece lock	11.671
Cracker index lock	5.169
Total	16.84
Average (Total by 8)	2.105

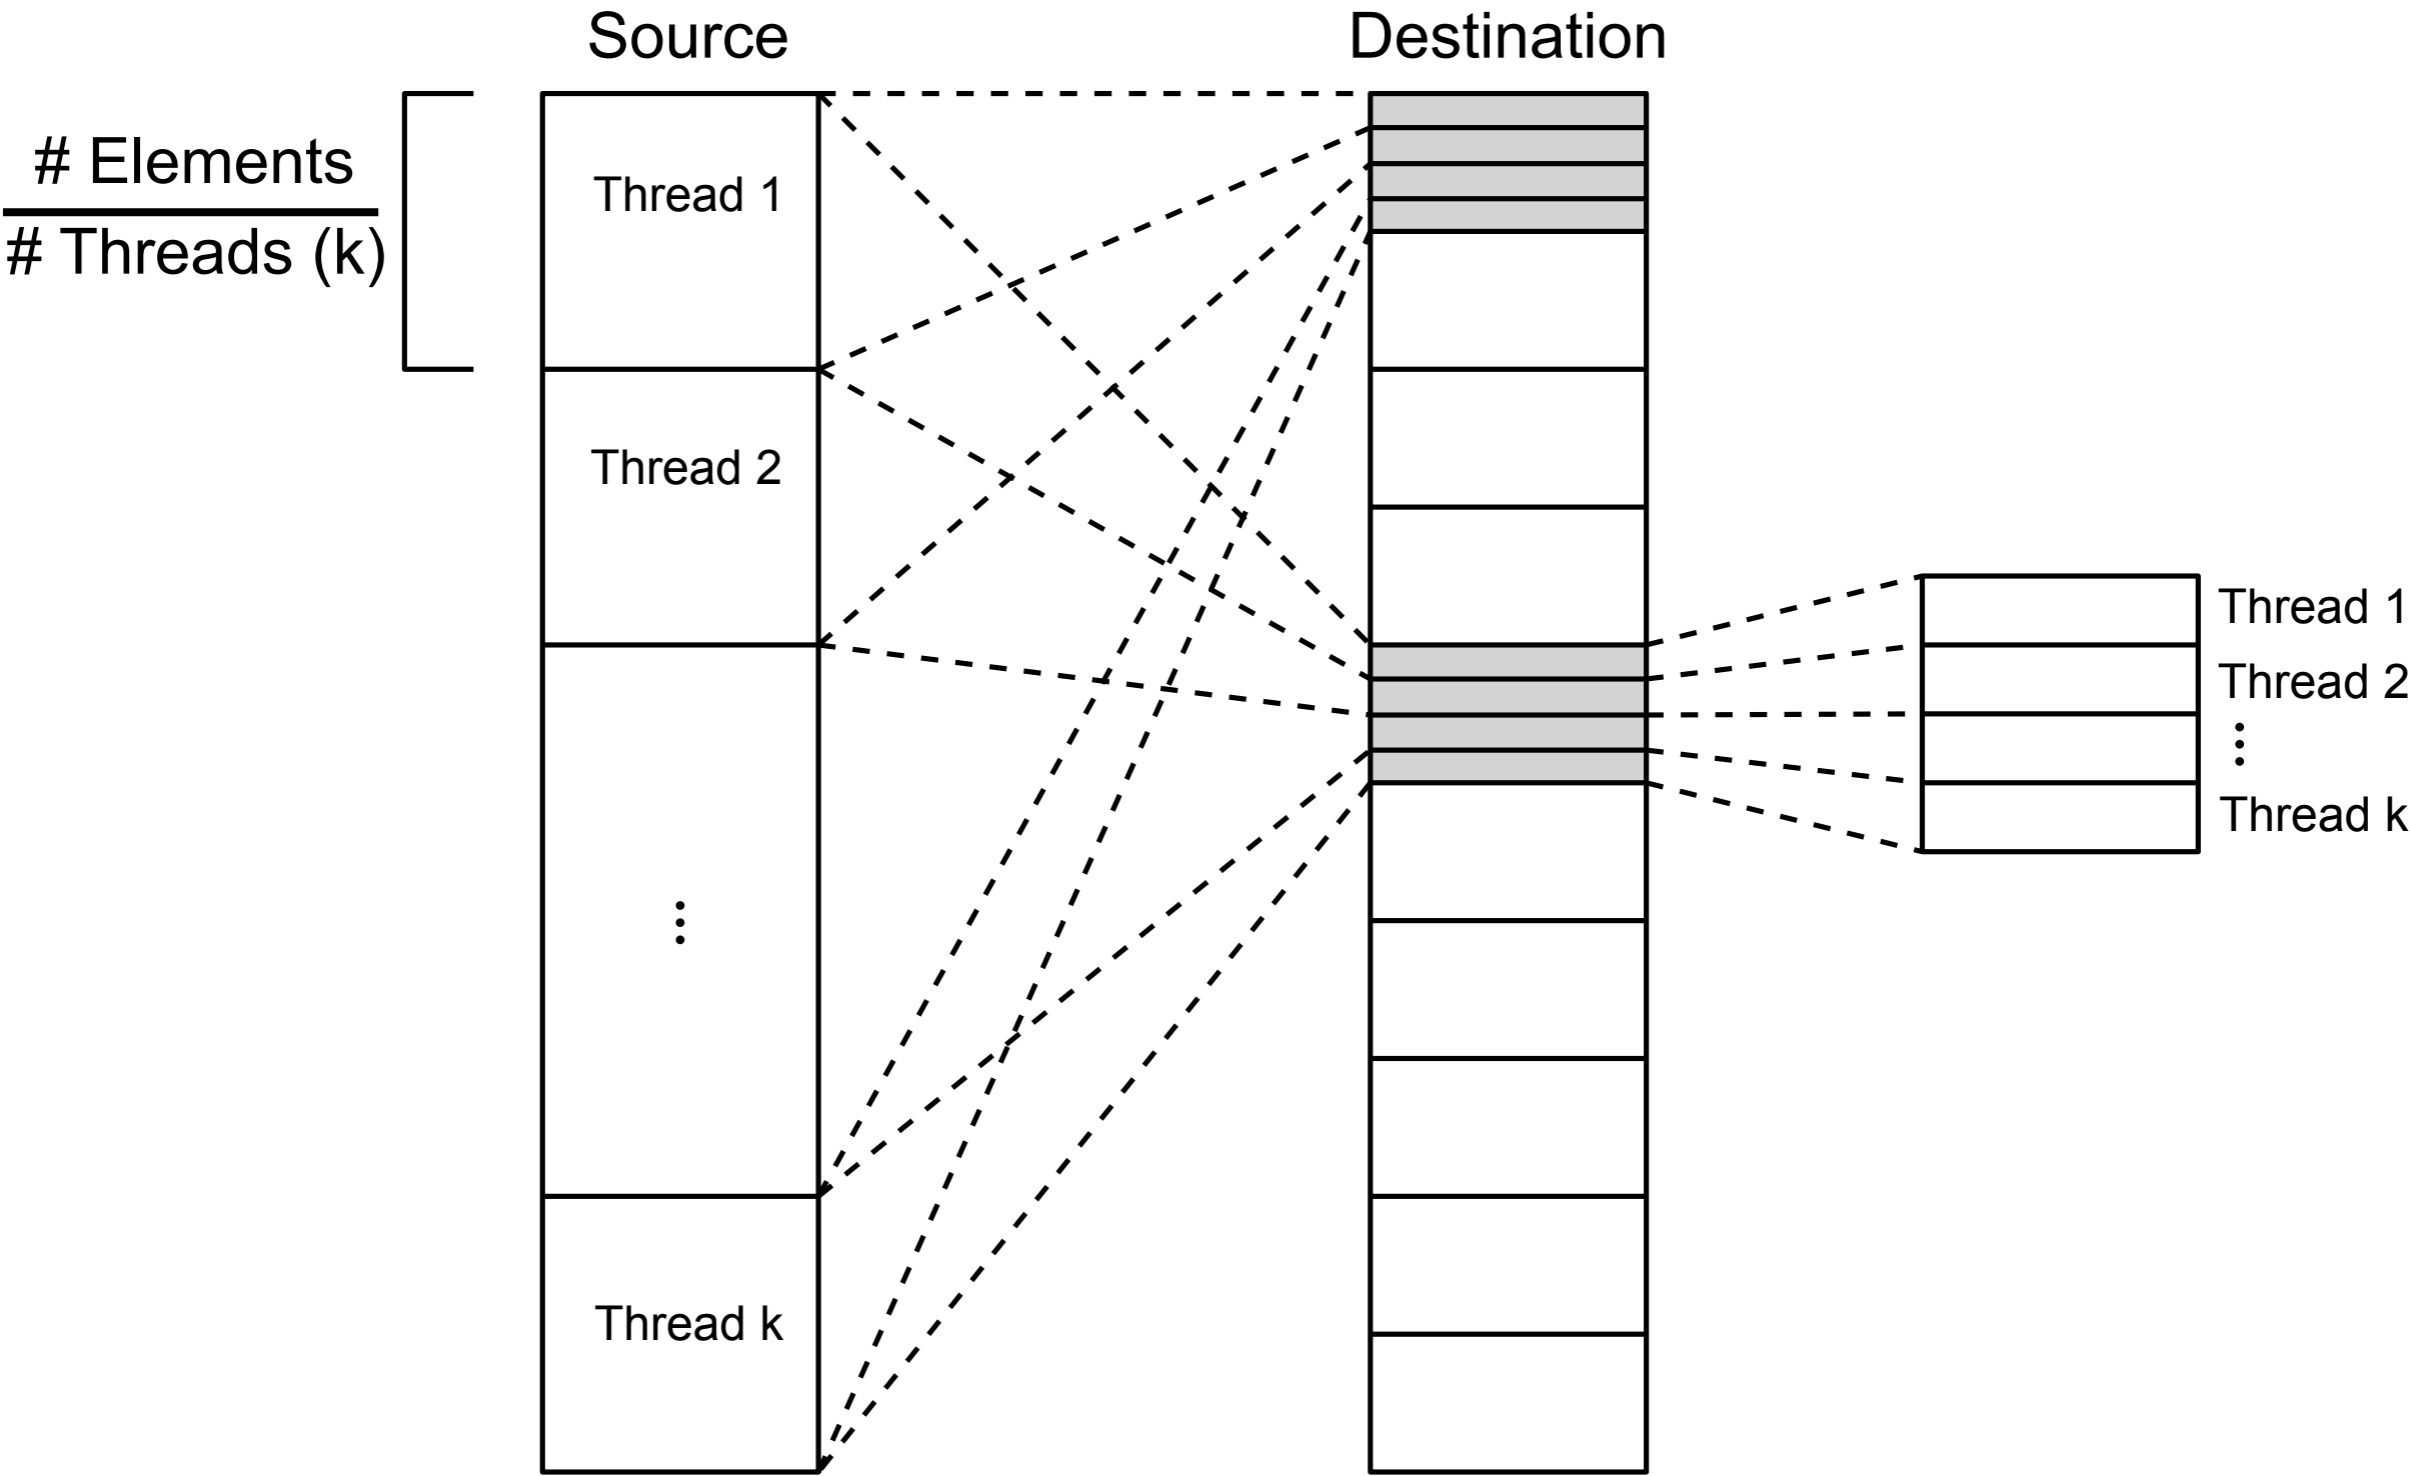
# Multi-threaded Results

## Factor Speedup from 1 to 8 Threads



