# I/O-Efficient Well-Separated Pair Decomposition and its Applications ${ }^{\star}$ 

(Extended Abstract)

Sathish Govindarajan ${ }^{1}$, Tamás Lukovszki ${ }^{2}$, Anil Maheshwari ${ }^{3}$, and Norbert Zeh ${ }^{3}$<br>${ }^{1}$ Duke University, gsat@cs.duke.edu<br>2 University of Paderborn, Heinz-Nixdorf-Institut, tamas@hni. upb. de<br>${ }^{3}$ Carleton University, \{maheshwa, nzeh\}@scs.carleton.ca


#### Abstract

We present an external memory algorithm to compute a well-separated pair decomposition (WSPD) of a given point set $P$ in $\mathbb{R}^{d}$ in $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory, where $N$ is the number of points in $P$, and $\operatorname{sort}(N)$ denotes the I/O complexity of sorting $N$ items. (Throughout this paper we assume that the dimension $d$ is fixed). We also show how to dynamically maintain the WSPD in $O\left(\log _{B} N\right)$ I/O's per insert or delete operation using $O(N / B)$ blocks of external memory. As applications of the WSPD, we show how to compute a linear size $t$-spanner for $P$ within the same I/O and space bounds and how to solve the $K$-nearest neighbor and $K$-closest pair problems in $O(\operatorname{sort}(K N))$ and $O(\operatorname{sort}(N+K))$ I/Os using $O(K N / B)$ and $O((N+K) / B)$ blocks of external memory, respectively. Using the dynamic WSPD, we show how to dynamically maintain the closest pair of $P$ in $O\left(\log _{B} N\right)$ I/O's per insert or delete operation using $O(N / B)$ blocks of external memory.


## 1 Introduction

Many geometric applications require computations involving the set of all distinct pairs of points (and their distances) from a set $P$ of $N$ points in $d$-dimensional Euclidean space (e.g. nearest neighbor for each point). Voronoi diagrams and multi-dimensional divide and conquer are the traditional techniques used for solving several distance based geometric problems, especially in two and three dimensions. But in $d$ dimensions, the worst case size of Voronoi diagrams can be $\Omega\left(N^{\lfloor d / 2\rfloor}\right)$, and divide and conquer will require an exponent in the polylog factor depending on the dimension. Callahan and Kosaraju [6] introduced the WSPD data structure to cope with higher dimensional geometric problems. It consists of a binary tree $T$ whose leaves are the points in $P$, with internal nodes representing the subsets of points within the subtree, and a list of "well-separated" pairs of subsets of $P$, each of which is a node in $T$. Intuitively a pair $\{A, B\}$ is well separated, if the distance between $A$ and $B$ is significantly greater than the distance between any two points within $A$ or $B$. It turns out that for many problems it is sufficient to perform only a constant number of operations on pair $\{A, B\}$ instead of performing $|A||B|$ separate operations on the corresponding pairs of points. Moreover for fixed $d$, a WSPD of $O(N)$ pairs of subsets can be constructed. This has resulted in fast sequential, parallel,

[^0]and dynamic algorithms for a number of problems on point-sets. Here we extend these results to external memory.

Previous Work In the Parallel Disk Model (PDM) [16], there is an external memory consisting of $D$ disks attached to a machine with internal memory size $M$. Each of the $D$ disks is divided into blocks of $B$ consecutive data items. Up to $D$ blocks, at most one per disk, can be transferred between internal and external memory in a single I/O-operation. The complexity of an algorithm in the PDM model is the number of I/O operations it performs. Relevant research work in the external memory (EM) setting can be found in the survey of Vitter [15]. In the PDM model it has been shown that sorting an array of size $N$ takes $\operatorname{sort}(N)=\Theta\left((N / D B) \log _{(M / B)}(N / B)\right)$ I/Os [16, 15]. Scanning an array of size $N$ takes $\operatorname{scan}(N)=\Theta(N / D B)$ I/Os.

For the geometric problems discussed in this paper, the best resources will be [10] for WSPD and its applications and [14] for results on proximity problems. We omit the discussion on the state of the art, importance and significance of these problems here due to the lack of space, and refer the reader to these references.

New Results In this paper we present external memory algorithms for the following problems for a set $P$ consisting of $N$ points in $d$-dimensional Euclidean space: In Sections 3 and 4, we present an algorithm to compute the WSPD data structure for $P$; it requires $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory. In Section 5, we present an algorithm to dynamically maintain the WSPD in $O\left(\log _{B} N\right)$ I/O's per insert/delete operation using $O(N / B)$ blocks of external memory. In Section 6.1, we present an algorithm to compute a linear size spanner for $P$; it requires $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory. Choosing spanner edges carefully, we can guarantee a diameter of $2 \log N$ for the constructed spanner. We also show that $\Omega(\min \{N, \operatorname{sort}(N)\})$ I/Os are required to compute any $t$-spanner of a given point set, for any $t>1$. In Section 6.2, we present an algorithm to compute $K$-nearest neighbors, i.e. for each point $a \in P$, compute the $K$-nearest points to $a$ in $P-\{a\}$; this takes $O(\operatorname{sort}(K N))$ I/Os using $O(K N / B)$ blocks of external memory. In Section 6.3, we present an algorithm to compute $K$-closest pairs, i.e. report the $K$ smallest interpoint distances in $P$; this takes $O(\operatorname{sort}(N+K))$ I/Os using $O((N+K) / B)$ blocks of external memory. In Section 6.4, we present an algorithm to dynamically maintain the closest pair of $P$ in $O\left(\log _{B} N\right)$ I/O's per insert or delete operation using $O(N / B)$ blocks of external memory.

In [5] an $O\left(\log _{B} N\right)$ algorithm for the dynamic closest pair problem in higher dimensions is given. Their approach uses the Topology B-tree data structure. For the remaining problems, our results are the only efficient external memory algorithms known in higher dimensions. For the $K$-nearest neighbor and $K$-closest pair problems, optimal algorithms were presented in [12] for the case where $d=2$ and $K=1$. In [2] it is shown that computing the closest pair of a point set requires $\Omega(\operatorname{sort}(N))$ I/Os, which implies the same lower bound for the general problems we consider in this paper. I/O-efficient construction of fault-tolerant spanners and bounded degree spanners for point sets in the plane and of spanners for polygonal obstacles in the plane has been discussed in [13].

