# Batched Dynamic Geometric Problems 

## Jeff Vitter

Duke University
Center for Geometric and Biological Computing and Department of Computer Science
 http://www.cs.duke.edu/CGBC/

July 2002

## Outline

设 Fundamental Techniques for batched problems.

- Merge sort, distribution sort.
$\Longrightarrow$ Techniques for solving batched geometric problems.
- Distribution sweeping, batched filtering, randomized incremental construction, parallel simulation.
- Red-blue orthogonal rectangle intersection, convex hull, range search, nearest neighbors.
- Empirical results (via TPIE programming environment).
\& Fundamental lower bounds.
- Sorting, permuting, FFT, matrix transposition, bundle sort.
- Dynamic memory allocation
- Hierarchical memory.
\& Parallel disks.
- Load balancing among disks is key issue.
- Duality: reading (prefetching) $\longleftrightarrow$ writing, merging $\longleftrightarrow$ distribution


## Review of Parallel Disk Model

[Aggarwal \& Vitter 88], [Vitter \& Shriver 90, 94], ...


$$
\begin{aligned}
N & =\text { problem data size. } \\
M & =\text { size of internal memory. } \\
B & =\text { size of disk block. } \\
D & =\text { number of independent disks. } \\
P & =\text { number of CPUs. } \\
Q & =\text { number of queries. } \\
Z & =\text { problem output size. }
\end{aligned}
$$

Notational convenience (in units of blocks):
$n=\frac{N}{B}, m=\frac{M}{B}, q=\frac{Q}{B}, z=\frac{Z}{B}$.

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

约 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: $\Theta\left(\min \left\{N, n \log _{m} n\right\}\right)$
i For other problems [CGGTVV95], [AKL95], .. .
- Graph problems $\asymp$ Permutation
- Computational Geometry $\asymp$ Sorting

设 Online problems:

- Searching and Querying: $\Theta\left(\log _{B} N+\frac{Z}{B}\right)=\Theta\left(\log _{B} N+z\right)$


## Batched Problems in Geometry

［GTVV93］，［AVV95］，［APRSV98a］，［APRSV98b］，［CFMMR98］
～Orthogonal rectangle intersection．
i Red－blue line segment intersection．
约 General line segment intersection．
解 All nearest neighbors．
论 2－D and 3－D convex hulls．
\＆Batched range queries．
\＆Trapezoid decomposition
纹 Batched planar point location．
\＆Triangulation．
Use of virtual memory $\Longrightarrow \Omega\left(N \log _{B} N+Z\right)$ I／Os．Bad ！！！
We can improve this to $O\left(n \log _{m} n+z\right)$ I／Os using
\＆Distribution sweep．
～Persistent B－trees and batched filtering．
该 Random incremental construction．
is Parallel simulation．

## Orthogonal Line Segment Intersection



Problem: Find all intersections of vertical segments with horizontal segments.

## Internal Memory Approach

设 Presort the endpoints in $y$－order．
i．Sweep the plane from top to bottom with a horizontal line．
4 When reaching a vertical segment，store its $x$ value in a balanced tree． When leaving a vertical segment，delete its $x$ value from the tree．

设 At any given time，the balanced tree stores the vertical segments hit by the sweep line．
i When reaching a horizontal segment，do a 1－d range query in the tree to find intersections with vertical segments．Time is $O\left(\lg N+Z^{\prime}\right)$ ， where $Z^{\prime}$ is number of intersections reported．

该 Total running time is $O(N \lg N+Z)$ ．

## External Solution?



ش. Internal plane-sweep solution runs in $O(N \log N+Z)$ time.
\& Using B-tree gives an $O\left(N \log _{B} n+z\right)$ I/O solution.
设 We want an $O\left(n \log _{m} n+z\right)$ I/O solution that takes advantage of batching!