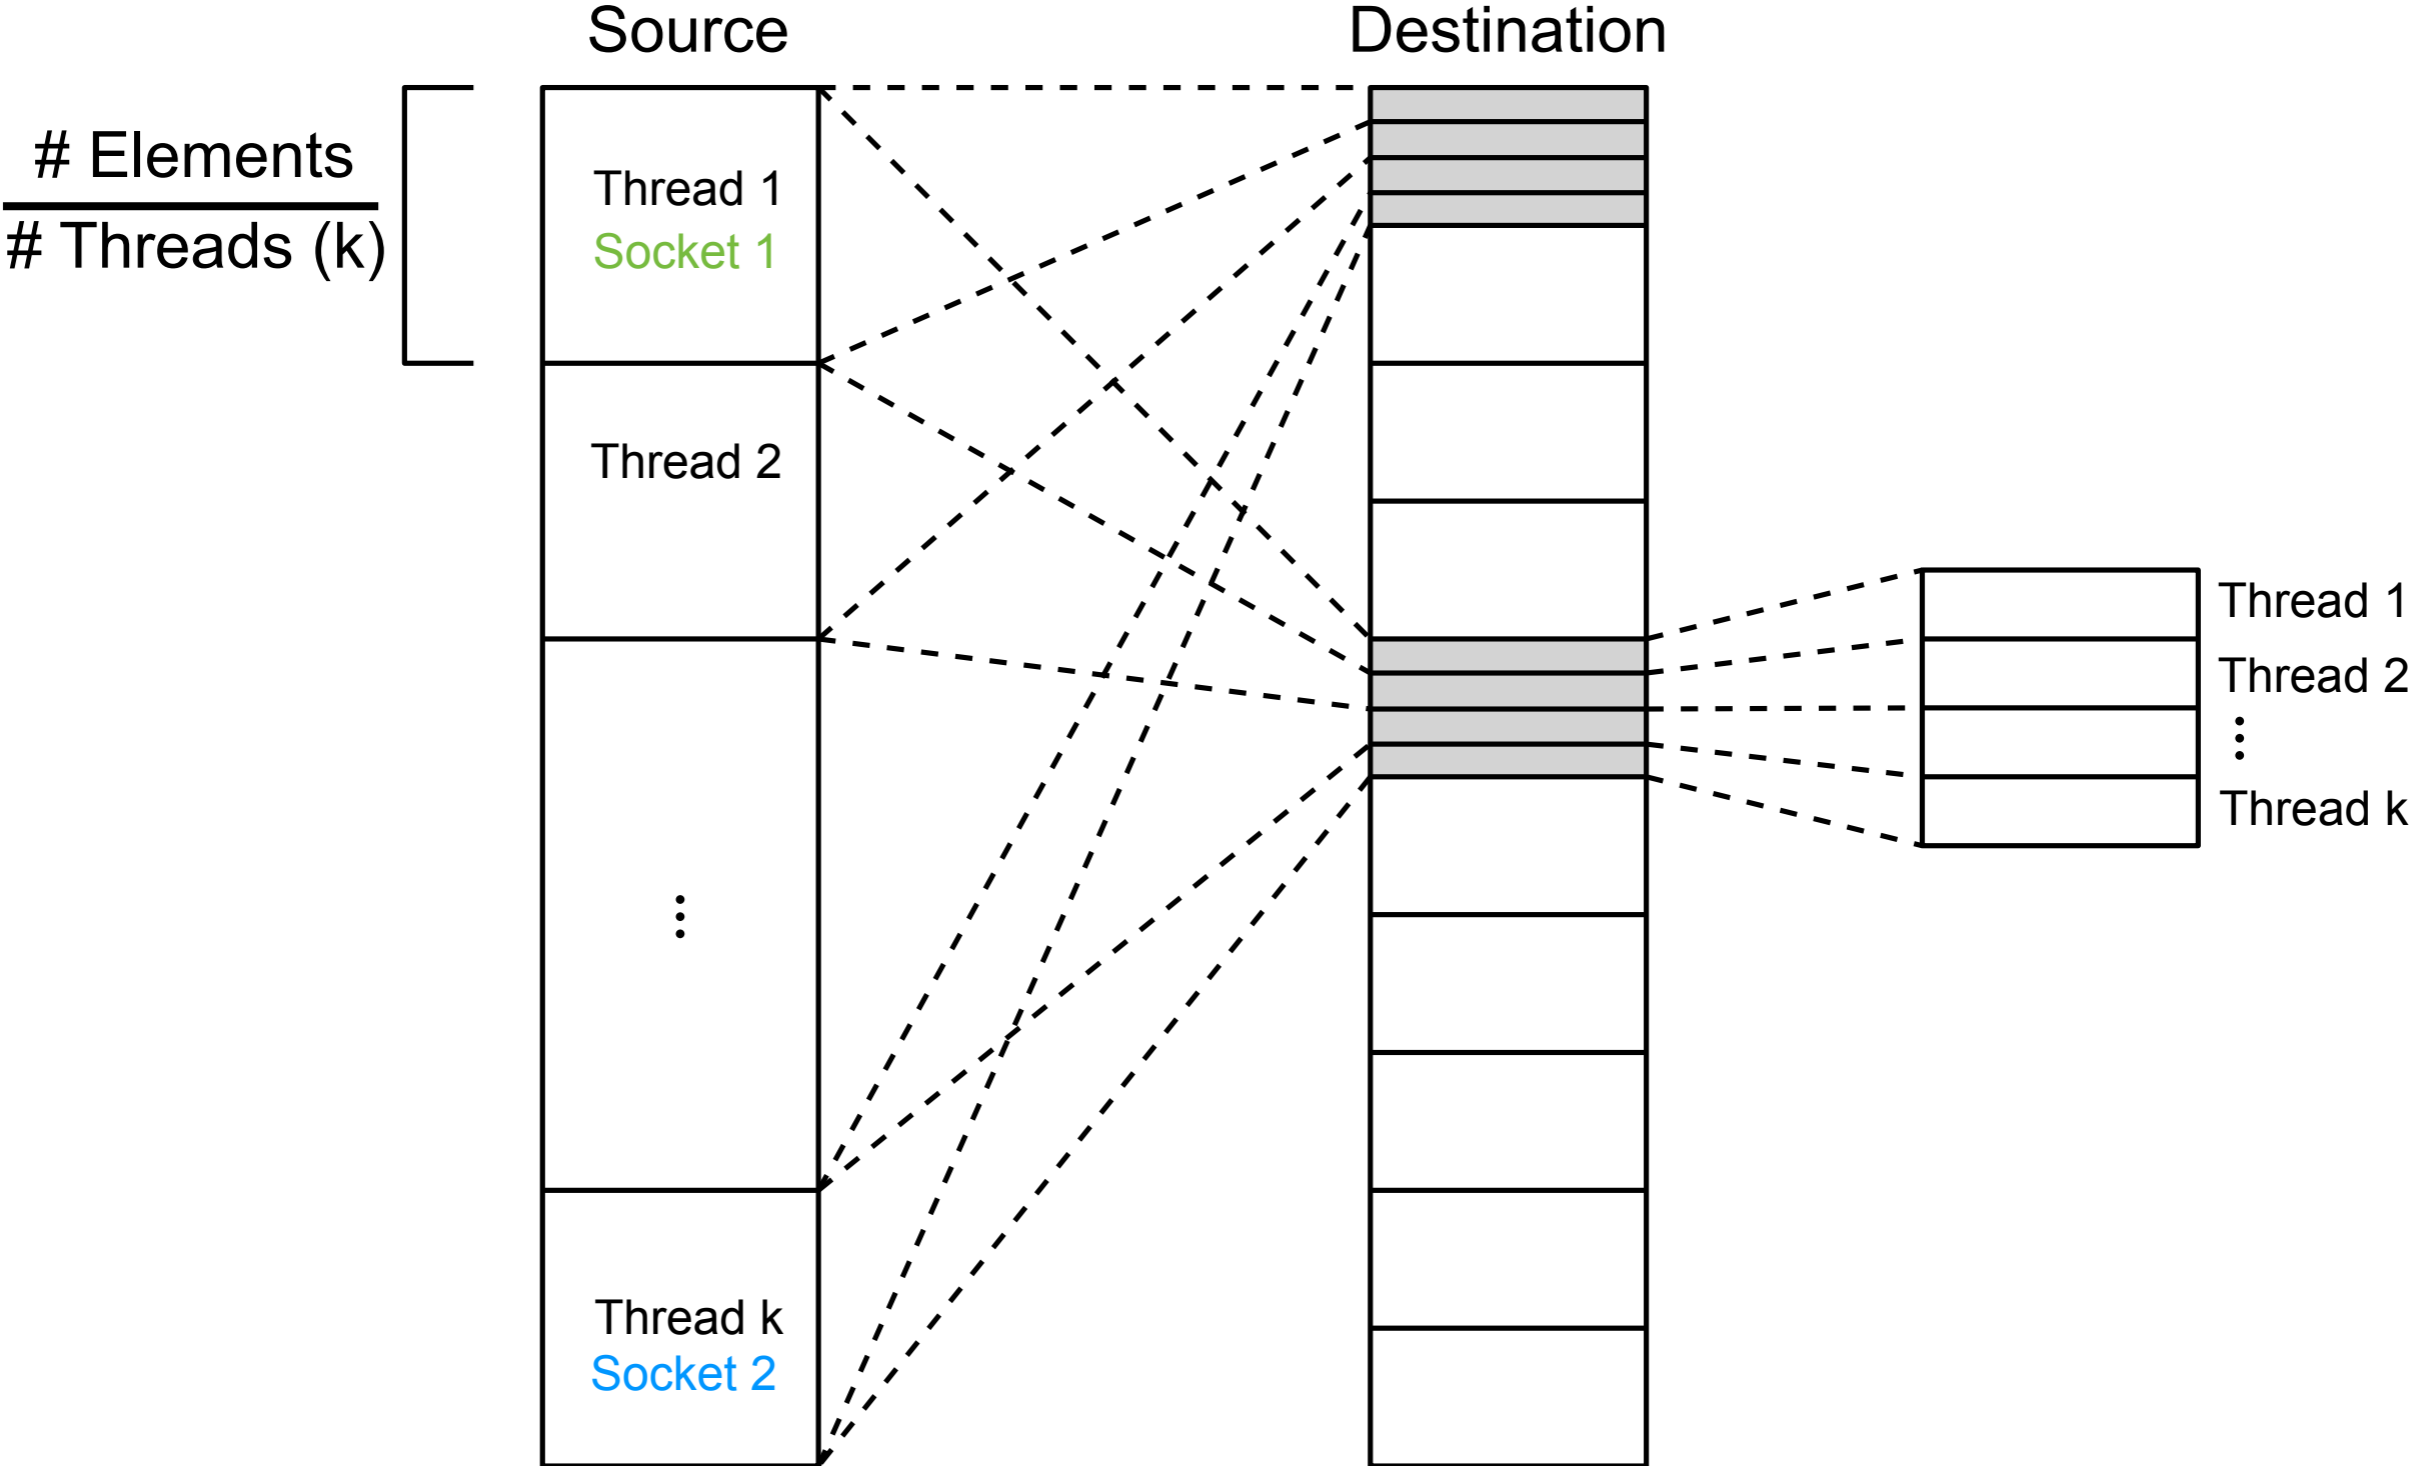
# Non-Chunked Algorithms: Analysis (P-RPRS)

Range Partitioning (RP) Phase:



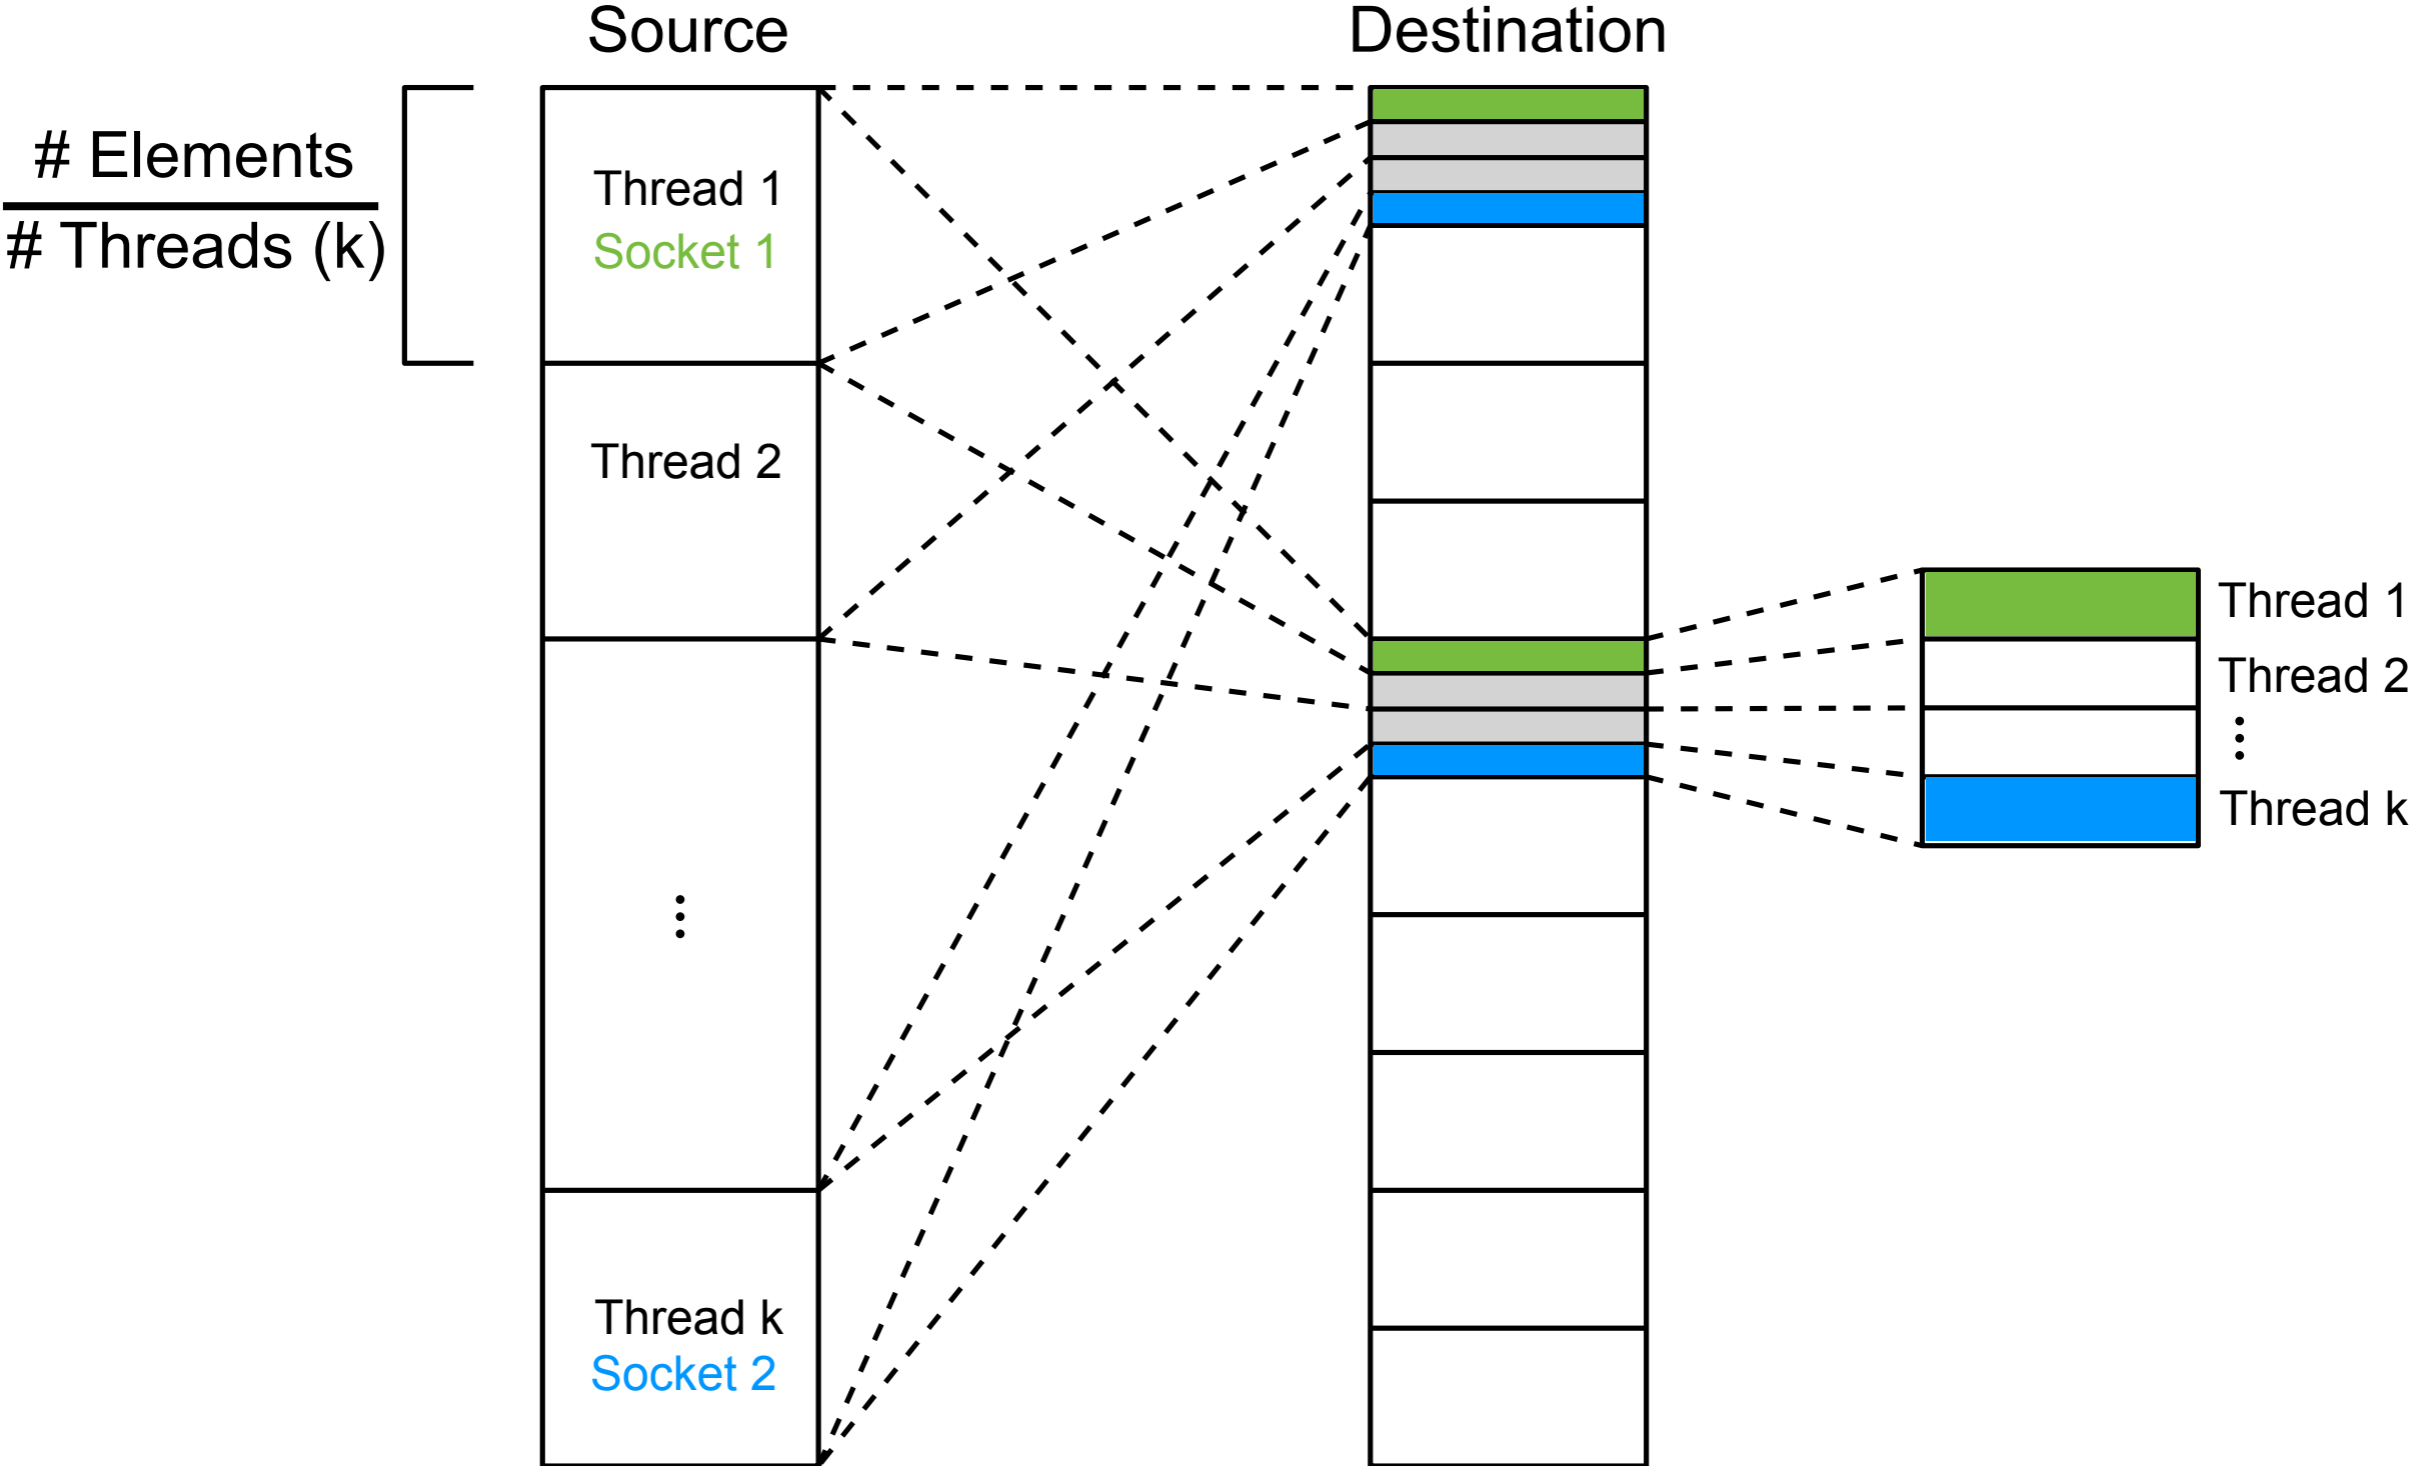
# Non-Chunked Algorithms: Analysis (P-RPRS)

