# Cascading Divide-and-Conquer: A Technique for Designing Parallel Algorithms 

Mikhail J. Atallah<br>Purdue University, mja@cs.purdue.edu

Richard Cole
Michael T. Goodrich
Report Number:
87-665

[^0]This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

# CASCADING DIVIDE-AND CONQUER: 

A TECHNIQUE FOR DESIGNING
PARALLEL ALGORITHMS

Mikhail J. Atallah
Richard Cole Michael T. Goodrich

CSD-TR-665
March 1987
(Revised June 1989)

# CASCADING DIVIDE-AND-CONQUER: A TECHNIQUE FOR designing parallel algorithms* 

MIKHAIL J. ATALLAHt, RICHARD COLE $\ddagger$, AND MICHAEL T. GOODRICH§


#### Abstract

Techniques for parallel divide-and-conquer are presented, resulting in improved parallel algorithms for a number of problems. The problems for which improved algorithms are given include segment intersection detection, trapezoidal decomposition, and planar point locution. Efficient parallel algorithms are algo given for fractional cascading, three-dimensional maxima, two-set dominance counting, and visibility from a point. All of the algorithms presented run in $O(\log n)$ time with either a linear or a sublinear number of processors in the CREW PRAM model.


Key words. parallel algorithms, parallel datal structures, divide-and-conquer, computational geometry, fractional cascading, visibility, planar point location, rrapezoidat decomposition, dominance, intersection delection

AMS(MOS) subject classifications. $68 \mathrm{E} 05,68 \mathrm{C} 05,68 \mathrm{C} 25$

1. Introduction. This paper presents a number of general techniques for parallel divide-and-conquer. These techniques are based on nontrivial generalizations of Cole's recent parallel merge sor result [13] and enable us to achieve improved complexity bounds for a large number of problems. In particular, our techniques can be applied to any problem solvable by a divide-and-conquer method such that the subproblem merging step can be implemented using a restricred, but powerful, set of operations, which include (i) merging sorted lists, (ii) computing the values of labeling functions on elements stored in sorted lists, and (iii) changing the identity of elements in a sored fist monotonically. The elements stored in such sorted lists need not belong to a total order, so long as the computation can be specified so that we will never try to compare two incomparable elements. We demonstrate the power of these techniques by using them to design efficient paraltel algorithms for solving a number of fundamental problems from computational geometry.

The general framework is one in which we want to design efficient parallel algorithms for the CREW PRAM or EREW PRAM models. Recall that the CREW PRAM model is the synchronous shared memory model in which processors may simultaneously read from any memory location but simultaneous writes are not allowed. The EREW PRAM model does not allow for any simultaneous access to a memory cell. Our goal is to find algorithms that run as fast as possible and are efficient in the following sense: if $p(n)$ is the processor complexity, $f(n)$ the parallel time complexity, and seq( $n$ ) the time complexity of the best-known sequential algorithm for the problem

[^1]under consideration, then $t(n) * p(n)=O(\operatorname{seq}(n))$. If the product $t(n) * p(n)$ achieves the sequential lower bound for the problem, then we say the algorithm is optimal. When specifying the processor complexity, we omit the "big oh," e.g., we say " $n$ processors" rather than " $O(n)$ processors"; this is justified because we can always save a constant factor in the number of processors at a cost of the same constant factor in the running time. In all of the problems listed below, we achieve $t(n)=O(\log n)$ and, simultaneously (except for planar point location), an optimal $t(n) * p(n)$.

Previous work on parallel divide-and-conquer has produced relatively few algorithms that are optimal in the above sense. Exceptions to this include some of the previous algorithms for the convex hull problem [1], [4], [6], [18], [27] and the problem of circumscribing a convex polygon with a minimum-area triangle [1]. Unfortunately, each of these approaches was very problem-specific. Thus, there is a need for techniques of wider scope.

This is in fact the motivation for our work, for we give a number of general techniques for efficiently solving problems in parallel by divide-and-conquer. We model the divide-and-conquer paradigm as a binary tree whose nodes contain sorted lists of some kind. The computation involves computing on this tree in a recursively defined bottom-up fashion using lists of items and labeling functions defined for each node in the tree. In Cole's scheme [13], the list at a node was defined to be the sorted merge of the two lists stored at its children. In our scheme, however, the lists at a node of the tree can depend on the lists of its children in more complex ways. For example, in our solution to the segment intersection detection problem, the lists at a node depend on computing, in addition to merges, set difference operations that are not directly solvable by the "cascading" method used by Cole [13]. Such operations arise here because the lists at a node contain segments ordered by their intersections with a vertical line (the so-called "above" relationship), which is obviously not a total order. One may be tempted to try to solve this problem by delaving the performance of these set difference operations until the end of the computation. Unfortunately, this is not feasible for many reasons, not the least of which is that this approach could lead to a situation in which a processor tries to compare two incomparable items. Nor does it seem possible to explicitly perform the set difference operations on-line without sacrificing the time-efficiency of the cascading method. Our solution avoids both of these problems by using an on-line "identity-changing" technique.

Another significant contribution of this paper is an optimal parallel construction of the "fractional cascading" data structure of Chazelle and Guibas [12]. This too is based on a generalization of Cole's method [13] in the sense that instead of having the computation proceeding up and down a tree, it now moves around a directed graph (possibly with cycles). Our solution to fractional cascading is quite different from the sequential method of Chazelle and Guibas (their method relies on an amortization scheme to achieve a linear running timel.

The following is a list of the problems for which our techniques result in improved complexity bounds. Unless otherwise specified, each performance bound is expressed as a pair $(t(n), p(n))$, where $t(n)$ and $p(n)$ are the time and processor complexities, respectively, in the CREW PRAM model.

Fractional cascading. Given a directed graph $G=(V, E)$, such that every node $v$ contains a sorted list $C(v)$, construct a data structure that, given a walk ( $v_{1}, v_{2}, \cdots, v_{m}$ ) in $G$ and an arbitrary element $x$, enables a single processor to locate $x$ quickly in each $C\left(v_{r}\right)$, where $n=|V|+|E|+\sum_{v \in V}|C(v)|$. In [12] Chazelle and Guibas gave an elegant $O(n)$ time, $O(n)$ space, sequential construction, where $n=\Sigma_{\text {cc }}|C(v)|$. We give a $(\log n, n / \log n)$ construction.

Trapezoidal decomposition. Given a set $S$ of $n$ line segments in the plane, determine for each segment endpoint $p$ the first segment "stabbed" by the vertical ray emanating upward (and downward) from $p$. A ( $\log ^{2} n, n$ ) solution to this problem was given by Aggarwal et al. in [1], later improved to ( $\log n \log \log n, n)$ by Atallah and Goodrich in [5]. We improve this to $(\log n, n)$.

Planar poins location. Given a subdivision of the plane into (possibly unbounded) polygons, construct, in parallel, a data structure that, once built, enables one processor to determine for any query point $p$ the potygonal face containing $p$. Let $Q(n)$ denote the time for performing such a query, where $n$ is the number of edge segments in the subdivision. A $\left(\log ^{2} n, n\right), Q(n)=O\left(\log ^{2} n\right)$ solution was given by Aggarwal et al. in [1], later improved to $(\log n \log \log n, n), Q(n)=O(\log n)$ by Atallah and Goodrich in [5]. In [14] Dadoun and Kirkpatrick further improved this to $\left(\log n \log ^{*} n, n\right)$, $Q(n)=O(\log n)$. We give a $(\log n, n), Q(n)=O(\log n)$ solution.

Segment intersection detection. Given a set $S$ of $n$ line segments it the plane, determine if any two segments in $S$ intersect. A $\left(\log ^{2} n, n\right)$ solution was given in [1], later improved to $(\log n \log \log n, n)$ in [5]. We improve this to $(\log n, n)$.

Three-dimensional maxima. Given a set $S$ of $n$ points in three-dimensional space, determine which points are maxima. A maximum in $S$ is any point $p$ such that no other point of $S$ has $x, y$, and $z$ coordinates that simultaneously exceed the corresponding coordinates of $p$. A. $(\log n \log \log n, n)$ solution was given in [5]. We improve this to ( $\log n, n$ ).

Tho-ses dominance counting. Given a set $A=\left\{q_{1}, q_{2}, \cdots, q_{l}\right\}$ and a set $B=$ $\left\{r_{1}, r_{2}, \cdots, r_{m}\right\}$ of points in the plane, determine for each point $r_{i}$ in $B$ the number of points in $A$ whose $x$ and $y$ coordinates are both less than the corresponding coordinates of $r_{i}$. The problem size is $n=l+m$. A $(\log n \log \log n, n)$ solution was given in [5]. We improve this to $(\log n, n)$.

Visibility from a point. Given $n$ line segments such that no two intersect (except possibly at endpoints) and a point $p$, determine that pan of the plane visible from $p$. if all the segments are opaque. A $(\log n \log \log n, n)$ solution was given in [5]. We improve this to $(\log n, n)$.

We recently learned that Reif and Sen [24] solved planar point location, trapezoidal decomposition, segment intersection and visibility in randomized $O(\log n)$ time using $O(n)$ processors in the CREW PRAM model. All of our algorithms are deterministic.

This paper is organized as follows. In $\S 2$ we present a generalized version of the cascading merge procedure and in $\S 3$ we give our method for doing fractional cascading in parallel. In § 4 we show how to apply the fractional cascading technique to a data structure we call the plane sweep tree, showing how to solve the trapezoidal decomposition and point location problems. In § 5 we show how to extend the cascading merge technique to allow for cascading in the "above" partial order of line segments, giving solutions to the problems of building the plane sweep tree and solving the intersection detection problem. In $\S 6$ we use the cascading divide-and-conquer technique to compute labeling functions and show how to use this approach to solve threedimensional maxima, two-set dominance counting, and visibility from a point. Finally, in §7, we briefly describe how most of our algorithms can be implemented in the EREW PRAM model with the same time and processor bounds as our CREW PRAM algorithms, and we conclude in $\S 8$.
2. A generalized cascading merge procedure. In this section we present a technique for a generalized version of the merge sorting problem. Suppose we are given a binary tree $T$ (not necessarily complete) with items, taken from some total order, placed at
the leaves of $T$, so that each leaf contains at most one item. For simplicity, we assume that the items are distinct. We wish to compute for each internal node $v \in T$ the soned list $U(v)$ that consists of all the items stored in descendant nodes of $v$. (See Fig. 1.) In this section we show how to constrict $U(v)$ for every node in the tree in $O($ height ( $T$ ) time using $|T|$ processors, where $|T|$ denotes the number of nodes in $T$. This is a generalization of the problem studied by Cole [13], because in his version the tree $T$ is complete. Without loss of generality, we assume that every internal node $v$ of $T$ has two children. For if $v$ has only one child then we can add a child to $v$ (a leaf node) that does not store any items from the total order. Such an augmentation will at most double the size of $T$ and does not change its height.


Fig. 1. An cxample of the generalized merge problem.
Let $a, b$, and $c$ be three items, with $a \leqq b$. We say $c$ is between $a$ and $b$ if $a<c \leqq b$. Let two sorted (nondecreasing) lists $A=\left(a_{1}, a_{2}, \cdots, a_{n}\right)$ and $B=\left(b_{1}, b_{2}, \cdots, b_{m}\right)$ be given. Given an element $a$, we define the predecessor of $a$ in $B$ to be the greatest element in $B$ that is less than or equal to $a$. If $a<b_{1}$, then we say that the predecessor of $a$ is $-\infty$. We define the rank of $a$ in $B$ to be the rank of the predecessor of $a$ in $B$ ( $-\infty$ has rank zero). We say that $A$ is ranked in $B$ if for every element in $A$ we know its rank in $B$. We say that $A$ and $B$ are cross-ranked if $A$ is ranked in $B$ and $B$ is ranked in $A$. We define two operations on sorted lists. We define $A \cup B$ to be the sorted merged list of all elements in $A$ or $B$. If $B$ is a subset of $A$, then we define $A-B$ to be the sorted list of the elemented in $A$ that are not in $B$.

Let $T$ be a binary tree. For any node $v$ in $T$ we let parent $(v)$, sibling $(v)$, Ichild ( $v$ ), rchild ( $v$ ), and depth( $v$ ) denote the parent of $v$, the sibling of $v$, the left child of $v$, the right child of $v$, and the depth of $v$ (the root is at depth zero), respectively. We also let $\operatorname{rool}(T)$ and heighr $(T)$ denote the root node of $T$ and the height of $T$, respectively. The altitude, denoted $\operatorname{aht}(v)$, is defined $\operatorname{alt}(v)=h e i g h t(T)-\operatorname{depth}(v)$. $\operatorname{Desc}(v)$ denotes the set of descendant nodes of $v$ (including $v$ itself).

Let a sorted list $L$ and a sorted list $J$ be given. Using the terminology of Cole [13], we say that $L$ is a c-cover of $J$ if between each two adjacent items in ( $-\infty, L, \infty$ ) there are at most $c$ items from $J$ (where $(-\infty, L, \infty)$ denotes the list consisting of $-\infty$, followed by the elements of $L$, followed by $\infty$ ). We let $S A M P_{r}(L)$ denote the sorted
list consisting of every $c$ th element of $L$, and call this list the $c$-sample of $L$. That is, $S A M P_{r}(L)$ consists of the cth element of $L$ followed by the ( $2 c$ )th element of $L$, and so on.

The algorithm for constructing $U(v)$ for each $v \in T$ proceeds in stages. Intuitively, in each stage we will be performing a portion of the merge of $U($ Ichild $(v))$ and $U(r c h i l d(v))$ to give the list $U(v)$. After performing a portion of this merge we will gain some insight into how to perform the merge at $v$ 's parent. Consequently, we will pass some of the elements formed in the merge at $v$ to $v$ 's parent, so we can begin performing the merge at $v$ 's parent.

Specifically, we denote the list stored at a node $v$ in $T$ at stage $s$ by $U_{s}(v)$. Initially, $U_{0}(v)$ is empty for every node except the lear nodes of $T$, in which case $U_{0}(v)$ contains the item stored at the leaf node $v$ (if there is such an item). We say that an internal node $v$ is acrive at stage $s$ if $[s / 3] \leqq a h(v) \leqq s$, and we say $v$ is full at stage $s$ if $\operatorname{alf}(v)=\lfloor s / 3\rfloor$. As will become apparent below, if a node $v$ is full, then $U_{s}(v)=U(v)$. For each active node $v \in T$ we define the list $U_{s+1}^{\prime}(v)$ as follows:

$$
U_{s+1}^{\prime}(v)= \begin{cases}S A M P_{a}\left(U_{s}(v)\right) & \text { if } \text { alt }(v) \geqq s / 3, \\ S A M P_{2}\left(U_{s}(v)\right) & \text { if } \operatorname{alh}(v)=(s-1) / 3 \\ S A M P_{1}\left(U_{s}(v)\right) & \text { if } \operatorname{alt}(v)=(s-2) / 3\end{cases}
$$

At stage $s+1$ we perform the following computation at each internal node $v$ that is currently active.

Per-stage computation $(v, s+1)$. Form the two lists $U_{s+1}^{\prime}($ child $(v))$ and $U_{i+1}^{\prime}($ rchild $(v))$, and compute the new list

$$
U_{s+1}(v):=U_{r+1}^{\prime}(\text { Ichild }(v)) \cup U_{s+1}^{\prime}(\text { rchild }(v))
$$

This formalizes the notion that we pass information from the merges performed at the children of $v$ in stage $s$ to the merge being performed at $v$ in stage $s+1$. Note that until $v$ becomes full, $U_{s+1}^{\prime}(v)$ will be the list consisting of every fourth element of $U_{s}(v)$. This continues to be true about $U_{s+1}^{\prime}(v)$ up to the point that $v$ becomes full. If $s_{v}$ is the stage at which $v$ becomes full (and $U_{s}(v)=U(v)$ ), then at stage $s_{v}+1$, $U_{s+1}^{\prime}(v)$ is the two-sample of $U_{s}(v)$, and, at stage $s_{v}+2, U_{s+1}^{\prime}(v)=U_{s}(v)(=U(v))$. Thus, at stage $s_{v}+3$, parent $(v)$ is full. Therefore, after $3 *$ height $(T)$ stages every node has become full and the algorithm terminates. We have yet to show how to perform each stage in $O(1)$ time using $n$ processors.

We begin by showing that the number of items in $U_{\mathrm{s}+1}(v)$ can be only a little more than twice the number of items in $U_{s}(v)$, a property that is essential to the construction.

