# I/O Lower Bounds for Sorting and Matrix Problems 

Jeff Vitter

Duke University
Department of Computer Science


EEF Summer School—July 2002

## Outline

设 Fundamental Techniques for batched problems.

- Merge sort, distribution sort.
i Techniques for solving batched geometric problems.
- Distribution sweeping, batched filtering, randomized incremental construction.
- Red-blue orthogonal rectangle intersection, convex hull, range search, nearest neighbors.
- Empirical results (via TPIE programming environment).
$\Longrightarrow$ Fundamental lower bounds.
- Sorting, permuting, FFT, matrix transposition, bundle sort.
- Dynamic memory allocation
- Hierarchical memory.

级 Parallel disks.

- Load balancing among disks is key issue.
- Duality: reading (prefetching) $\longleftrightarrow$ writing, merging $\longleftrightarrow$ distribution



## Fundamental I/O Bounds (with $D=1$ disk)

~ Batched problems [AV88], [VS90], [VS94]:

- Scanning (touch problem): $\Theta\left(\frac{N}{B}\right)=\Theta(n)$
- Sorting:

$$
\Theta\left(\frac{N}{B} \frac{\log \frac{N}{B}}{\log \frac{M}{B}}\right)=\Theta\left(\frac{N}{B} \log _{M / B} \frac{N}{B}\right)=\Theta\left(n \log _{m} n\right)
$$

- Permuting: $\Theta\left(\min \left\{N, n \log _{m} n\right\}\right)$
\& For other problems [CGGTVV95], [AKL95], ...
- Graph problems $\asymp$ Permutation
- Computational Geometry $\asymp$ Sorting
~ Online problems:
- Searching and Querying: $\Theta\left(\log _{B} N+\frac{Z}{B}\right)=\Theta\left(\log _{B} N+z\right)$
\& What if there are $D$ parallel disks ???


## Fundamental I/O Bounds (with $D=1$ disk)

in Batched problems [AV88], [VS90], [VS94]:

- Scanning (touch problem): $\Theta\left(\frac{N}{B}\right)=\Theta(n)$
- Sorting:

$$
\Theta\left(\frac{N}{B} \frac{\log \frac{N}{B}}{\log \frac{M}{B}}\right)=\Theta\left(\frac{N}{B} \log _{M / B} \frac{N}{B}\right)=\Theta\left(n \log _{m} n\right)
$$

- Permuting: $\Theta\left(\min \left\{N, n \log _{m} n\right\}\right)$

触 For other problems [CGGTVV95], [AKL95], ...

- Graph problems $\asymp$ Permutation
- Computational Geometry $\asymp$ Sorting
( ) Online problems:
- Searching and Querying: $\Theta\left(\log _{B} N+\frac{Z}{B}\right)=\Theta\left(\log _{B} N+z\right)$
~ $D$ parallel disks: Saves factor of $D$ for batched problems, Replace $B$ by $D B$ in online problems (disk striping).


## I/O Lower Bound for Permuting

Permuting problem: Given $N$ distinct items from $\{1,2, \ldots, N\}$, rearrange the $N$ items into sorted order.
~ We will show the lower bound that permuting requires $\Omega\left(\min \left\{N, n \log _{m} n\right\}\right)$ I/Os.
iर Typically the min term is $n \log _{m} n$.
is Permuting is a special case of sorting.
出 I/O lower bound also applies to sorting. It is based only upon routing considerations, since the order is already known.

约 For the pathological case when $N<n \log _{m} n$, we can show that sorting requires $\Omega\left(n \log _{m} n\right) \mathrm{I} / \mathrm{Os}$ in comparison model.
\& In the RAM model, permutation takes only $O(N)$ time. But in I/O model, it (and most interesting problems) require sorting complexity (except for pathological case)!

## I/O Lower Bound for Permuting

Goal: See how many I/O steps $T$ are needed so that any of the $N$ ! permutations of the $N$ items can be realized.

We say that a permutation is realizable if it appears in extended memory in the required order.


Tactic: Determine how much the $t$ th I/O step can increase the number of possible realizable permutations.