Range Partitioning (RP) Phase:



# Non-Chunked Algorithms: Analysis (P-RPRS)

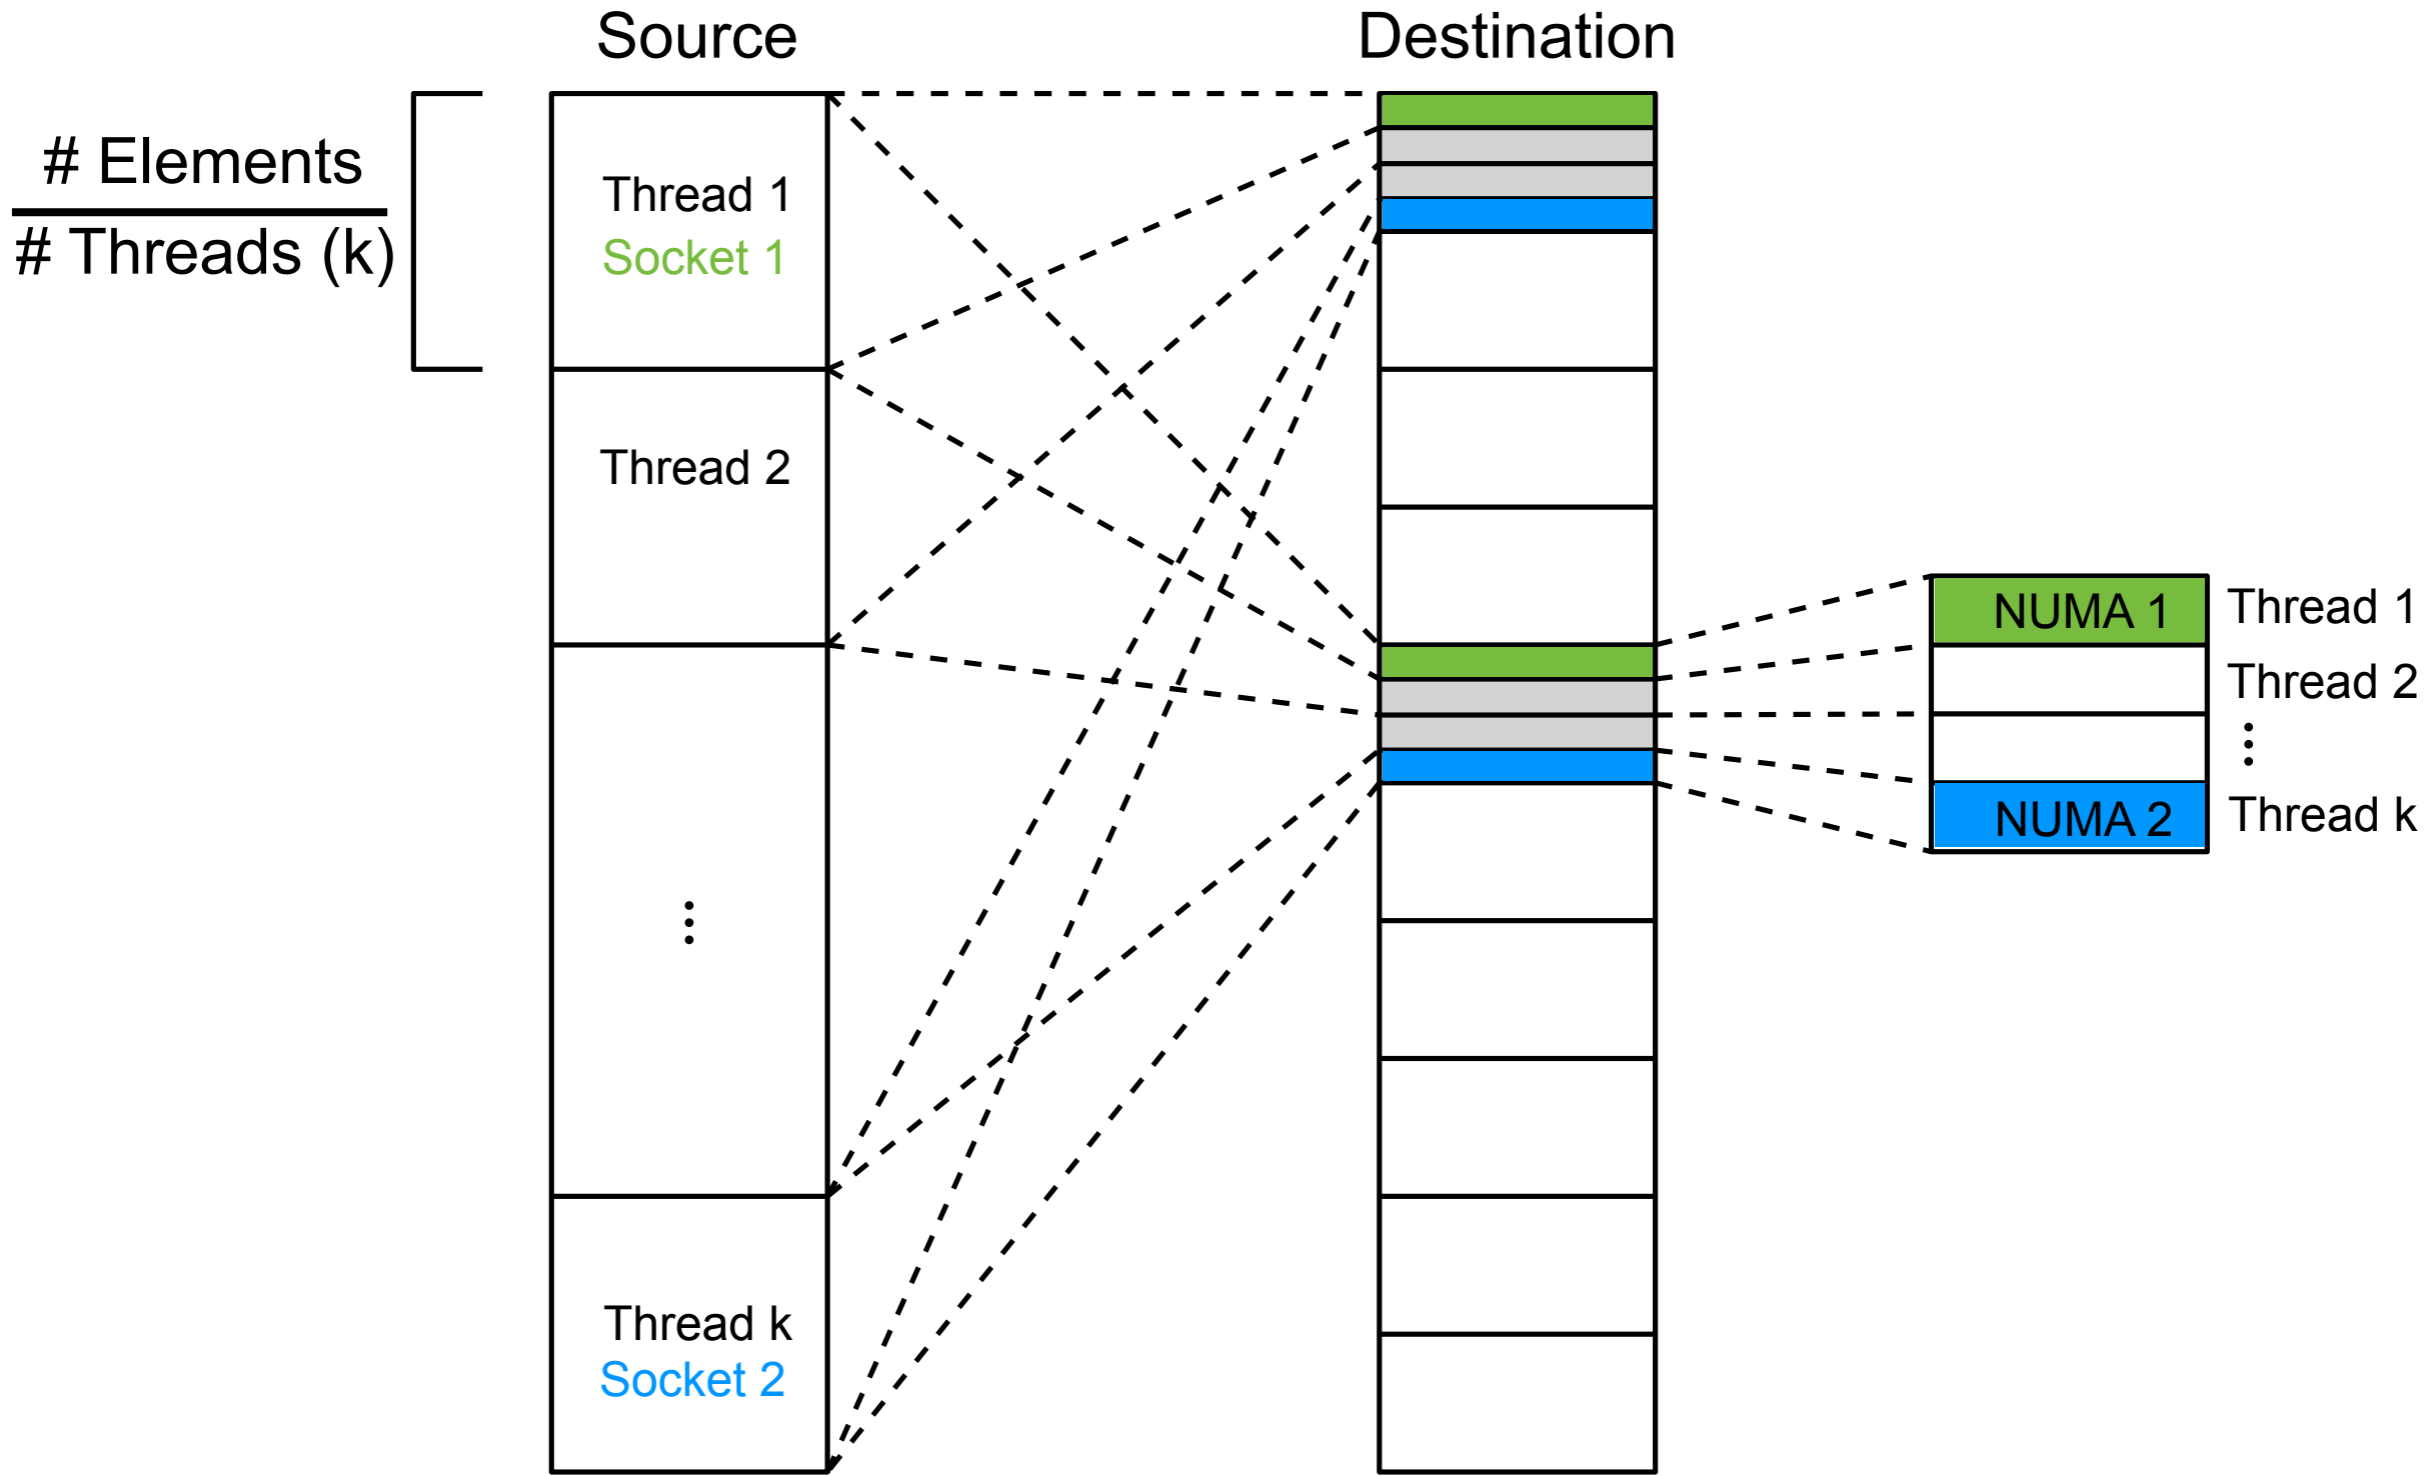
Range Partitioning (RP) Phase:





# Non-Chunked Algorithms: Analysis (P-RPRS)

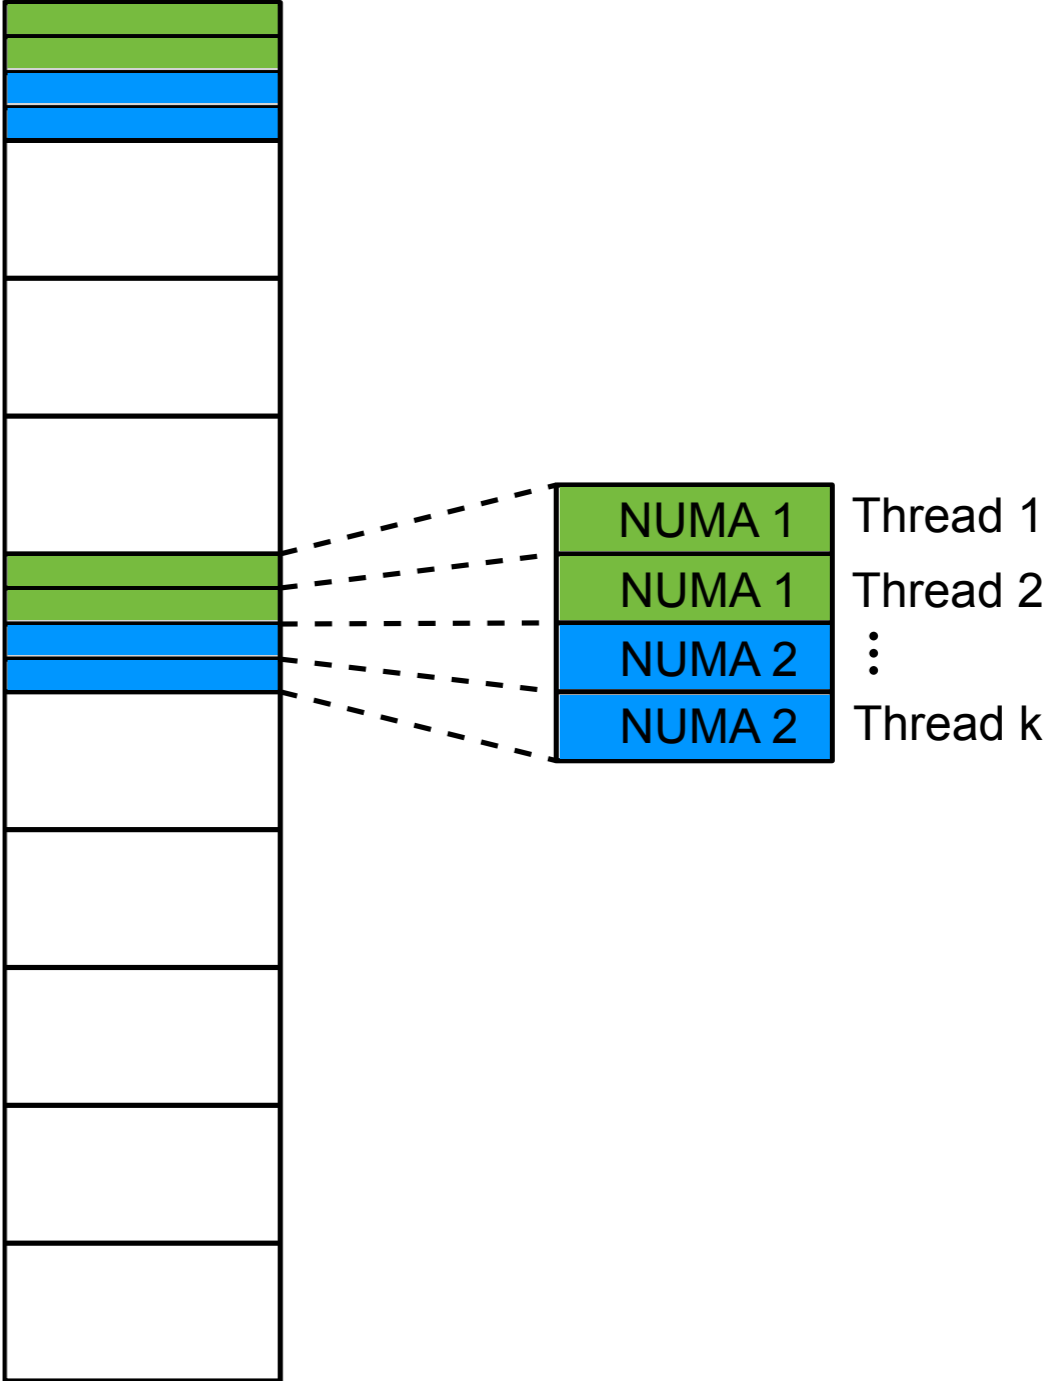
Range Partitioning (RP) Phase:



# Non-Chunked Algorithms: Analysis (P-RPRS)

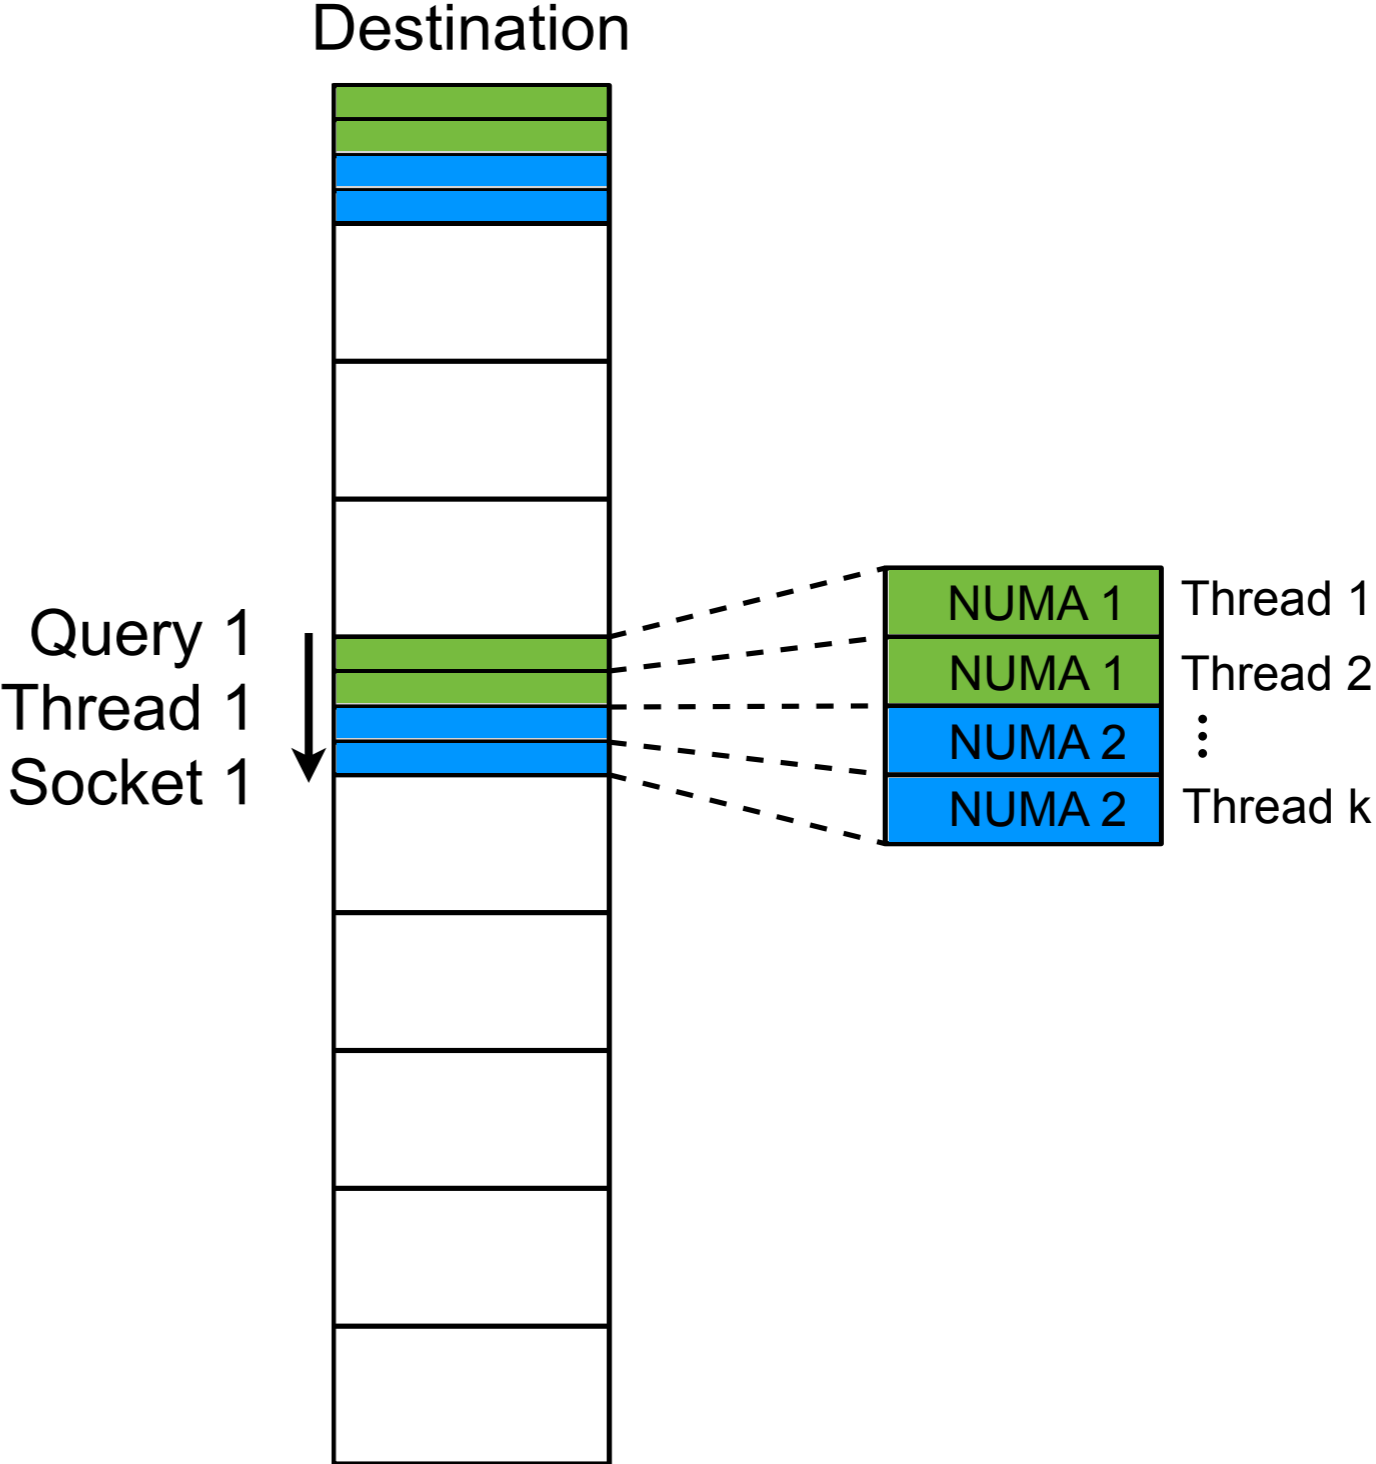
Query Phase:

Destination



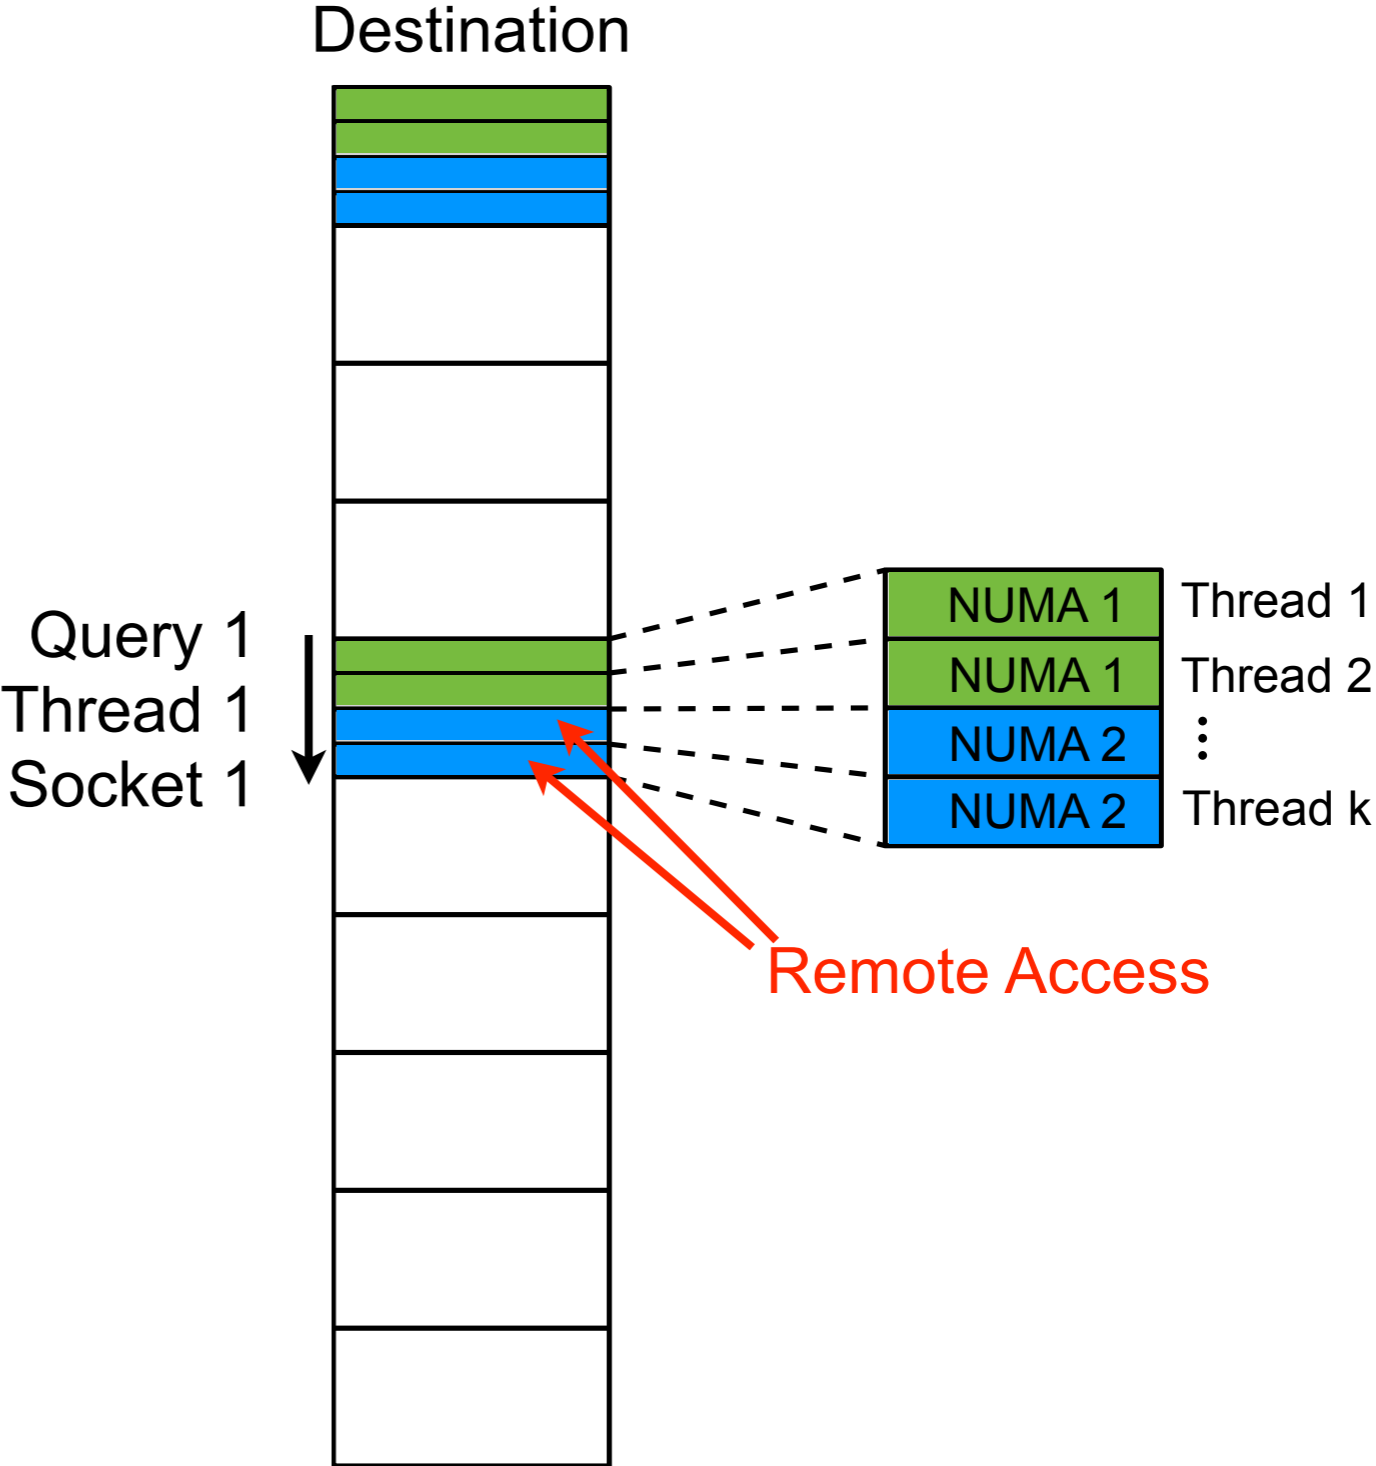
# Non-Chunked Algorithms: Analysis (P-RPRS)