Lemma 2.1. For any srage $s \geqq 0$ and any node $v \in T,\left|U_{s+1}(v)\right| \leqq 2\left|U_{s}(v)\right|+4$.
Proof. The proof is by induction on $s$.
Basis ( $s=0$ ). The claim is clearly true for $s=0$.
Induction step $(s>0)$. Assume the claim is true for stage $s-1$. If $v$ is full (i.e., $a l(v)=\lfloor s / 3\rfloor)$, then the claim is obviously true, since $U_{\mathrm{s}+1}(v)=U_{s}(v)=U(v)$. Consider the case where either the children of $v$ were not full at stage $s$ or had just become full at stage $s$. We know that $U_{\mathrm{r}+1}(v)=U_{\mathrm{c}+1}^{\prime}(x) \cup U_{\mathrm{s}+1}^{\prime}(y)$, where $x=1$ child $(v)$ and
$y=$ rchild $(u)$. In addition, we have the following:

$$
\begin{aligned}
\left|U_{x+1}(v)\right| & =\left\lfloor\frac{\left|U_{1}(x)\right|}{4}\right\rfloor+\left\lfloor\frac{\left|U_{r}(y)\right|}{4}\right\rfloor \quad \text { (from definitions) } \\
& \leqq\left\{\frac{2\left|U_{t-1}(x)\right|+4}{4}\right\rfloor+\left\lfloor\frac{2\left|U_{s-1}(y)\right|+4}{4}\right\rfloor \quad \text { (by induction hypothesis) } \\
& \leqq 2\left(\left\lfloor\frac{\left|U_{s-1}(x)\right|}{4}\right\rfloor+\left\lfloor\frac{\left|U_{r-1}(y)\right|}{4}\right\rfloor\right)+4 \\
& =2\left|U_{s}(v)\right|+4 .
\end{aligned}
$$

The case when the children of $v$ are full at stage $s-1$ is similar (except that one divides by 2 or 1 instead of 4). Actually, it is simpler, since in this case the children of $v$ were full in stage $s-1$; hence, the step using the induction hypothesis can be replaced by a simple algebraic substitution step.

In the next lemma we show that the way in which the $U_{s}(v)$ s grow is "well behaved."

Lemma 2.2. Let $[a, b]$ be an interval with $a, b \in\left(-\infty, U_{s}^{\prime}(v), \infty\right)$. If $[a, b]$ intersects $k+1$ items in $\left(-\infty, U_{s}^{\prime}(v), \infty\right)$, then it intersects at most $8 k+8$ items in $U_{s}(v)$ for all $k \geqq 1$ and $s \geqq 1$.

Proof. The proof is by induction on $s$. The claim is initially true (for $s=1$ ). Actually, for any stage $s$, if $U_{s}^{\prime}(v)$ is empty, then $U_{s-1}(v)$ contains at most three items, hence, $U_{s}(v)$ contains at most ten elements, by the previous lemma. Also, if $U_{s}^{\prime}(v)$ contains one item, then $U_{s-1}(v)$ contains at most seven items, hence, $U_{s}(v)$ contains at most 18 items, by the previous lemma. At most 1.5 of these items can be between any two adjacent items in $\left(-\infty, U_{r}^{\prime}(v), \infty\right)$, since the item in $U_{s}^{\prime}(v)$ was the fourth item in $U_{s-1}(v)$ by definition.

Inductive step (assume true for stage $s$ ). Let $[a, b]$ be an interval with $a, b$ both in the list $\left(-\infty, U_{s+1}^{\prime}(v), \infty\right)$, and suppose $[a, b]$ intersects $k+1$ items in $\left(-\infty, U_{s+1}^{\prime}(v), \infty\right)$. The lemma is immediately true if $v$ was full stage $s$, since the smallest sample we take is a four-sample. So, next, suppose that either the children of $v$ are not full or have just become full in stage $s$. Let $g$ be the number of items in $\left(-\infty, U_{s}(v), \infty\right)$ intersected by $[a, b]$. Recall that $U_{s}(v)=$ $U_{s}^{\prime}($ lchild $(v)) \cup U_{r}^{\prime}(r c h i l d(v))$. Let $\left[a_{1}, b_{1}\right]$ (respectively, $\left.\left[a_{2}, b_{2}\right]\right)$ be the smallest interval containing $[a, b]$ such that $a_{1}, b_{1} \in\left(-\infty, U_{s}^{\prime}(\right.$ Ichild $\left.(v)), \infty\right)$ (respectively, $a_{2}, b_{2} \in(-\infty), U_{s}^{\prime}($ rchild $\left.\left.(v)), \infty\right)\right)$. Suppose the interval $\left[a_{1}, b_{1}\right]$ intersects $h+1$ items in the list $\left(-\infty, U_{s}^{\prime}(\operatorname{lchild}(v)), \infty\right)$ and $\left[a_{2}, b_{2}\right]$ intersects $j+1$ items in $\left(-\infty, U_{s}^{\prime}(\right.$ rchild $\left.(\nu)), \infty\right)$. Note that $h+j=g$. By the induction hypothesis, $\left[a_{1}, b_{1}\right]$ intersects at most $8 h+8$ items in $U_{s}(\operatorname{lchild}(v))$, and hence at most $(8 h+8) / 4=2 h+2$ items in $U_{i+1}^{\prime}(\operatorname{lchild}(v))$. Likewise, $\left[a_{2}, b_{2}\right]$ intersects at most $2 j+2$ items in $U_{i+1}^{\prime}\left(\right.$ rchild $(v)$ ). The definition of $U_{s+1}^{\prime}(v)$ implies that $g \leqq 4 k+1$. Therefore, since $U_{s+1}(v)=U_{s+1}^{\prime}($ lchild $(v)) \cup U_{s+1}^{\prime}(r c h i l d(v)), \quad[a, b]$ intersects at most $(2 h+2)+$ $(2 j+2)$ items in $U_{1+1}(v)$, where $(2 h+2)+(2 j+2) \leqq(2 h+2)+(2(4 k-h+1)+2) \leqq$ $8 k+8$.

The proof for the case when the children of $v$ were full in stage $s-1$ is similar. Actually, it is simpler, since the induction steps can be replaced by algebraic substitution steps in this case.

Corollary 2.3. The list $\left(-\infty, U_{s}^{\prime}(v), \infty\right)$ is a four-cover for $U_{s+1}^{\prime}(v)$, for all $s \geqq 0$.

This corollary is used in showing that we can perform each stage of the merge procedure in $O(1)$ time. In addition to this corollary, we also need to maintain the following rank information at the start of each stage $s$ :
(1) For each item in $U_{i}^{\prime}(v)$ : its rank in $U_{5}^{\prime}($ sibling $(v))$.
(2) For each item in $U_{5}^{\prime}(v)$ : its rank in $U_{s}(v)$ (and hence, implicitly, its rank in $\left.U_{r+1}^{\prime}(v)\right)$.

The lemma that follows shows that the above information is sufficient to allow us to merge $U_{s+1}^{\prime}($ Ichild $(v))$ and $U_{s+1}^{\prime}($ rchild $(v))$ into the list $U_{t+1}(v)$ in $O(1)$ time using $\left|U_{\mathrm{s}+1}(v)\right|$ processors.

Lemma 2.4 (The Merge Lemma) [13]. Suppose we are given sorted lisis $A_{5}, A_{5+1}^{\prime}$, $B_{5}^{\prime}, B_{s+1}^{\prime}, C_{s}^{\prime}$, and $C_{s+1}^{\prime}$, where the following (input) conditions are true:
(1) $A_{1}=B_{s}^{\prime} \cup C_{s}^{\prime}$;
(2) $A_{r+1}^{\prime}$ is a subset of $A_{3}$;
(3) $B_{s}^{\prime}$ is a $c_{1}$-cover for $B_{++1}^{\prime}$;
(4) $C_{5}^{\prime}$ is a $c_{2}$-cover for $C_{s+1}^{\prime}$;
(5) $B_{s}^{s}$ is ranked in $B_{s+1}^{\prime}$;
(6) $C_{5}^{\prime}$ is ranked in $C_{s+1}^{\prime}$;
(7) $B_{s}^{\prime}$ and $C_{s}^{\prime}$ are cross-ranked.

Then in $O(1)$ time using $\left|B_{++1}^{\prime}\right|+\left|C_{s+1}^{\prime}\right|$ processors in the CREW PRAM model, we can compute the following (outpur compurations):
(1) the sorted list $A_{s+1}=B_{s+1}^{\prime} \cup C^{\prime}+1$;
(2) the ranking of $A_{s+1}^{\prime}$ in $A_{s+1}$;
(3) the cross-ranking of $B_{x+1}^{\prime}$ and $C_{s+1}^{\prime}$.

We apply this lemma by setting $A_{s}=U_{s}(v), A_{s+1}^{\prime}=U_{s+1}^{\prime}(v), A_{s+1}=U_{s+1}(v), B_{s}^{\prime}=$ $U_{r}^{\prime}(x), B_{s+1}^{\prime}=U_{s+1}^{\prime}(x), C_{s}^{\prime}=U_{s}^{\prime}(y)$, and $C_{s+1}^{\prime}=U_{s+1}^{\prime}(y)$, where $x=\operatorname{lchild}(v)$ and $y=$ rchild $(v)$. Note that assigning the lists of Lemma 2.4 in this way satisfies input conditions (1)-(4) from the definitions. The ranking information we maintain from stage to stage satisfies input conditions (5)-(7). Thus, in each stage $s$, we can construct the list $U_{s+1}(v)$ in $O(1)$ time using $\left|U_{s+1}(v)\right|$ processors. Also. the new ranking information (of output computations (2) and (3)) gives us the input conditions (5)-(7) for the next stage. By Corollary 2.3 we have that the constants $c_{1}$ and $c_{2}$ (of input conditions (3) and (4)) are both equal to four. Note that in stage $s$ it is only necessary to store the lists for $s-1$; we can discard any lists for stages previous to that.

The method for performing all these merges with a total of $|T|$ processors is basically to stan out with $O(1)$ virtual processors assigned to each leaf node, and each time we pass $k$ elements from a node $v$ to the parent of $v$ (to perform the merge at the parent), we also pass $O(k)$ virtual processors to perform the merge. When $v$ 's parent becomes full, then we no longer "store" any processors at $v$. (See [17] for details.) There can be at most $O(n)$ elements present in active nodes of $T$ for any stage $s$ (where $n$ is the number of leaves of $T$ ), since there are $n$ elements present on the full level, at most $n / 2$ on the level above that, $n / 8$ on the level above that, and so on. Thus, we can perform the entire generalized cascading procedure using $O(n)$ virtual processors, or $n$ actual processors (by a simple simulation argument). This also implies that we need only $O(n)$ storage for this computation, in addition to that used for the output, since once a node $v$ becomes full we can consider the space used for $U(v)$ to be part of the output. Equivalently, if we are using the generalized merging procedure in an algorithm that does not need a $U(v)$ list once $v$ 's parent becomes full, then we can implement that algorithm in $O(n)$ space by deallocating the space for a $U(v)$ list once it is no longer needed (chis is in fact what we will be doing in $\S 6$ ).

It will often be more convenient to relax the condition that there be at most one item stored at each leaf. So, suppose there is an unsorted set $A(v)$ (which may be empty) stored at each leaf. In this case we can construct a tree $T^{\prime}$ from $T$ by replacing each leaf $v$ of $T$ with a complete binary tree with $|A(v)|$ leaves, and associating each item in $A(v)$ with one of these leaves. $T^{\prime}$ would now satisfy the conditions of the method outlined above. We incorporate this observation in the following theorem, which summarizes the discussion of this section.

Theorem 2.5. Suppose we are given a binary tree $T$ such that there is an unsorted set $A(v)$ (which may be empty) stored at each leaf. Then we can compute, for each node $v \in T$, the list $U(v)$, which is the union of all items stored at descendents of $v$, sorted in an array. This computation can be implemented in $O\left(\right.$ height $\left.(T)+\log \left(\max _{v}|A(v)|\right)\right)$ time using a toral of $n+N$ processors in the CREW PRAM computational model, where $n$ is the number of leaves of $T$ and $N$ is the total number of items stored in $T$.

Proof. The complexity bounds follow from the fact that the tree $T^{\prime}$ described above would have height at most $O\left(\right.$ height $\left.(T)+\log \left(\max _{v}|A(v)|\right)\right)$ and $|T|$ is $O(|T|+N)$.

The above method comprises one of the main building blocks of the algorithms presented in this paper. We present another important building block in the following section.
3. Fractional cascading in parallel. Given a directed graph $G=(V, E)$, such that every node $v$ contains a sorted list $C(v)$, the fractional cascading problem is to construct an $O(n)$ space data structure that, given a walk ( $v_{1}, v_{2}, \cdots, v_{m}$ ) in $G$ and an arbitrary element $x$, enables a single processor to locate $x$ quickly in each $C\left(v_{i}\right)$, where $n=$ $|V|+|E|+\sum_{\mathrm{rc} \cdot}|C(v)|$. Fractional cascading problems arise naturally from a number of computational geometry problems. As a simple example of a fractional cascading problem, suppose we have five different English dictionaries and would like to build a data structure that would allow us to look up a word $w$ in all the dictionaries. Chazelle and Guibas [12] give an elegant $O(n)$ time sequential method for constructing a fractional cascading data structure from any graph $G$, as described above, achieving a search time of $O(\log n+m \log d(G))$, where $d(G)$ is the maximum degree of any node in $G$. However, their approach does not appear to be "parallelizable."

In this section we show how to construct a data structure achieving the same performance as that of Chazelle and Guibas in $O(\log n)$ time using $\lceil n / \log n\rceil$ processors. Our method begins with a preprocessing step similar to one used by Chazelle and Guibas where we "expand" each node of $G$ into two binary trees-one for its in-edges and one for its out-edges-so that each node in our graph has in-degree and out-degree at most 2 . We then perform a cascading merge procedure in stages on this graph. Each catalogue $C(v)$ is "fed into" the node $v$ in samples that double in size with each stage and these lists are in turn sampled and merged along the edges of $G$. Lists continue to be sampled and "pushed" across the edges of $G$ (even in cycles) for a logarithmic number of stages, at which time we stop the computation and add some links between elements in adjacent lists. We conclude this section by showing that this gives us a fractional cascading data structure, and that the computation can be implemented in $O(\log n)$ time and $O(n)$ space using $\lceil n / \log n\rceil$ processors.

We show below how to perform the computations in $O(\log n)$ time and $O(n)$ space using $n$ processors. We will show later how to get the number of processors down to $\lceil n / \log n\rceil$ by a careful application of Brent's theorem [11].

Define $\operatorname{In}(v, G)$ (respectively, $\operatorname{Out}(v, G)$ ) to be the set of all nodes $w$ in $V$ such that $(u, v) \in E$ (respectively, $(v, w) \in E)$. The degree of a vertex $v$, denoted $d(v)$, is
defined as $d(v)=\max \{|\operatorname{In}(v, G)|,|O u f(v, G)|\}$. The degree of $G$, denoted $d(G)$, is defined as $d(G)=\max _{v \in v}\{d(v)\}$. A sequence $\left(v_{1}, v_{2}, \cdots, v_{m}\right)$ of verices is a walk if $\left(v_{1}, v_{i+1}\right) \in E$ for all $i \in\{1,2, \cdots, m-1\}$.

As mentioned above, we begin the construction by preprocessing the directed graph $G$ to convert it into a directed graph $\hat{G}=(\hat{V}, \hat{E})$ such that $d(\hat{G}) \leqq 2$ and such that an edge $(v, w)$ in $G$ corresponds to a path in $\hat{G}$ of length at most $O(\log d(G))$. Specifically, for each node $v \in V$ we construct two complete binary trees $T_{v}^{\text {in }}$ and $T_{v}^{\text {out }}$. Each leal in $T_{n}^{\text {in }}$ (respectively, $T_{v}^{\text {out }}$ ) corresponds to an edge coming into $u$ (respectively, going out of $v$ ). So there are $|\operatorname{In}(v, G)|$ leaves in $T_{v}^{\text {in }}$ and $|O u t(v, G)|$ leaves in $T_{v}^{\text {out }}$. (See Fig. 2.) We call $T_{v}^{\text {in }}$ the fan-in tree for $v$ and $T_{v}^{\text {out }}$ the fan-out tree for $v$. An edge $e=(v, w)$ in $G$ corresponds to a node $e$ in $\hat{G}$ such that $e$ is a lear of the fan-out tree for $v$ and $e$ is also a leaf of the fan-in tree for $w$. The edges in $T_{v}^{\text {in }}$ are all directed up towards the root of $T_{v}^{i n}$, and the edges in $T_{c}^{\text {out }}$ are all directed down towards the leaves of $T_{v}^{\text {out }}$. For each $v \in V$ we create a new node $v^{\prime}$ and add a directed edge from $v^{\prime}$ to $v$, a directed edge from the root of $T_{v}^{\text {in }}$ to $v^{\prime}$, and an edge from $v^{\prime}$ to the root of $T_{v}^{\text {out }}$. We call $v^{\prime}$ the gateway for $v$. (See Fig. 2). Note that $d(\hat{G})=2$. We assume that for each node $v$ we have access to the nodes in $\operatorname{In}(v, \hat{G})$ as well as those in $\operatorname{Out}(v, \hat{G})$. We structure fan-out trees so that a processor needing to go from $v$ to $w$ in $\hat{G}$, with $(0, w) \in E$, can correctly determine the path down $T_{\mathrm{V}}^{\text {out }}$ to the lear corresponding to ( $u, w$ ). More specifically, the leaves of each fan-out tree are ordered so that they are listed from left to right by increasing destination name, i.e., if the leaf in $T_{v}^{\text {out }}$ for $e=(v, u)$ is to the left of the leaf for $f=(v, w)$, then $u<w$. (The leaves of $T_{v}^{u n}$ need not be sorted, since all edges are directed towards the root of that tree.) If we are not given the $\operatorname{Out}(\nu, G)$ sets in sorted order, then we must perform a sort as a part of the $T_{\mathrm{t}}^{\text {out }}$ construction, which can be done in $O(\log d(G))$ time using $n$ processors using Cole's merge sorting algorithm [13]. We also store in each internal node $z$ of $T_{v}^{\text {out }}$ the leaf node $u$ that has the smallest name of all the descendants of $z$.