## Distribution Sweeping

[Goodrich, Tsay, Vengroff \& Vitter 93]


## Distribution Sweeping

4．Presort endpoints by $x$ and $y$ coordinates．
设 Divide the $x$－range into $\Theta(m)$ slabs，so that each slab contains the same number of $x$ values of vertical segments．

级 Sweep all slabs simultaneously from top to bottom，keeping the vertical segments of a slab in a stack．
i For each slab spanned by a horizontal segment，output all＂living＂ vertical segments in the slab＇s stack and delete all＂dead＂vertical segments from stack．

次 For the left and right＂endpieces＂of a horizontal segment，that stick out into a slab but don＇t completely span it，handle those intersections recursively for each slab．

## Implementing a Stack

Various stack operations:

1. Push element onto top
2. Read top entry,
3. Pop entry from top.

Variants: We can read the top $k$ entries from the stack by iterating operation $3 k$ times and then operation $1 k$ times.

Keep current block and one other in internal memory (using LRU).

It takes $O(B)$ pushes or pops to require one I/O. $\Longrightarrow$ \# I/Os per operation $=O\left(\frac{1}{B}\right)$ amortized.

|  |
| :---: |
| $\vdots$ |
| $i_{1}$ |
| $i_{2}$ |
| $i_{3}$ |
| $i_{4}$ |
| $\vdots$ |
| $i_{k}$ |
| $i_{k+1}$ |
| $i_{k+2}$ |
| $\vdots$ |

## Analysis of External Distribution Sweeping

约 Each of the $\Theta(m)$ stacks can use $O(1)$ blocks in internal memory.
स Therefore, each push, pop, or read uses $\left(\frac{1}{B}\right)$ I/Os amortized.
设 In each pass, the $O(N)$ vertical segments are inserted into the stack in $O(n)$ I/Os.
id For each of the $O(N)$ horizontal segments, we report intersections in the slabs it completely spans. If the total number of intersections reported in this pass is $Z^{\prime}$, the number of I/Os is $O(n)$ plus the cost of $Z^{\prime}$ stack push, pop, or read operations, which is $O\left(n+Z^{\prime} / B\right)$.

## Analysis of External Distribution Sweeping

设 We recurse on each of the $\Theta(m)$ slabs to handle the left endpieces and right endpieces of the horizontal segments．
（ Note that the total number of endpieces at every level of recursion is at most $2 \times \#$ horizontal segments．
It doesn＇t double at each level．
㭥 Number levels of recursion is $O\left(\log _{m} n\right)$ ．
设 Final result：$O\left(n \log _{m} n+z\right)$ I／Os．

## Class Quiz

What about batched range searching?


We want to be able to do $Q$ range queries on $N$ points in $O\left((n+q) \log _{m} n+z\right)$ I/Os.

Ideas???

## Distribution Sweeping to the Rescue

Sweep
Line

Slab 1 Slab 2
Slab 3
Slab 4
Slab 5

## Distribution Sweeping

i Presort points on $x$ and $y$ coordinates．
纹 Presort the bottom horizontal sides of the query rectangles by their $y$ coordinate．

该 Sweep all slabs simultaneously from top to bottom，keeping the points of each slab in a stack．

纹 For each slab spanned by a bottom horizontal side，traverse its stack．
㕸 Recursively handle the left endpiece and the right endpiece．

## Analysis of Distribution Sweeping

, Each sweep uses $O\left(n+q+\frac{Z^{\prime}}{B}\right)$ I/Os.
2 In each pass, the points are inserted into the stacks in $O(n)$ I/Os.
\% For each query rectangle, we report the points that are both inside the rectangle and inside the slab spanned by the rectangle. If the total number of points reported in this pass is $Z^{\prime}$, the number of I/Os is $O\left(q+Z^{\prime} / B\right)$.

4 We recurse in each of the $\Theta(m)$ slabs to handle the left endpieces and right endpieces of the query rectangles.

该 The total number of endpieces at every level of recursion is at most $2 Q$.