Query Phase:




# Non-Chunked Algorithms: Analysis (P-RPRS)

Query Phase:



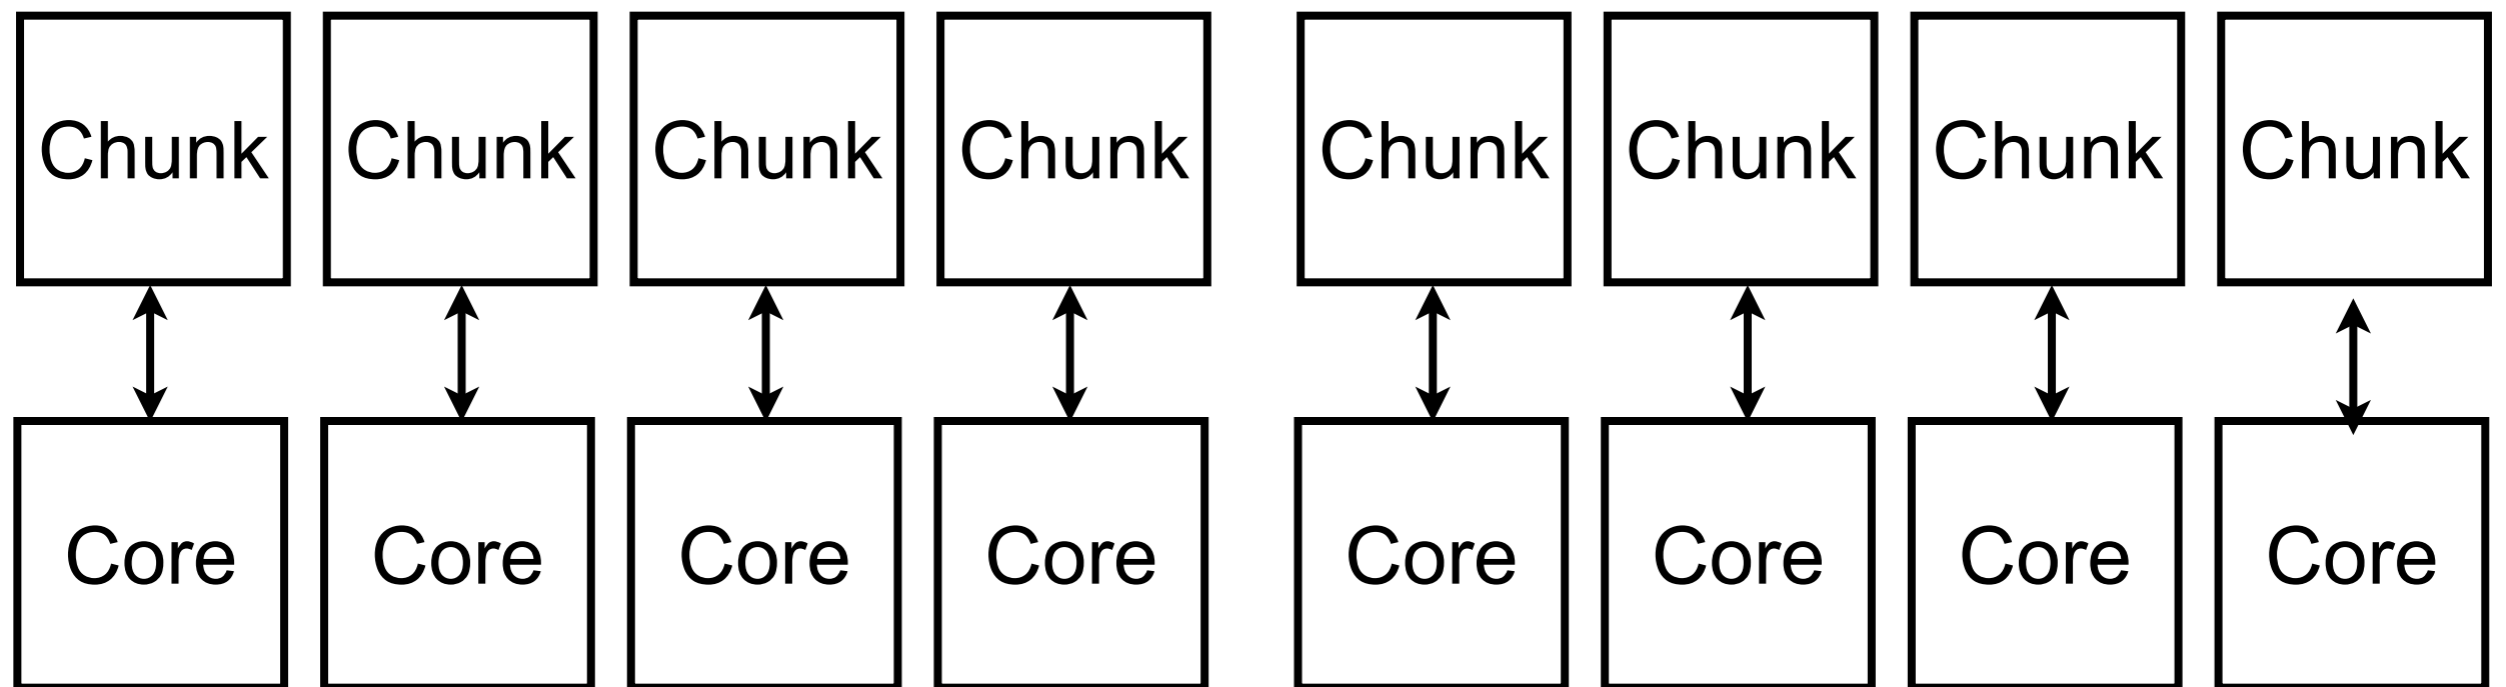
# Multi-threaded Results

## Factor Speedup from 1 to 8 Threads



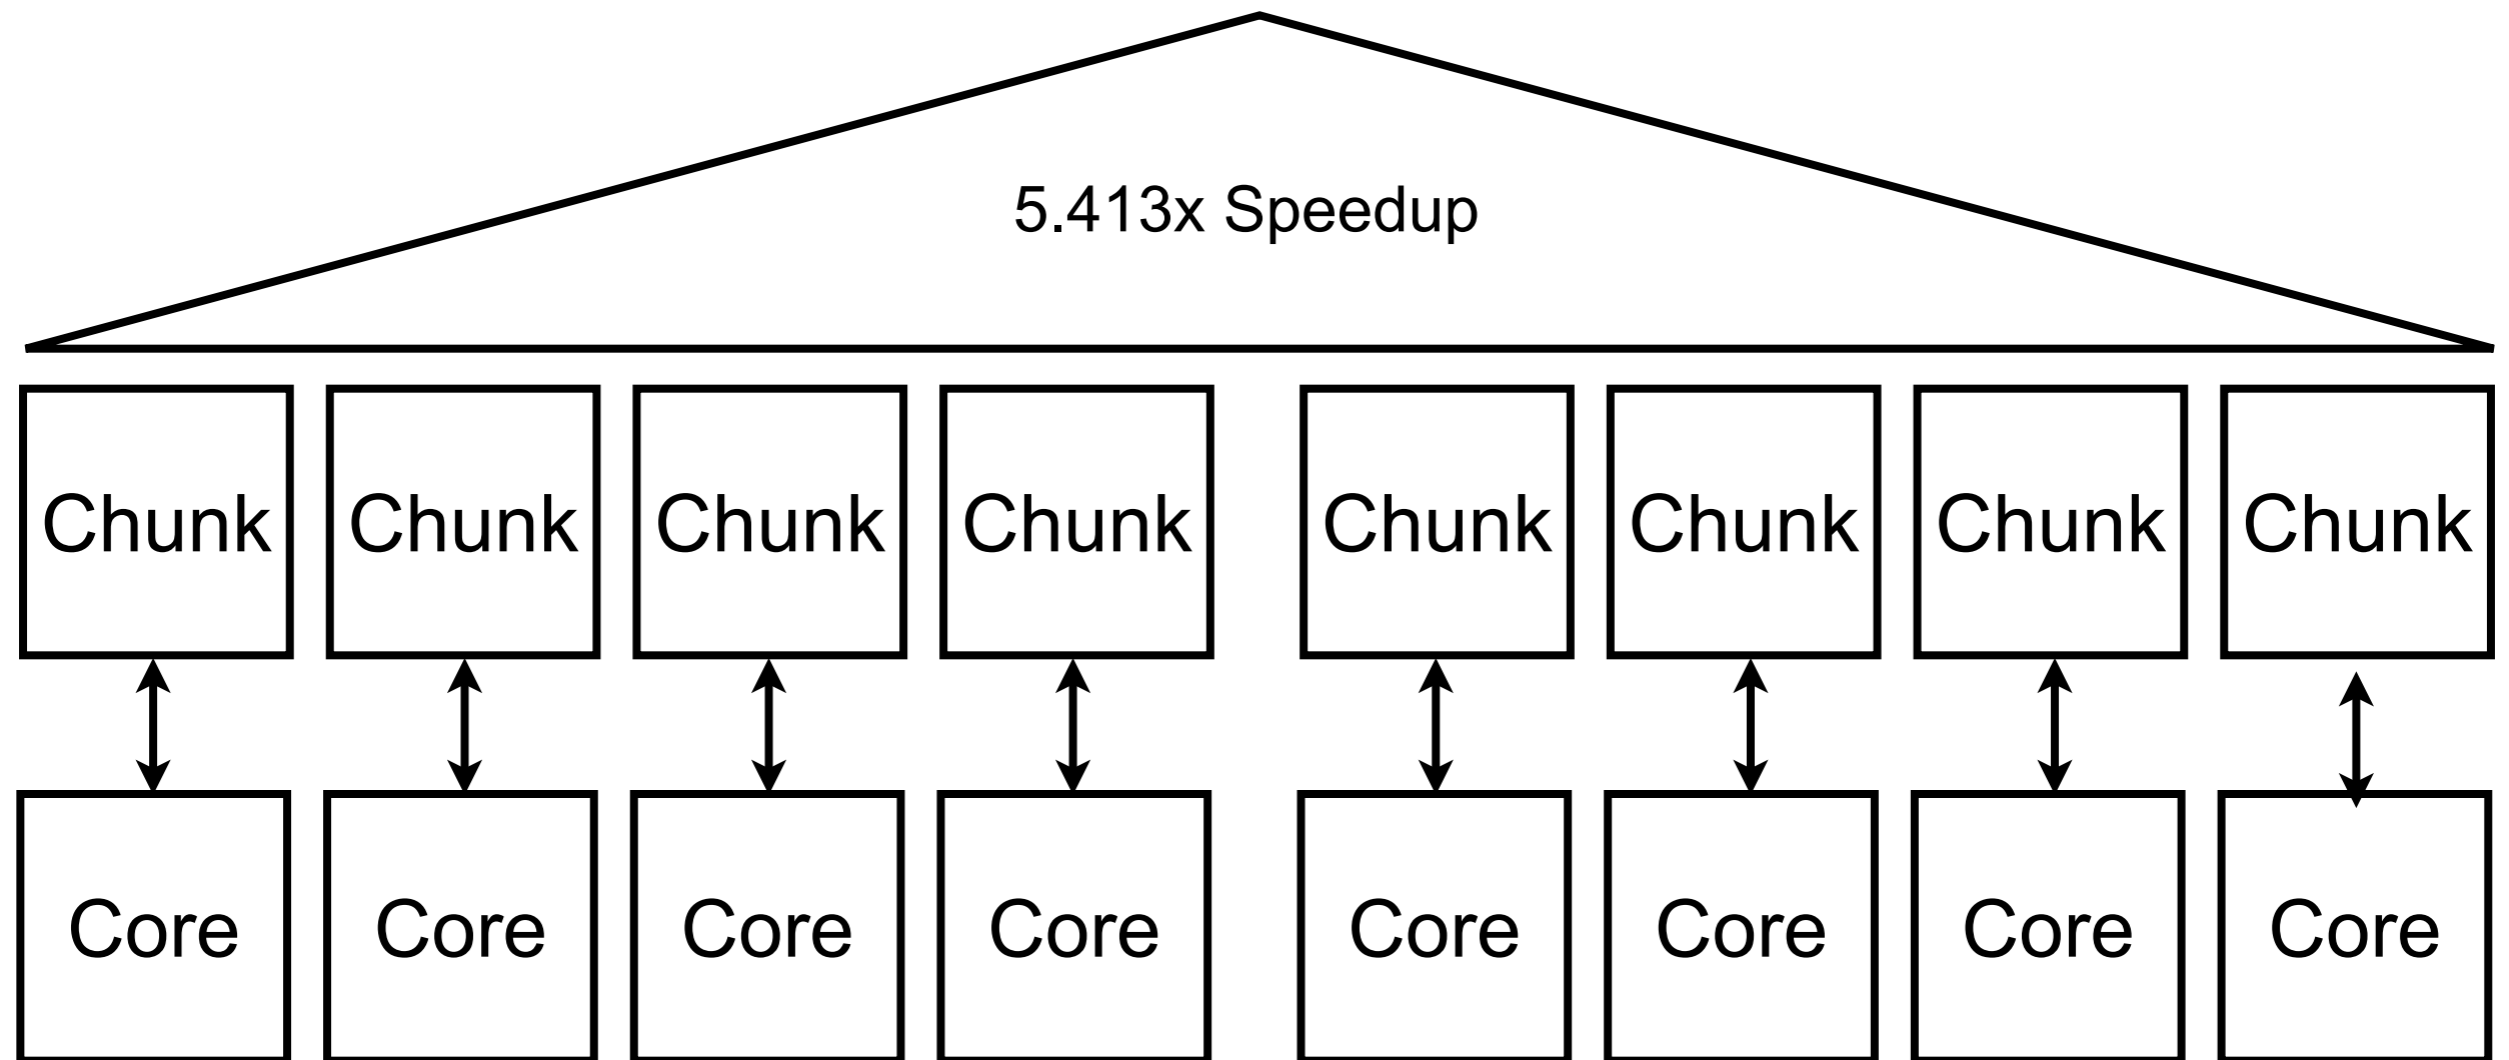
# Chunked Algorithms: Analysis (P-CRS)

All chunks are completely independent - 8x Speedup?



