Reason 1: Less Storage Space

compression goal: less storage space

Reason 2: Bandwidth!

compression goal: less bandwidth

Example: Scanning from Disk

large sequential scan of 1 GB **uncompressed** data

100 MB/s sequential read bandwidth

=> 10 sec read time

3 GHz CPU (full overlap)

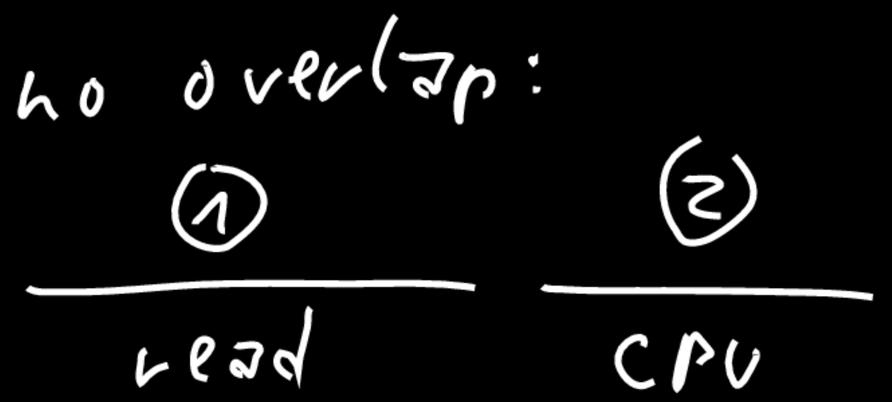

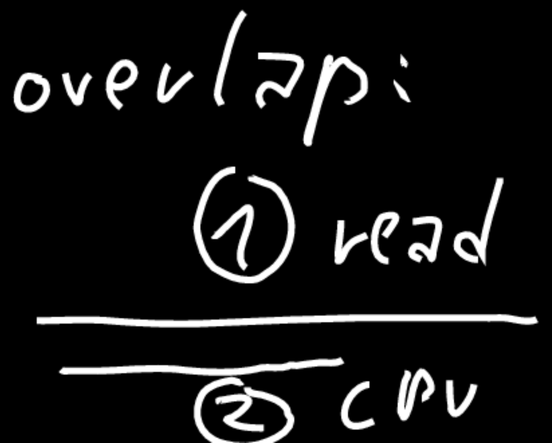

Example: Scanning from Disk

large sequential scan of 1 GB **uncompressed** data

100 MB/s sequential read bandwidth

=> 10 sec read time

3 GHz CPU (full overlap)

=> 30 clock ticks to burn **for every single uncompressed byte** in the input

Let‘s compress it:

1:4 compression ratio => 0.25 GB **compressed** data

=> 2.5 sec read time

=> factor 4 faster

=> up to 2.5*3G = 7.5G clock ticks to burn

=> on average up to 7.5G/0.25G = 30 clock ticks to **uncompress and process** for each **compressed byte**!!

=> 7.5 clock ticks per **uncompressed byte**

40 => 10G clock ticks

10/3 = 3.3̄

Example: Scanning from DRAM

large sequential scan of 1 GB **uncompressed** data

10 GB/s sequential read bandwidth

=> 0.1 sec read time

3 GHz CPU (full overlap)

=> 0.3 clock ticks to burn **for every single uncompressed byte** in the input

Let‘s compress it:

1:4 compression ratio => 0.25 GB **compressed** data

=> 0.025 sec read time

=> factor 4 faster

=> up to 0.025*3G = 0.075G clock ticks to burn

=> on average up to 0.075G/0.25G = 0.3 clock ticks to **uncompress and process** for each **compressed byte**!!

=> 0.075 clock ticks per **uncompressed byte**

Lightweight Compression

goal:

compression + write < write

decompression + read **compressed** < read **uncompressed**

CPU I/O I/O

Lightweight Compression

goal:

decompression + read **compressed** < read **uncompressed**

features:

“CPU-friendly“

lossless vs. lossy

Compression Granularities

	Accessibility	Compression Ratio
attribute values		
tuples		
pages		
horizontal partitions of a table		
tables		
databases