क Recursion levels: $O\left(\log _{m} n\right)$.
纹 Final result: $O\left((n+q) \log _{m} n+z\right)$ I/Os.

## Output-Sensitive Convex Hulls

i Goal: Compute the convex hull in
$T(N, H)=O\left(n \log _{m}\lceil h\rceil+n\right)$ I/Os, where $H=h B$ is the size of the convex hull.

约 Motivation: $H$ is often $\ll N$.
h. Follow internal memory approach of [Kirkpatrick-Seidel].
\& We no longer have time to sort by $x$ coordinate for a distribution sweep.
L. We can avoid the need to presort by $x$ coordinate and can instead do the partitioning into slabs using the partitioning method described earlier.
~ Cost is $O(n)$ I/Os to do the partitioning.
设 But the number of slabs needs to be smaller: $O(\sqrt{m})$. But that's OK: \# levels of recursion is still $O\left(\log _{m} h\right)$.

## Output-Sensitive Convex Hulls

Main Ideas:

1. Apply Partitioning Lemma. Each of the $S=\sqrt{m}$ slabs has between $\frac{3 N}{4 S}$ and $\frac{5 N}{4 S}$ points.
2. Find hull edges crossing dividers in $O(n)$ I/Os (à la [Goodrich]).
3. Recurse only when needed.

Result: $O\left(n \log _{m}\lceil h\rceil+n\right)$ I/Os.
Analysis: Assuming Step 2 requires $O(n)$ I/Os, each recursive call
次 either finds more than $\sqrt{m} / 2$ edges

ch or it eliminates $N / 2$ points.


## Proof that $T(N, H) \leq c n \log _{m}\lceil h\rceil+n$

Divide-and-conquer gives $T(N, H)=\sum_{i} T\left(N_{i}, H_{i}\right)+n$.
By convexity, the worst case is when each $H_{i}=\frac{N_{i}}{N} H_{i}$, which is between $\frac{3}{4} \frac{H}{\sqrt{m}}$ and $\frac{5}{4} \frac{H}{\sqrt{m}}$.
Case 1: $\left\lceil\frac{5}{4} \frac{H}{B \sqrt{m}}\right\rceil \leq \frac{2 H}{B \sqrt{M}}$. By D-and-C and induction hypothesis,

$$
\begin{aligned}
T(N, H) & \leq \sum_{i}\left(\frac{c n}{\sqrt{m}} \log _{m}\left[\frac{2 H}{B \sqrt{m}}\right\rceil+n_{i}\right)+n \\
& \leq c n \log _{m} \frac{2 H}{B \sqrt{m}}+\sqrt{m}+n+n \\
& \leq c n \log _{m} \frac{H}{B}+c n \log _{m} 2-\frac{c n}{2}+2 n+\sqrt{m} \\
& \leq c n \log _{m} h+n,
\end{aligned}
$$

assuming $m>4$ and $c$ is large enough s.t.
$c n \log _{m} 2-c n / 2+2 n+\sqrt{m} \leq n$.

## Proof that $T(N, H)=O\left(n \log _{m}\lceil h\rceil+n\right)$

Case 2: $\frac{5}{4} \frac{H}{B \sqrt{m}} \leq 1$.
By divide-and-conquer and induction hypothesis,

$$
\begin{aligned}
T(N, H) & \leq \sum_{i}\left(\frac{c n}{\sqrt{m}}(0)+n_{i}\right)+n \\
& =2 n
\end{aligned}
$$

## 3－d Convex Hull［Goodrich－Tsay－Vengroff－Vitter］

纹 Plane sweep and disribution sweep don＇t seem applicable．
纹 Instead we use externalization of randomized construction of ［Reif－Sen］to compute 3－d convex hulls ．
i．Idea：Use random sampling in the dual problem（intersecting half－spaces containing origin）．
i Take $O\left(\log _{m} n\right)$ samples of $S=N^{\epsilon}$ half－spaces and recursively compute intersection of each sample．

埌 For each sample，construct（triangulated）＂cones＂formed from origin to faces and find cones hit by the $N$ input half－spaces．


## 3-d Convex Hull

4. Eliminate redundant half-spaces.

识 Poll to find a sample that gives a well-balanced partition.
in With high probability, there will be a sample such that the subproblem sizes add up to $O(N)$ and the largest is at most $\log N$ times the smallest.