Ftg. 2. Converting $G$ into a bounded degrec graph $\dot{\vec{G}}$. A node $v$ in $G$ (a) corresponds imo a node vadjacent 10 its gateway $v^{\circ}$. which is connected to the fan in tree and she fan-out tree for $v$ (b). An edge $c$ in $G$ (c) is converted into a node in $\dot{G}$ which corresponds to a leaf node of two trees (d).

The above preprocessing step is similar to a preprocessing step used in the sequential fractional cascading atgorithm of Chazelle and Guibas [12]. This is where the resemblance to the sequential algorithm ends, however.

The goal for the rest of the computation is to construct a special sorted list $B(v)$, which we call the bridge list, for every node $v \in \hat{V}$. We shall define these bridge lists so that $B(v)=C(v)$ if $v$ is in $V$; if $v$ is in $\hat{V}$ but not in $V$, then for every $(v, w) \in \hat{E}$, if a single processor knows the position of a search item $x$ in $B(v)$, it can find the position of $x$ in $B(u)$ in $O(1)$ time.

The construction of the $B(v)$ 's proceeds in stages. Let $B_{r}(v)$ denote the bridge list stored at node $v \in \hat{V}$ at the end of stage $s$. Initially, $B_{0}(v)=\varnothing$ for all $v$ in $\hat{V}$. Intuitively, the per-stage computation is designed so that if $v$ came from the original graph $G$ (i.e.. $v \in V$ ), then $v$ will be "reeding" $B_{s}(v)$ with samples of the catalogue $C(v)$ that double in size with each stage. These samples are then cascaded back into the gateway $v^{\prime}$ for $v$ and from there back through' the fan-in tree for $v$. We will also be merging any samples "passed back" from the fan-out tree for $v$ with $B_{r}(v$ '), and cascading these values back through the fan-in tree for $v$ as well. We iterate the per-stage computation for $\lceil\log N\rceil$ stages, where $N$ is the size of the largest catalogue in $G$. We will show that after we have completed the last stage, and updated some ranking pointers, $\hat{G}$ will be a fractional cascading data structure for $G$. The details follow.

Recall that $B_{0}(v)=\emptyset$ for all $v \in \hat{V}$. For stage $s \geqq 0$, we define $B_{s+1}^{\prime}(v)$ and $B_{r+1}(v)$ as follows:

$$
\begin{aligned}
& B_{s+1}^{\prime}(v)= \begin{cases}\operatorname{SAMP}_{s}\left(B_{s}(v)\right) & \text { if } v \in \hat{V}-V, \\
\operatorname{SAMP}_{c(s)}(C(v)) & \text { if } v \in V,\end{cases} \\
& B_{s+1}(v)= \begin{cases}B_{s+1}^{\prime}\left(w_{1}^{\prime}\right) \cup B_{r+1}^{\prime}\left(w_{2}\right) & \text { if } \operatorname{Out}(v, \hat{G})=\left\{w_{1}^{\prime}, w_{2}\right\}, \\
B_{s+1}^{\prime}\left(w^{\prime}\right) & \text { if } \operatorname{Out}(v, \hat{G})=\left\{w^{\prime}\right\}_{1}, \\
\varnothing & \text { if } \operatorname{Out}(v, \hat{G})=\varnothing,\end{cases}
\end{aligned}
$$

where $c(s)=2^{\text {flog } N i-s}$ and $N$ is the size of the largest catalogue. The per-stage computation, then, is as follows.

Per-stage computation ( $v, s+1$ ). Using the above definitions, construct $B_{s+1}(v)$ for all $v \in \hat{V}$ in parallel (using $\left|B_{\mathrm{x+1}}(v)\right|$ processors for each $v$ ).

The function $c(s)$ is defined so that if $v \in V$, then as the computation proceeds the list $B_{\mathrm{s}+1}^{\prime}(v)$ will be empty for a while. Then at some stage $s+1$, it will consist of a single element of $C(v)$ (the ( $2^{\text {hog } N 1-s}$ )th element), in stage $s+2$ at most three elements (evenly sampled), in stage $s+3$ at most five elements, in stage $s+4$ at most nine elements, and so on. This continues until the final stage (stage $\lceil\log N\rceil$ ), when $B_{s+1}^{\prime}(v)=C(v)$. Intuitively, the $c(s)$ function is a mechanism for synchronizing the processes of "feeding" the $C(v)$ lists into the nodes of $\hat{G}$ so that all the processes complete at the same time. We show below that each stage can be performed in $O(1)$ time, resulting in a running time of the cascading computations that is $O(\log N)$ (plus the time it takes time to compute the value of $N$, namely, $O(\log n)$ ). The following important lemma is similar to Lemma 2.1 in that it guarantees that the bridge lists do not grow "too much" from one stage to another.

Lemma 3.1. For any sfage $s \leqq 0$ and any node $v \in T,\left|B_{\mathrm{c}+1}(v)\right| \leqq 2\left|B_{s}(v)\right|+4$.
Proof. The proof is by induction on $s$.
Basis ( $s=0$ ). The claim is clearly true for $s=0$.
Induction step ( $s>0$ ). Assume the claim is true for stage $s-1$. If $v \in V$, then the claim follows immediately from the definition of $c(s)$, since in this case $B_{s+1}(v)$ and
$B_{s}(v)$ are both samples of $C(v)$ with $B_{1+1}(v)$ being twice as finc as $B_{1}(v)$, i.e., $\left|B_{s+1}(v)\right| \leqq 2\left|B_{r}(v)\right|+1$.

Consider the case when $v \in \hat{V}-V$, and $\operatorname{Out}(v, \hat{G})=\left\{u_{1}^{\prime}, w_{2}\right\}$. We know in this case $B_{r+1}(v)=B_{r+1}^{\prime}\left(w_{1}\right) \cup B_{r+1}^{\prime}\left(w_{2}\right)$. Thus, we have the following:

$$
\begin{aligned}
\left|B_{\mathrm{s}+1}(v)\right| & =\left[\frac{\left|B_{\mathrm{s}}\left(w_{1}\right)\right|}{4}\right\rfloor+\left\lfloor\frac{\left|B_{\mathrm{s}}\left(w_{2}\right)\right|}{4}\right] \quad \text { (from definitions) } \\
& \leqq\left\lfloor\frac{2\left|B_{r-1}\left(w_{1}\right)\right|+4}{4}\right\rfloor+\left\lfloor\frac{2\left|B_{r-1}\left(w_{2}\right)\right|+4}{4}\right] \quad \text { (by induction hypothesis) } \\
& \leqq 2\left(\left\lfloor\frac{\mid B_{s-1}\left(w_{1}\right)}{4}\right\rfloor+\left\lfloor\frac{\left|B_{s-1}\left(w_{2}\right)\right|}{4}\right]\right)+4 \\
& =2\left|B_{s}(v)\right|+4 .
\end{aligned}
$$

For the case when $v \in \hat{V}-V$ and $\operatorname{Our}(v, \hat{G})$ contains only one node, $w$, the argument is similar and, in fact, simpler. We simply repeat the above argument, replacing $w_{1}$ with $w$ and eliminating those terms that contain $w_{2}$.

In the next lemma we show that the way in which the $B_{s}(v)$ 's grow is "well behaved," much as we did in Lemma 2.2.

Lemma 3.2. Let $[a, b]$ be an interval with $a, b \in\left(-\infty, B_{( }^{\prime}(v), \infty\right)$. If $[a, b]$ intersects $k+1$ items in $\left(-\infty, B_{s}^{\prime}(v), \infty\right)$, then it intersects at most $8 k+8$ items in. $B_{s}(v)$ for all $k \geqq 1$ and $s \geqq 1$.

Proof. The proof is structurally the same as that of Lemma 2.2, since that lemma was based on a merge definition similar to that for $B_{s+1}(v)$.

Corollary 3.3. The list $\left(-\infty, B_{s}^{\prime}(v), \infty\right)$ is a four-cover for $B_{s+1}^{\prime}(v)$, for $s \geqq 0$.
Corollary 3.4. The list $\left(-\infty, B_{s}(v), \infty\right)$ is a 16 -cover for $B_{s}(w)$, for $s \geqq 0$ and $(v, w) \in \hat{E}$.

The first of these two corollaries implies that we can satisfy all the c-cover input conditions for the Merge Lemma (Lemma 2.4) for performing the merge operations for the computation at stage $s$ in $O(1)$ time using $n_{s}$ processors, where $n_{s}=\sum_{\nu \in \dot{\nu}}\left|B_{s}(v)\right|$. We use the second corollary to show that when the computation is completed we will have a fractional cascading data structure (after adding the appropriate rank pointers). We maintain the following rank information at the start of each stage $s$.
(1) For each itern in $B_{s}^{\prime}(v)$ : its rank in $B_{s}^{\prime}(w)$ if $\operatorname{In}(v, \hat{G}) \cap \ln (w, \hat{G})$ is nonempty, i.e., if there is a vertex $u$ such that $(u, v) \in \hat{E}$ and $(u, w) \in \hat{E}$.
(2) For each item in $B_{s}^{\prime}(v)$ : its rank in $B_{s}(v)$ (and thus, implicitly, its rank in $\left.B_{5+1}^{\prime}(\nu)\right)$.

By having this rank information available at the start of each stage $s$, we satisfy all the ranking input conditions of the Merge Lemma. Thus, we can perform each stage in $O(1)$ time using $n_{s}$ processors. Moreover, the output computations of the Merge Lemma allow us to maintain all the necessary rank information into the next stage. Note that in stage $s$ it is only necessary to store the lists for $s-1$; we can discard any lists for stages previous to that, as in the generalized cascading merge.

Recall that we perform the computation for $[\log N]$ stages, where $N$ is the size of the largest catalogue. When the computation completes, we take $B(v)=B_{s}(v)$ for all $v \in \hat{V}$, and for each $(v, w) \in \hat{E}$ we rank $B(v)$ in $B(w)$. We can perform this ranking step by the following method. Assign a processor to each element $b$ in $B(v)$ for all $v \in \hat{V}$ in parallel. The processor for $b$ can find the rank of $b$ in each $B_{s}^{\prime}(w)$ such that $w \in \operatorname{Out}(v, \hat{G})$ in $O(1)$ time because $B_{r}(v)$ contains $B_{r}^{\prime}(w)$ as a proper subset ( $B_{s}^{\prime}(w)$ was one of the lists merged to make $\left.B_{s}(v)\right)$. This processor can then determine the
rank of $b$ in $B(w)=B_{s}\left(w^{-}\right)$for eacin $w \in O u t(v, \hat{G})$ in $O(1)$ time by using the ranking information we maintained (from $B_{:}^{\prime}\left(w^{\prime}\right)$ to $B_{s}\left(w^{\prime}\right)$ ) for stage $s$ (rank condition (2) above).

Given a walk $W=\left(v_{1}, \cdots, v_{m}\right)$, and an arbitrary element $x$, the query that asks for locating $x$ in every $C\left(v_{t}\right)$ is called the multilocation of $x$ in $\left(v_{t}, \cdots, v_{m}\right)$. To perform a multilocation of $x$ in a walk ( $v_{1}, \cdots, v_{m i}$ ), we extend the walk $W$ in $G$ to its corresponding walk $\hat{W}=\left(\hat{v}_{1}, \cdots, \hat{v}_{m}^{\prime}\right)$ in $\hat{G}$ and perform the corresponding multilocation in $\hat{G}$, similar to the method given by Chazelle and Guibas [12] for performing multilocations in their data structure. The multilocation begins with the location of $x$ in $B\left(\hat{v}_{1}\right)=B\left(v_{1}^{\prime}\right)$, the gateway bridge list for $v_{1}$, by binary search. For each other verex in this walk we can locate the position of $x$ in $B\left(\hat{v}_{i}\right)$ given its position in $B\left(\hat{v}_{i-1}\right)$ in $O(1)$ time. The method is to follow the pointer from $x$ 's predecessor in $B\left(\hat{v}_{-1}\right)$ to the predecessor of that element in $B\left(\hat{v}_{\mathrm{i}}\right)$ and then locate $x$ in $B\left(\hat{v}_{i}\right)$ by a linear search from that position (which will require at most 15 comparisons by Corollary 3.4). In addition, if $\hat{v}$, corresponds io a gateway $v^{\prime}$, then we can locate $x$ in $C(v)$ in $O(1)$ time given its position in $B\left(v^{\prime}\right)$ by a similar argument. (See Fig. 3.) Since each edge in the walk $W$ corresponds to a path in $\hat{W}$ of length at most $O(\log d(G))$, this implies that we can perform the multilocation of $x$ in $\left(v_{1}, \cdots, v_{m}\right)$ in $O\left(\log \left|B\left(v_{1}^{\prime}\right)\right|+m \log d(G)\right)$ time. In other words, $\hat{G}$ is a fractional cascading data structure. We show that $\hat{G}$ uses $O(n)$ space in the following lemma.
in $G$ :


Fig. 3. Multilacating an element $x$ in $\left(v_{1}, v_{2}, v_{3}\right)$.

Lemma 3.5. Let $n_{v}$ denote the amount of space that is added to $\hat{G}$ because of the presence of a particular casalogue $C(v), v \in V$. Then $n_{\mathrm{v}} \leqq 2|C(v)|$.

Proof. Recall that while constructing the bridge lists in $\hat{G}$ we copy one-fourth of the elements in each bridge list to at most two of its neighbors. Thus, we have the
following:

$$
\begin{aligned}
n_{\mathrm{r}} & \leqq|C(v)|+2[|C(v)| / 4]+2^{2}\left[|C(v)| / 4^{2}\right\rfloor+2^{3}\left[|C(v)| / 4^{3}\right]+\cdots \\
& \leqq 2|C(v)| .
\end{aligned}
$$

(This is obviously an overestimate, but it is good enough for the purposes of the analysis.)

Corollary 3.6. The total amount of space used by the fractional cascading data structure is $O(n)$, where $n=|V|+|E|+\sum_{v \in V}|C(v)|$.

Proof. The total amount of space used by the fractional cascading data structure is $O\left(|\hat{V}|+|\hat{E}|+\sum_{U^{\prime \prime} \in \dot{V}}|B(v)|\right)$. Since all the bridge lists stan out empty, $\left.\sum_{u \in \dot{v}}|B(v)|\right)=\sum_{v \in V} n_{v}$. The previous lemma implies that $\sum_{v \in v} n_{v} \leqq \sum_{u c v} 2|C(v)|$. Therefore, since $|\hat{V}|+|\hat{E}|$ is $O(|V|+|E|)$ by the definition of $\hat{G}$, the total amount of space used by the fractional cascading data structure is $O(n)$.

Note that the upper bound on the space of the fractional cascading data structure holds even if $\hat{G}$ contains cycles. This corollary, then, implies that we can construct a fractional cascading data structure $\hat{G}$ from any catalogue graph $G$ in $O(\log n)$ time and $O(n)$ space using $n$ processors, even if $G$ contains cycles. We have not shown, however, how to assign these $n$ processors to their respective jobs.

The method for performing the processor allocation is as follows. Initially, we assign $2|C(v)|$ vinual processors to each node $v \in V$ and no processors to each node $\nu \in \hat{V}-V$. This requires at most $2 n$ virtual processors; hence, can be easily simulated with $n$ actual processors. Each time we pass $k$ elements from a node $v$ to a node $w$ (in performing the merge at node $w^{\prime}$ ) we also pass along (exactly) $k$ virtual processors to go with them. When we say that we are passing a virtual processor from some node $o$ to some node $n$, all we are actually changing is the node to which that processor is assigned. Since, by Lemma $3.5, n_{t} \leqq 2|C(v)|$, we know that there are enough virtual processors assigned to $v \in V$ to do this. To see that this also suffices for $v \in \hat{V}-V$ note that at the beginning of stage $s$ node $v$ has $\left|B_{s-1}(v)\right|$ elements (and processors). We "give away" at most $2\left\lfloor\left|B_{s-1}(v)\right| / 4\right\rfloor$ elements (and processors) from $B_{s-1}(v)$ in stage $s$ and receive $\left|B_{s}(v)\right|$ elements (and processors). Consequently, there are enough processors to perform the merge to construct $B_{s}(v)$ and repeat the give-away procedure for the next stage. In addition, since we pass a processor for each item we pass to another node, each processor $p_{i}$ can maintain not only which node it is assigned but $p_{1}$ can also maintain $m_{\nu}$, the number of other processors that are assigned to that node, as well as maintaining a unique integer identification for itself in the range [ $1, m_{4}$ ]. Thus, we have the following lemma.

Lemma 3.7. Given any catalogue graph $G$, we can construct a fracrional cascading data structure for $G$ in $O(\log n)$ time and $O(n)$ space using $n$ processors in the CREW PRAM model.

Thus, we can solve the fractional cascading problem in $O(\log n)$ time using $n$ processors. For the applications we study in this paper, however, we can do even betuer. The following lemma enumerates two important situations where the method just described can be improved.

Lemma 3.8. Given any catalogue graph $G$, if $d(G)$ is $O(1)$ or if we are given $\operatorname{Out}(v, G)$ in sorted order for each $v \in V$, then the total number of operations performed by the fractional cascading algorithm is $O(n)$.

