New constructions of SSPDs and their applications ☆

Mohammad A. Abam [a], Sariel Har-Peled [b],[*],[1]

[a] *Department of Computer Engineering, Sharif University of Technology, Tehran, Iran*
[b] *Department of Computer Science, University of Illinois, 201 N. Goodwin Avenue, Urbana, IL 61801, USA*

ABSTRACT

We present a new optimal construction of a semi-separated pair decomposition (i.e., SSPD) for a set of n points in $\mathbb{R}^d$. In the new construction each point participates in a few pairs, and it extends easily to spaces with low doubling dimension. This is the first optimal construction with these properties.

As an application of the new construction, for a fixed $t > 1$, we present a new construction of a t-spanner with $O(n)$ edges and maximum degree $O(\log^2 n)$ that has a separator of size $O(n^{1-1/d})$.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

For a point-set P, a pair decomposition of P is a set $\mathcal{W}$ of pairs of subsets of P, such that for every pair of points of p, q ∈ P there exists a pair $\{\mathcal{X}, \mathcal{Y}\} \in \mathcal{W}$ such that p ∈ $\mathcal{X}$ and q ∈ $\mathcal{Y}$—see Section 2 for the formal definition. A well-separated pair decomposition (WSPD) of P is a pair decomposition of P such that for every pair $\{\mathcal{X}, \mathcal{Y}\}$, the distance between $\mathcal{X}$ and $\mathcal{Y}$ is large when compared to the maximum diameter of the two point sets. The notion of WSPD was developed by Callahan and Kosaraju [7], and it provides a compact representation of the quadratic pairwise distances of the point-set P, since there is a WSPD with a linear number of pairs.

The total weight of a pair decomposition $\mathcal{W}$ is the total size of the sets involved; that is $\omega(\mathcal{W}) = \sum_{\{\mathcal{X},\mathcal{Y}\}\in\mathcal{W}}(|\mathcal{X}| + |\mathcal{Y}|)$. Naturally, a WSPD with near linear weight is easier to manipulate and can be used in applications where the total weight effects the overall performance. Unfortunately, it is easy to see that, in the worst case, the total weight of any WSPD is $\Omega(n^2)$. Callahan and Kosaraju [7] overcame this issue by generating an implicit representation of the WSPD using a tree. Indeed, they build a tree $\mathcal{T}$ (usually a compressed quadtree, or some other variant) that stores the points of the given point set in the leaves. Pairs are reported as $\{u, v\}$ where u and v are nodes in the tree such that their respective pair $\mathcal{S}(u) \otimes \mathcal{S}(v)$ is well separated, where $\mathcal{S}(u)$ denotes the set of points stored at the subtree of u.

1.1. SSPDs

To overcome this obesity problem, a weaker notion of semi-separated pair decomposition (SSPD) had been suggested by Varadarajan [22]. Here, an SSPD $\mathcal{S}$ of a point set has the property that for each pair $\{\mathcal{X}, \mathcal{Y}\} \in \mathcal{S}$ the distance between $\mathcal{X}$
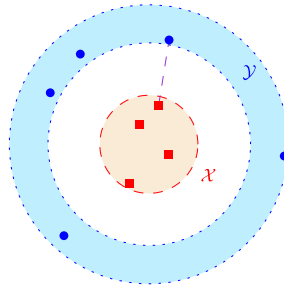

Fig. 1.

and $\mathcal{Y}$ is large when compared to the *minimum* diameter of the two point sets (in the WSPD case this was the *maximum* of the diameters). See Fig. 1 for an example of a semi-separated pair that is not well-separated.

By weakening the separation, one can get an SSPD with near-linear total weight. Specifically, Varadarajan [22] showed how to compute an SSPD of weight $O(n \log^4 n)$ for a set of n points in the plane, in $O(n \log^5 n)$ time (for a constant separation factor). He used the SSPD for speeding up his algorithm for the min-cost perfect-matching in the plane. Recently, Abam et al. [3] presented an algorithm which improves the construction time to $O(\varepsilon^{-2} n \log n)$ and the weight to $O(\varepsilon^{-2} n \log n)$, where $1/\varepsilon$ is the separation required between pairs. It is known that any pair decomposition has weight $\Omega(n \log n)$ [6], implying that the result of Abam et al. [3] is optimal. This construction was generalized to $\mathbb{R}^d$ with the construction time being $O(\varepsilon^{-d} n \log n)$ and the total weight of the SSPD being $O(\varepsilon^{-d} n \log n)$, see [2].