कर Polling uses random sampling to find the good sample in $O(n)$ I/Os.
i Recurse in each cone.

## Batched Persistent B-trees

(4. Problem: given $\sigma_{1}, \sigma_{2}, \ldots, \sigma_{N}$, where $\sigma_{i}=\operatorname{insert}(x)$ or delete $(x)$, construct a data structure that allows a "B-tree search" in the past.

纹 We will apply distribution sweeping to construct a structure with $\sqrt{m}$-way branching.
~ We achieve $O\left(n \log _{m} n\right)$.
~ Online method takes $O\left(N \log _{m} n\right)$.

## Batched Persistent B－trees



谈 Online property doesn＇t hold for batched persistent B－trees．
设 Online Property：For any time $t$ ，a root to leaf search or range search w．r．t．time $t$ traverses only blocks that are $\geq$ half－full．

设 Important for output－sensitivity in time－stamped 1－d range search （3－sided range search）．

## Batched Persistent B-trees

设 Online property not important for applications like batched planar point location.
( Applications:

- $K$ simultaneous point location queries.
- $K$ ray-shooting queries in CSG model.
- $K$ range queries.
- Graph drawing.


## Persistent B-trees and Batch Filtering


~ Outdegree $\leq m$
i. Search a layered planar dag in

$$
O(n+(q+1) \text { height }) \text { I/Os, }
$$

where $Q=q B$ is the number of queries.

## Persistent B－trees and Batch Filtering

约 Start by sending all queries to the root node．
i Proceed level by level，sending all $Q$ queries to level $i$ before sending any to level $i+1$ ．

该 To do this I／O－efficiently，maintain a FIFO queue of queries that flow through the edges between current level and next level．
－If less than B queries traverse an edge，store edges in queue．
－Otherwise，store a pointer to a linked list of blocks．
设 The queue for the next level is produced from the current one I／O－efficiently．

## Map Overlay / Spatial Join

Spatial Data:

- Maps
- Terrains
- CAD models
- VLSI models

Traditionally, spatial data is stored in layers.

Overlaying layers (map overlay)

is a fundamental operation in geographical information systems
(GIS).

## Geographical Information Systems

A typical GIS might store the following layers:

- Roads
- Rivers and lakes
- Railroads

Example: roads in Triangle Area.


## Geographical Information Systems

Query: "Find all bridges in Triangle Area"
Requires map overlay (the roads map with the rivers/lakes map), a type of spatial join.


## Spatial Join



令 In database literature often solved in two steps：
－Filter step：Compute minimal bounding rectangles for each region and compute intersections between rectangles from different maps （red－blue rectangle intersection）．
－Refinement step：Validate intersections．
该 We consider filter step：intersecting the two sets of rectangles．
资 Issues：
－\＃I／Os，
－Indexed vs．non－indexed structures for storing the rectangles．
－Skewed data

## Spatial Join



令 In database literature often solved in two steps：
－Filter step：Compute minimal bounding rectangles for each region and compute intersections between rectangles from different maps （red－blue rectangle intersection）．
－Refinement step：Validate intersections．
该 We consider filter step：intersecting the two sets of rectangles．
资 Issues：
－\＃I／Os，
－Indexed vs．non－indexed structures for storing the rectangles．
－Skewed data

## Case I: No Indexes

## Previous Algorithm: PBSM [PD96]

Partitions data into tiles

## Drawbacks:

Reports duplicate intersections
A tile may not fit in memory


Robust to skewed data

## Red-Blue Rectangle Intersection



该 Sweep plane while maintaining two active lists of red and blue rectangles intersecting vertical sweep line [BW80]:

- When top of blue rectangle is reached:
(i) Insert blue rectangle in blue active list.
(ii) Find intersections with rectangles in red active list.
- When bottom of blue rectangle is reached:
(i) Remove rectangle from blue active list.