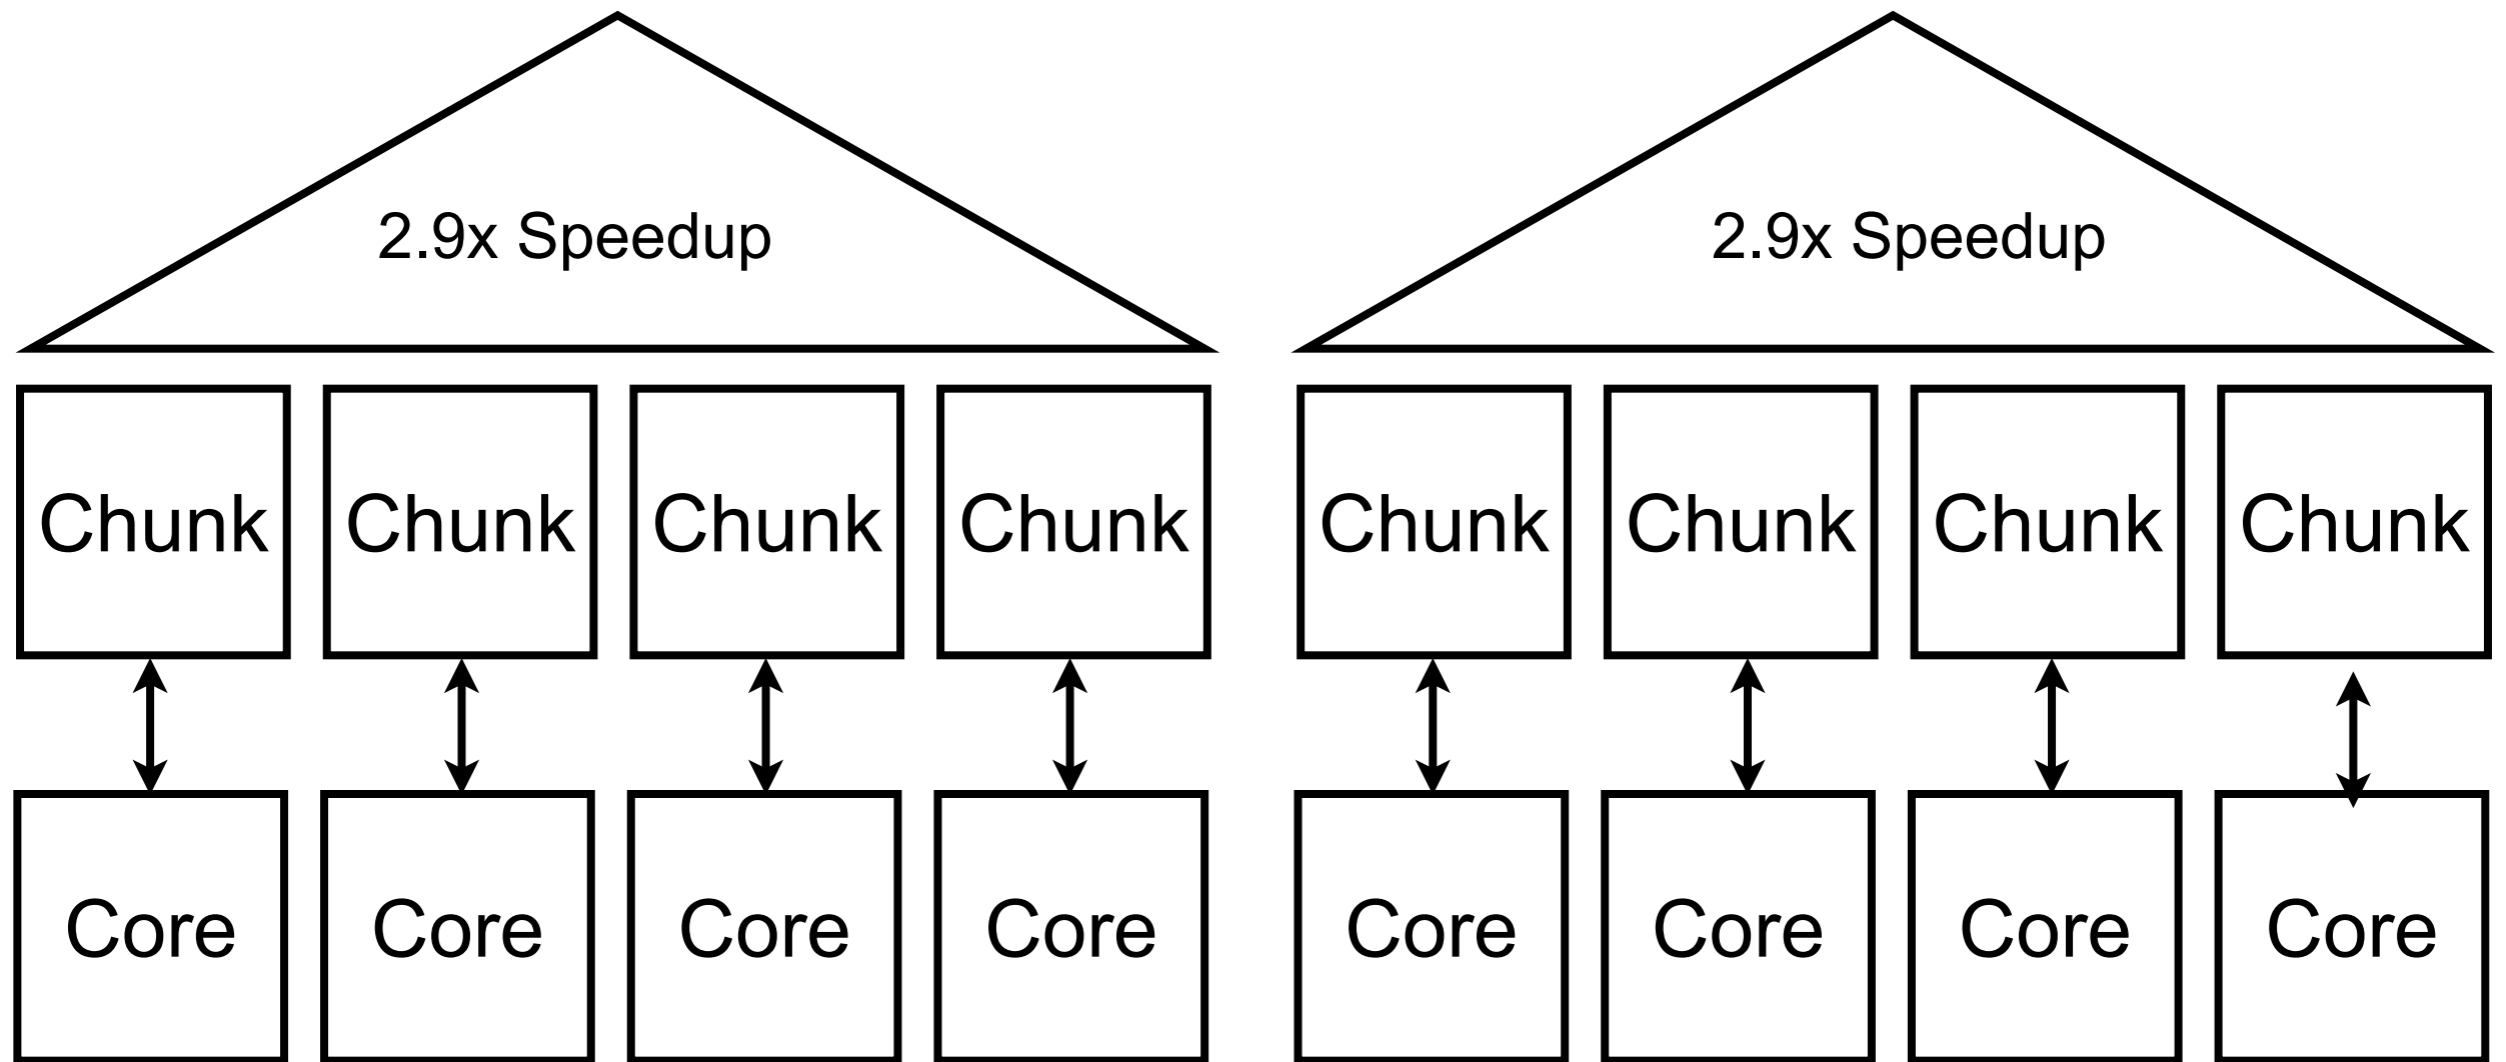
# Chunked Algorithms: Analysis (P-CRS)

All chunks are completely independent - 8x Speedup?



# Chunked Algorithms: Analysis (P-CRS)

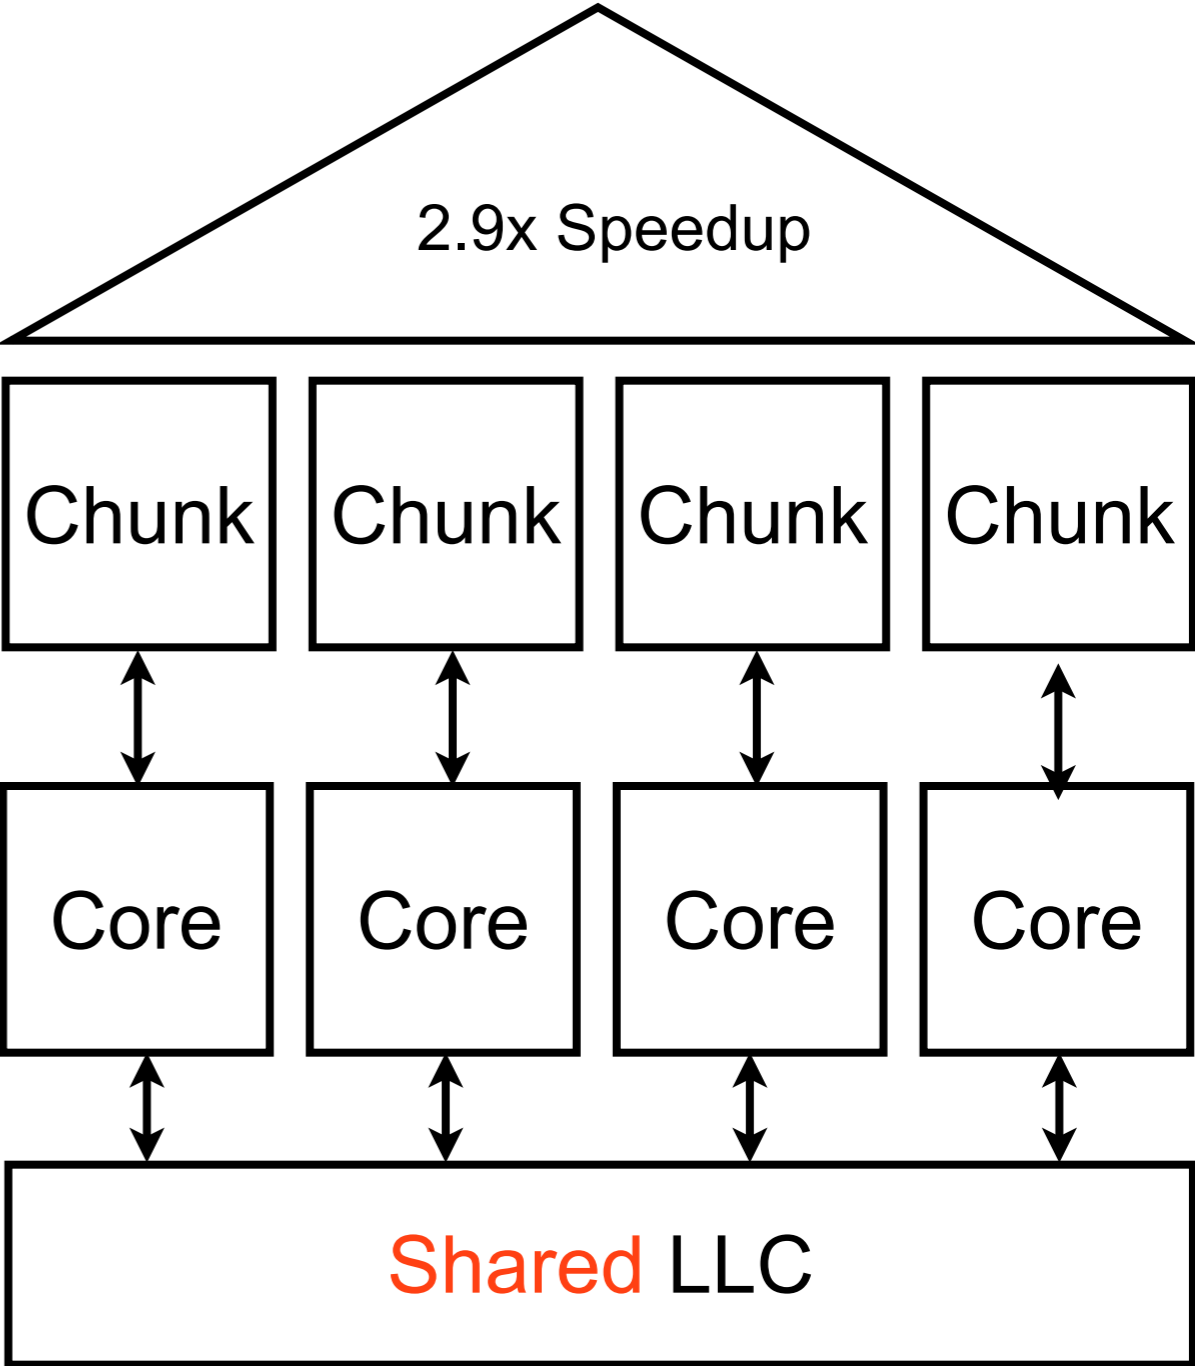
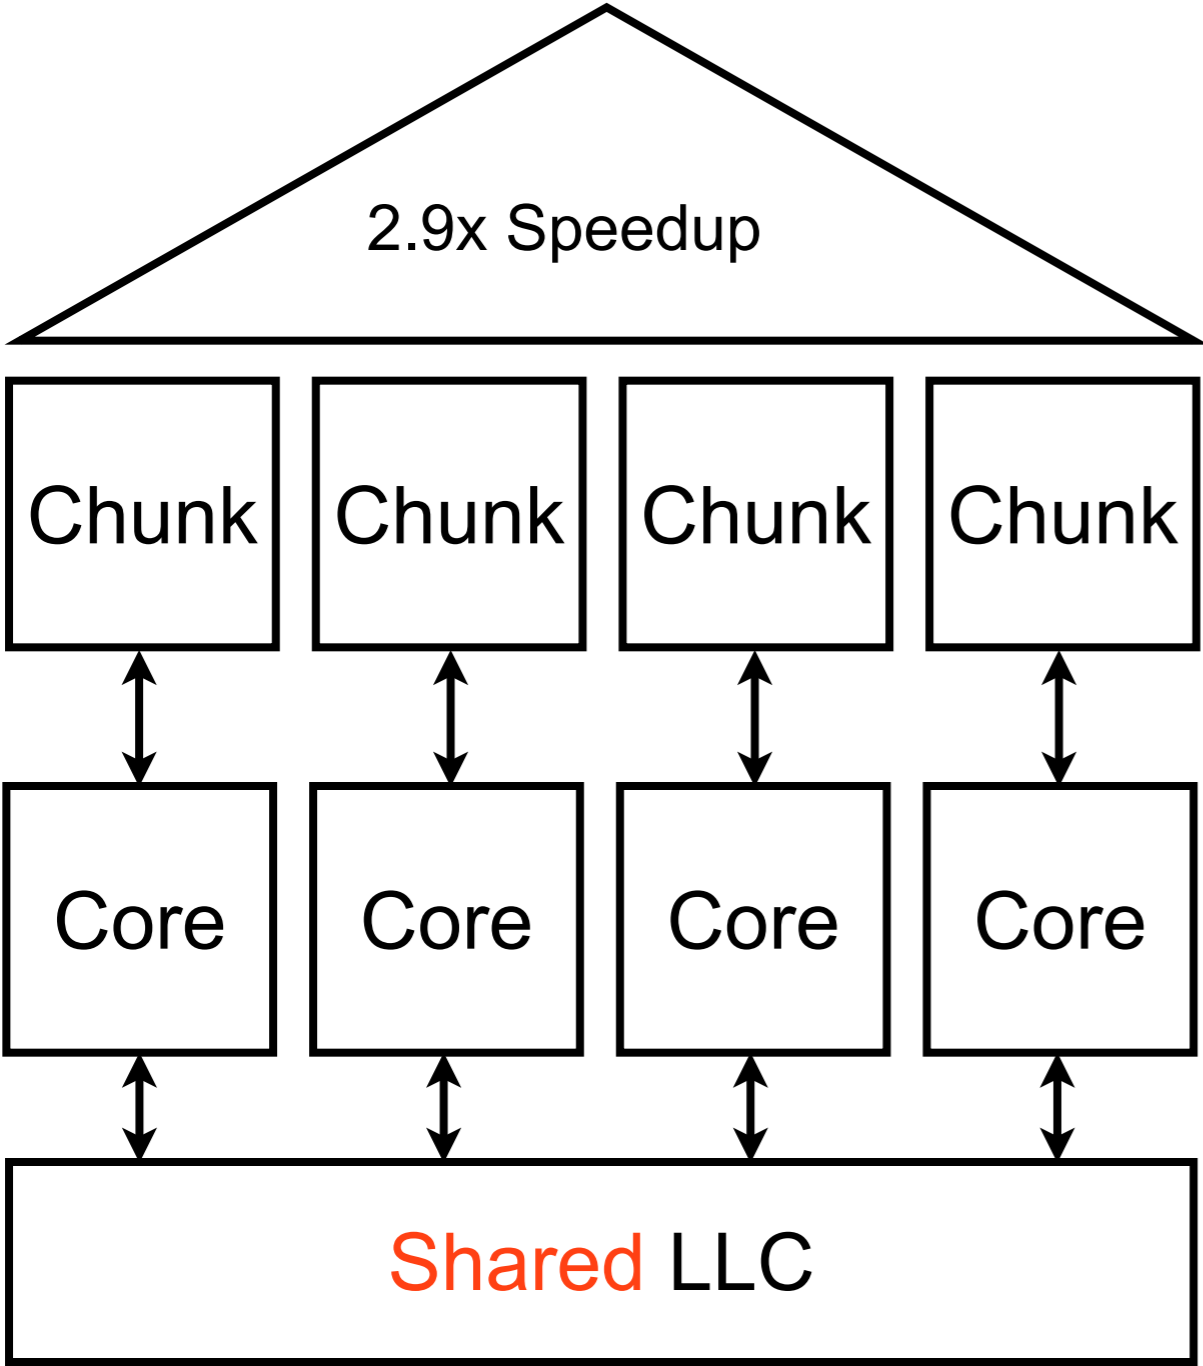
All chunks are completely independent - 8x Speedup?