## 2 Preliminaries

For a given point set $P$, the bounding rectangle $R(P)$ is the smallest rectangle containing all points in $P$, where a rectangle is the Cartesian product $\left[x_{1}, x_{1}^{\prime}\right] \times\left[x_{2}, x_{2}^{\prime}\right] \times \cdots \times$ $\left[x_{d}, x_{d}^{\prime}\right]$ of a set of closed intervals. The length of $R$ in dimension $i$ is $l_{i}(R)=x_{i}^{\prime}-x_{i}$. The maximum and minimum lengths of $R$ are $l_{\max }(R)=\max \left\{l_{i}(R): 1 \leq i \leq d\right\}$ and $l_{\text {min }}(R)=\min \left\{l_{i}(R): 1 \leq i \leq d\right\}$. When all lengths of $R$ are equal, $R$ is a cube; we denote it's side length by $l(R)=l_{\max }(R)=l_{\min }(R)$. Let $i_{\max }(R)$ be the dimension such that $l_{i_{\max }}(R)=l_{\max }(R)$. For a point set $P$, let $l_{i}(P)=l_{i}(R(P)), l_{\max }(P)=l_{\max }(R(P))$, $l_{\min }(P)=l_{\min }(R(P))$, and $i_{\max }(P)=i_{\max }(R(P))$. Let $d(x, y)$ denote the Euclidean distance between points $x$ and $y$. For point sets $X$ and $Y$, let $d(X, Y)=\min \{d(x, y): x \in$ $X \wedge y \in Y\}$. Given a separation constant $s$, we say that two point sets $A$ and $B$ are wellseparated if $R(A)$ and $R(B)$ can be enclosed in two $d$-balls of radius $r$ such that the distance between the two balls is at least $s r$. We define the interaction product $A \otimes B$ of two point sets $A$ and $B$ as $A \otimes B=\{\{a, b\}: a \in A \wedge b \in B \wedge a \neq b\}$. A well-separated realization of $A \otimes B$ is a set $\left\{\left\{A_{1}, B_{1}\right\}, \ldots,\left\{A_{k}, B_{k}\right\}\right\}$ with the following properties [9]:
(R1) $A_{i} \subseteq A$ and $B_{i} \subseteq B$, for $1 \leq i \leq k$,
(R2) $A_{i} \cap B_{i}=\emptyset$, for $1 \leq i \leq k$,
(R3) $\left(A_{i} \otimes B_{i}\right) \cap\left(A_{j} \otimes B_{j}\right)=0$, for $1 \leq i<j \leq k$,
(R4) $A \otimes B=\bigcup_{i=1}^{k} A_{i} \otimes B_{i}$.
(R5) $A_{i}$ and $B_{i}$ are well-separated, for $1 \leq i \leq k$.
Intuitively, this means that for every pair $\{a, b\}$ of points $a \in A$ and $b \in B$, there is a unique pair $\left\{A_{i}, B_{i}\right\}$ such that $a \in A_{i}$ and $b \in B_{i}$ and that for any pair $\left\{A_{i}, B_{i}\right\}$ the distance between the points in one of the sets is small compared to the distance between the sets. We can use a binary tree $T$ to define a partition of a point set $P$ into subsets. In particular, there is a leaf in $T$ for each point in $P$. An internal node represents the set of points associated with its descendant leaves. We refer to a node representing a subset $A \subseteq P$ as node $A$. A leaf representing point $a \in P$ is referred to as leaf $a$. The parent of a node $A$ in $T$ is denoted by $p(A)$. We say that a realization of $A \otimes B$ uses a tree $T$ if all sets $A_{i}$ and $B_{i}$ in the realization are nodes in $T$. A well-separated pair decomposition of a point set $P$ is a binary tree $T$ associated with $P$ and a well-separated realization of $P \otimes P$ that uses $T$. A split of a point set $P$ is a partition of $P$ into two non-empty point sets lying on either side of a hyperplane perpendicular to one of the coordinate axes, and not intersecting any points in $P$. A split tree $T$ of $P$ is a binary tree whose nodes are associated with subsets of $P$, defined as follows: If $P=\{x\}, T$ contains a single node $x$. Otherwise, we use a split to partition $P$ into two subsets $P_{1}$ and $P_{2} ; T$ consists of two split trees for point sets $P_{1}$ and $P_{2}$ whose roots are the children of the root node $P$ of $T$. For a node $A$ in $T$, we define the outer rectangle $\hat{R}(A)$ as follows: For the root $P$, let $\hat{R}(P)$ be an open $d$-cube centered at the center of $R(P)$ with $l_{i}(\hat{R}(P))=l_{\max }(P)$. For all other nodes $A$, the hyperplane used for the split of $p(A)$ divides $\hat{R}(p(A))$ into two open rectangles. Let $\hat{R}(A)$ be the one that contains $A$. A fair split of $A$ is a split of $A$ where the hyperplane splitting $A$ is at distance at least $l_{\max }(A) / 3$ from each of the two boundaries of $\hat{R}(A)$ parallel to it. A split tree formed using only fair splits is called a fair split tree. A partial fair split tree is defined in the same way; but the leaves may represent point sets instead of single points.

## 3 Constructing a WSPD

In this section we assume that we are given a fair split tree $T=(V, E)$ for a point set $P$. Also, let every interior node $A$ of $T$ be labeled with its bounding rectangle $R(A)$. Every leaf is labeled with the point it represents. Given this input, we show how to find a wellseparated pair decomposition of $P$ I/O-efficiently. In the output, every pair $\left\{A_{i}, B_{i}\right\}$ in the WSPD will be represented by the IDs of the two nodes $A_{i}$ and $B_{i}$ in $T$.