设 Red rectangles are handled similarly.

## Red-Blue Rectangle Intersection



H Algorithm performs badly ( $>N$ I/Os)
if size of active lists $>M$.

## Red-Blue Rectangle Intersection



设 Algorithm performs badly ( $>N$ I/Os)
if size of active lists $>M$.
设 Solved in optimal $O\left(n \log _{m} n+z\right)$ I/Os using general method for solving Batched Dynamic Problems.
i Sequence of operations $a_{1}, a_{2}, \ldots, a_{N}$ known beforehand.
( $a_{i}$ is Insert, Delete or Query.)
ش Key point: Updates and queries are batched!

## Sketch of External Solution [APRSV98]:



1. Divide plane into $\sqrt{m}$ slabs, each with $O(N / \sqrt{m})$ endpoints.
2. Break rectangles into three pieces: left endpiece, centerpiece, and right endpiece.
3. Find $Z^{\prime}$ intersections involving at least one centerpiece.
4. Recursively solve problem in each slab for endpieces.
is $O\left(\log _{\sqrt{m}} n\right)=O\left(\log _{m} n\right)$ levels of recursion.
स Performing Step 3 in $O\left(n+\frac{Z^{\prime}}{B}\right)$ I/Os

$$
\Longrightarrow \quad O\left(n \log _{m} n+z\right) \text { I/Os total. }
$$

## Sketch of External Solution [APRSV98]:



1. Divide plane into $\sqrt{m}$ slabs, each with $O(N / \sqrt{m})$ endpoints.
2. Break rectangles into three pieces: left endpiece, centerpiece, and right endpiece.
3. Find $Z^{\prime}$ intersections involving at least one centerpiece.
4. Recursively solve problem in each slab for endpieces.

क $O\left(\log _{\sqrt{m}} n\right)=O\left(\log _{m} n\right)$ levels of recursion.
स Performing Step 3 in $O\left(n+\frac{Z^{\prime}}{B}\right)$ I/Os

$$
\Longrightarrow \quad O\left(n \log _{m} n+z\right) \text { I/Os total. }
$$

## Sketch of External Solution [APRSV98]:



1. Divide plane into $\sqrt{m}$ slabs, each with $O(N / \sqrt{m})$ endpoints.
2. Break rectangles into three pieces: left endpiece, centerpiece, and right endpiece.
3. Find $Z^{\prime}$ intersections involving at least one centerpiece.
4. Recursively solve problem in each slab for endpieces.
is $O\left(\log _{\sqrt{m}} n\right)=O\left(\log _{m} n\right)$ levels of recursion.
स Performing Step 3 in $O\left(n+\frac{Z^{\prime}}{B}\right)$ I/Os

$$
\Longrightarrow \quad O\left(n \log _{m} n+z\right) \text { I/Os total. }
$$

## Key Idea



Consider intersections of red centerpieces and tops of blue rects.:
\& Use $\sqrt{m}$ slabs
设 $\Longrightarrow O(m)$ multislabs (continuous ranges of slabs)
( Store each red centerpiece in a multislab, implemented as a stack.
~ Stack effectively keeps the first $B$ rectangles of each multislab in internal memory.
~ Perform top down sweep:

- Maintaining active list for each multislab.


## Sketch of Sweep


in Intersections between red centerpieces and tops of blue rects.:

- At red rectangle: Insert into relevant multislab list (stack).
- At blue rectangle: Scan through all relevant multislab lists of red rectangles.
(i) Report intersection with "non-expired" red rectangles.
(ii) Remove "expired" red rectangles ("lazy" deletion). (Combine block with neighbor if $<B / 2$ living rectangles.)
f. Other cases handled similarly-in one sweep!


