


External Memory Algorithms

Cache-efficient Algorithms

nvm

NVMe


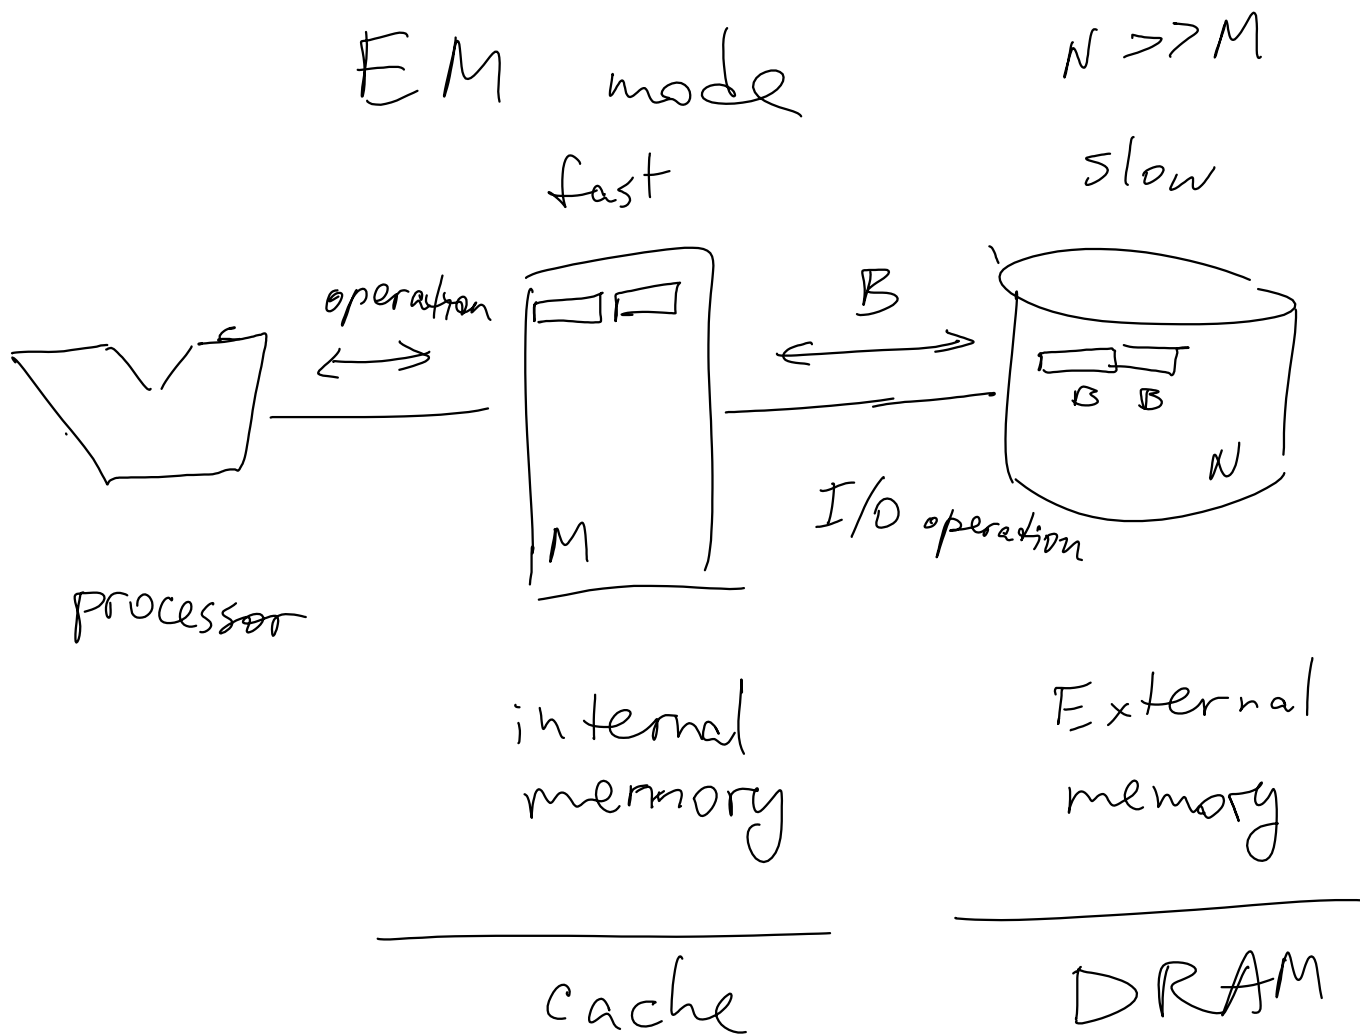


seek time

Disk Access Memory (DAM) model

External Memory Model

I/O Model





E.g. On hard drive $B \approx 4kB$ $M \approx 4GB$

cache line $B \approx 64B$ $M \approx 4-8MB$



Complexity Metrics

only I/O-complexity = # of I/Os
(ignore actual operations)

for $i = 1$ to N $O(n)$ time/operations
 $a[i] = a[i] + 1$
 $scan(n) = \underline{\underline{O\left(\frac{n}{B}\right)}}$ I/Os