Proof. If $d(G)$ is $O(1)$ or we are given $\operatorname{Our}(v, G)$ in sorted order, then the construction of the graph $\hat{G}$ (without any bridge lists) requires only $O(n)$ operations, since we do not have to perform any sorting. Let us account for the total work performed
by computing the total number of other operations that are performed because of the fact that the catalogue for each node $u$ contains $|C(v)|$ elements (we will only charge vertices in $V$ ). Let $s_{1}$ be the first stage that $B_{1}\left(v^{\prime}\right)$ becomes nonempty. In this stage $B_{r}\left(v^{\prime}\right)$ receives one element of $C(v)$ from $v$, and hence we charge one operation in stage $s_{r}$. for the node $v$. In stage $s_{r}+1$ we will then perform at most 3 operations, at most 7 in stage $s_{1}+2$, at most 15 in stage $s_{t}+3$, and so on. As soon as $B_{s}\left(u^{\prime}\right)$ contains at least four elements from $v$ (as early as stage $s_{v}+2$ ), then we will perform one more operation, passing one element to the fan-in tree for $v$. In the next stage, $s_{x}+3$, we will perform at most two additional operations, then at most four additional operations in stage $s_{r}+4$, and so on. This pattern will "ripple" back through the fan-in tree for $v$ and on through the graph $\hat{G}$ for as long as the computation proceeds. Specifically, the number of operations charged to a node $v \in V$ is, at most, the sum of the following $k_{r}=\left\lceil\log _{4}|C(v)|\right\rceil$ rows:

| 1 | 3 | 7 | 15 | 31 | 63 | 127 | $\cdots$ | $\|C(v)\|$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $2 * 1$ | $2 * 3$ | $2 * 7$ | $2 * 15$ | $2 * 31$ | $\cdots$ | $2\lfloor\|C(v)\| / 4\rfloor$ |  |
|  |  |  | $2^{2} * 1$ | $2^{2} * 3$ | $2^{2} * 7$ | $\cdots$ | $2^{2}\left\lfloor\mid C(v) \\| / 4^{2}\right\rfloor$ |  |
|  |  |  |  |  | $\vdots$ |  | $\vdots$ |  |
|  |  |  |  |  |  |  | $2^{*} * 1$ |  |

where the number in row $i$ and column $j$ corresponds to the maximum number of operations performed in stage $s_{v}+j-1$ at nodes at distance $i$ from $v$ because of the fact that the catalogue at node $v$ contains $|C(v)|$ elements. (This is actually an overestimate, since not all nodes in $\hat{G}$ have out-degree 2). Summing the number of operations for each row, and then summing the rows, we get that the number of operations charged to $v \in V$ is at most $2\left(|C(v)|+2\lfloor|C(v)| / 4\rfloor+2^{2}\left\lfloor|C(v)| / 4^{2}\right\rfloor+\cdots+\right.$ $\left.2^{h_{1}}\right)$, which is at most $\left.4 \mid C(v)\right]$. Thus, the total number of operations performed by the fractional cascading algorithm is $O(n)$.

This lemma immediately suggests that we may be able to apply Brent's theorem to the fractional cascading algorithm so that it runs in $O(\log n)$ time using $\lceil n / \log n\rceil$ processors.

Theorem 3.9 ([11]). Any synchronous parallel algorithm taking time $T$ that consists of a toral of $N$ operations can be simulated by $P$ processors in $O(\lfloor N / P\rfloor+T)$ time.

Proof of Brent's theorem. Let $N$, be the number of operations performed at step $i$ in the paraltel algorithm. The $P$ processors can simulate step $i$ of the algorithm in $O\left(\left\lceil N_{1} / P\right\rceil\right)$ time. Thus, the total running time is $O(\lfloor N / P\rfloor+T)$ :

$$
\sum_{i=1}^{T}\left\lceil N_{i} / P\right\rceil \leqq \sum_{i=1}^{T}\left(\left\lfloor N_{i} / P\right\rfloor+1\right) \leqq\lfloor N / P\rfloor+T .
$$

There are two qualifications we must make to Brent's theorem before we can apply it in the PRAM model, however. The first is that we must be able to compute $N_{i}$ at the beginning of step $i$ in $O\left(\left\{N_{i} / P\right\rceil\right)$ time using $P$ processors. And, second, we must know how to assign each processor to its job. Thus, in order to apply Brent's theorem to our problem of doing fractional cascading, we must deal with these processor allocation problems.

Let $\Gamma=\left\{p_{1}, p_{2}, \cdots, p_{m}\right\}$ be the set of virtual processors used in the fractional cascading algorithm (with $m \leqq 2 n$ ), and let $\Gamma^{\prime}=\left\{p_{1}^{\prime}, p_{2}^{\prime}, \cdots, p_{\{n / l o g n\}}^{\prime}\right\}$ be the set of processors we will be using to simulate the fractional cascading algorithm. Assuming that $d(G)$ is constant or we are given the list of vertices in $\operatorname{Out}(\nu, G)$ in sorted order, we can compute the graph $\hat{G}$ and the initial assignment of processors from $\Gamma$, so that
we assign $2|C(v)|$ virtual processors to each node $v \in V$, in $O(\log n)$ time using the processors in $\Gamma^{\prime}$ by a parallel prefix computation. (Recall that the problem of computing all prefix sums $c_{h}=\sum_{i-1}^{h}, a$ of a sequence of integers $\left(a_{1}, a_{2}, \cdots, a_{n}\right)$ can be done in $O(\log n)$ time using $\lceil n / \log n\rceil$ processors [21], [22].) Let $v\left(p_{i}\right)$ denote the vertex in $\hat{G}$ to which $p_{1} \in \Gamma$ is assigned. Recall that we will be "passing" the processor $p_{\text {, around }}$ $\hat{G}$ during the computation, so the value of $v\left(p_{1}\right)$ can change from one stage to the next. Once a processor $p$, becomes active, it stays active for the remainder of the computation. So, the only thing left to show is how to compute the number of processors active in stage $s$, and to assign the processors in $\Gamma^{\prime}$ to their respective tasks of simulating the processors in $\Gamma$. We do this by sorting the set of processors in $\Gamma$ by the stage in which they become active. It is easy to compute the stage in which a processor $p_{i}$ becomes active in $O(\mathrm{I})$ time, because this depends only on the initial value of $v\left(p_{i}\right)$ and the size of $C\left(v\left(p_{i}\right)\right)$ relative to $N$ (the size of the largest catalogue). We can son the processors in $\Gamma$ by the stage in which they become active in $O(\log n)$ time using the $[n / \log n\}$ processors in $\Gamma^{\prime}$, by using an algorithm from Reif [23] (since the stage numbers fall in the range $[1,\lceil\log N \mid]$ ). Thus, by performing a parallel prefix computation on this ordered list of processors, we can determine the number of processors active in each stage $s$, and also know how to assign the processors in $\Gamma^{\prime}$ so that they optimally simulate the activities of the processors in $\Gamma$ during stage $s$. We thus have established the following theorem.

Theorem 3.10. Given a catalogue graph $G=(V, E)$, such that $-d(G)$ is $O(1)$ or given each $\operatorname{Out}(\nu, G)$ set in sorted order, we can build a fractional cascading data structure for $G$ in $O(\log n)$ time and $O(n)$ space using $[n / \log n\}$ processors in the CREW PRAM model, where $n=|V|+|E|+\sum_{v \in v}|C(v)|$. This bound is optimal.
4. The plane-sweep tree data structure. In this section we define a data structure, which we call the plane-sweep tree, and show how to use it and the fractional cascading procedure of the previous section to solve the trapezoidal decomposition problem and the planar-point location problem in $O(\log n)$ time using $n$ processors. Since the construction of this data structure is quite involved, we merely define the data structure now, and show how to construct it in these same bounds in $\S 5$.

Let $S=\left\{s_{1}, s_{2}, \cdots, s_{n}\right\}$ be a set of nonintersecting line segments in the plane, and let $X^{\prime}(S)=\left(\alpha_{1}, \alpha_{2}, \cdots, \alpha_{2 n}\right)$ be the (nondecreasing) sorted list of the $x$-coordinates of the endpoints of the segments in $S$. To simplify the exposition, we assume that no two endpoints in $S$ have the same $x$-coordinate, i.e., $\alpha_{i}<\alpha_{i+1}$. Let $X^{\prime}=\left(x_{1}, x_{2}, \cdots, x_{m}\right)$ be some subsequence of $X(S)$ and let $T$ be the complete binary tree whose $m+1$ leaves, in left to right order, correspond to the intervals $\left(-\infty, x_{1}\right],\left[x_{1}, x_{2}\right],\left[x_{2}, x_{3}\right]$, $\cdots,\left[x_{m-1}, x_{m}\right],\left[x_{m},+\infty\right)$, respectively. Associated with each internal node $v \in T$ is the interval $I_{v}$ which is the union of the intervals associated with the descendants of v. Let $\Pi_{v}$ denote the verical strip $I_{\mathrm{r}} \times(-\infty,+\infty)$. We say a segment $s_{1}$ covers a node $v \in T$ i ${ }^{\text {it }}$ spans $\Pi_{v}$ but not $\Pi_{\text {marentir) }}$. No segment covers more than two nodes of any level of $T$; hence, every segment covers at most $O(\log m)$ nodes of $T$. For each node $v \in T$ we let $\operatorname{Cover}(v)$ denote the set of all segments in $S$ that cover $v$.

The idea of using a tree data structure such as this to parallelize plane-sweeping is due to Aggarwal et al. [1] and is itself based on the "segment tree" of Bentley and Wood [8]. The data structure of Aggarwal et al. consists of the tree $T$ described above with $X^{\prime}=X^{\prime}(S)$ (i.e., it has $2 n+1$ leaves). Aggarwal et al. store the list Cover $(v)$ at each node $v$ sorted by the "above" relation for tine segments. They construct these lists by first collecting the segments in each Coveri $v$ ) and then sorting all the Cover $(v)$ 's in parallel, an operation that requires $\Theta\left(\log ^{2} n\right)$ time using $n$ processors [13], since
there are a total $\Theta(n \log n)$ items to sort. Once these lists are constructed the data structure can then be used to solve various problems by performing certain searches on the nodes of $T$. These searches are of the following nature: given a set of $O(n)$ input points, for each point $p$ locate the segment in $\operatorname{Cover}(v)$ that is directly above (or below) $p$, for all $v \in T$ such that $p \in \Pi_{r}$. Notice that for the leaf-to-root walk starting with the leaf $v$ such that $p \in \Pi_{r}$, this search can be solved by the multilocation of $p$ in that walk. Aggarwal et al. [1] perform all $O(n)$ multilocations in $O\left(\log ^{2} n\right)$ time using $n$ processors by assigning a processor to each point $p$ and doing a binary search for $p$ in all the $\operatorname{Cover}(v)$ lists such that $p \in \Pi_{v}$ (there are $O(\log n)$ such lists for each $p$ ).

Although based on the structure of Aggarwal et al., the plane-sweep tree differs from it in some important ways. One such difference is that the plane-sweep tree allows us to perform $O(n)$ multilocations in $O(\log n)$ time using $n$ processors, after a preprocessing step that takes $O(\log n)$ time using $n$ processors. Also, instead of taking $X^{\prime}$ to be the entire $X^{\prime}(S)$ list, we define $X^{\prime}$ to be the list consisting of every $\lceil\log n\rceil$ th element of $X(S)$, i.c., $X^{\prime}=S A M P_{\text {「log } n \mid}(X(S))$. Thus, each vertical strip $\Pi_{c}$ associated with a leaf of $T$ in our construction contains $O(\log n)$ segment endpoints. Like Aggarwal et al., we also store each Cover $(v)$ list sorted by the "above" relation. In addition, for every node $v$ of $T$ we define the set End $(v)$ as follows:

$$
\text { End }(v)=\left\{s_{t} \mid s_{1} \in S \text {, has an endpoint in } \Pi_{r}, \text { and does not span } \Pi_{r}\right\} \text {. }
$$

Although End $(v)$ is defined for each node of $T$ we only construct a copy of End ( $v$ ) if $v$ is a leaf node. We do not store the elements of any End $(v)$ in any particular order. This is due to the fact that $\operatorname{End}(v)$ contains $O(\log n)$ segments for any leaf node; hence a single processor can search the entire list in $O(\log n)$ time.

Note that all the segments in the Cover( $b$ l's of any root-to-leaf path in $T$ are comparable by the "above" relation. Thus, if we direct all the edges in $T$ so that each edge goes from a child to its parent, then the elements stored in any directed walk in $T$ are all comparable by the "above" relationship. Therefore, we can apply the fractional cascading technique of the previous section to $T$ (with each $\operatorname{Cover}(v)$ playing the role of the catalogue $C(v)$ ). Since $T$ has bounded degree and has $O(n \log n)$ space, we can, by Theorern 3.10, construct a fractional cascading data structure $\hat{T}$ for $T$ in $O(\log n)$ time and $O(n \log n)$ space using $n$ processors. This data structure allows us to perform the multilocation of any point $p$ (in a leaf-to-root walk) in $O(\log n)$ time $(O(\log n)$ for the binary search at the leaf, and an additional $O(1)$ for each internal node on the path to the root). We also store the set End $(v)$ in each leal $v$ of $\hat{\gamma}$. The plane-sweep tree data structure, then, consists of the tree $\hat{T}$ constructed from $T$ by fractional cascading, where $T$ is defined with $X^{\prime}=\operatorname{SAMP}_{\left[\log { }_{n]}\right.}(X(S))$, has $\operatorname{Cover}(v)$ stored in sorted order for every node $v \in T$, and the set End ( $v$ ) stored (unsorted) for each leaf node $v \in T$ (see Fig. 4).

In $\S 5$ we show how to construct this data structure efficiently in parallel. Since the construction is rather involved, before giving the details of the construction, we give two applications of this data structure. We begin with the trapezoidal decomposition problem.
4.1. The trapezoidal decomposition problem. Let $S=\left\{s_{1}, s_{2}, \cdots, s_{n}\right\}$ be a set of nonintersecting line segments in the plane. For any endpoint $p$ of a segment in $S$ a trapezoidal segment for $p$ is a segment of $S$ that is directly above or below $p$ such that the vertical line segment from $p$ to this edge is not intersected by any other segment in $S$. The trapezoidal decomposition problem is to find the trapezoidal segment(s) for each endpoint of the segments in $S$. Even in the parallel setting, this problem is often


Ftg. 4. A portion of a planc-sweep tree. The segments are numbered in this example by embedding the "above" relation of $\$ 2$ in the fotal order 1, 2. $\cdots, 11$. For simplicity we denote the list Covert b by parentheses and the set End (v) br set braces.
used as a building block to solve other problems, such as polygon triangulation [1], [19], [28] or shortest paths in a polygon [16].

Theorem 4.1. A trapezoidal decomposition of a set $S$ of $n$ nonintersecting segments in the plane can be constructed in $O(\log n)$ time using $n$ processors in the CREW PRAM model, and this is optimal.

Proof. Construct the plane-sweep tree data structure $T$ for $S$. Theorem 5.2 (to be given later, in §5) shows that this structure can be constructed in $O(\log n)$ time using $n$ processors. And we already know that $T$ can be made into a fractional cascading data structure $\hat{T}$ in these same bounds. We assign a single processor to every segment endpoint (there are $2 n$ such points). Let us concentrate on computing the trapezoidal segment below a single segment endpoint $p$. Let $\left(v_{1} \cdots, \operatorname{roor}(T)\right.$ ) be the leaf-to-root path in $\hat{T}$ that starts with the leaf $v$ such that $p \in \Pi_{r .}$. We first search through End $(v)$ to see if there are any segments in this set that are below $p$, and take the one that is closest to $p$ (recall that $\operatorname{End}(v)$ contains $O(\log n)$ segments). We then perform the multilocation of $p$ in the leaf-to-root walk starting at $v$, giving us for each $w$ such that $p \in \Pi_{n}$, the segment in $\operatorname{Cover}\left(u^{\prime}\right)$ directly below $p$. We choose among these $\lceil\log n\rceil$ segments the segment that is closest to $p$. Comparing this segment to the one (possibly) found in End (v), we get the segment in $S$, if there is one, that is directly below $p$. Since the length of the walk from $v$ to $\operatorname{root}(T)$ is at most $\lceil\log n\rceil$, by the method outined at the end of $\S 3$ [12], this computation can be done in $O$ (log $n$ ) time using $n$ processors. Since the two-dimensional maxima problem can be reduced to trapezoidal decomposition in $O(1)$ time using $n$ processors [17], and the two-dimensional maxima problem has a sequential lower bound of $\Omega(n \log n)$ in the algebraic computation tree model [7], [20], we cannot do better than $O(\log n)$ time using $n$ processors.

Solving the trapezoidal decomposition problem efficiently in parallel has proven to be an important step in triangulating a polygon efficiently in parallel [1]. [2], [5], [17], [28]. In fact, Theorem 4.1 is used in the algorithms of Goodrich [19] and Yap [28] to achieve an $O(\log n)$ time solution to polygon triangulation using only $n$ processors. We next point out that the plane-sweep tree can also be used to solve the planar point location problem.
4.2. The planar point location problem. The planar point location problem is the following: Given a planar subdivision $S$ consisting of $n$ edges, construct a data structure that, once constructed, enables one processor to determine for a query point $p$ the face in $S$ containing $p$. This problem has applications in several other parallel computational geometry problems, such as Voronoi diagram construction.