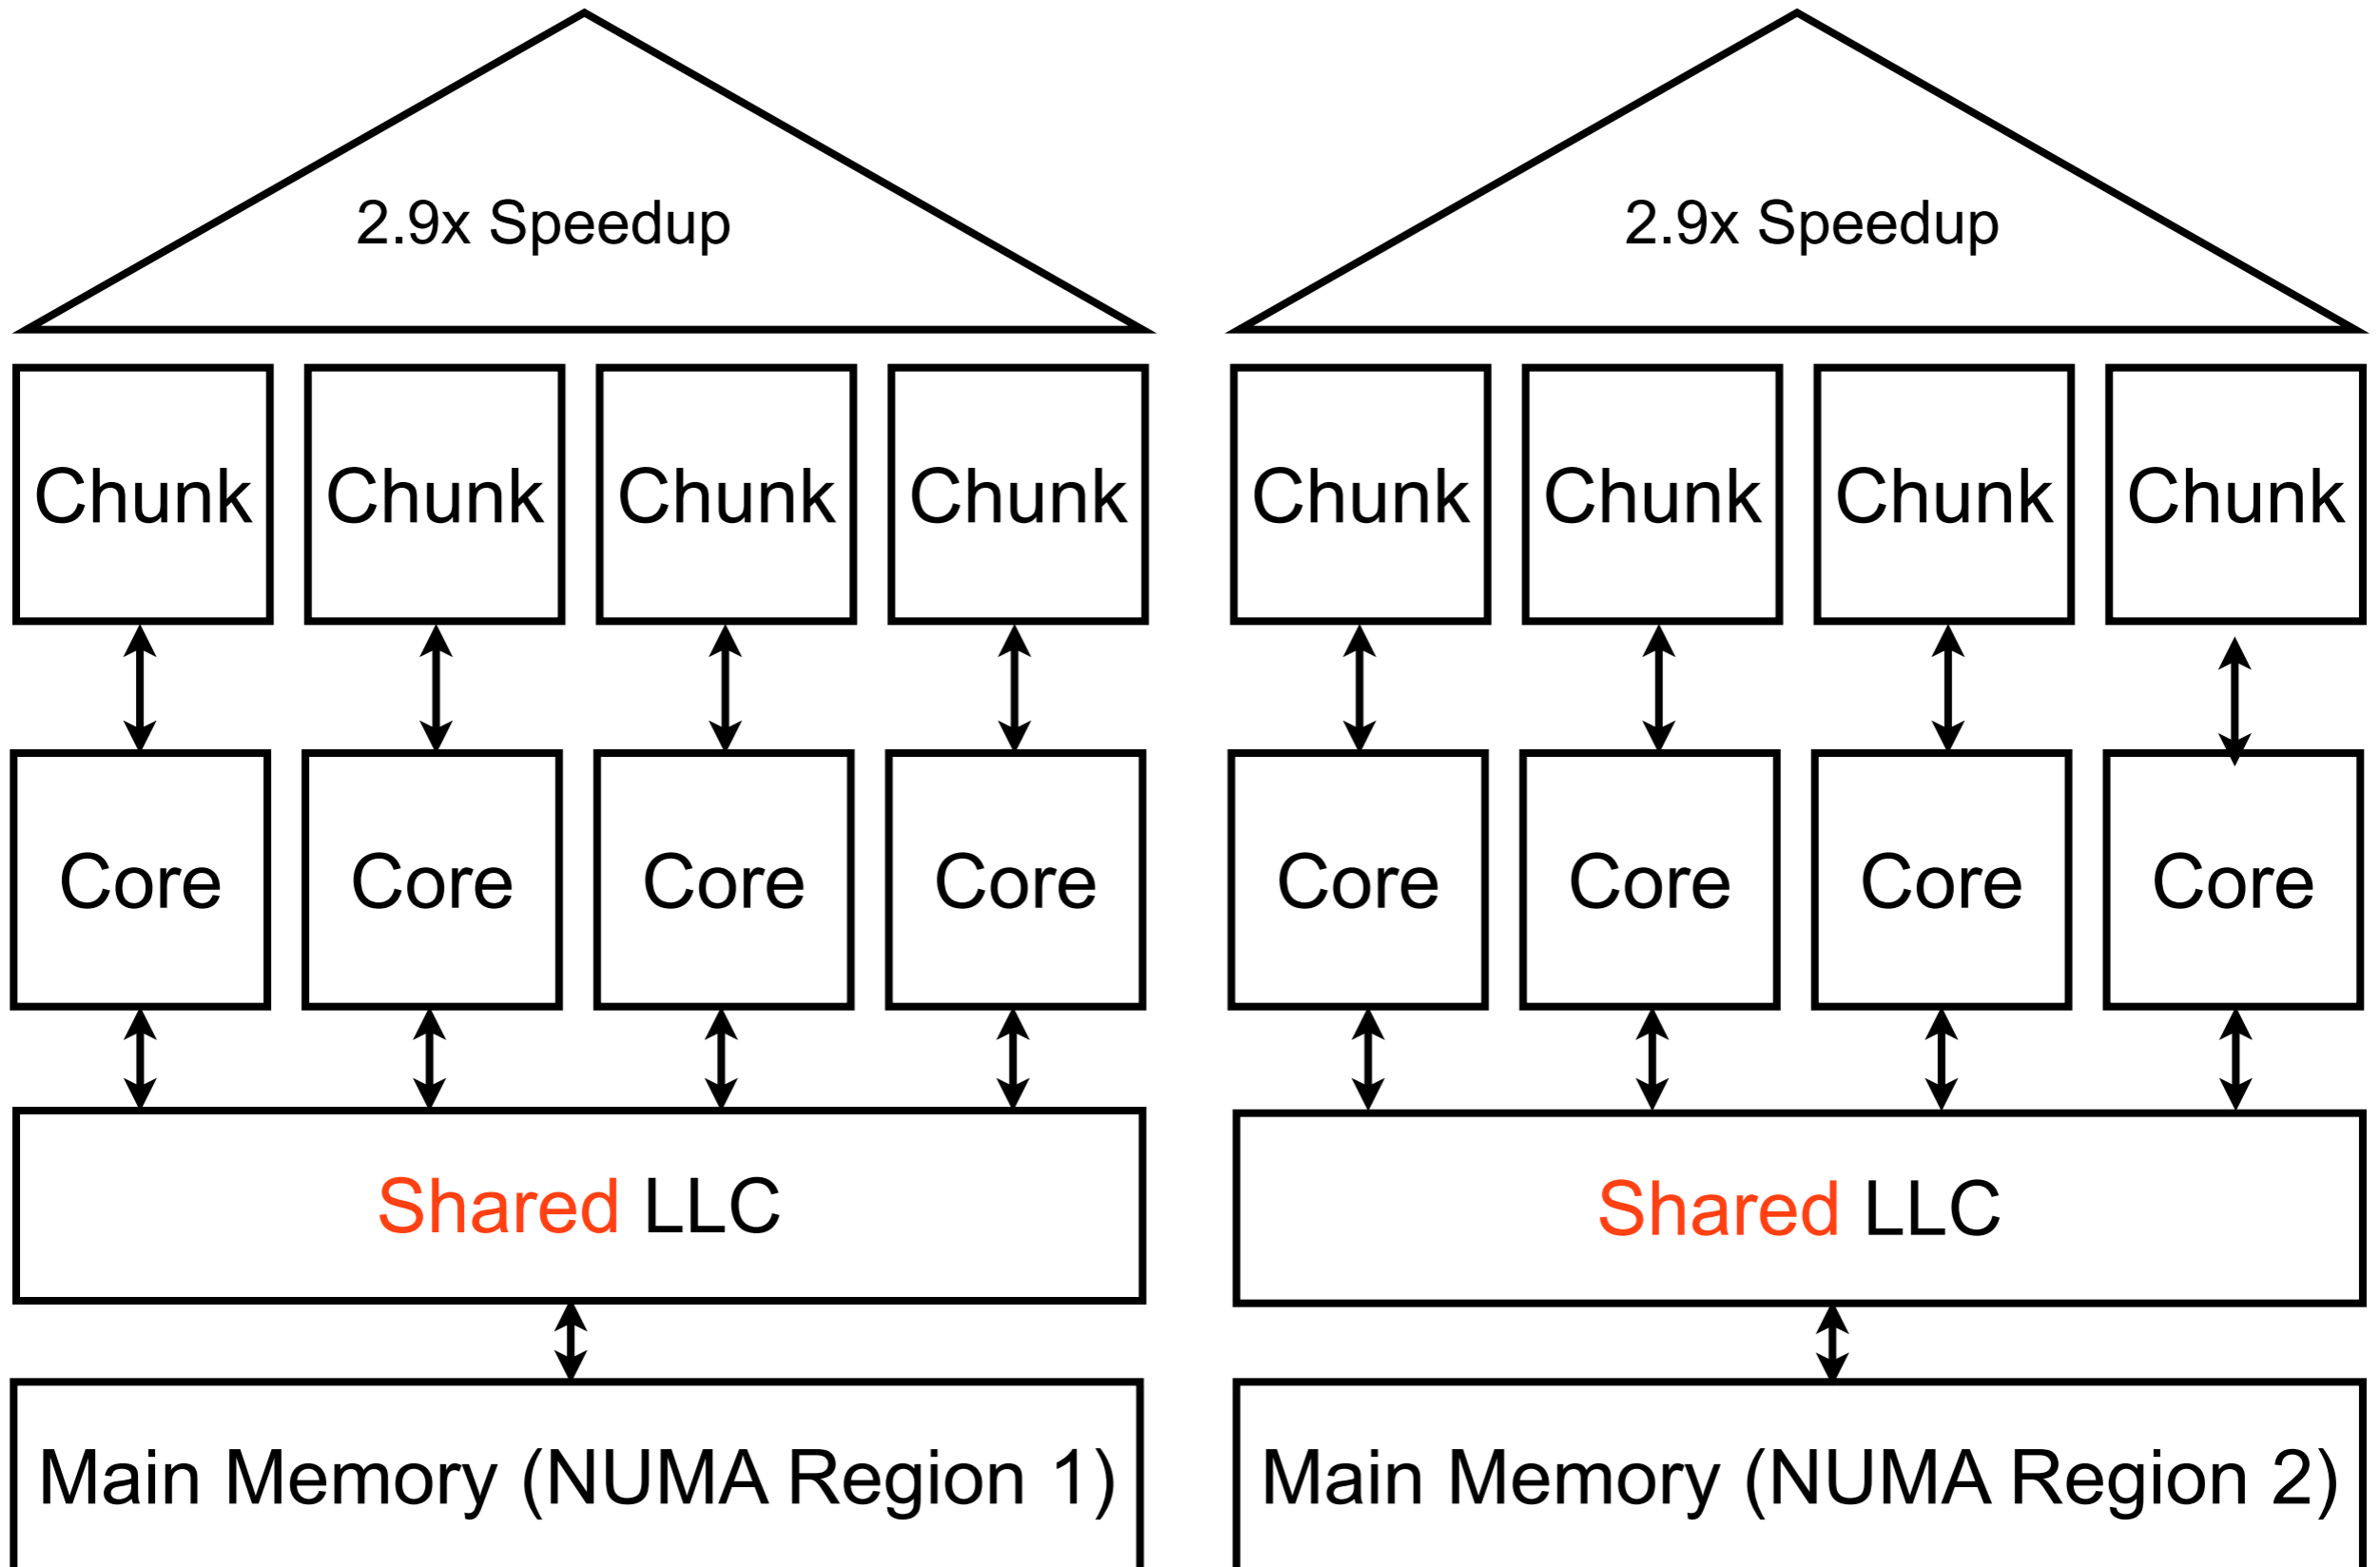
# Chunked Algorithms: Analysis (P-CRS)

All chunks are ~~completely~~ independent - 8x Speedup?



