Balanced Aspect Ratio Trees: An Introduction

Christian A. Duncan*
Center for Geometric Computing
The Johns Hopkins University
duncan@jhu.edu

Michael T. Goodrich†
Center for Geometric Computing
The Johns Hopkins University
goodrich@jhu.edu

Stephen Kobourov‡
Center for Geometric Computing
The Johns Hopkins University
kobourov@cs.jhu.edu

1 Introduction

Binary space partition (BSP) trees such as k-d trees and octrees have uses influencing a wide range of disciplines. Areas such as geographic information systems, natural language processing, speech recognition, computer vision, and computer graphics rely heavily on the need to find nearest and farthest neighbors to given positions in some high-dimensional space or to report all points in a data set that may lie within some given range. BSP trees and their various derivatives can typically handle these queries with both ease and efficiency, yielding vast improvements over exhaustive searches of the entire point set. There are, unfortunately, some serious theoretical and practical drawbacks with most BSP trees. Many of the standard techniques employed by these data structures fall apart when the dimensionality of the point set becomes extremely high, i.e. non-constant, resulting in performances which are no better than exhaustive searches even in practice. Another drawback is that most results are based on empirical rather than theoretical performance. It is easy to construct examples, in any dimension, which cause problems with many of these types of trees.

One measure of performance, a weak spot in most BSP trees and the focus of this talk, is the aspect ratio bound, or "fatness", of the regions produced in the tree. The aspect ratio of a region is roughly equivalent to the ratio of the region's longest side to its smallest side. If a tree has several "skinny" regions close together, a single query may require traversing all of these skinny regions prior to finding the region containing the nearest point. All variations of the popular k-d tree, for example, are unable to maintain any bounded aspect ratio bounds. Imagine a cube with nearly all points concentrated at one of the corners. On the other hand, octrees, with regions that are perfect cubes, maintain a balanced aspect ratio but only by sacrificing balanced tree depth.

Recently, some major progress has been achieved in maintaining aspect ratio performance of BSP-like trees. Arya et al [2, 1] describe a related data structure, the balanced box-decomposition (BBD) tree, based on the fair-split tree of Callahan and Kosaraju [3, 4], which successfully achieves both balanced, $O(\log n)$, depth, and bounded aspect ratios, but sacrifices the convexity of the regions by introducing holes.

In our talk, we generalize an earlier result in the plane [5] to higher dimensions and introduce both a general framework for the balanced aspect ratio (BAR) tree, the first BSP tree structure to combine the advantages of k-d trees, octrees, and BBD trees, and a particular construction of such a BAR tree. We also describe some of the various uses and properties BAR trees have. For example, BAR trees can efficiently find approximate nearest and farthest neighbors as well as answer basic approximate range queries.

*This research partially supported by ARO under grant DAAH04-96-1-0013.
†This research partially supported by NSF under Grant CCR-9625289, and ARO under grant DAAH04-96-1-0013.
‡This research partially supported by ARO under grant DAAH04-96-1-0013.

2 General Framework

Before describing the general framework, we need to introduce some standard definitions.

Definition 2.1 A region R in $\mathbb{R}^d$ has aspect ratio $\mathtt{asp}(R) = O_R/I_R$ with respect to some underlying metric, where O_R is the radius of the smallest circumscribed hypersphere in $\mathbb{R}^d$ and I_R is the radius of the largest inscribed d-hypersphere. R has *balanced aspect ratio* with maximum aspect ratio $\alpha \geq 1$, if $\mathtt{asp}(R) \leq \alpha$. We say, a region R is an *α-balanced region* if it has maximum aspect ratio α. A collection of regions, $\mathcal{R}$, has balanced aspect ratio with balancing factor α if each region $R \in \mathcal{R}$ is an α-balanced region.

Definition 2.2 A *canonical cut set*, $\mathcal{C} = \{\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_\gamma\}$, is a collection of γ, not necessarily independent, vectors that span $\mathbb{R}^d$ (thus, $\gamma \geq d$). A *canonical cut* is any hyperplane, H, in $\mathbb{R}^d$ with a normal in $\mathcal{C}$. A *canonical region* is a convex polyhedron in $\mathbb{R}^d$ with every facet having a normal in $\mathcal{C}$.

The canonical cut set varies depending on the type of BSP tree constructed. A general BSP tree uses the entire range of vectors possible, whereas the k-d tree uses only vectors parallel to one of the d axes. The canonical set that we use in our construction of a BAR tree is a union of axis-parallel cuts and corner cuts, cuts whose intersections with each of the coordinate axes form 45° angles.

Definition 2.3 An α-balanced canonical region, R, is *one-cuttable* with reduction factor β, where $1/2 \leq \beta < 1$, if there is a cut $s_1 \in \mathcal{C}$, called a *one-cut*, dividing R into two subregions R_1 and R_2 such that

1. R_1 and R_2 are α-balanced canonical regions,
2. $|R_1| \leq \beta|R|$ and $|R_2| \leq \beta|R|$.

Definition 2.4 An α-balanced canonical region, R, is *k-cuttable* with reduction factor β, for $k > 1$, if there is a cut $s_k \in \mathcal{C}$, called a *k-cut*, dividing R into two subregions R_1 and R_2 such that

1. R_1 and R_2 are α-balanced canonical regions,
2. $|R_2| \leq \beta|R|$,
3. Either $|R_1| \leq \beta|R|$ or R_1 is $(k-1)$-cuttable with reduction factor β.

In other words, the sequence of cuts, $s_k, s_{k-1}, \ldots, s_1$, results in $k+1$ balanced canonical regions each containing no more than βn points. If β is understood, we simply say R is k-cuttable.

Theorem 2.5 For a canonical set, $\mathcal{C}$, if every possible α-balanced canonical region is k-cuttable with reduction factor β, then a BAR tree with maximum aspect ratio α can be constructed with depth $O(k \log_{1/\beta} n)$, for any set S of n points.

The difficulty in using the above theorem is in proving that such a canonical set exists. In our talk, we describe how the canonical set formed by the union of all axis-parallel and corner cuts guarantees that every possible α-balanced canonical region is $(d+1)$-cuttable for $\alpha \geq 3d$ and $\beta \geq d/(d+1)$.

References

[1] S. Arya and D. M. Mount. Approximate range searching. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 172–181, 1995.

[2] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching. In *Proc. 5th ACM-SIAM Sympos. Discrete Algorithms*, pages 573–582, 1994.

[3] P. B. Callahan and S. R. Kosaraju. Algorithms for dynamic closest-pair and n-body potential fields. In *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, pages 263–272, 1995.

[4] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *J. ACM*, 42:67–90, 1995.

[5] C. A. Duncan, M. T. Goodrich, and S. G. Koboruov. Balanced aspect ratio trees and their use for drawing very large graphs. In *Sixth Symposium on Graph Drawing*, to be published in 1998.