Theorem 4.2. Given a planar subdivision $S$ consisting of $n$ edges, we can construct a data structure that can be used to devermine for any query poimt $p$ the face in $S$ comaining $p$ in $O(\log n)$ serial time. This construction takes $O(\log n)$ time using n processors in the CREW PRAM model.

Proof. The solution to this problem is to build the plane-sweep tree data structure for $S$ (with fractional cascading) and associate with each edge $s_{i}$ the name of the face above $s_{i}$. As already mentioned, Theorem 5.2 (to be given later, in $\S 5$ ) shows that the tree $T$ can be constructed in $O(\log n)$ time using $n$ processors. Also recall that $T$ can be made a fractional cascading data structure $\hat{T}$ in these bounds. Let a query point $p$ be given. A planar point location query for $p$ can be solved in $O(\log n)$ serial time by performing a multilocation like that used in the proof of Theorem 4.1 to find the segment in $S$ directly below $p$. After we have determined the segment $s_{1}$ in $S$ that is directly below $p$, we then can read off the face of $S$ containing $p$ by looking up which face is directly above $s_{i}$.

Incidentally, Theorem 4.2 immediately implies that the running time of the Voronoi diagram algorithm of Aggarwal et al. [1] can be improved from $O\left(\log ^{3} n\right)$ to $O\left(\log ^{2} n\right)$, still using only $n$ processors. (We have recently learned that in the final version of their paper [2], they reduce the time bound of their algorithm to $O\left(\log ^{2} n\right)$ using a substantially different technique.)

The results of $\S \S 4.1$ and 4.2 are conditional: they hold if we can construct the plane-sweep tree data structure efficiently in parallel. We next show how to construct the plane-sweep tree in $O(\log n)$ time using only $n$ processors.
5. Cascading with line segment partial orders. In this section we show how to modify the cascading divide-and-conquer technique of $\S 2$ to solve some geometric problems in which the elements being merged belong to the partial order defined by a set of nonintersecting line segments. Recall that in this partial order a segment $s_{1}$ is "above" a segment $s_{2}$ if there is a vertical line that intersects both segments, and its intersection with $s_{1}$ is above its intersection with $s_{2}$. We apply this technique to the problems of constructing the plane-sweep tree data structure and of detecting if any two of $n$ segments in the plane intersect.

We now give a brief overview of the problems encountered and our solutions to them. The essential computation is as follows: we have a binary tree with lists stored in its leaves, and we wish to combine them in pairs (up the tree) to construct lists at internal nodes. The main difficulty is that the list stored at some node $v$ is not defined as a simple merge of the lists stored at the children of $v$. Instead, its definition involves deleting elements from lists stored at children nodes before performing a merge. These deletions are quite troublesome, because if we try to perform these deletions while cascading, then the rank information will become corrupted, and the cascade will fail.

On the other hand, if we try to postpone the deletions to some postprocessing step, then there will be nondeleted elements that are not comparable to others at the same node; hence, there will be instances when processors try to compare two elements that are not comparable, and the cascade will fail. The main idea of our method for getting around these problems is to embed partial orders in total orders "on the fly" while we are cascading up the tree. That is, we change the identity of segments as they are being passed up the tree, so that the segments in any list are always linearly ordered. To be able to do this, however, we must do some preprocessing that involves simultaneously performing a number of cascading merges in parallel. We complete the computation by performing a purging postprocessing step to remove the segments that "changed identity" (as an alternative to being deleted).

For the intersection detection problem, we need to dovetail the detection of intersections with the cascading. That is, we cascade the results of intersection checks along with the segments being passed up the tree. The complication here is that if we should ever detect an intersection on the way up the tree we cannot stop and answer "yes" as this would require $O(\log n)$ time (to "fan-in" all the possible answers). Thus we are forced to proceed with the merging until we reach the root, even though in the case of an intersection the segments being merged no longer even belong to a partial order. We show that in this case we can replace the segment with a special place holder symbol so that the cascades can proceed. After the cascading merge completes we perform some postprocessing to then check if any intersections are present.

The next two subsections give the details.
5.1. Plane-sweep tree construction. In this subsection we describe how to construct the Cover $(v)$ lists for each node $v$ in the plane-sweep tree $T$. We begin by making a few definitions and observations. We let left $\left(\Pi_{v}\right)$ (respectively, right $\left(\Pi_{v}\right)$ ) denote the teft (right) vertical boundary line for $\Pi_{\mathrm{c}}$. . We define the dominator node of a segment $s_{i}$, denoted $\operatorname{dom}\left(s_{i}\right)$, to be the deepest node $v$ (i.e., farthest from the root) in $T$ such that $s_{i}$ is completely contained in $\Pi_{\text {}}$. That is, the dominator of $s_{i}$ is the node $v$ such that $s_{i}$ does not intersect left $\left(\Pi_{v}\right)$ or right $\left(\Pi_{v}\right)$, but $s_{i}$ does intersect the vertical boundary separating $\Pi_{\text {icturd }(\mathbb{F})}$ and $\Pi_{\text {rchildtor }}$. In addition, we define the following sets for each node $v \in T$ :

$$
\begin{aligned}
& L(v)=\left\{s_{i} \mid s_{i} \in E n d(v) \text { and } s_{i} \cap \operatorname{left}\left(\Pi_{v}\right) \neq \varnothing\right\}, \\
& R(v)=\left\{s_{i} \mid s_{i} \in E n d(v) \text { and } s_{i} \cap \text { right }\left(\Pi_{v}\right) \neq \varnothing\right\}, \\
& I(v, d)=\left\{s_{i} \mid s_{i} \in L(v) \text { and } d=\operatorname{depth}\left(\operatorname{dom}\left(s_{i}\right)\right)\right\}, \\
& r(v, d)=\left\{s_{i} \mid s_{i} \in R(v) \text { and } d=\operatorname{depth}\left(\operatorname{dom}\left(s_{i}\right)\right)\right\} .
\end{aligned}
$$

Note that $I(v, d)$ and $r(v, d)$ are only defined for $0 \leqq d<\operatorname{depth}(v)$. Any time we construct one of these sets it will be ordered by the "above" relation, so for the remainder of this section we represent these sets as soned lists. In the following lemma we make some observations concerning the relationships between the various lists defined above.

Lemma 5.1. Lef v be a node in $T$ with left child $x$ and right child $r$. Then we have the following (see Fig. 5):
(1) $I(v, d)=I(x, d) \cup l(y, d)$ for $d<\operatorname{deph}(v)$,
(2) $r(v, d)=r(x, d) \cup r(\rho, d)$ for $d<\operatorname{depth}(v)$,
(3) $L(v)=I(v, 0) \cup I(v, 1) \cup \cdots \cup I(v$, depth $(v)-1)$,
(4) $R(v)=r(v, 0) \cup r(v, 1) \cup \cdots \cup r(v, \operatorname{deph} h(v)-1)$,
(5) $L(v)=L(x) \cup(L(y)-I(y, \operatorname{dept}(v)))$,


Fig.5. The plane-sweeptree equations. $(\mathrm{a}) /(v, d)=I(x, d) \cup I(1, d) ;(b) L(v)=\{(\mathrm{c}, 0) \cup I(c . \operatorname{deph}(v)-1)$ : (c) $L(v)=L(x) \cup(L(y)-I(y, d e p t h(v)))$ (d) Cover $(x)=L(y)-4(y, d e p h(v))$.
(6) $R(v)=(R(x)-r(x, \operatorname{depth}(v))) \cup R(y)$,
(7) $\operatorname{Cover}(x)=L(y)-1(y, \operatorname{dep} t h(v))$,
(8) Cover $(y)=R(x)-r(x, \operatorname{deph}(v))$.

Proof. The proof follows from the definitions. $]$
Lemma 5.1 essentially states that the lists $l, r, L, R$, and Cover for the nodes on a particular level of $T$ can be defined in terms of lists for nodes on the next lower level of $T$. We could use this lemma and the parallel merge technique of Valiant [26], as implemented by Borodin and Hopcroft [10], to construct a sorted copy of each Cover $(v)$ list in $O(\log n \log \log n)$ time using $n$ processors, improving on the previous bound of $O\left(\log ^{2} n\right)$ time using the same number of processors, due to Aggarwal et al. [1]. We can do even better, however, by exploiting the structure of the $L$ and $R$ lists. We describe how to do this below, in order to achieve a running time of $O(\log n)$ still using $n$ processors. Before going into the details of the plane-sweep tree construction, we give a brief overview of the algorithm.

## High-Level Description of Plane-Sweep Tree Construction.

The construction consists of the following four steps:
Step 1. Construct $I(v, d)$ and $r(v, d)$ for every $v \in T$. To implement this step, we perform $\{\log n\}$ generalized cascading merges in parallel (one for each $d$ ) based on (I) and (2) of Lemma 5.1 (starting with the leaf nodes of $T$ ). We implement this step in $O(\log n)$ time using $n$ processors in total for all the merges.

Step 2. Let $d_{v}=\operatorname{depth}(\operatorname{parent}(v))$. Compute for each segment in $l\left(v, d_{v}\right)$ (respectively, $r\left(v, d_{v}\right)$ ) its predecessor segment in $L(v)-I\left(v, d_{v}\right)$ (respectively, $R(v)-r\left(v, d_{v}\right)$ ) based on (3) and (4). We do this, for each $v \in T$, by making $d_{v}$ copies of $l\left(v, d_{v}\right)$ and $r\left(v, d_{v}\right)$, and merging $l\left(v, d_{v}\right)$ (respectively, $r\left(v, d_{v}\right)$ ) with all the $l(v, d)$ (respectively, $r(v, d)$ ) such that $d<d_{v}$. Note: we perform this step without actually constructing $L(v)$ or $R(v)$.

Step 3. Construct $L(v)$ and $R(v)$ for every $v \in T$. To implement this step we perform a generalized cascading merge procedure based on (5) and (6) and the information computed in Step 2 (starting with the leaf nodes of $T$ ). We never actually perform the set difference operations of (5) and (6), however. Instead, at the point in the merge that a segment in, say, $I\left(v, d_{v}\right)$, should be deleted we "change the identity"
of that segmem to its predecessor in $L(c)-/(t, d$,$) (which we know from Step 2). That$ is, from this point on in the cascading merge this segment is indistinguishable from its predecessor in $L(c)-/\left(c, d_{1}\right)$. We show below that (i) the cascading merge will not be corrupted by doing this, (ii) the lists never contain too many duplicate entries (that would require us to use more than $n$ processors), and (iii) after the merge completes, we can construct $L(v)$ and $R(v)$ for each node by removing duplicate segments in $O(\log n)$ time using $n$ processors.

Step 4. Construct Cover ( $v$ ) for every $v \in T$ using (7) and (8) and the lists constructed in Step 3. The implementation of this step amounts to compressing each $L(v)$ (respectively, $R(v)$ ) so as to delete all the segments in $/\left(u, d_{v}\right)$ (respectively, $r\left(v, d_{\mathrm{t}}\right)$ ), and then copying the list of segments so computed to the sibling node in $T$.

End of High-Level Descrirtion.
We now describe how to perform each of these high-level steps.
5.2. Step I: Constructing $I(v, d)$ and $r(v, d)$. We construct the $I(v, d)$ and $r(c, d)$ lists as follows. We make $\lceil\log n\rceil$ copies of $T$, and let $T(d)$ denote tree number $d$. Note that by our definition of $T$ the space needed to store the "skeleton" of each $T(d)$ is $O(n / \log n)$. This of course results in a total of $O(n)$ space for all the $T(d)$ s. For each node $v$ of $T(d)$ such that $d e p h(v)>d$ we wish to construct the lists $l(c, d)$ and $r(c, d)$, as given by (1) and (2) of Lemma 5.I. This implies that if we store $\|(v, d)$ ( respectively, $r(c, d)$ ) in every leaf node $v$ of $T(d)$, then for any node $v \in T(d), I(v, d)$ is precisely the sorted merge of the lists stored in the descendants of $v$. We start with the elements belonging to $I(v, d)$ (respectively, $r(u, d)$ ) stored (unsorted) in a list $A(v)$ for each leaf $v$ in $T(d)$, and construct each $l(b, d)$ and $r(c, d)$ by the generalized cascading merge technique of Theorem 2.5 (using the $A(v)$ 's as in the theorem). Note: since $I(c, d)$ and $r(L, d)$ are only defined for $d<\operatorname{depth}(v)$, we only proceed up any tree $T(d)$ as far as nodes at depth $d+1$, terminating the cascading merge at that point. We allocate $[n / \log n\}+N_{a}$ processors to each tree $T(d)$, where $N_{d}$ denotes the number of segments stored initially in the leaves of $T(d)$. Thus, since $\sum_{d=1}^{\text {flog } " i} N_{d}=n$, we have shown how to construct all the $I(c, d)$ and $r(v, d)$ lists in $O(\log n)$ time and $O(n \log n)$ space using $n$ processors.
5.3. Step 2: Computing predecessors. In Step 2 we wish to compute for each segment in the list $l\left(v, d_{\mathrm{v}}\right)$ (respectively, $r\left(v, d_{\mathrm{v}}\right)$ ) its predecessor segment in $L(v)$ $\|\left(v, d_{v}\right)$ (respectively, $R(v)-r\left(v, d_{r}\right)$ ), where $d_{r}=\operatorname{dep} t h($ parent $(v))$. Without loss of generality, we restrict our attention to the segments in $/\left(b, d_{r}\right)$ (the treatment for the segments in $r\left(v, d_{c}\right)$ is similar). Recall that (3) and (4) state that $L(v)=$ $\|(v, 0) \cup\|(u, 1) \cup \cdots \cup \|\left(v, d_{\mathrm{c}}\right)$ and that $R(v)=r(v, 0) \cup r(v, 1) \cup \cdots \cup r\left(v, d_{v}\right)$. We make $d_{x}$ copies of $l\left(v, d_{v}\right)$ and, using the merging procedure of Shiloach and Vishkin [25] or that of Bilardi and Nicolau [9], we merge a copy of $I\left(u_{1}, d_{r}\right)$ with each of $l(v, 0), \cdots, l\left(v, d_{r}-1\right)$. This takes $O(\log n)$ time using $\lceil|L(v)| / \log n\rceil+$ $\left\{d_{\mathrm{v}}\left|/\left(v_{1}, d_{\mathrm{t}}\right)\right| / \log n\right\}$ processors for each $v \in T$. Since (i) there are $O(n / \log n)$ nodes in each $T(d)$; (ii) each segment appears exactly once in some $l\left(v_{1} d_{\mathrm{v}}\right)$; and (iii) $\sum_{v \in T}|L(v)|$ is $O(n \log n)$, we can implement all these merges in parallel using $n$ processors. Once we have completed all the merges, we assign a single processor to each segment $s_{i}$ and compare the predecessors of $s$, in $l(v, 0), \cdots, l\left(v, d_{\mathrm{r}}-1\right)$ so as to find the predecessor of $s_{i}$ in $L(v)-/\left(c, d_{i}\right)\left(=\|(v, 0) \cup \cdots \cup l\left(v, d_{\mathrm{v}}-1\right)\right)$. This amounts to $O(\log n)$ additional work for each $s_{1}$; thus Step 2 can be implemented in $O(\log n)$ time using $n$ processors.
5.4. Step 3: Constructing $L(\nu)$ and $R(v)$. In this step we perform another cascading merge on $T$; this time to construct $L(v)$ and $R(v)$ for each $v \in T$ based on (5) and (6)
of Lemma 5.1. Initially, we have $L(v)$ and $R(v)$ constructed only for the leaves. We then merge these lists up the tree based on (5) and (6) as in Theorem 2.5. The computation for this step differs from the cascading merge of Step 2, however, in that we need to be performing set-difference operations as well as list merges as we are cascading up the tree. Unfortunately, it is not clear how to perform these difference operations on-line any faster than $O(\log n)$ time per level, which would result in a running time that is $O\left(\log ^{2} n\right)$. We get around this problem by never actually performing the difference operations. That is, we do not actually delete segments from any lists. Instead, we change the identity of a segment $s_{r}$ in say, $l\left(l, d_{1}\right)$, to its predecessor in $L(v)-l\left(y, d_{v}\right)$ when we are performing the merge as node $v$, where $y=r c h i l d(v)$ (see Fig. 6). We do this instead of simply marking $s_{i}$ as "deleted" in $L(v)$, because segments in $I\left(y, d_{v}\right)$ may not be comparable to segments in $L(x)$ (the list with which we wish to merge $L\left(y^{\prime}\right)-l\left(y, d_{r}\right)$. Simply marking a segment as being "deleted" could thus result in a processor attempting to compare two incomparable segments.