Our approach is based on [9], which we sketch next. Given a postorder traversal of $T$, denote the postorder number of a node $A$ in $T$ by $\eta(A)$. Define an ordering $\prec$ on the nodes of $T$ as $A \prec B$ if and only if either $l_{\max }(A)<l_{\max }(B)$ or $l_{\max }(A)=l_{\max }(B)$ and $\eta(A)<\eta(B)$. Given the sequence of nodes in $T$ sorted by $\succ$, define $v(A)$ to be the position of node $A$ in this sequence.

Denote the internal memory algorithm of [9] to construct a well-separated realization of $P \otimes P$ by Computewsr. It takes the tree $T$ as an input. For every internal node $X$ of $T$, let $A$ and $B$ be the two children of $X$. Then ComputeWSR calls a procedure FIndPairs with argument $(A, B)$. Given a pair $\left(A^{\prime}, B^{\prime}\right)$ as an input, procedure FindPairs does the following: First it ensures that $B^{\prime} \prec A^{\prime}$, by swapping $A^{\prime}$ and $B^{\prime}$ if necessary. If $A^{\prime}$ and $B^{\prime}$ are well-separated, the pair $\left\{A^{\prime}, B^{\prime}\right\}$ is added to the well-separated realization. Otherwise, let $A_{1}^{\prime}$ and $A_{2}^{\prime}$ be the children of $A^{\prime}$ in $T$. The procedure invokes itself recursively with pairs $\left(A_{1}^{\prime}, B^{\prime}\right)$ and $\left(A_{2}^{\prime}, B^{\prime}\right)$ as arguments.

The recursive invocation pattern of procedure FindPairs can be described using the concept of computation trees. The root of such a tree $T^{\prime}$ is a pair $\{A, B\}$ such that $A$ and $B$ are the children of an internal node $X$ of $T$. A node $\left\{A^{\prime}, B^{\prime}\right\}$ in $T^{\prime}$ is a leaf if $A^{\prime}$ and $B^{\prime}$ are well-separated. Otherwise, let $B^{\prime} \prec A^{\prime}$ and $A_{1}^{\prime}$ and $A_{2}^{\prime}$ be the two children of $A^{\prime}$; then node $\left\{A^{\prime}, B^{\prime}\right\}$ has two children $\left\{A_{1}^{\prime}, B^{\prime}\right\}$ and $\left\{A_{2}^{\prime}, B^{\prime}\right\}$ in $T^{\prime}$. Tree $T^{\prime}$ represents the recursive invocations of procedure FINDPAIRS made to compute a realization of $A \otimes B$. Every leaf of $T^{\prime}$ corresponds to a pair $\left\{A_{i}, B_{i}\right\}$ in the well-separated realization of $P \otimes P$. In [9] it is shown that there are only $O(N)$ pairs in this realization, which implies that the total size of all computation trees is $O(N)$.

We simulate COMPUTEWSR in external memory by applying time-forward processing [11] to the forest of computation trees. The difficulty is that this forest is not known in advance, and we have to generate it on the fly while processing it.

We use the split tree $T$ to guide the processing of $F$. Assume that we are given the forest $F=\left(V_{F}, E_{F}\right)$ of computation trees and the fair split tree $T=(V, E)$. We sort the nodes of $T$ by the $\succ$-relation and define a mapping $\mu: V_{F} \rightarrow V$ as $\mu(\{A, B\})=A$, assuming w.l.o.g. that $B \prec A$. This mapping naturally defines an edge set $\mu\left(E_{F}\right)$ between the vertices in $V$ as $\mu\left(E_{F}\right)=\left\{(\mu(v), \mu(w)):(v, w) \in E_{F}\right\}$. Note that $\mu\left(E_{F}\right)$ may be a multiset. Given this mapping $\mu$ we process $F$ by applying time-forward processing to the multigraph $\mu(F)$. In particular, we process the nodes of $\mu(F)$ sorted by $\succ$, which guarantees that we process the source vertex of every edge before the target vertex. Indeed, for any edge $\left(A, A^{\prime}\right)$ in $\mu\left(E_{F}\right)$, either $\left\{A, A^{\prime}\right\}$ is a vertex in $F$ with $A^{\prime} \prec A$ or $A^{\prime}$ is a child of $A$, which also implies that $A^{\prime} \prec A$.

As $F$ is not given beforehand, we have to show how to generate the edges of $\mu(F)$ while processing $\mu(F)$. For every node $A \in V$ with children $A_{1}$ and $A_{2}$, we store a label $\lambda(A)$ consisting of $v\left(A_{1}\right), v\left(A_{2}\right), R\left(A_{1}\right)$, and $R\left(A_{2}\right)$. Given an edge $\left(\{A, B\},\left\{A^{\prime}, B\right\}\right) \in$ $E_{F}$, where $A^{\prime}$ is a child of $A$ in $T$, the information $\phi(\mu(\{A, B\}))$ sent along edge
$\left(\mu(\{A, B\}), \mu\left(\left\{A^{\prime}, B\right\}\right)\right)$ consists of $v\left(A^{\prime}\right), v(B), R\left(A^{\prime}\right)$, and $R(B)$. Now consider a node $A \in V$ receiving some tuple describing pair $\{A, B\}$ as input. We have to show how to generate the two children $\left\{A_{1}, B\right\}$ and $\left\{A_{2}, B\right\}$ of $\{A, B\}$ in $F$ (if any) as well as the two edges $\left(A, \mu\left(\left\{A_{1}, B\right\}\right)\right)$ and $\left(A, \mu\left(\left\{A_{2}, B\right\}\right)\right)$ from this input and the label $\lambda(A)$. From the received information, decide whether $A$ and $B$ are well-separated. If they are, add $\{A, B\}$ to the well-separated realization, and proceed to the next input received by $A$. Otherwise, as the input has been sent to $A, B \prec A$. Hence, $\left\{A_{1}, B\right\}$ and $\left\{A_{2}, B\right\}$ are indeed the two children of $\{A, B\}$ in $F$. Generate the description of these two pairs from the received information about $B$ and the information about $A_{1}$ and $A_{2}$ stored in $\lambda(A)$. Compare $v\left(A_{1}\right)$ and $v\left(A_{2}\right)$ with $v(B)$ to compute $\mu\left(\left\{A_{1}, B\right\}\right)$ and $\mu\left(\left\{A_{2}, B\right\}\right)$. Queue $\left\{A_{1}, B\right\}$ with priority $v\left(\mu\left(\left\{A_{1}, B\right\}\right)\right)$ and $\left\{A_{2}, B\right\}$ with priority $v\left(\mu\left(\left\{A_{2}, B\right\}\right)\right)$.