## More Refined Analysis to Get Leading Coefficient

Assuming that $M / B$ is an increasing function, \# I/Os required to sort or permute $n$ items is at least

$$
\frac{2 N}{D} \frac{\log n}{B \log m+2 \log N} \sim \begin{cases}\frac{2 n}{D} \log _{m} n & \text { if } B \log m=\omega(\log N) \\ \frac{N}{D} & \text { if } B \log m=o(\log N)\end{cases}
$$

. WLOG, we can assume that each I/O is simple: at any time there is only one copy of each item-on disk or in memory. No copying!

* We need to do enough write I/Os to keep up with read I/Os.
\& The problem is that read I/Os may have fewer than $B$ items.
L Let $b_{i}=\#$ items read in $i$ th read I/O.
出 Let $R=$ \# read I/Os, and $W=\#$ write I/Os.
设 $W \geq \frac{1}{B}\left(\sum_{1 \leq i \leq R} b_{i}\right)$.
~ Each read I/O boosts \# realizable permutations by a factor of $N(1+\log N)\binom{M}{b_{i}}$.
\& Each write I/O boosts \# realizable permutations by a factor of $N(1+\log N)$.


## More Refined Analysis to Get Leading Coefficient

$$
\Longrightarrow(N(1+\log N))^{R+W} \prod_{1 \leq i \leq R}\binom{M}{b_{i}} \geq \frac{N!}{(B!)^{N / B}}
$$


出 By convexity argument，LHS is maximized by setting each $b_{i}:=\widetilde{b}$ ．
设 $W \geq \frac{1}{B}\left(\sum_{1 \leq i \leq R} b_{i}\right)=\frac{1}{B}(R \widetilde{b}) \Longrightarrow R \leq(R+W) /(1+\widetilde{b} / B)$ ．
出 $(N(1+\log N))^{R+W}\binom{M}{\widetilde{b}}^{(R+W) /(1+\widetilde{b} / B)} \geq \frac{N!}{(B!)^{N / B}}$ ．
～Maximize LHS by setting $\widetilde{b}=B$ ，so we get

$$
(N(1+\log N))^{R+W}\binom{M}{B}^{(R+W) / 2} \geq \frac{N!}{(B!)^{N / B}}
$$

which gives desired lower bound on the total number $R+W$ of I／Os．





## Recursive Matrix Multiplication



设 $\mathrm{I} / \mathrm{O}$ complexity for $K \times K$ matrices:

$$
\begin{align*}
T(K) & =8 T\left(\frac{K}{2}\right)+6 \frac{K^{2}}{B}  \tag{1}\\
& =9 \sqrt{3} \frac{K^{3}}{B \sqrt{M}} . \tag{2}
\end{align*}
$$

## Iterative Matrix Multiplication


~ Rather than do partitioning at each level of recursion, do the partitioning all at once, up front.
~ Preprocess by reblocking row-major $K \times K$ input matrices into blocks of size $\sqrt{M / 3} \times \sqrt{M / 3}$.
\& Do matrix multiplication on blocks.
i Reblock output into row-major order.


## The Need for Memory-Adaptive EM Algorithms.

h Traditional EM algorithms assume fixed memory allocation.
~ Problem:

- OS/DBMS can dynamically change memory allocation.
- EM applications exhibit thrashing.
is Solution:
EM algorithms that adapt online to memory fluctuations.
d All prior work has been exclusively empirical:
- Memory-Adaptive Hash Join (Zeller\& Gray, Pang et al.)
- Pang et al., 1995: Non-optimal memory-adaptive sort.
- Zhang and Larson, 1997: Memory-adaptive sort, works only for very restricted kinds of fluctuations.



## Why Traditional EM Algorithms Thrash.



Merging 8 runs using 5 internal memory blocks: Leading blocks of 4 runs are out of memory

设 If $m$ drops to less than 8 but merge-order remains 8 , worst case cost is one I/O per element output by merge.
~ Solution: Reorganize computation; ie, change merge-order in response to change in $m$.

## Dynamic Memory Environment

i EM algorithm is allocated $m$ memory blocks by the OS/DBMS for an unspecified amount of time.

设 When OS/DBMS wants to change the allocation of $m$, it first allows EM algorithm to carry out $m$ I/Os ("Reaction time"). Then it changes $m$.

设 We use a simplified "constant factor approximation" of this model.