Fig. 6. Segnent identity changing during the cascading merge. We illustrate the way segment names change identity to that of their predecessor as we are performing the cascading merge. In this case we are construcing the L(ul's. We denote the predecessor of each segment by a dotted arrow:

Clearly, the fact that we change the identity of a segment in $l\left(y, d_{y}\right)$ to its predecessor in $L(y)-l\left(r, d_{y}\right)$ means that there will be multiple copies of some segments. This will not corrupt the cascading merge, however, because one of the properties of the "above" relation for segments is that all duplicate copies of a segment will be contiguous. Moreover, they will remain contiguous as the cascading merge proceeds up the tree. In addition, even though we will have multiple copies of segments in lists as they are merging up the tree, we can still implement this step with a total of $n$ processors, because there will never be more items present in any $L(b)$ than the total number of items stored in the (leaf) descendants of $c$. At the end of this step we assign $\lceil|L(v)| / \log n\rceil$ processors to each $v$ and compress out the duplicate entries in $L(v)$ in $O(\log n)$ rime. Thus, we can construct $L(v)$ and $R(v)$ (compressed and sorted) for each $v \in T$ in $O(\log n)$ time using $n$ processors.
5.5. Step 4: Constructing Cover(v). In this step we construct Cover(v) for every $v$ in $T$, based on (7) and (8) of Lemma 5.1. We impiement this step by first compressing each $L(v)$ (respectively, $R(v)$ ) so as to delete all the segments in $l\left(v, d_{v}\right.$ ) (respectively,
$r\left(v, d_{\mathrm{r}}\right)$ ), and then by copying the list of segments so computed to the sibling of $v$ in $T$. This can all be done in $O(\log n)$ time using $n$ processors.

Thus, summarizing the entire previous section, we have the following theorem.
Theorem 5.2. Given a set $S$ of nonintersecting line segments in the plane, we can construct the plane-sweep tree $T$ for $S$ in $O(\log n)$ time using $n$ processors in the CREW PRAM model, and this is optimal.

Proof. We have already established the correctness and complexity bounds. To see that our construction is optimal, note that the plane-sweep tree requires $\Omega(n \log n)$ space.

In the previous sections we assumed that segments did not intersect. Indeed, $T$ is defined only if they do not intersect. We show in the next section that we can detect an intersection, if there is one, by constructing $T$ while simultaneously checking for intersections.
5.6. The segment intersection detection problem. The problem we solve in this section is the following: given a set $S$ of $n$ line segments in the plane, determine if any two segments in $S$ intersect. We begin by stating the conditions that we use to test for an intersection.

Lemma 5.3 [1]. The segments in $S$ are nonintersecting if and only if we have the following for the plane-sweep tree $T$ of $S$ :
(1) For every $v \in T$ all the segments in Cover $(v)$ intersect left $\left(\Pi_{v}\right)$ in the same order as they intersect right ( $\Pi_{r}$ ).
(2) For every $v \in T$ no segment in End (v) intersects any segment in Cover (v).

Aggarwal et al. [1] used this lemma and their data structure to solve the intersection detection problem in $O\left(\log ^{2} n\right)$ time using $n$ processors. Their method consisted of constructing the Cover $(v)$ lists independently of one another, basing comparisons on segment intersections with left ( $\Pi_{\bullet}$ ), and then testing for condition (1) by checking if each list Cover ( $v$ ) would be in the same order if they based comparisons on segment intersections with right $\left(\Pi_{v}\right)$. If no intersection was detected by this step, then they tested for condition (2) by performing $O(n)$ multilocations of segment endpoints. This entire process took $O\left(\log ^{2} n\right)$ time using $n$ processors.

We use this lemma by testing for condition (1) while we are constructing the plane-sweep tree for $S$ (instead of waiting until after it has been built) and in so doing we achieve an $O(\log n)$ time bound for this test (since our construction takes only $O(\log n)$ time $)$. We test condition (2) in the same fashion as Aggarwal et al., that is, by doing $O(n)$ multilocations after the plane-sweep tree has been built. Since with our data structure the multiplications can all be performed in $O(\log n)$ time, the entire intersection-detection process takes $O(\log n)$ time using $n$ processors.

Since we do not construct the $\operatorname{Cover}(v)$ lists independently of one another, but instead construct them by performing several cascading merges, we must be very careful in how we base segment comparisons, and in how we test for condition (1). For if two segments intersect, then determining which segment is above the other depends on the vertical line upon which we base the comparison.

We consider each step of the construction in turn, beginning with Step 1. Recall that in Step 1 we construct all the $l(v, d)$ and $r(v, d)$ lists for each $v \in T$. In the following lemma we show that if we base segment comparisons on appropriate vertical lines, Step 1 can be performed just as before.

Lemma 5.4. Let $v \in T$ and $0 \leqq d<d e p t h(v)$ be given, and let $s_{1}$ and $s_{2}$ be two segmems such that $s_{1} \in I(w, d)$ and $s_{2} \in I(z, d)$ (or $s_{1} \in r(n, d)$ and $s_{2} \in r(z, d)$, where $w_{1}, z \in \operatorname{Desc}(v)$. Then $\operatorname{dom}\left(s_{1}\right)=\operatorname{dom}\left(s_{2}\right)$.

Proof. Let $v \equiv T$ and $0 \leqq d \leqq\lceil\log n\rceil$ be given. Recall that $\|(d)$ (respectively, $r(v, d)$ is defined to be the list of all segments in $L(v)(R(v))$ that have a dominator node at depth $d$ in $T$. Note that the dominator node for any segment $s_{i}$ in $l(u, d)$, $r(w, d), l(=, d)$, or $r(\Sigma, d)$, where $w, z E \operatorname{Desc}(v)$, must be an ancestor of $v$, since $d<d c p h(v)$ and, by definition, $s, \in \operatorname{End}(v)$ and $s, \in \operatorname{End}(d o m(s)$,$) . There is only one$ node that is an ancestor of $c$ and is at depth $d$ in $T$.

Thus, we can perform the merges based on (1) and (2) of Lemma 5.1 (e.g., $\|(v, d)=\|(x, d) \cup l(y, d))$ by basing all segment comparisons on the intersection of the segments with the vertical boundary separating the two children of their dominator node. That is, if $s_{1}$ and $s_{2}$ are two segments to be compared in Step 1, then we say that $s_{1}$ is "above" $s_{2}$ if and only if the intersection of $s_{1}$ with $L$ is above the intersection of $s_{2}$ with $L$, where $L$ is the vertical boundary line separating the two children of $\operatorname{dom}\left(s_{1}\right)\left(=\operatorname{dom}\left(s_{2}\right)\right)$.

In Step 2 we computed for each segment in $\left(t v, d_{r}\right)$ (respectively, $r\left(v, d_{r}\right)$ ) its predecessor segment in $L(v)-\|\left(v, d_{\mathrm{c}}\right)$ (respectively, $R(v)-r\left(v, d_{1}\right)$ ), where $d_{\mathrm{r}}=$ depth(parem $(v))$. Recall that we did this by merging $\|\left(v, d_{x}\right)$ with each of $l(u, 0), \cdots, l\left(u, d_{\mathrm{r}}-1\right)$. A similar computation was done for $r\left(u, d_{\mathrm{r}}\right)$; without loss of generality, we concentrate on the computation involving $\left(L_{1}, d_{\mathrm{t}}\right)$. Also recall that all the segments in $/(c, 0), \cdots, /\left(c, d_{\mathbf{c}}\right)$ belong to $L(E)$; hence they intersect left $\left(\Pi_{r}\right)$. After Step 1 finishes, each list $/(c, d)$ will be sorted based on segment intersections with the vertical boundary line separating the two children of the ancestor of $c$ at depth $d$ (the dominator of all the segments in $((c, d) \|$. In $O(\log n)$ time we can cleck if this order is preserved in each of $l(c, 0), \cdots, l\left(c, d_{r}\right)$ if we base segment intersections on left $\left(\Pi_{r}\right)$, instead. If the order changed in any $l(v, d)$, then we have detected an intersection, and we are done. Otherwise, we proceed with Step 2 just as before, basing comparisons on segment intersections with le $\int\left(\Pi_{\mathrm{r}}\right)$.

In Step 3 we performed a cascading merge up the tree $T$, constructing $L(v)$ and $R(v)$ for every node $v \in T$. Recall that this cascading merge was based on ( 5 ) and (6) of Lemma 5.1 (e.g., $L(v)=L(x) \cup(L(y)-I(1, d e p t h(v)))$ ). Let us concentrate on the testing procedure for the $L(v)$ s, since the method for the $R(v)$ s is similar. Initially, let us start with each $L(v)$ constructed at the leaves of $T$ sorted by segment intersections with left ( $\Pi_{r}$ ). Thus, before we perform the merge based on the equation $L(v)=L(x) \cup$ ( $L(y)-I(y, \operatorname{depth}(v)))$, we must first check to see if the segments in the sample of $L(y)-I(y, \operatorname{dep} h(v))$ (to be merged with the sample of $L(x)$ ) have the same order independent of whether comparisons are based on segment intersections with left ( $\Pi_{y}$ ) or left $\left(\Pi_{2}\right)$. Unfortunately, to do this completely would require $O(\log n)$ time at every level of the tree, resutting in an $O\left(\log ^{-} n\right)$ time algorithm. So, instead of broadcasting at each level whether an intersection has occurred or not, we cascade that information $u p$ along with the merges. More precisely, before doing the merge at a node $v$, we test if every consecutive pair of items in the sample of $L(y)-l(y, d e p t h(v))$ would remain in the same order independent of whether comparisons were based on segment intersections with left $\left(\Pi_{v}\right)$ or with left $\left(\Pi_{v}\right)$. If we detect that an intersection has occurred, then we will have two elements that are out of order. If this should occur, we replace both items by the distinguished symbol $S$. Then, as the merges continue up the tree, any time we compare an item with $S$, we replace that item with $\$$ and proceed just as before. This keeps the merging process consistent, and after the cascading merge completes we can then in $O(\log n)$ time test if any of the items in any $L(v)$ or $R(v)$ contain a $S$ symbol, by assigning $\lceil|L(v)| / \log n\rceil$ processors to each $v \in T$.

In Step 4 we consiructed $\operatorname{Cover}(v)$ for each $v \in T$. Recall that we did this by simply performing compressing and copying operations on lists constructed in Step 3. Thus,
assuming that no intersection was detected in Step 3, we can perform Step 4 iust as before. After Step 4 completes we can assign $[\mid$ Cover $(b) \mid / \log n]$ processors to each $v \in T$ and test condition (1) directly in $O(\log n)$ time, checking if the items in Cover $v$ ) would be in the same order independent of whether comparisons were based on Icft ( $\Pi_{.}$) or on right ( $\Pi_{1}$ ).

If we have not discovered an intersection after Step 4, then the only computation left is to perform fractional cascading on the plane-sweep tree $T$, constructing a fractional cascading data structure $\hat{T}$. In directing all the edges in $T$ to the root, and performing the fractional cascading preprocessing on $T$ to construct $\hat{T}$, we associate a vertical strip with each node in $\hat{T}$. Since $T$ is a tree then $\hat{T}$ is also a tree (recall the preprocessing step of the (ractional cascading algorithm). For each node $v$ in $\hat{T}$ if $v$ is also in $T_{\text {, then }}$ we take $H_{r}$ for $v$ in $\hat{T}$ to be the same as $I_{r}$. for $v$ in $T$. Then, for any $v$ that is in $\dot{T}$ but not in $T$ (i.e., $v$ is a gateway or a node in a fan-in or fan-out tree), we take Il, to be the union of all the vertical strips that are descendents of $e$. Every time we perform the per-stage merge computation we compare adjacent entries in each bridge list $B(v)$ to see if they would be in the same order independent of whether we base comparisons on segment intersections with left ( $\Pi_{r}$ ) or right ( $\Pi_{i}$ ). If we detect that two adjacent segments intersect, then we replace both with the special symbol $S$. Then, as before, any time we compare a segment with $S$ we replace that segment by S. Finally, when we complete the computation for Step 5, we assign $\lceil|B(b)| / \log n\rceil$ processors to each node $v$ and check if there are any $S$ symbols present in any $B(b)$ iist.

If there are no intersections detected during the fractional cascading, then we perform $O(n)$ multilocations of all the segment endpoints as in [1] to test condition (2). Let $p$ be an endpoint of some segment $s_{t}$. We perform the multilocation of $p$ in the plane-sweep tree for $S$, and check if $s_{i}$ intersects the segment directly above $p$ or the segment directly below $p$ in each $\operatorname{Cover}(v)$ list such that $p \in \mathrm{II}_{\text {.. This }}$. Thest is sufficient, since if $s$, intersects any segment in $\operatorname{Cover}(v)$, it must intersect the segment directly above $p$ in Cover ( $c$ ) or the segment directly below $p$ in $\operatorname{Cover}(v)$. Thus, by performing a multilocation for $p$, we can test for condition (2) in $O(\log n)$ time using $n$ processors. We summarize this discussion in the following theorem.

Theorem 5.5. Given a set of $n$ line segments in the plane, we can detect if any two imersect in $O(\log n)$ time using n processors in the CREW PRAM model.

So far in this paper we have restricted ourselves to applications involving line segments. In the next section we show how to apply the cascading divide-and-conquer technique to other geometric problems as well.
6. Cascading with labeling functions. In this section we show how to solve several different geometric problems by combining the merging procedure of $\$ 2$ with divide-and-conquer strategies based on merging lists with labels defined on their elements. For most of these problems our divide-and-conquer approach gives an efficient sequential alternative to the known sequential algorithms (which use the plane-sweeping paradigm) and gives rise to efficient parallel algorithms as well. We begin with the three-dimensional maxima problem.
6.1. The three-dimensional maxima problem. Let $V=\left\{p_{1}, p_{2}, \cdots, p_{n}\right\}$ be a set of points in $: 1^{3}$. For simplicity, we assume that no two input points have the same $x$ (respectively,,$y=$ ) coordinate. We denote the $x, y$, and $z$ coordinates of a point $p$ by $x(p), r(p)$, and $z(p)$, respectively. We say that a point $p_{r}$ one-dominares another point $p_{\text {, if }} x\left(p_{1}\right)>\left(p_{i}\right)$, mo-dominares $p_{i}$ if $x\left(p_{i}\right)>x\left(p_{i}\right)$ and $r\left(p_{i}\right)>r\left(p_{j}\right)$, and three dominates
$p_{1}$ if $x\left(\rho_{1}\right)>x\left(p_{i}\right), r\left(p_{1}\right)>r\left(p_{i}\right)$, and $z\left(p_{1}\right)>z\left(p_{i}\right)$. A point $p_{1} \in V$ is said to be a maximum if it is not threc-dominated by any other point in $V$. The three-dimensional maxima problem, then, is to compute the set, $M$, of maxima in $V$. We show how to solve the three-dimensional maxima problem efficiently in parallel in the following algorithm.

Our method is based on cascading a divide-and-conquer strategy in which the subproblem merging step involves the computation of two labeling functions for each point. The labels we use are motivated by the optimal sequential plane-sweeping algorithm of Kung, Luccio, and Preparata [20]. Specifically, for each point $p_{i}$ we compute the maximum $z$-coordinate from among all points that one-dominate $p_{1}$ and use that label to also compute the maximum $z$-coordinate from among all points that two-dominate $p_{i}$. We can then test if $p_{i}$ is a maximum point by comparing $z\left(p_{\mathrm{r}}\right)$ to this latter label. The details follow.

Without loss of generality, we assume the input points are given sorted by increasing $y$-coordinates, i.e., $P\left(p_{i}\right)<y\left(p_{r+1}\right)$, since if they are not given in this order we can sor them in $O(\log n)$ time using $n$ processors [13]. Let $T$ be a complete binary tree with leaf nodes $v_{1}, v_{2}, \cdots, v_{n}$ (in this order). In each leaf node $v_{i}$ we store the list $B\left(v_{i}\right)==$ $\left(-\infty, p_{i}\right)$, where $-\infty$ is a special symbol such that $x(-\infty)<x\left(p_{j}\right)$ and $y(-\infty)<y\left(p_{j}\right)$ for all points $p_{j}$ in $V$. Initializing $T$ in this way can be done in $O(\log n)$ time using $n$ processors. We then periorm a generalized cascading merge from the leaves of $T$ as in Theorem 2.5, basing comparisons on increasing $x$-coordinates of the points (not their $y$-coordinates). Using the notation of $\S 2$, we let $U(v)$ denote the sorted array of the points stored in the descendants of $v \in T$ soned by increasing $x$-coordinates. For each point $p_{1}$ in $U(v)$ we store two labels: $\operatorname{zod}\left(p_{i}, v\right)$ and $\operatorname{ztd}\left(p_{i}, v\right)$, where $\operatorname{zod}\left(p_{i}, v\right)$ is the largest $z$-coordinate of the points in $U(v)$ that one-dominate $p_{i}$, and $z t d\left(p_{i}, v\right)$ is the largest $z$-coordinate of the points in $U(v)$ that two-dominate $p_{i}$. Initially, zod and zrd labels are only defined for the leaf nodes of $T$. That is, $\operatorname{sod}\left(p_{i}, v_{1}\right)=z t d\left(p_{i}, v_{i}\right)=$ $-\infty$ and $\operatorname{zod}\left(-\infty, v_{i}\right)=\operatorname{zid}\left(-\infty, v_{i}\right)=z\left(p_{i}\right)$ for all leal nodes $v_{i}$ in $T$ (where $U\left(v_{i}\right)=$ ( $-\infty, p_{r}$ )). In order to be more explicit in how we refer to various ranks, we let pred $\left(p_{i}, v\right)$ denote the predecessor of $p_{i}$ in $U(v)$ (which would be $-\infty$ if the $x$ coordinates of the points in $U(v)$ are all larger than $x\left(p_{i}\right)$ ) (see Fig. 7). As we are performing the cascading merge, we update the labels zod and ztd based on the equations in the following lemma.