To bound the I/O-complexity of this procedure, observe that we queue a constant amount of information per edge in $F$. As there are $O(N)$ edges in $F$, applying known results $[11,1]$ about time-forward processing, we obtain the following lemma.

Lemma 1. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$, a fair split tree $T$ of $P$ with $O(N)$ nodes and a separation constant $s>0$, a well-separated realization of $P$ can be computed in $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory. Together with $T$ this gives a well-separated pair decomposition of $P$.

## 4 Constructing a Fair Split Tree

Our algorithm to construct a fair split tree $T$ for a given point set $P$ follows the framework of the optimal PRAM algorithm in [4]. ${ }^{1}$ The idea is to construct $T$ recursively. First we construct a partial fair split tree $T^{\prime}$ whose leaves have size $O\left(N^{\alpha}\right)$ for some constant $1-\frac{1}{6 d} \leq \alpha<1$. Then we recursively build split trees for the leaves, proceeding with an optimal internal memory algorithm for every leaf whose size is at most $M$. We will show how to implement one such recursive step in $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory, which leads to the following result.
Lemma 2. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$, a fair split tree for $P$ can be computed in $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory.

From Lemmas 1 and 2 we obtain the following result.
Theorem 1. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$ and a separation constant $s>0$, a wellseparated pair decomposition for $P$ can be computed in $O(\operatorname{sort}(N))$ I/Os using $O(N / B)$ blocks of external memory.

To construct a partial fair split tree $T^{\prime}$, we first construct a compressed version $T_{c}$ as in [4]. First each dimension is partitioned into slabs containing $N^{\alpha}$ points each. Every rectangle $R$ associated with a node in $T_{c}$ satisfies the following three invariants: (1) In each dimension at least one side of $R$ lies on a slab boundary; (2) If $R^{\prime}$ is the largest rectangle contained in $R$ all of whose sides fall onto slab boundaries, then either $l_{i}\left(R^{\prime}\right)=$ $l_{i}(R)$ or $l_{i}\left(R^{\prime}\right) \leq \frac{2}{3} l_{i}(R) ;(3) l_{\min }(R) \geq \frac{1}{3} l_{\max }(R)$. Now if we want to split rectangle $R$ associated with an internal node $v$ into smaller rectangles, we split $R$ along its longest

[^1]side according to the following three cases: (1) $l_{\max }(R)=l_{i_{\max }}\left(R^{\prime}\right)$ : We find the slab boundary that comes closest to splitting $R$ into equal-sized pieces. If this slab boundary is at distance at least $\frac{1}{3} l_{\text {max }}(R)$ from either side of $R$, we split $R$ along this slab boundary. Otherwise, we split $R$ into two equal-sized pieces. This case produces the two resulting rectangles $R_{1}$ and $R_{2}$. (2) $l_{\max }\left(R^{\prime}\right) \geq \frac{8}{81} l_{\max }(R)$ : If $l_{i_{\max }}\left(R^{\prime}\right) \geq \frac{1}{3} l_{\max }(R)$, we split $R$ along the side of $R^{\prime}$ that is not shared with $R$. Otherwise, we find the unique integer $j$ such that $y=\frac{2}{3}\left(\frac{4}{3}\right)^{j} l_{\max }\left(R^{\prime}\right)$ lies in the interval $\left(\frac{1}{2}, \frac{2}{3}\right] l_{\max }(R)$ and split along a hyperplane that is at distance $y$ from the slab boundary shared by $R^{\prime}$ and $R$. This case produces the rectangle containing $R^{\prime}$. The other rectangle is being ignored for the time being; we reattach it later. (3) $l_{\max }\left(R^{\prime}\right)<\frac{8}{81} l_{\max }(R)$ : In this case, $R$ and $R^{\prime}$ share a unique corner. We construct a $d$-cube sharing the same corner with $R$ and with side length $\frac{3}{2} l_{\max }\left(R^{\prime}\right)$. This case produces the cube $C$. Later we have to construct a sequence of fair splits cutting $R$ down to $C$.

The construction of $T_{c}$ takes $O(\operatorname{sort}(N))$ I/Os using standard external-memory graph techniques. The reattachment of the missing Case 2 leaves takes sorting and scanning. Next we describe how to expand compressed edges corresponding to a Case 3 split.

Such a Case 3 edge corresponds to a rectangle $R$ that has been "shrunk" to a $d$ cube $C$ with the following properties: ( P 1 ) $C$ and $R$ share exactly one corner, and $C$ is contained in $R$. (P2) Any rectangle contained in $R$ and not intersecting $C$ contains at most $O\left(N^{\alpha}\right)$ points. (P3) $l(C)<\frac{8}{81} l_{\max }(R)$ and $l_{\min }(R) \geq \frac{1}{3} l_{\max }(R)$.

Properties P1 and P3 guarantee that there exists a sequence of fair splits that produces $C$ from $R$. Property P2 guarantees that for every split in this sequence, the rectangle not containing $C$ can be made a leaf of $T^{\prime}$. We show how to construct these sequences of fair splits for all Case 3 edges in $O(\operatorname{sort}(N))$ I/Os. We also have to assign the points in $P$ to the containing leaves of $T^{\prime}$, as this information is needed for the next level of recursion in the fair split tree construction algorithm.

Every point of $P$ is contained in some leaf rectangle or in some region $R \backslash C$, where $(R, C)$ is a compressed edge produced by Case 3. We compute an assignment of the points in $P$ to these regions similarly to the corresponding step in the optimal PRAMalgorithm in [4]. This takes a constant number of sorting and scanning steps.