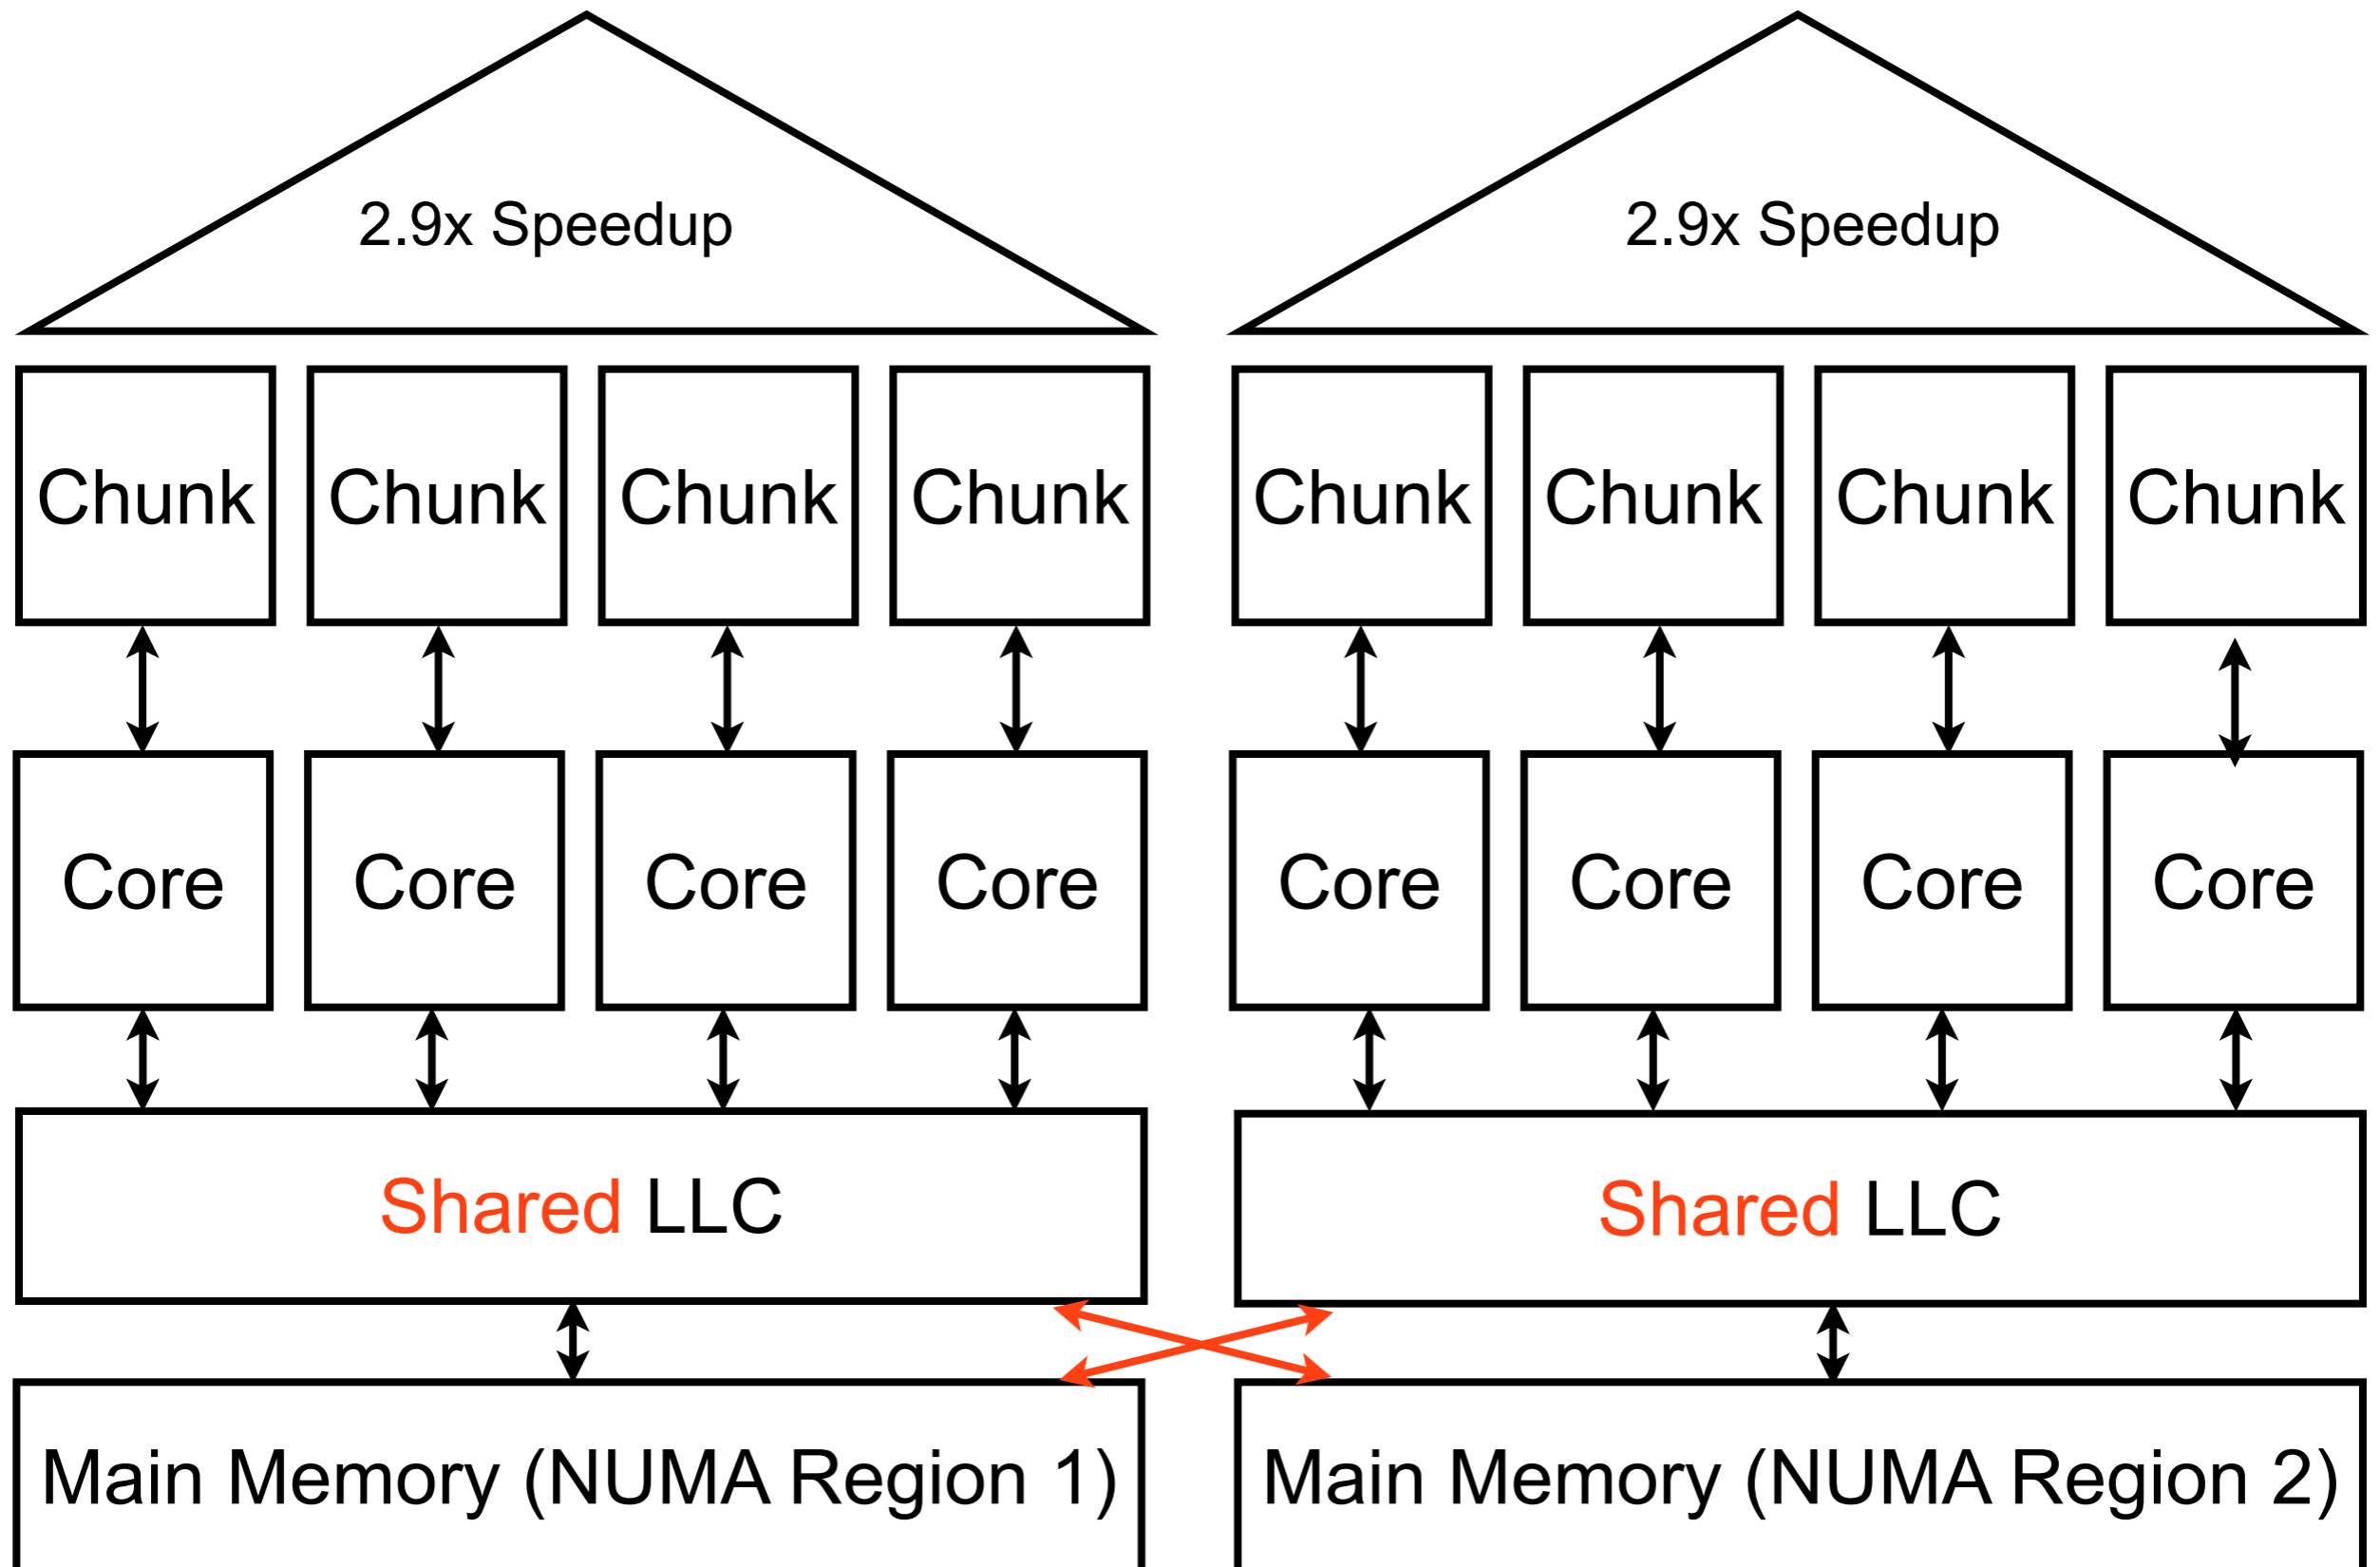
# Chunked Algorithms: Analysis (P-CRS)

All chunks are ~~completely~~ independent - 8x Speedup?




# Chunked Algorithms: Analysis (P-CRS)

All chunks are ~~completely~~ independent - 8x Speedup?




# Conclusion

# Conclusion





# Conclusion

■ RS ■ P-RPRS ■ P-CRS






# Conclusion

■ RS   ■ P-RPRS   ■ P-CRS



# Conclusion




■ RS ■ P-RPRS ■ P-CRS





# Conclusion

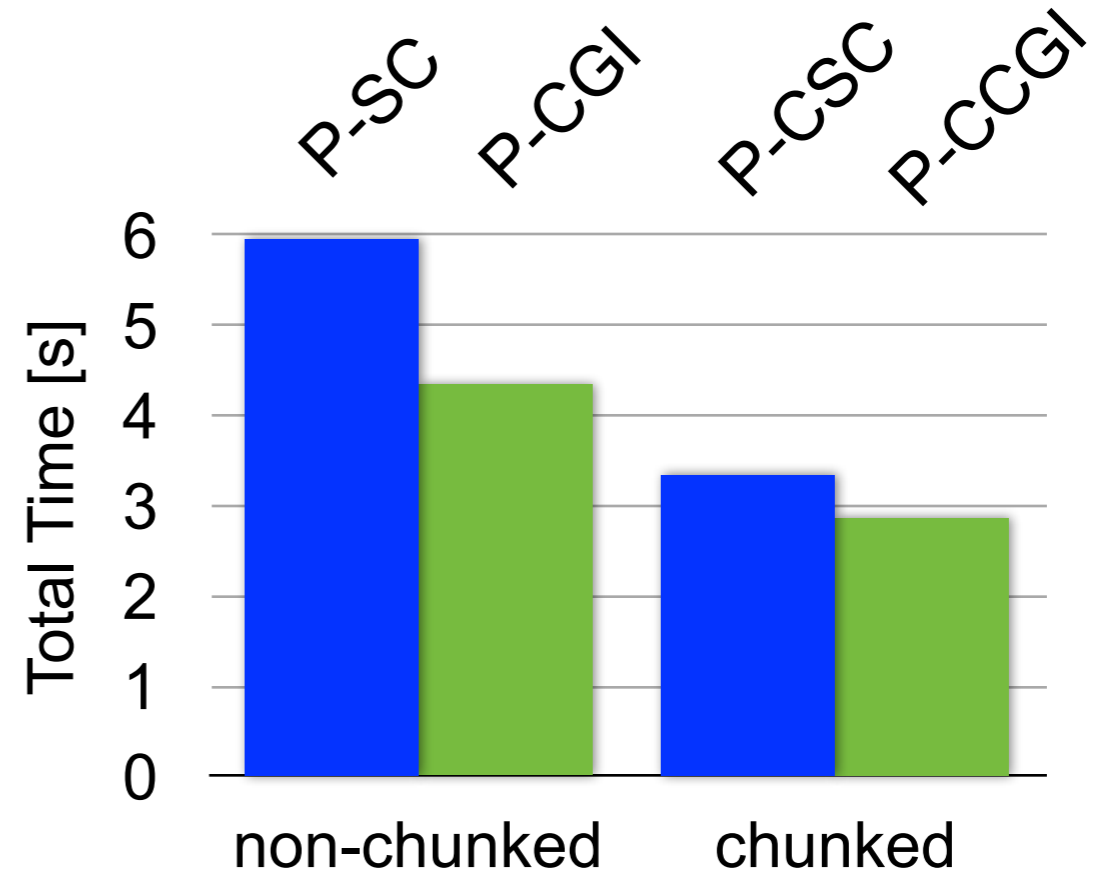

■ RS ■ P-RPRS ■ P-CRS



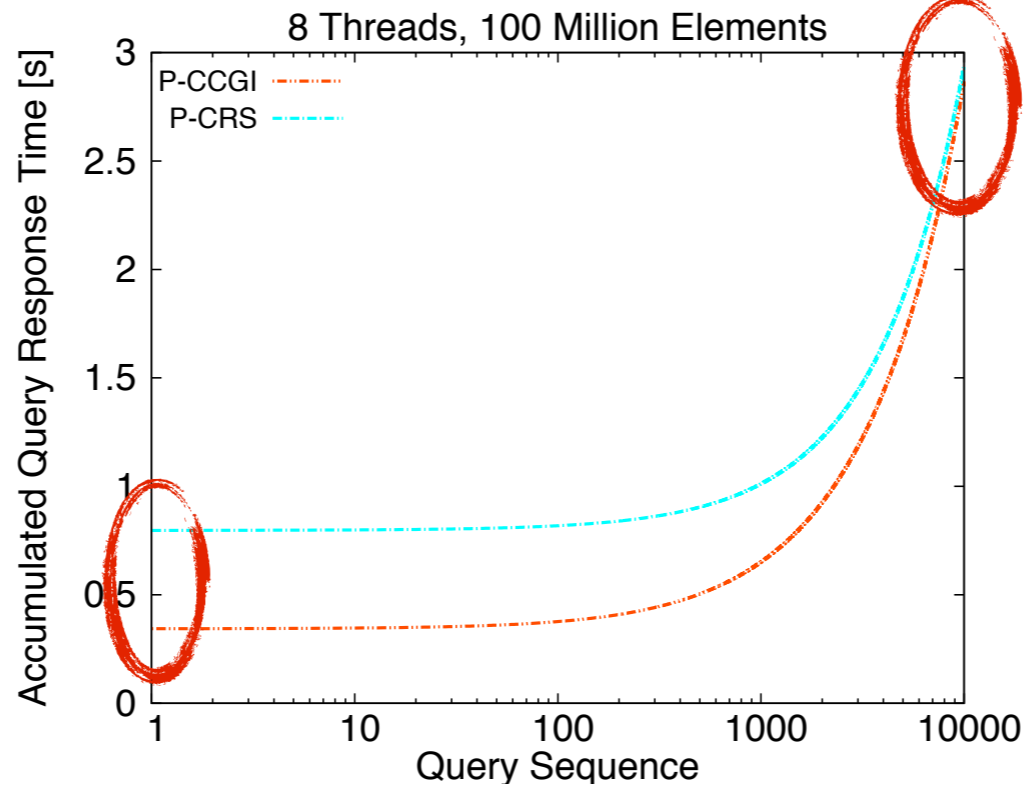
> 10000 queries to win over best cracking

# Conclusion

■ RS   
 ■ P-RPRS   
 ■ P-CRS




gap decreased from 5 seconds (1T) to 0.5 seconds (8T)




> 10000 queries to win over best cracking

# Upcoming


# Upcoming



# Upcoming



# Upcoming



Thank you!