Lemma 6.1. Let $p$, be an element of $U(v)$ and let $u=l$ child $(v)$ and $w=$ rchild (v). Then we have the following:

$$
\begin{align*}
& \operatorname{zod}\left(p_{i}, v\right)= \begin{cases}\max \left\{\operatorname{zod}\left(p_{i}, u\right), \operatorname{zod}\left(\operatorname{pred}\left(p_{r}, w\right), w\right)\right\} & \text { if } p_{i} \in U(u), \\
\max \left\{\operatorname{zod}\left(\operatorname{pred}\left(p_{i}, u\right), u\right), \operatorname{zod}\left(p_{i}, w\right)\right\} & \text { if } p_{i} \in U(w)\end{cases}  \tag{9}\\
& \operatorname{zrd}\left(p_{i}, v\right)= \begin{cases}\max \left\{\operatorname{ztd}\left(p_{i}, u\right), \operatorname{zod}\left(\operatorname{pred}\left(p_{t}, w\right), w\right)\right\} & \text { if } p_{i} \in U(u) \\
\operatorname{ztd}\left(p_{i}, w\right) & \text { if } p_{i} \in U(w)\end{cases}
\end{align*}
$$

Proof. Consider (9). If $p_{\mathrm{r}} \in U(u)$, then every point that one-dominates $p_{i}$ 's predecessor in $U\left(w^{\prime}\right)$ also one-dominates $p_{1}$, since $p_{t}$ 's predecessor in $U\left(w^{\prime}\right)$ is the point with largest $x$-coordinate less than $x\left(p_{i}\right)$ (or $-\infty$ if every point in $U(w)$ has targer $x$-coordinate than $p_{i}$ ). Thus $\operatorname{zod}\left(p_{i}, v\right)$ is the maximum of $\operatorname{zod}\left(p_{i}, u\right)$ and $\operatorname{zod}\left(\operatorname{pred}\left(p_{i}, w\right), w\right)$ in this case. The case when $p_{r} \in U(w)$ is similar. Next, consider (10). We know that every point in $U\left(w^{\prime}\right)$ has $y$-coordinate greater than every point in $U(u)$, by our construction of $T$. Therefore, if $p_{1} \in U(u)$, then every point in $U(w)$ that one-dominates $p_{\mathrm{\prime}}$ 's predecessor in $U(w)$ must two-dominate $p_{i}$. Thus, zid $\left(p_{i}, v\right)$ is the


Fig. 7. The combening stepfor three-dimensional maxima. Points to the right of the dotied line one-dominate $p_{r}$ (respectivel!, $p_{i}$ ), and poinrs enclosed in the dashed lines wo-dominare $p_{i}\left(p_{i}\right)$.
maximum of $z t\left(p_{i}, u\right)$ and $z o d\left(p r e d\left(p_{r}, w^{\prime}\right), w\right)$. On'the other hand, if $p_{\mathrm{i}} \in U(w)$ then no point in $U(u)$ can two-dominate $p_{i}$; thus, $z r d\left(p_{i}, v\right)=\operatorname{ztd}\left(p_{i}, w\right)$.

We use these equations during the cascading merge to maintain the labels for each point. By Lemma 6.1, when $v$ becomes full (and we have $U(u), U(w)$, and $U(u) \cup U(w)$ available), we can determine the labels for all the points in $U(v)$ in $O(1)$ additional time using $|U(v)|$ processors. Thus, the running time of the cascading merge algorithm, even with these additional label computations, is still $O(\log n)$ using $n$ processors. Moreover, after $c$ 's parent becomes full we no longer need $U(v)$, and can deallocate the space it occupies. resulting in an $O(n)$ space algorithm, as outlined in $\$ 2$. After we complete the merge, and have computed $U(\operatorname{root}(T))$, along with all the labels for the points in $U(\operatorname{rool}(T))$, note that a point $p_{i} \in U(\operatorname{root}(T))$ is a maximum if and only if $z t d\left(p_{t}, \operatorname{roor}(T)\right) \leqq z\left(p_{i}\right)$ (there is no point that two-dominates $p_{i}$ and has z-coordinate greater than $z\left(p_{i}\right)$ ). Thus, after completing the cascading merge we can construct the set of maxima by compressing all the maximum points into one contiguous list using a simple parallel prefix computation. We summarize in the following theorem.

Theorem 6.2. Given a set $V$ of $n$ points in $9^{3}$, we can construct the set $M$ of maxima points in $V$ in $O(\log n)$ time and $O(n)$ space using $n$ processors in the CREW PRAM model, and this is optimal.

Proof. We have established the correctness and complexity bounds for parallel three-dimensional maxima finding in the discussion above. Kung, Luccio, and Preparata [20] have shown that this problem has an $\Omega(n \log n)$ sequential lower bound (in the comparison model). Thus, we can do no better than $O(\log n)$ time using $n$ processors.

It is worth noting that we can use roughly the same method as that above as the basis step of a recursive procedure for solving the general $k$-dimensional maxima problem. The resulting time and space complexities are given in the following theorem. We state the theorem for $k \geqq 3$ (since the two-dimensional maxima problem can easily be solved in $O(\log n)$ time and $O(n)$ space by a sorting step followed by a parallel prefix step).

TitaOREM 6.3. For $k \equiv 3$ the h-dimensional maxima problem can be solued in O(log $n)^{h-3}$ ) fime using $n$ processors in the CREW PRAM model.

Proof. The method is a straightforward parallelization of the algorithm by Kung, Luccio, and Preparata [20]. using a procedure very similar to that described above as the basis for the recursion. We leave the decails to the reader.

Next, we address the two-set dominance counting problem. We also show how the mutiple range-counting problem and the rectitinear segment intersection counting problem can be reduced to two-set dominance problems efficiently in parallet.
6.2. The two-set dominance counting probiem. In the two-set dominance counting problem we are given a set $A=\left\{q_{1}, q_{2}, \cdots, q_{m}\right\}$ and a set $B=\left\{r_{1}, r_{2}, \cdots, r_{t}\right\}$ of points in the plane, and wish to know for each point $r_{r}$ in $B$ the number of points in $A$ that are two-dominated by $r_{i}$. For simplicity, we assume that the points have distinct $x$ (respectively, ,) coordinates. Our approach to this problem is simitar to that of the previous subsection, in that we will be performing a cascading merge procedure while maintaining two labeling functions for each point. In this case the labels maintain for each point $p_{1}$ (from $A$ or $B$ ) how many points of $A$ are one-dominated by $p_{i}$ and also how many points of $A$ are two-dominated by $p_{i}$. As in the previous solution, the first label is used to maintain the second. The details follow.

Let $Y=\left\{p_{1}, p_{2}, \cdots, p_{1+m}\right\}$ be the union of $A$ and $B$ with the points listed by increasing $y$-coordinate, i.e., $r\left(p_{i}\right)<r\left(p_{i+1}\right)$. We can construct $Y$ in $O(\log n)$ time using $n$ processors [13], where $n=1+m$. Our method for solving the two-set dominance counting problem is similar to the method used in the previous subsection. As before, we let $T$ be a complete binary tree with leaf nodes $v_{1}, v_{2}, \cdots, v_{n}$, in this order, and in each leal node $v_{\text {, }}$ we store the list $U\left(v_{i}\right)=\left(-\infty, p_{i}\right)(-\infty$ still being a special symbol such that $x(-\infty)<x\left(p_{i}\right)$ and $y(-\infty)<y^{\prime}\left(p_{i}\right)$ for all points $p_{i}$ in $Y$. We then perform a generalized cascading merge from the leaves of $T$ as in Theorem 2.5, basing comparisons on increasing $x$-coordinates of the points (not their $y$-coordinates). We let $U(v)$ denote the sorted array of the points stored in the descendants of $v \in T$ sorted by increasing $x$-coordinate. For each point $p_{1}$ in $U(v)$ we store two labels: nod $\left(p_{i}, v\right)$ and $n d\left(p_{i}, v\right)$. The label $\operatorname{nod}\left(p_{i}, v\right)$ is the number of points in $U(v)$ that are in $A$ and are one-dominated by $p_{\mathrm{r}}$, and the label $m d\left(\rho_{r}, v\right)$ is the number of points in $U(v)$ that are in $A$ and are two-dominated by $p_{i}$. Initially, the nod and nid labels are only defined for the leaf nodes of $T$. That is, $\operatorname{nod}\left(p_{i}, v_{i}\right)=\operatorname{nod}\left(-\infty, v_{i}\right)=\operatorname{md}\left(p_{i}, v_{i}\right)=\operatorname{md}\left(-\infty, v_{i}\right)=$ 0 . For each $p_{i} \in Y$ we define the function $\chi_{A}\left(p_{r}\right)$ as follows: $\chi_{A}\left(p_{i}\right)=1$ if $p_{r} \in A$, and $\chi_{A}\left(p_{1}\right)=0$ otherwise. (We also use pred $\left(p_{1}, v\right)$ to denote the predecessor or $p$ in $U(v)$. As we are performing the cascading merge, we update the labels nod and nrd based on the equations in the following lemma (see Fig. 8).

Lemma 6.4. Let $p_{i}$ be an element of $U(v)$ and let $u=$ lchild $(v)$ and $w=$ rchild $(v)$. Then we have the following:
(11) $\operatorname{nod}\left(p_{i}, v\right)= \begin{cases}\operatorname{nod}\left(p_{1}, u\right)+\operatorname{nod}\left(\operatorname{pred}\left(p_{1}, u^{\prime}\right), u^{\prime}\right)+\chi_{A}\left(\operatorname{pred}\left(p_{1}, u^{\prime}\right)\right) & \text { if } p_{1} \in U(u), \\ \operatorname{nod}\left(\operatorname{pred}\left(p_{i}, u\right), u\right)+\operatorname{nod}\left(p_{i}, u^{\prime}\right)+\chi_{A}\left(\operatorname{pred}\left(p_{i}, u\right)\right) & \text { if } p_{r} \in U\left(w^{\prime}\right),\end{cases}$

$$
\operatorname{md}\left(p_{r}, v\right)=\left\{\begin{array}{l}
n d\left(p_{i}, u\right) \quad \text { if } p_{i} \in U(u)  \tag{12}\\
\operatorname{nod}\left(\operatorname{pred}\left(p_{i}, u\right), u\right)+\operatorname{nrd}\left(p_{i}, w\right)+\chi_{A}\left(\operatorname{pred}\left(p_{i}, u\right)\right) \quad \text { if } p_{i} \in U(u) .
\end{array}\right.
$$

Proof. Consider (11). For any point $p_{1} \in U(u)$ the number of points one-dominated by $p_{1}$ is equal to the number of points in $U(u)$ that are in $A$ and one-dominated by


Figi. 8 . The combinmg step for dominance couming. Perints to the leff of the doned the are one-dominated

$p_{1}$. plus the number of points in $U(w)$ that are in $A$ and one-dominated by pred $\left(p_{1}, w\right)$, plus one if pred ( $p_{r}, w$ ) is in $A$ (since the predecessor of $p_{i}$ is one-dominated by $p_{i}$ ). Thus, we have the equation for the case when $p_{1} \in U(u)$. The case when $p_{i} \in U(w)$ is similar. Next, consider (12). By our construction, every point in $U(u)$ has 1 -coordinate less than the $y$-coordinate of every point in $U(u)$. So if $p, \in U(u)$, then the number of points in $U(0)$ that are in $A$ and are two-dominated by $p_{i}$ is preciscly $m d\left(p_{1}, u\right)$, since $p_{1}$ cannot two-dominate any points in $U(w)$. I $p_{\mathrm{r}} \in U(w)$, on the other hand, then the number of points in $U(v)$ that are in $A$ and two-dominated by $p$, is the number of points in $U(u)$ that are in $A$ and one-dominated by pred ( $p_{i}, u$ ), plus the number of points in $U(w)$ that are in $A$ and two-dominated by $p_{i}$, plus one if pred ( $p_{i}, u$ ) is in $A$. This is exactly (12) in this case.

By Lemma 6.4, when $v$ becomes full (and we have $U(u), U\left(u^{\prime}\right)$, and $U(v)=U(u) \cup$ $U(w)$ available), we can determine the labels for all the points in $U(v)$ in $O(1)$ additional time using $|U(v)|$ processors. Thus, the running time of the cascading merge algorithm, even with these additional label computations, is still $O(\log n)$ using $n$ processors. After we complete the merge, and have computed $U$ (rool ( $T$ ), along with all the labels for the points in $U(\operatorname{rool}(T))$, then we are done. We summarize in the following theorem.

Theorem 6.5. Given a set $A$ of $I$ points in the plane and a set $B$ of $m$ points in the plane, we can compute for each point $p$ in $B$ the number of points in $A$ who-dominated $b y$ $p$ in $O(\log n)$ time and $O(n)$ space using $n$ processors in the CREW PRAM model, where $n=1+m$, and this is optimal.

Proof. The correctness and complexity bounds should be apparent from the discussion above. To prove the lower bound note that the wo-dimensional maxima problem can be reduced to dominance counting in $O(1)$ time using $n$ processors (see [17]). Since the maxima problem has an $\Omega(n \log n)$ lower bound [20] in the comparison model, we conclude that we can do no better than $O(\log n)$ time using n processors in the CREW PRAM model.

There are a number of other problems that can be reduced to two-set dominance counting. We mention two here, the first being the multiple range-counting problem:
given a set $V$ of $/$ points in the plane and a set $R$ of $m$ isothetic rectangles (ranges) the multiple range-counting problem is to compute the number of points interior to each rectangle.

Corollary 6.6. Given a set $V$ of $I$ points in the plane and a set $R$ of $m$ isothetic rectangles, we can solve the multiple range-counting problem for $V$ and $R$ in $O(\log n)$ time and $O(n)$ space using $n$ processors, where $n=1+m$.

Proof. Let $d(p)$ be the number of points in $V$ two-dominated by a point $p$. Edelsbrunner and Overmars [15] have shown that counting the number of points interior to a rectangle can be reduced to dominance counting. That is, given a rectangle $r=\left(p_{1}, p_{2}, p_{3}, p_{4}\right)$ (where vertices are listed in counterclockwise order starting with the upper righthand corner), the number of points in $V$ interior to $r$ is $d\left(p_{1}\right)-d\left(p_{2}\right)+$ $d\left(p_{3}\right)-d\left(p_{4}\right)$. Therefore, it suffices to solve the two-set dominance counting problem.

Another problem that reduces to two-set dominance counting is rectilinear segment intersection counting: given a set $S$ of $n$ rectilinear line segments in the plane, determine for each segment the number of other segments in $S$ that intersect it.

Corollary 6.7. Given a set $S$ of $n$ rectilinear line segments in the plane, we can determine for each segment the number of other segments in $S$ that intersect it in $O(\log n)$ time and $O(n)$ space using n processors in the CREW PRAM model.

Proof. Let $U_{1}\left(U_{2}\right)$ be the set of left (right) endpoints of horizontal segments, and fet $d_{1}(p)\left(d_{2}(p)\right)$ denote the number of points in $U_{1}\left(U_{2}\right)$ two-dominated by $p$. For any vertical segment $s$, with upper endpoint $p$ and lower endpoint $q$, the number of horizontal segments that intersect $s$ is $d_{1}(p)-d_{1}(q)+d_{2}(q)-d_{2}(p)$. This is because $d_{1}(p)-d_{1}(q)$ (respectively, $\left.d_{2}(p)-d_{2}(q)\right)$ counts the number of horizontal segments with a left (right) endpoint to the left of $s$ and $y$-coordinate in the interval $[y(q), y(p)]$. Thus, $d_{1}(p)-d_{1}(q)-\left(d_{2}(p)-d_{2}(q)\right)$ counts the number of horizontal segments with left endpoint to the left of $s$, right endpoint to the right of $s$, and $y$-coordinate in the interval $[p(q), p(p)]$ (i.e., the set of horizontal segments that intersect $s$ ).

The final problem we address at is visibility from a point.
6.3. The visibility from a point problem. Given a set of line segments $S=$ $\left\{s_{1}, s_{2}, \cdots, s_{n}\right\}$ in the plane that do not intersect, except possibly at endpoints, and a point $p$, the visibility from a point problem is to determine the part of the plane that is visible from $p$ assuming every $s_{i}$ is opaque. Intuitively, we can think of the point $p$ as a specular light source, the segments as walls, and the problem to determine all the parts of the plane that are illuminated. We can use the cascading divide-and-conquer technique to solve this problem in $O(\log n)$ time and $O(n)$ space using $n$ processors. Without loss of generality, we assume that the point $p$ is at negative infinity below all the segments. The algorithm is essentially the same if $p$ is a finite point, except that the notion of segment endpoints being ordered by $x$-coordinate is replaced by the notion that they are ordered radially around $p$. In other words, it suffices to compute the lower envelope of the $n$ segments to give a method for computing the visibility from a point. For simplicity of expression, we also assume that the $x$-coordinates of the endpoints are distinct.

In the previous two subsections the set of objects consisted of points, but in the visibility problem we are dealing with line segments. The method is slightly different in this case. In this case, we store the segments in the leaves of a binary tree and perform a cascading merge of the $x$-coordinates of intervals of the $x$-axis determined by segment endpoints. We maintain a single label for each interval which represents the segment which is visible from $-\infty$ on that interval. The details follow.