To replace each edge corresponding to Case 3 by a sequence of fair splits, we simulate one phase of the sequential algorithm of [9]. We sort the points in $R \backslash C$ in each dimension, producing a sorted list $L_{i}^{R}$ of points for each dimension $i$. Consider the current rectangle $R$ which we want to split in dimension $i_{\max }$, producing rectangles $R_{1}$ and $R_{2}$, where $R_{1}$ contains the cube $C$. If $l_{\max }(R)>3 l(C)$, split $R$ in half. Otherwise, choose a hyperplane containing a side of $C$ for the split. We use list $L_{i_{\max }}^{R}$ to decide whether there are any points in $R_{2}$. If so, we perform a split producing two new nodes $R_{1}$ and $R_{2}$ of $T^{\prime \prime}$ and assigning the points in $R_{2}$ to the leaf $R_{2}$. Otherwise, shrink $R$ to $R_{1}$, not producing any new nodes. Shrinking $R$ does not cost any I/Os as the coordinates of $R$ are maintained in internal memory. This approach ensures that a linear number of splits are performed in the number of points in $R \backslash C$. Thus, in total only $O(N)$ splits for all edges $(R, C)$ to be expanded are performed. We have to show how to maintain the lists $L_{i}^{R}$ representing the set of points in the current rectangle $R$.

Assume that we split $R$ in dimension $i$ and that this produces two new rectangles $R_{1}$ and $R_{2} . R_{1}$ contains the cube $C$ and $R_{2}$ does not. We scan $L_{i}^{R}$ from the tail until we
find the first point that is in $R_{1}$. All points after this position in $L_{i}^{R}$ are not in $R_{1}$ and are therefore removed and put into the point list of leaf $R_{2}$. As $R_{1}$ is going to act as rectangle $R$ for the next split, we need to delete all points in $R_{2}$ from all lists $L_{j}^{R}$. The above scan takes care of deleting these points from $L_{i}^{R}$. However, we cannot afford to delete these points from all the other lists $L_{j}^{R}, j \neq i$, because this would cause many random accesses or would require a scan over all these lists per split. Instead we delete points lazily. When splitting the rectangle $R_{1}$ in dimension $j$, which produces rectangles $R_{3}$ and $R_{4}$, where $R_{3}$ contains $C$, we scan $L_{j}^{R}$ from the tail until we find the first point in $R_{3}$. For each point visited by this scan we test whether it is in $R_{4}$. If it is, it is appended to the point set of leaf $R_{4}$. Otherwise, as it is not in $R_{3}$ and not in $R_{4}$, it must be contained in a leaf produced by a previous split and can therefore be discarded. In total, we scan every list at most once, for a total of $O(\operatorname{scan}(N))$ I/Os.

Up to this point, the construction algorithm has in fact computed a "super-tree" $T^{\prime \prime}$ of $T^{\prime}$, as some leaves produced by the algorithm may be empty. Even though some leaves of $T^{\prime \prime}$ are empty, the size of $T^{\prime \prime}$ is $O(N)$ [4]. Thus, given the assignment of point sets to the leaves of $T^{\prime \prime}$ as computed above, it takes $O(\operatorname{sort}(N))$ I/Os using standard external memory graph techniques to remove empty leaves and all internal nodes such that all descendant leaves of at least one of its two children are empty. The result is the desired partial fair split tree $T^{\prime}$.

## 5 Dynamic Well-Separated Pair Decomposition

Dynamic Fair-Split Tree We present an I/O-efficient algorithm for dynamically maintaining the fair-split tree using $O\left(\log _{B} N\right)$ I/Os per insert or delete operation. We follow the approach of [8].

The main idea is to maintain a rectangle $\tilde{R}(A)$, such that $R(A) \subseteq \tilde{R}(A)$, for each node $A \in T$ that is not modified during updates. A fair cut of any rectangle $\bar{R}$ is defined as partitioning the rectangle using an axis-parallel hyperplane that is at distance at least $\frac{1}{3} l_{\text {max }}(\bar{R})$ from its nearest parallel sides. Note that a fair cut of $\tilde{R}(A)$ satisfies the fair split condition on $R(A)$, since $l_{\max }(R(A)) \leq l_{\max }(\tilde{R}(A))$ and $R(A) \subseteq \tilde{R}(A)$ at any moment in time. Let $R \rightarrow R^{\prime}$ indicate that $R^{\prime}$ can be constructed from $R$ by a sequence of fair cuts.

We maintain a binary tree $T$ under insertions and deletions in which each node satisfies the following invariants:
(I1) For all internal nodes $A$ with children $A_{1}$ and $A_{2}$, there exists a fair cut that partitions $\tilde{R}(A)$ into rectangles $R_{1}$ and $R_{2}$ such that $R_{1} \rightarrow \tilde{R}\left(A_{1}\right)$ and $R_{2} \rightarrow \tilde{R}\left(A_{2}\right)$.
(I2) For all leaves $a, \tilde{R}(a)=a$.
It is easy to see that $T$ is a fair-split tree. Invariant I1 ensures the existence of a fair cut, which satisfies the fair split property.

Insertion: When we want to insert a point $a$, we find the deepest node $A$ in $T$ such that $a \in \tilde{R}(A)$. Let $R_{1}$ and $A_{1}$ have the same meaning as in Invariant I1 and assume that $a \in R_{1}$. We insert a new node $B$, which replaces $A_{1}$ as a child of $A$. We make $A_{1}$ and $a$ children of $B$. We show how to construct $\tilde{R}(B)$ satisfying Invariant I1.

Consider Invariant I1 for node $A$. Since $R_{1} \rightarrow \tilde{R}\left(A_{1}\right)$, there exists a sequence of fair cuts that construct $\tilde{R}\left(A_{1}\right)$ from $R_{1}$. Clearly the last rectangle $R$ in the sequence of cuts that satisfies $a \in R$ can be assigned as $\tilde{R}(B)$. However, this does not give us an efficient way to find $\tilde{R}(B)$, since the sequence of cuts may be long. In [8] it is shown that given the deepest node $A, \tilde{R}(B)$ can be computed in constant time.

We can compute $A$ using $O\left(\log _{B} N\right)$ I/Os by posing it as a deepest-intersect query on the Topology B-tree [5].