Matrix Transpose



$$\sqrt{N} \Rightarrow \frac{M}{B}$$

row-major layout

column-major layout



for $i = 1$ to \sqrt{N}
 for $j = 1$ to \sqrt{N}
 $A[i][j]++$

↑ row ↑ column

$O(n)$ time

$O(\frac{n}{B})$ I/Os

for $row = 1$ to \sqrt{n}

$O(n)$ time



for $col = 1$ to \sqrt{n}

$A[col][row]++$

$O(n)$ I/Os if $\sqrt{n} > \frac{M}{B}$


$O(\frac{n}{B})$ I/Os if $\sqrt{n} \leq \frac{M}{B}$

Morton Z-order




$A[i][j]$


$$\text{addr} = \text{offset} + i \cdot \sqrt{n} + j$$



Morton 2-order Space-filling curves



Matrix transposition



⇒

for $i = 1$ to \sqrt{n}

for $j = 1$ to \sqrt{n}


$\rightarrow C[i, j] = A[j, i]$

$O\left(\frac{n}{B}\right)$

$\Rightarrow O(n) + O\left(\frac{n}{B}\right) = O(n)$
 I/Os

$$M \geq B^2$$

A

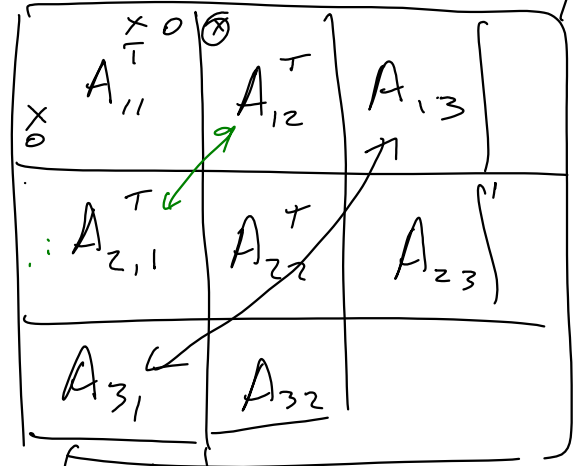


of submatrices = $\frac{n}{n'} = \frac{n}{B^2}$


each takes B I/Os

#A^T

⇒



$$\Rightarrow O\left(\frac{n}{B^2} \cdot B\right) = O\left(\frac{n}{B}\right) \text{ I/Os}$$



$n' = B^2$ elements

for $i = 1$ to B

for $j = 1$ to B

$$A^T[j][i] = A[i][j]$$

$O(B)$ I/Os.



$$\Rightarrow O\left(\frac{B^2}{B}\right) \text{ I/Os}$$

$$\Rightarrow O\left(\frac{n'}{B}\right) \text{ I/Os}$$

for $i = 1$ to $\frac{\sqrt{n}}{B}$

for $j = 1$ to $\frac{\sqrt{n}}{B}$

$$A_{j_i} \Leftrightarrow A_{i_j}$$



copy (A, C):

for $i = 1$ to B

for $j = 1$ to B

$$C[i][j] = A[i][j]$$

$$O\left(\frac{B^2}{B}\right) = O(B)$$

$$\Rightarrow O\left(\frac{|A|}{B}\right) = O\left(\frac{\sqrt{n}}{B}\right)$$

of submatrices = $\frac{n}{n'}$

Total I/os for copying = $O\left(\frac{n}{n'}\right) O\left(\frac{n'}{B}\right) = O\left(\frac{n}{B}\right)$

without any assumptions

to transpose $k \times m$ matrix

takes $\Theta\left(\frac{N}{B} \cdot \log(m, m, \{k, m, M, \frac{N}{B}\})\right)$

Cache-oblivious Algorithms.

Q_{opt} - I/o complexity of an algorithm
under optimal cache-
replacement policy

LRU = least recently used

Q_{LRU} - I/O complexity of an algorithm
under LRU replacement policy

Theorem: $Q_{LRU}(N, M, B) \leq$
 $\leq 2 \cdot Q(N, \underline{2M}, B)$


Flush-when-full policy

Cache-oblivious Matrix
transposition.

$M \geq B^2$

A

$$C = A^T$$



Each submatrix is transposed
recursively

$$\text{Transpose}(A, 1, 1, \sqrt{n}, \sqrt{n})$$



$$\text{Transpose}(A, 1, 1, \frac{\sqrt{n}}{2}, \frac{\sqrt{n}}{2})$$

$$\text{Transpose}(A, 1, \frac{\sqrt{n}}{2} + 1, \frac{\sqrt{n}}{2}, \sqrt{n})$$

;

⋮


$$Q(n, m, B) = \begin{cases} 4Q(\frac{n}{4}, m, B) + O(\frac{n}{B}) \\ O(1) \end{cases} \text{ if } n \leq 1$$




$$Q(n) = \begin{cases} 4Q\left(\frac{n}{4}\right) + O\left(\frac{n}{B}\right) & \text{if } n > M \\ O\left(\frac{n}{B}\right) & \text{if } n \leq M \end{cases}$$

$$\Rightarrow O\left(\frac{n}{B}\right) \quad M \geq B^2$$


Binary Search





$$Q(n) = \begin{cases} Q\left(\frac{n}{2}\right) + O(1) & ; \text{if } n \geq 1 \\ O(1) & ; \text{if } \underline{n \leq B} \end{cases}$$

$$= O\left(\log \frac{N}{B}\right) = O(\log N - \log B)$$



B-tree - a I/O-efficient search tree

$(2, 4)$ -tree - each node has between 2 & 4 children

(a, b) -tree - each node - - -
a & b children



$$2 \leq a \leq \frac{b+1}{2}$$





all leaves are on the same level

B-tree is an (a, b) -tree, where

$$a = \frac{B}{4} \quad \& \quad b = B$$

$\left(\frac{B}{4}, B\right)$ -tree



⇒ Each node has $\Theta(B)$ keys



$$Q(N) = \begin{cases} Q\left(\frac{N}{B}\right) + \Theta(1) & \text{if } N > B \\ \Theta(1) & \text{if } N \leq B \end{cases}$$

$$= \underline{\underline{\Theta\left(\log_B \frac{N}{B}\right)}}$$

van Ende Boas layout



$$\Rightarrow O(\log_B N)$$