Let $T$ be a complete binary tree with leal nodes $v_{1}, v_{2}, \cdots, v_{\text {, }}$ ordered from lefl 10 right. We associate the segment $s_{i}$ with the leal $v_{i}$ and at $v_{i}$ store the list $U\left(v_{1}\right)=$ $\left(-\infty, p_{1}, p_{2}\right)$, where $p_{1}$ and $p_{2}$ are the two endpoints of $s_{1}$, with $x\left(p_{1}\right)<x\left(p_{2}\right)$, and $-\infty$ is defined such that $x(-\infty)<x(p)$ and $v(-\infty)<\underline{r}(p)$ for all points $p$. We then perform a generalized cascading merge from the teaves of $T$ as in Theorm 2.5, basing comparisons on increasing $x$-coordinates of the points. For each internal node $v$ we let $U(v)$ denote an array of the points stored in the descendants of $v \in T$ sorted by increasing $x$-coordinates. For each point $p_{\mathrm{t}}$ in $U(v)$ we store a label vis $\left(p_{1}, v\right)$ which stores the segment with endpoints in $U(v)$ that is visible in the interval $\left(x\left(p_{1}\right), x\left(\operatorname{succ}\left(p_{i}, v\right)\right)\right)$, where $\operatorname{succ}\left(p_{i}, v\right)$ denotes the successor of $p_{\mathrm{r}}$ in $U(v)$ (based on $x$-coordinates). Initially, the vis labels are only defined for the leaf nodes of $T$. That is, if $U(v)=\left(-\infty, p_{1}, p_{2}\right)$, where $s_{i}=p_{1} p_{2}$, then vis $(-\infty)=+\infty$, vis $\left(p_{1}\right)=s_{t}$, and $\operatorname{vis}\left(p_{2}\right)=+\infty$. We use pred $\left(p_{1}, v\right)$ to denote the predecessor of $p_{i}$ in $U(v)$. As we are performing the cascading merge, we update the vis tabels based on the equation in the following lemma (see Fig. 9).

Lemma 6.8. Let $p_{i}$ be an element of $U(v)$ and let $u=/ c h i l d$ ( $v$ ) and $w=$ rchild ( $v$ ). Then we have the following (if wo segments $s_{i}$ and $s_{j}$ are comparable by the "above"


After merge: $U(v)$


Fig. 9. An example of visibility merging. The dashied segments correspond to the tisible region for $X^{\prime}($ u) and the solid segnents correspond so the visible region for $X(x)$. For simplicity, we store the pointers pred ( $p, u$ ) and pred ( $\rho_{1}, w^{\prime}$ ) in arrars and denote each point $\rho_{1}$ br its mdex. Note that poms are never removed. even if the same segment defines the visible region for many consecume intervals (e.g., $p_{3}$, hrough $p$ ).
relation, then we let min $\left\{s_{1}, s_{1}\right\}$ denote the lower of the two ):

$$
\text { vis }(p, v)= \begin{cases}\min \left\{\operatorname{vis}\left(p_{1}, u\right), \text { vis }\left(\operatorname{pred}\left(p_{1}, w^{\cdot}\right), w^{\prime}\right)\right\} & \text { if } p_{1} \in U(u), \\ \min \left\{\operatorname{vis}\left(\operatorname{pred}\left(p_{1}, u\right), u\right), \text { vis }\left(p_{1}, w^{\prime}\right)\right\} & \text { if } p_{1} \in U(w)\end{cases}
$$

Proof. If we restrict our attention to the segments with an endpoint in $U(n)$, then for any point $\Gamma_{r} \in U(u)$ the segment visible (from $-\infty$ ) on the interval $\left(x\left(p_{i}\right), x(\operatorname{succ}(p, v))\right)$ is the minimum of the segment visible on the interval $\left(x\left(p_{i}\right)\right.$, $\left.x\left(\operatorname{succ}\left(p_{i}, u\right)\right)\right)$ and the segment that is visible on the interval ( $x$ (pred $\left(p_{1}, w^{\prime}\right)$ ), $x\left(\operatorname{succ}\left(\operatorname{pred}\left(p_{i}, w\right), w\right)\right)$. This is because the interval $\left(x\left(p_{i}\right), x\left(\operatorname{succ}\left(p_{i}, v\right)\right)\right)$ is exactly the intersection of the interval $\left(x\left(p_{i}\right), x\left(\operatorname{succ}\left(p_{1}, u\right)\right)\right.$ ) and the interval ( $x$ (pred $\left(p_{i}, w\right)$ ), $x\left(\operatorname{succ}\left(\operatorname{pred}\left(p_{1}, w\right), w\right)\right)$, and there is no segment in $U(v)$ with an endpoint interior to the interval $\left(x\left(p_{1}\right), x\left(\operatorname{succ}\left(p_{1}, c^{\prime}\right)\right)\right.$. Thus, $\operatorname{vis}\left(p_{i}, v\right)$ is equal to minimum of vis $\left(p_{1}, u\right)$ and vis ( $\left.\operatorname{pred}\left(p_{i}, w\right), w\right)$. The case when $p, \in U(v)$ is similar.

By Lemma 6.8, after merging the lists $U(u)$ and $U(w)$, we can determine the labels for all the points in $U(v)$ in $O(1)$ additional time using $|U(v)|$ processors. Thus, the running time of this generalized cascading merge algorithm is still $O$ ( $\log n$ ) using $n$ processors. After we complete the merge and have computed $U(r o o t(T))$, along with all the vis labels for the points in $U(\operatorname{roof}(T))$, then we can compress out duplicate entries in the list ( $\operatorname{vis}\left(p_{1}, \operatorname{root}(T)\right)$, $\operatorname{vis}\left(p_{2}, \operatorname{root}(T)\right), \cdots, \operatorname{vis}\left(p_{2 n}, \operatorname{root}(T)\right)$ ) using a paraliel prefix computationt to construct a compact representation o $\Gamma$ the visible portion of the plane. We summarize in the following theorem.

Theorem 6.9. Given a set $S$ of $n$ nonimersecting segments in the plane, we can find the lower envelope of $S$ in $O(\log n)$ time and $O(n)$ space using n processors in the CREW PRAM model, and this is optimal.

Proof. The correctness and complexity bounds follow from the discussion above. Since we require that the points in the description of the lower envelope be given by increasing $x$-coordinates, we can reduçe sorting to this problem, and thus can do no better than $O(\log n)$ time using $n$ processors.
7. EREW PRAM implementations. In this section we briefly note that the same techniques as employed by Cole in [13] to implement his merging procedure in the EREW PRAM model (no simultaneous reads) can be applied to our algorithms for generalized merging, fractional cascading, constructing the plane-sweep tree. threedimensional maxima, two-set dominance counting, and visibility from a point, resulting in EREW PRAM algorithms for these problems. Apparently, we cannot apply his techniques to our algorithms for trapezoidal decomposition and segment intersection detection, however. since our algorithms for these problems explicitly use concurrent reads (in the multilocation steps).

Applying his techniques to our algorithms results in EREW PRAM algorithms with the same asymptotic bounds as the ones presented in this paper, except that the space bounds for the problems addressed in $\S 6$ all become $O(n \log n)$. The reason that his tectniques increase the space complexity of these problems is because of our use of labeling functions. Specifically, it is not clear how to perform the merges on-line and still updace the labels in $O(1)$ lime after a node becomes full. This is because a label whose value changes on level $/$ may have to be broadcasted to many elements in level $l-1$ to update their labels, which would require $\Omega(\log n)$ time in this model if there were $O(n)$ such elements.

We can get around the problem arising from the labeling functions, however. For the three-dimensional maxima problem and the two-set dominance counting problem,
we separate the computation of the $U\left(t^{\prime}\right)$ lists and computation of the labeiing functions into two separate steps, rather than "dovetailing" the two computations as before. Each of the labeling functions we used for these two problems can be redefined so as to be EREW-computabie. Specificatly, the tabel for an element $p$ in $U(w)$, on level $l$, can be expressed in terms of a label $\operatorname{prcf}(p, t)$ and a label $u p(p, v)$, where pref( $p, v)$ can be computed by performing a parallel prefix computation [21], [22] in $U(v)$ and $u p(p, v)$ can be defined in terms of pref(pred ( $p$, lchild ( $b$ ) ). Ichild ( $a$ )), pref ( $p r e d$ ( $p$, refild $(v)$ ), rchild $(b)$ ), and the up label $p$ had on level $1+1$ (say, in $U(r c h i l d(v))$ if $p \in U($ rchild ( $v)$ )). In particular, for the three-dimensional maxima problem $\operatorname{pref}(p, v)=\operatorname{sod}(p, v)$ and $u p(p, v)=z t d(p, v)$, and for the two-set dominance counting problem pref $(p, v)=\operatorname{nod}(p, v)$ and $u p(p, v)=n d(p, v)$. We can compute all the pref( $n, v)$ labels in $O(\log n)$ time using $n$ processors by assigning $\{|U(v)| / \log n\rceil$ processors to each node $\mathfrak{c}$ [21]. We can then broadcast each pref( $p, v$ ) label to the successor of $v$ in sibling $(b)$. which takes $O(\log n)$ time using $n$ processors by assigning $\lceil|U(v)| / \log n\rceil$ processors to each node $v$. Finally, we can compute all the up $(p, v)$ labels in $O(\log n)$ additional time by assigning a single processor to each point $p$ and tracing the path in the tree from the leal node that contains $p$ up to the root. This is an EREW operation because computing all the $u p(p, v)$ labels only depends upon accessing memory locations associated with the point $p$.

The EREW solution to the visibility from a point problem requires $O(n \log n)$ space for a different reason, namely, because we can solve it by constructing the plane-sweep tree for the segments (we need not have the Cover (u)'s in sorted order, howeverl, computing the lowest segment in each Cover( $v$ ), and then performing a top-down parallel min-finding computation to find the segment visible on each interval ( $p_{i}, p_{i 1}$ ). Since these are all straightforward computations, given the discussion presented earlier in this paper, we leave the details to the reader.
8. Conclusion. In this paper we gave several general techniques for solving problems efficiently using parallel divide-and-conquer. Our techniques are based on nontrivial generalizations of the merge-sorting approach of Cole [13]. It is interesting to note that Cole's algorithm improved the previous results by a constant factor, whereas our algorithms improve the previous results asymptotically.

Two of our techniques involved methods for performing fractional cascading and a generalized version of the merge-sorting problem optimally in parallel. Our method for doing fractional cascading runs in $O(\log n)$ time using $\lceil n / \log n\rceil$ processors, and, if implemented as a sequential algorithm, results in a sequential alternative to the method of Chazelle and Guibas [12] for Sractional cascading.

We also showed how to apply the generalized merging procedure and fractional cascading to efficiently solve several problems by "cascading" the divide-and-conquer paradigm. For three of the problems-trapezoidal decomposition, planar point location, and segment intersection detection-the method involved merging in the line segment partial order, and required considerable care to avoid situations in which the algorithm would halt because it attempted to compare two incomparable segments. All three of these algorithms ran in $O(\log n)$ time using $n$ processors, which is optimal for all but the point location problem. In addition, since our algorithm for doing planar point location results in a query time of $O(\log n)$, our result immediately implies an $O\left(\log ^{2} n\right)$ time, $n$ processor solution to the problem of constructing the Voronoi diagram of $n$ planar points, using the algorithm of Aggarwal et al. [1].

We showed how to apply the cascading divide-and-conquer technique to problems that can be solved by merging with labeling functions. We used this approach to solve
the three-dimensional maxima problem, the two-set dominance counting problem, the rectilinear segment intersection counting problem, and the visibility from a point problem. Our algorithms for these problems all ran in $O(\log n)$ time using $n$ processors, which is optimal.

## REFERENCES

[1] A. Aggarwal, B. Chazelle, l. Guibas, C. Ó Dúnlaing. and C. Y'ap. Parallel computational geontetry: Proc. 26th JEEE Symposium on Foundations of Computer Science. 1985, np. 468-477.
[2] ——Parallel computational geometry, Algorithmical. 3 (1988), pp. 293-328.
[3] M. J. Atallah, R, Cole, and M. T. Goodrich. Cascading divide-and-conquer: A iechniquc for designing parallef algorithms. Proc. 2sth IEEE Symposium on Foundations of Computer Science, 1987. pp. 151-160.
[4] M. J. Atallah and M. T. Goodrich, Efficient parallel sohuions to some geonerric problems, J. Paralled and Distributed Computing. 3 (1986). pp. 492-507.
[5] —. Efficient plane swecpug in parallel. Proc. 2nd ACM Symposium on Computational Geometry, 1986, pp. 216-225.
[6] ——. Parallel algorithms for some functions of wo conecx polygons. Algorithmica. 3 (19881, pp. 535-548.
[7] M. Ben-Or, Lower bounds for algebraic computasion irees, Proc. 15th ACM Symposium on Theory of Computing. 1983, pp. 80-86.
[8] J. L. BENTLEY AND D. WOOD, An opimal worst case algorithm for repornng mersections of rectangles, IEEE Trans. Comput. C-29 (1980), pp. 571-576.
[9] G. Bilarid and A. Nicolau, ddaptive bironic sorting: An optimal parallel algortlim for shared memory machines. TR 86-i69. Deparment of Computer Science, Corneli University, August 1986.
[10] A. Boroutn and J. E. Hopcroft. Ruuring, merging, and soring on parallel modets of computation, J. Comput. System Sci., 30 (1985). pp. 130-145.
[11] R. P. BRENT, The parallel evaluation of general arithmetic expressions, J. ACM, 21 (1974), pp. 201-206.
[12] B. Chazelle and L. J. Guibas. Fracrional cascading: I. A data strucruring technique. Algorithmica, I (1986). pp. 133-162.
[13] R. Cole, Parallel merge sort, Proc. 27th IEEE Symposium on Foundations of Computer Science. 1986, pp. 51I-516; SIAM J. Compus. 17 (1988). pp. 770-785.
[14] N. DADOUN ANDID. Kirkpatrick, Parallel pracessing for efficient subditision search. Proc. 3rd ACM Symposium on Computational Geometry, 1987, pp. 205-214.
[15] H. Edelshrunner and M. H. Overmars. On the equivalence of some reciangle problems. Inform, Process. Lett., 14, (1982), pp. 124-127.
[16] H. ElGindy and M..T. Goodrich, Parallet algorithms for shoriesi parh problems in polygons, The Visual Computer: Internat. J. Comput. Graphics, 3 (1988), pp. 371-378.
[17] M. T. Goodrich, Efficient parallel techniques for compuational geometr: Ph.D. thesis. Department of Computer Science. Purdue University, W. Lafayctte, IN. 1987.
[18] -. Finding the convex hull of a sorted point ser in parallel, Inform. Process. Lett., 26 (1987). pp. 173-179.
[19] -, Triangulating a polygon in parallel. J. Algorithms, 10 appear.
[10] H. T. Kung, F. luccio. and F. P. Preparata, On finding the maxima of a ser of vectors. J. ACM, 22 (1975), pp. 469-476.
[21] C. P. Kruskal, L. Runolph, and M. SNir, The Power of Parallel Prefix, Proc. 1985 iEEE Intemat. Conference on Parallel Processing, St. Charles, IL, 1985. pp. 180-185.
[22] R. E. Ladner and M. J. Fischer. Parallel prefix computaion, J. ACM (1980), pp. 831-838.
[23] J. H. ReIF, An optimal parallel algoridlm for integer soring, Proc. 26th IEEE Symposium on Foundations of Computer Science. 1985, pp. 496-504.
[24] 1. H. Reif and S. SEN, Optimal Randomized Parallel Algorithms for Computational Geometṛ; Proc. 1987 IEEE Intermat. Conference on Parallel Processing, 1987, pp. 270-277.
[25] Y. Shiloach and U. Vishkin, Finding the maximum merging, and sorting in a parallel computation model, J. Algorithms. 2 (1981), pp. 88-102.
[26] L. Valiant, Parallelism in comparison prohlems. SIAM J. Comput., 4 (1975), pp. 348-355.
[27] H. Wagener. Optimally parallel algorithms for concex hudl determination. manuscript. 1985.
[28] C. K. Yap, Parallel triangulation of a polygon in rwo calls to the trapezoidal map, Algorithmica, 3 (1988). pp. 279-288.


[^0]:    Atallah, Mikhail J.; Cole, Richard; and Goodrich, Michael T., "Cascading Divide-and-Conquer: A Technique for Designing Parallel Algorithms" (1987). Department of Computer Science Technical Reports. Paper 576.
    https://docs.lib.purdue.edu/cstech/576

[^1]:    * Received by the editors September 14, 1987: accepted for publication (in revised form) August 12, 1988. This paper appeared in preliminary form as [3] and as porions of [17].
    ; Department of Computer Science, Purdue University, West Lafayette. Indiana 47907. The research of this author was suppored by the Office of Naval Research under grants N00014.84-K-0502 and N00014.86-K-0689, and by the National Science Foundation under grant DCR-84-51393, with matching funds from AT\&T.
    $\doteqdot$ Courant Institute. New York University, New York, New York 10012. The research of this author was supported in part by National Science Foundation grants DCR.84-01633 and CCR-8702271, and by an Ofice of Naval Research grant N00014-85-K-0046.
    § Deparment of Computer Science. Johns Hopkins University, Baltimore. Maryland 21218. The research of this author was supported by Office of Naval Research grants N00014-84.K-0502 and N000-86.K-0689, National Science Foundation grants DCR-84-51393, with matching funds from AT\&T. and CCR-88-10568. and a David Rass grant from the Purdue Research Foundation.