## Analysis of I／O Performance in each Pass

Intersections of red centerpieces and tops of blue rects．
i．Centerpieces of red rectangles are scanned in $O(n)$ I／Os．
设 For each top of a blue rectangle，we report intersections with non－expired red centerpieces in all relevant multislab lists．

埌 Since the first block of each multlislab list（stack）is in internal memory，if a multislab list has $k$ centerpieces，\＃I／Os $=\left\lfloor\frac{k}{B}\right\rfloor \leq \frac{k}{B}$ ．
纹 Each centerpiece is deleted in lazy manner at most once．
\＆Sum of $\frac{k}{B}$ over all reportings is thus at most $\frac{Z^{\prime}+N^{\prime}}{B}$ ， where $N^{\prime}$ is number of red centerpieces in the current pass．
ㅅ．Over the $\log _{m} n$ passes，summing $O\left(n+\frac{Z^{\prime}+N^{\prime}}{B}\right)$ gives a total of $O\left(n \log _{m} n+z\right)$ I／Os．

## Avoiding redundant reportings of intersections

设 Example：A given blue rectangle could intersect the centerpiece of a red rectangle，and the blue rectangle＇s endpiece could intersect the red rectangle＇s endpiece．

约 Two intersections would be reported at different levels of recursion．
级 How to fix this without sorting all intersections？ （Technically，sorting would require $O\left(z \log _{m} z\right) \mathrm{I} / \mathrm{Os}$ ，which is too much theoretically，and inefficient in practice．）

## Avoiding redundant reportings of intersections

iर Example: A given blue rectangle could intersect the centerpiece of a red rectangle, and the blue rectangle's endpiece could intersect the red rectangle's endpiece.
H Two intersections would be reported at different levels of recursion.
स How to fix this without sorting all intersections? (Technically, sorting would require $O\left(z \log _{m} z\right) \mathrm{I} / \mathrm{Os}$, which is too much theoretically, and inefficient in practice.)
约 Solution: Avoid redundant reportings of intersections by adopting a convention as to when to report an intersection.
: For example, each intersection could be reported only at the first available opportunity. At each potential reporting time, the two rectangles must be examined to determine if the intersection has already been reported.
is Charge each non-reporting to the actual intersection. Each intersection is non-reported at most $O(1)$ times.

## Higher Dimensions

~ Technique can be used recursively in dimension $d>2$ by decreasing number of slab boundaries to $m^{1 / 2(d-1)}$ in each of the $d-1$ dimensions orthogonal to sweep.
\& For $d=3$, consider a checkerboard of slabs, $m^{1 / 4} \times m^{1 / 4}$.
\& There are at most $m^{1 / 2} \times m^{1 / 2}=m$ multislabs.
Rectangles are partitioned in $x$ dimension and then a sweep is done in the $z$ dimension simultaneously for all $x$-slabs to solve the $y, z$-dimension subproblems. COMPLICATED!

设 I/O performance using technique:

- $d$-dim. batched range searching: $O\left(n \log _{m}^{d-1} n+t\right)$ I/Os, $O(n)$ space.
- $d$-dim. rectangle intersection:

$$
O\left(n \log _{m}^{d-1} n+t\right) \mathrm{I} / \mathrm{Os}, O(n) \text { space. }
$$

- Batched semidynamic planar point location:

$$
O\left((n+k) \log _{m}^{2}(n+k)\right) \text { I/Os, } O(n+k) \text { space. }
$$

## TPIE, http://www.cs.duke.edu/TPIE/

Many problems can be solved using small number of paradigms.
OS often provides inadequate support for I/O and internal memory management.