Deletion: To delete a point $a$, delete the leaf $a$ and compress the parent node $p(a)$. Note that this preserves the invariants.

Theorem 2. The fair-split tree of a point set $P$ can be maintained using $O(N / B)$ disk blocks and $O\left(\log _{B} N\right)$ I/Os per insert or delete operation.

Maintaining the pairs dynamically We follow the approach of [8]. We use the following characterization of well-separated pairs. An ordered pair of nodes $(A, B), A, B \in T$ is well-separated if and only if it satisfies the following conditions:

1. The point sets $A$ and $B$ are well-separated,
2. Nodes $p(A)$ and $B$ are not well-separated, and
3. $l_{\text {max }}(R(B))<l_{\text {max }}(R(p(A))) \leq l_{\text {max }}(R(p(B)))$.

It is easy to construct examples where the insertion of a new point can add a linear number of well-separated pairs. Thus, we cannot do a trivial update of the pairs upon insertion or deletion. The main idea of the approach is to anticipate all but a constant number of pairs that would be added when a new point is inserted. So, when a point is inserted or deleted, we need to update only a constant number of pairs. In [8] it is shown that this invariant can be maintained only if the point distribution of $P$ is uniform. If it is not uniform, we add dummy points to make it uniform. In [8] it is shown that it is necessary to add only $O(1)$ dummy points per well-separated pair.

When a point $p$ is inserted, we compute a $c$-approximation of the distance between $p$ and its nearest neighbor in $P$, for some constant $c=f(s)$. In [5] it is shown how to compute the approximate nearest neighbor in $O\left(\log _{B} N\right)$ I/Os using a Topology B-tree.

Let $d_{p}$ be the approximate distance computed above. It can be shown that the interactions between $p$ and points $\left\{q \in P \mid d(q, p)>d_{p}\right\}$ are already covered in the current WSPD. We need to compute well-separated pairs that account for interactions between $p$ and points within distance $d_{p}$ from $p$. The number of such pairs needed is $O(1)$. We can compute these pairs as follows:

Extract all nodes $A \in T$ whose bounding rectangles $R(A)$ overlap the cube $C$ centered at $p$ with length $l(C)=d_{p}$ and satisfy the conditions of a WSPD defined earlier in this subsection. This can be done by filtering $C$ through a Topology-B-tree that corresponds to the fair-split tree. Since the number of pairs extracted is $O(1)$, we can perform this extraction using $O\left(\log _{B} N\right)$ I/O's.

We need to add dummy points to make sure that the region covering the newly computed pairs $(A, B)$ is uniform. Note that for the existing pairs, we have already added the required dummy points. A straightforward way is to add a mesh of suitably spaced points covering the region around $R(A)$ and $R(B)$. We insert the dummy points
and compute the pairs corresponding to them as described above. The total number of dummy points added is $O(1)$ and the number of pairs added to the WSPD is $O(1)$. Thus, the entire process requires $O\left(\log _{B} N\right) \mathrm{I} / \mathrm{Os}$. The deletion of a point $p$ can be taken care of using a similar procedure.

Theorem 3. Let $P$ be a set of $N$ points in $\mathbb{R}^{d}$. The well-separated pair decomposition of $P$ with respect to a fair-split tree $T$ can be maintained under insertions and deletions using $O\left(\log _{B} N\right)$ I/Os per update and $O(N / B)$ disk blocks.

## 6 Applications of the WSPD

## 6.1 t-Spanners

Given a point set $P$, let $G=(P, E)$ be a graph whose edges have weights corresponding to the Euclidean distance between their two endpoints. $G$ is called a $t$-spanner for $P$ if for any two points $p$ and $q$ in $P$, the shortest path from $p$ to $q$ in $G$ is at most $t$ times longer than the Euclidean distance between $p$ and $q$. In [9] it is shown that the following graph $G=(P, E)$ is a $t$-spanner of linear size for the point set $P$ : For every node $A \in T$ choose a representative $q(A) \in A$. For every pair $\left\{A_{i}, B_{i}\right\}$ in the given WSPD of $P$, add an edge $\left\{q\left(A_{i}\right), q\left(B_{i}\right)\right\}$ to $E$. In [3] it is shown that one can compute a spanner of diameter at most $2 \log N$ by choosing representatives carefully. In particular, we partition the edges in $T$ into heavy and light edges. Given a node $A$ with children $A_{1}$ and $A_{2}$, edge $\left(A, A_{1}\right)$ is heavy and edge $\left(A, A_{2}\right)$ is light if $\left|A_{1}\right| \geq\left|A_{2}\right|$. Otherwise, edge $\left(A, A_{2}\right)$ is heavy and edge $\left(A, A_{1}\right)$ is light. Every node of $T$ is contained in a unique chain of heavy edges starting at a leaf. For node $A$, let $q(A)$ be the leaf whose chain contains $A$.

This assignment of representatives corresponds to sending representatives up the tree; at every node $A$ receiving representative $q\left(A_{1}\right)$ and $q\left(A_{2}\right)$ from its children, we choose $q(A)$ as the one sent along the heavy edge. The decision which edge is heavy can be made by sending sizes $\left|A_{1}\right|$ and $\left|A_{2}\right|$ along with the representative $q\left(A_{1}\right)$ and $q\left(A_{2}\right)$. This can be done using time-forward processing.

Theorem 4. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$, it takes $O(\operatorname{sort}(N))$ I/Os and $O(N / B)$ blocks of external memory to compute a $t$-spanner $G$ of linear size and diameter at most $2 \log N$ for $P$.

The next theorem states that the I/O-complexity of the spanner construction is practically optimal.

Theorem 5. It takes $\Omega(\min \{N, \operatorname{sort}(N)\})$ I/Os to compute a $t$-spanner, $t>1$, for a point set $P$ of $N$ points in $\mathbb{R}^{d}$.