1.2. Spanners

More recently SSPDs were used in constructing certain geometric spanners [3,2,1]. Let $\mathcal{G} = (P, E)$ be a geometric graph on a set P of n points in $\mathbb{R}^d$. That is, $\mathcal{G}$ is an edge-weighted graph where the weight of an edge $pq \in E$ is the Euclidean distance between p and q. The distance in $\mathcal{G}$ between two points p and q, denoted by $d_{\mathcal{G}}(p, q)$, is the length of the shortest (that is, minimum-weight) path between p and q in $\mathcal{G}$. A graph $\mathcal{G}$ is a (geometric) *t-spanner*, for some $t \geqslant 1$, if for any two points $p, q \in P$ we have $d_{\mathcal{G}}(p, q) \leqslant t \cdot \|pq\|$. Note that the concept of geometric spanners can be easily extended to any metric space. Geometric spanners have received considerable attention in the past few years—see [19] and references therein. Obviously, the complete graph is a *t-spanner*, but a preferable spanner would provide short paths between its nodes, while having few edges. Other properties considered include (i) low total length of edges, (ii) low diameter, (iii) low maximum degree, and (iv) having small separators.

1.3. Separators

A graph $\mathcal{G}$ is said to have a *k-separator* if its vertices can be decomposed into three sets $\mathcal{X}$, $\mathcal{Y}$ and $\mathcal{Z}$ such that both $|\mathcal{X}|$ and $|\mathcal{Y}|$ are $\Omega(n)$ and $|\mathcal{Z}| \leqslant k$ (i.e., the set $\mathcal{Z}$ is the *separator*). Furthermore, there is *no* edge in $\mathcal{G}$ connecting a vertex in $\mathcal{X}$ to a vertex in $\mathcal{Y}$. Graphs with good separators are the best candidates to applying the divide-and-conquer approach—see [16,18,21] and references therein for more results and applications. Lipton and Tarjan [16] showed that any planar graph has an $O(\sqrt{n})$-separator. The Delaunay triangulation of a planar point-set P is an $O(1)$-spanner [15]. Furthermore, it is a planar graph and is thus an $O(1)$-spanner with an $O(\sqrt{n})$-separator. Fürer and Kasiviswanathan [9] recently presented a *t-spanner* for a ball graph which is an intersection graph of a set of n balls in $\mathbb{R}^d$ with arbitrary radii which has an $O(n^{1-1/d})$-separator. Since complete Euclidean graphs are a special case of unit ball graphs, their results yield a new construction of *t*-spanner for geometric graphs with small separators.

1.4. Metric spaces with low doubling dimension

Recently the notion of doubling dimension [5,11,13] which is a generalization of the Euclidean dimension, has received considerable attention. The *doubling constant* of a metric space $\mathcal{M}$ is the maximum, over all balls b in the metric space $\mathcal{M}$, of the minimum number of balls needed to cover b, using balls with half the radius of b. The logarithm of the doubling constant is the *doubling dimension* of the space—note that $\mathbb{R}^d$ has $\Theta(d)$ doubling dimension. Constructions of WSPD, spanners, and a data-structure for approximate nearest neighbor for metric spaces with fixed doubling dimension were provided by Har-Peled and Mendel [14].

1.5. Limitations of known constructions of SSPDs

The optimal construction of SSPDs of [3] uses BAR-trees [8], and as such it is not applicable to metric spaces with constant doubling dimension. Moreover, in all previously known optimal constructions of SSPDs a point might appear in many (i.e., linear number of) pairs, see Appendix B for an example demonstrating this. Note, that in Varadarajan's original