## TPIE, http://www.cs.duke.edu/TPIE/

(1) Many problems can be solved using small number of paradigms.
( $\because$ OS often provides inadequate support for I/O and internal memory management.
h TPIE originally designed by former student Darren Vengroff:

- Make implementation easy (and portable). I/O-efficient (and portable) programs.
- Framework oriented: Implements a number of high-level paradigms on streams (C++) -Scanning, merging, distribution, sorting, permuting, ...
- Access-Oriented: For index structures.


## TPIE's Distribution Access Method



## TIGER/Line Data

is TIGER/Line data from U.S. Census Bureau (standard benchmark data for spatial databases)

| State | Category | Size | Objects |
| :---: | :---: | ---: | ---: |
| Rhode Island (RI) | Roads | 4.3 MB | 68,278 |
|  | Hydrography | 0.4 MB | 7,013 |
| Connecticut (CT) | Roads | 12.0 MB | 188,643 |
|  | Hydrography | 1.8 MB | 28,776 |
| New Jersey (NJ) | Roads | 26.5 MB | 414,443 |
|  | Hydrography | 3.2 MB | 50,854 |
| New York (NY) | Roads | 55.7 MB | 870,413 |
|  | Hydrography | 10.0 MB | 156,568 |
| All | Roads | 98.5 MB | 1541,777 |
|  | Hydrography | 15.4 MB | 243,211 |

## Performance Comparison with PBSM [DP96]



Sun SparcStation 20 (Solaris 2.5), 32MB memory (TPIE 12MB)

## Performance Comparison with PBSM [DP96]




## Case II: Indexes Exist

## Previous Algorithm: ST [BKS93]

Carefully synchronized depth-first traversal.

Our Algorithm: PQ [APRSVV00]
R-tree $\rightarrow$ Priority Queue $\rightarrow$ Sweep

## Related Results

. External segment tree used in conjunction with batched filtering [GTVV93] and external fractional cascading to solve large number of problems with GIS applications [AVV95]:

- Red-blue line segment intersection in $O\left(n \log _{m} n+t\right)$ I/Os.

设 Persistent B-trees [GTVV93] to solve batched point location in $O\left(n \log _{m} n+t\right)$ I/Os.
~ Random incremental construction [CFMMR98] to get optimal $O\left((n+q) \log _{m} n+z\right)$ I/Os for general line segment intersection.

## Parallel Simulation Paradigm [CGGTVV95]

Let $A$ be an $N$-processor PRAM algorithm such that

- $A$ reduces a problem of size $N$ to one of size $\alpha N$ in constant time.
- Parallel running time of $A$ is $\Theta(\log N)$.
if For each PRAM statement, sort the $N$ operands so that they are contiguous.
4 Simulate $N$ operations via a linear pass through the data.
设 $\mathrm{I} / \mathrm{O}$ Complexity for $D=1$ :

$$
\begin{aligned}
T(N) & =O(\operatorname{sort}(N))+T(\alpha N) \\
& =O(\operatorname{sort}(N)) .
\end{aligned}
$$

设 Gives optimal EM algorithms for list ranking, Euler tours, expression tree evaluation, connected components of sparse graph.
3. Sometimes the sorting can be done in $O(N)$ I/Os because of constraints and assumptions [DDH97, SK97].
h Some problems like topological sorting, BFS, DFS are hard.

## Conclusions and Open Problems

$\dot{\mu}$ 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
L. Lots of open problems in the design and analysis of external memory algorithms and data structures. Stay tuned!
- TPIE, see http://www.cs.duke.edu/TPIE/
- Handling many disks, large merge orders, many partition elements, large fanouts. (Don't use square root trick.)
- GIS applications (e.g. practical red-blue line segment intersection, nearest neighbor, spatial join, terrain processing).
- Image processing (indexing images, analyzing images).
- Fundamental graph problems (e.g. topological sorting, BFS, DFS, connectivity).


## Conclusions and Open Problems

- Online dynamic data structures (e.g. dynamic point location, range search in higher dimensions, clustering, similarity search).
- String processing, molecular databases.
- Typical-case behavior of popular data structures (e.g., R-trees).
- . .