Proof. We prove the lemma for $d=1$. The proof generalizes to higher dimensions by placing all points on a straight line in $\mathbb{R}^{d}$. Given a list $X$ of items $x_{1}, \ldots, x_{N}$ and a permutation $\sigma:[1, N] \rightarrow[1, N]$, we are finally interested in computing the sequence $x_{\sigma(1)}, \ldots, x_{\sigma(N)}$. Let $\sigma$ be given as a sequence $S=\langle\sigma(1), \ldots, \sigma(N)\rangle$. We map data item $x_{i}$ to the point $i+\frac{1}{t}$ and $\sigma(i)$ to the point $\sigma(i)$. Any $t$-spanner, $t>1$, of the resulting point set must contain edges $\left(\sigma(i), x_{\sigma(i)}\right)$. We first remove all edges that are not of this type and
then reverse the permutation performed on elements $\sigma(i)$ by the spanner construction algorithm to arrange the remaining edges in the order $\left(\sigma(1), x_{\sigma(1)}\right), \ldots,\left(\sigma(N), x_{\sigma(N)}\right)$. This reversal can be done by recording the I/Os performed by the spanner construction algorithm and playing this I/O-sequence backward, exchanging the meaning of read and write operations.

The construction of the point set takes $O(\operatorname{scan}(N))$ I/Os. The reversal of the permutation takes $O\left(T_{S}(N)\right)$ I/Os, where $T_{S}(N)$ is the I/O-complexity of the spanner algorithm. In total, we can compute any permutation in $O\left(T_{S}(N)\right.$ ) I/Os, so that $T_{S}(N)=$ $\Omega(\operatorname{perm}(N))=\Omega(\min \{N, \operatorname{sort}(N)\})$.

## 6.2 $K$-Nearest Neighbors

In this section we show how to compute the $K$-nearest neighbors for every point $p \in P$. The construction follows the sequential algorithm in [9] for this problem.

Lemma 3 ([9]). Let $\{A, B\}$ be a pair in a well-separated realization of $P \otimes P$ with separation $s>2$. If there is a point $b \in B$ that is a $K$-nearest neighbor of a point $a \in A$, then $|A| \leq K$.

Given a set $B$, let $O_{B}$ be the center of $R(B)$. Then we divide the space around $B$ into a constant number of cones with apex $O_{B}$ such that for any two points $a$ and $a^{\prime}$ in the same cone, $\angle a O_{B} a^{\prime}<\frac{s}{s+1}$.

Lemma 4 ([9]). Let a point set $X$ and a cone $c$ with apex $O_{B}$ be given, such that for any two points $a$ and $a^{\prime}$ in $c, \angle a O_{B} a^{\prime}<\frac{s}{s+1}$. Let $X_{c}=\left(x_{1}, \ldots, x_{l}\right)$ be the set of points in $X$ that lie in $c$, sorted by distance from $O_{B}$. For $i>K$, no point in $B$ can be a $K$-nearest neighbor of $x_{i}$.

Based on Lemma 3 the algorithm in [9] first extracts all pairs $\left\{A_{i}, B_{i}\right\}$ with $\left|A_{i}\right| \leq K$. This can be done in $O(\operatorname{sort}(N)+\operatorname{scan}(K N))$ I/Os using time-forward processing and a reverse preorder traversal of $T$. For a node $B$ in $T$, let $\left\{A_{1}^{\prime}, B\right\}, \ldots,\left\{A_{q}^{\prime}, B\right\}$ be the set of pairs such that $\left|A_{i}^{\prime}\right| \leq K$. Note that the sets $A_{i}^{\prime}$ are pairwise disjoint. We store the set $f(B)=\bigcup_{i=1}^{q} A_{i}^{\prime}$ as a candidate set at node $B$ in $T$. The construction of sets $f(B)$ takes sorting and scanning and thus $O(\operatorname{sort}(K N))$ I/Os. Then the set of candidates is narrowed down by a top-down procedure in $T$, which can be realized using time-forward processing. At the root node $B$ of $T$, the algorithm partitions the space around $O_{B}$ into cones as described above and uses $K$-selection in each cone to find the $K$ points in $f(B)$ that fall into this cone and are closest to $O_{B}$. By Lemma 4, all other points in $f(B)$ cannot have a $K$-nearest neighbor in $B$. Call the resulting set of points $N(B)$. As $N(B)$ contains at most $K$ points per cone and there are $O(1)$ cones, $|N(B)|=O(K)$. The set $N(B)$ is passed on to the children $B_{1}$ and $B_{2}$ of $B$. The points possibly having a closest neighbor in $B_{i}$ are in the set $f\left(B_{i}\right) \cup N(B)$. Thus, the algorithm applies the same $K$-selection algorithm to this set to construct $N\left(B_{i}\right)$ and pass it on to $B_{i}$ 's children. As we send only $O(K)$ points along every edge of $T$, this takes $O(\operatorname{sort}(K N))$ I/Os.

Finally every leaf $b$ stores a set $N(b)$ of size $O(K)$ such that $b$ is a potential $K$-nearest neighbor for the points in $N(b)$. The total size of all these sets is $O(K N)$. Now we build sets $N^{\prime}(a)$ for all points $a \in P$ such that $b \in N^{\prime}(a)$ if and only if $a \in N(b)$ and apply
$K$-selection to each set $N^{\prime}(a)$ separately to find the $K$-nearest neighbors of point $a$. This takes sorting of $K N$ points and a standard $K$-selection algorithm, hence, $O(\operatorname{sort}(\operatorname{KN}))$ I/Os.

Theorem 6. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$, it takes $O(\operatorname{sort}(K N)) I / O s$ and $O(K N / B)$ blocks of external memory to find the $K$-nearest neighbors for all points in $P$.

### 6.3 K-Closest Pairs

Given a well-separated realization of $P \otimes P$, let the pairs $\left\{A_{i}, B_{i}\right\}$ be sorted by increasing distances $d\left(R\left(A_{i}\right), R\left(B_{i}\right)\right)$. For any such pair, $\left|A_{i} \times B_{i}\right|=\left|A_{i}\right|\left|B_{i}\right|$. To find the $K$ closest pairs, we first determine the smallest $i$ such that $\sum_{j=1}^{i}\left|A_{j}\right|\left|B_{j}\right| \geq K$ and retrieve all pairs $\{A, B\}$ such that $d(R(A), R(B)) \leq(1+4 / s) r$, where $r=d\left(R\left(A_{i}\right), R\left(B_{i}\right)\right)$. We extract the set of pairs $\{a, b\}$ such that $a \in A$ and $b \in B$ for some pair $\{A, B\}$ that we retrieved.