## Simple Model for Memory－Adaptive EM Algorithms

（ EM algorithm $\mathcal{A}$ is allocated memory in an
allocation sequence $\sigma=m_{1}, m_{2}, m_{3}, \ldots$ of allocation phases．
该 OS／DBMS determines $\sigma$ in an online adversarial manner．
设 $i$ th phase：Algorithm owns $m_{i}$ blocks of memory for $2 m_{i} \mathrm{I} /$ Os．
纹 EM algorithm must adapt to allocation sequence．
iर Suppose that $\mathcal{A}$ solves problem $\mathcal{P}$ during $\sigma$ ．
i． $\mathcal{A}$ is dynamically optimal for $\mathcal{P}$ iff
－No other algorithm $\mathcal{A}^{\prime}$ can solve problem $\mathcal{P}$ more than a constant number of times during $\sigma$ ．

## Dynamic Memory Lower Bound for Sorting

$i$ th phase:


Use comparison model:

$$
\begin{aligned}
& \text { \# possible outcomes } \\
& \text { to comparisons per I/O }
\end{aligned}=\left\{\begin{array}{ll}
B!\times\binom{ M_{i}}{B} & \text { reading unread block. } \\
\binom{M_{i}}{B} & \text { reading dirty block. }
\end{array} \quad \begin{array}{l}
(B!)^{N / B} \prod_{i}\binom{M_{i}}{B}^{2 m_{i}} \geq N!\quad \Longrightarrow \quad \sum_{i} 2 m_{i} \log m_{i}=\Omega(n \log n) .
\end{array}\right.
$$

## Resource Consumption of Sorting

Sorting algorithm completes in $\ell$ phases
$\Longrightarrow \quad \sum_{i=1}^{\ell} 2 m_{i} \log m_{i}=\Omega(n \log n)$.
iर Resource Consumption of an $\mathrm{I} / \mathrm{O}$ in phase $i$ is

$$
\log m_{i}
$$

(4) Algorithm is dynamically optimal iff

Total Resource Consumption $(\mathrm{RC})=O(n \log n)$.

## A Framework for Memory-Adaptive Mergesort

\& Run Formation

- Phase $i \Longrightarrow$ Generate a run of length $m_{i}$ blocks.
- Number of runs in $\mathcal{Q}$ is $n_{0} \leq n$. (Very often, $n_{0} \ll n$.)
- Total Resource Consumption

$$
\begin{aligned}
\mathrm{RC}_{\text {run_formation }} & =O(\# \mathrm{I} / \mathrm{Os} \times \text { Max cost of each } \mathrm{I} / \mathrm{O}) \\
& =O\left(n \log m_{\max }\right)
\end{aligned}
$$

\& Merging Stage

- Memory-adaptive merging routine $\mathcal{M}$.
- Repeat: Merge $R$ runs from $\mathcal{Q}$, append output run to $\mathcal{Q}$.


## Resource Consumption Requirement for Merging

$$
\begin{aligned}
\mathrm{RC}_{\text {sort }} & =O\left(\mathrm{RC}_{\text {run_formation }}+\frac{\log n_{0}}{\log R} \mathrm{RC}_{\text {pass }}\right) \\
& =O\left(n \log m_{\max }+\frac{\log n_{0}}{\log R} \mathrm{RC}_{\text {pass }}\right)
\end{aligned}
$$

For dynamic optimality,

$$
\begin{aligned}
& \text { ) } \mathrm{RC}_{\mathrm{pass}}=O(n \log R) \text {. } \\
& \text { ( } R=\Omega\left(m_{\max }^{c}\right) \text {. }
\end{aligned}
$$

设 Various external memory data structures and techniques are required for the scheme to work efficiently．
\％Lower Bounds for problems related to sorting and matrix multiplication（and related problems）．

设 Sorting algorithm was used to get dynamically optimal algorithms for permuting，permutation networks，FFT．
～Dynamically Optimal memory－adaptive version of a buffer tree．
i Techniques applicable via sorting and buffer trees to many other applications．

次 Dynamically optimal matrix multiplication algorithm．

## Conclusions and Open Problems

स Répertoire of useful paradigms (distribution, merging, distribution sweeping, persistence, parallel simulation, B-trees, external interval tree, external priority search tree) for important problems.

- Worst-case optimality requires overhead.
- Simpler versions are practical!
- Building blocks for external data structures
i Lots of interesting open problems!
- Lower bounds without indivisibility assumption.
- [Adler] showed that removing the indivisibility assumption for an artificial problems related to transposition can lead to faster algorithms.
- New models: hierarchical memory, oblivious caching, dynamic memory allocation, MEMS, optical storage, ....