In [7] it is shown that the $K$ closest pairs are among the extracted point pairs and that the total number of extracted point pairs is $O(N+K)$. Given this set of point pairs, we apply $K$-selection again to find the $K$ closest ones. Given all point pairs, this takes $O(\operatorname{scan}(N+K))$ I/Os. We have to show how to extract all candidate pairs $\{a, b\}$.

Sorting the pairs $\left\{A_{i}, B_{i}\right\}$ by their distances $d\left(R\left(A_{i}\right), R\left(B_{i}\right)\right)$ between the bounding rectangles takes $O(\operatorname{sort}(N))$ I/Os. Then it takes a single scan to extract all pairs with $d(R(A), R(B)) \leq(1+r / s) r$. We extract the points in sets $A_{i}$ and $B_{i}$ in the same way as in the previous section and sort these at most $O(N+K)$ points to guarantee that the points in the same pair $\left\{A_{i}, B_{i}\right\}$ are stored consecutively. We can now construct the set of all candidate point pairs by scanning the resulting point list. We obtain the following result.
Theorem 7. Given a set $P$ of $N$ points in $\mathbb{R}^{d}$, it takes $O(\operatorname{sort}(N+K)) I / O s$ and $O((N+$ $K) / B)$ blocks of external memory to compute the $K$ closest pairs in $P$.

### 6.4 Dynamic Closest Pair

We show how to maintain the closest pair of points in $P$ under insertion and deletion. We make the following simple observation:
Lemma 5. Let $P$ be a point set of $n$ points in $\mathbb{R}^{d}$. Let $(a, b), a, b \in P$ be the closest pair in $P$. Then, the pair $\{\{a\},\{b\}\}$ belongs to the well-separated pair decomposition of $P$.

Proof. Let $(A, B)$ be the pair in the well-separated decomposition such that $a \in A, b \in B$ and $b$ is a nearest neighbor of $a$. Applying Lemma 3, we have $|A|=1$. Similarly, $a$ is a nearest neighbor of $b$. By Lemma 3, we have $|B|=1$.

Denote the well-separated pairs of the form $\{\{a\},\{b\}\}$ as singleton pairs. We maintain the singleton pairs in a B-tree based on the distance between them. The closest pair is the singleton pair with minimum distance.

During an insert or delete operation, update the well-separated pairs and insert or delete the singleton pairs in the B-tree, respectively. Since only $O(1)$ well-separated pairs change during an insert or delete operation, we obtain the following theorem.
Theorem 8. Let $P$ be a point set of $n$ points in $\mathbb{R}^{d}$. The closest pair in $P$ can be maintained using $O\left(\log _{B} N\right)$ I/Os per insert or delete operation.

## Acknowledgments

We would like to thank Pankaj Agarwal, Pat Morin, and Michiel Smid for helpful discussions.

## References

1. Lars Arge. The buffer tree: A new technique for optimal I/O-algorithms. In Proceedings of the Workshop on Algorithms and Data Structures, volume 955 of Lecture Notes in Computer Science, pages 334-345, 1995.
2. Lars Arge and P. B. Miltersen. On showing lower bounds for external-memory computational geometry problems. In J. Abello and J. S. Vitter, editors, External Memory Algorithms and Visualization. AMS, 1999.
3. S. Arya, D. M. Mount, and M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In 35th IEEE Symposium on Foundations of Computer Science (FOCS'94), pages 703-712, 1994.
4. P. B. Callahan. Optimal parallel all-nearest-neighbors using the well-separated pair decomposition. In Proc. FOCS'93, pages 332-341, 1993.
5. P. B. Callahan, M. Goodrich, and K. Ramaiyer. Topology B-trees and their applications. In Proc. WADS'95, LNCS 955, pages 381-392, 1995.
6. P. B. Callahan and S. R. Kosaraju. A decomposition of multi-dimensional point sets with applications to $k$-nearest-neighbors and $n$-body potential fields. In Proc. STOC'92, pages 546-556, 1992.
7. P. B. Callahan and S. R. Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In Proc. SODA'93, pages 291-300, 1993.
8. P. B. Callahan and S. R. Kosaraju. Algorithms for dynamic closest pair and $n$-body potential fields. In Proc. SODA'95, pages 263-272, 1995.
9. P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to $k$-nearest neighbors and $n$-body potential fields. Journal of the ACM, 42:6790, 1995.
10. Paul B. Callahan. Dealing with Higher Dimensions: The Well-Separated Pair Decomposition and Its Applications. PhD thesis, Johns Hopkins University, Baltimore, Maryland, 1995.
11. Yi-Jen Chiang, Michael T. Goodrich, Edward F. Grove, Roberto Tamassia, Darren Erik Vengroff, and Jeffrey Scott Vitter. External-memory graph algorithms. In Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms, January 1995.
12. Michael T. Goodrich, Jyh-Jong Tsay, Darren Erik Vengroff, and Jeffrey Scott Vitter. External-memory computational geometry. In Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science, November 1993.
13. T. Lukovszki, A. Maheshwari, and N. Zeh. I/O-efficient spanner construction in the plane. http://www.scs.carleton.ca/~nzeh/Pub/spanners00.ps.gz, 2000.
14. Michiel Smid. Closest-point problems in computational geometry. In J.-R. Sack and J. Urrutia, editors, Handbook of Computational Geometry, pages 877-936. North-Holland, 2000.
15. J. S. Vitter. External memory algorithms. In Proceedings of the 17 th Annual ACM Symposium on Principles of Database Systems, June 1998.
16. J.S. Vitter and E.A.M. Shriver. Algorithms for parallel memory I: Two-level memories. Algorithmica, 12(2-3):110-147, 1994.

[^0]:    * Research supported by NSERC, NCE GEOIDE, and by DFG-SFB376.

[^1]:    ${ }^{1}$ We do not simulate the PRAM algorithm as this would lead to a much higher I/O complexity.

