Lehrstuhl Algorithmen & Datenstrukturen, Institut für Informatik, Fakultät für Angewandte Wissenschaften, Albert-Ludwigs-Universität Freiburg
IP Network Algorithms
Algorithms & Data Structures Research Group
Georges-Köhler-Allee 51, D-79110, Freiburg im Breisgau, Germany. Fax: +49-761-206-8162.
URL: http://netalg.ad.informatik.uni-freiburg.de/

Level Intended: Bachelor **Project Type:** Thesis (12 weeks)
Collaborative Theme: Spatial data structures, conflict detection
Supervisor(s): Khaireel A. Mohamed (khaireel@informatik.uni-freiburg.de)

Studying the *Priority* R-Tree in Preparation for IP-Lookup

Preamble

The original R-Tree by Gutmann [1], is a hierarchical index-structure derived from the B-Tree [2]. The structure is designed to perform intersection queries of *d*-dimensional spatial objects efficiently, so that a spatial search requires visiting only a small number of nodes. The R-Tree offers a solution for combined dimensions that store *d*-dimensional spatial objects in their primitive forms while maintaining versatility of its internal structure. That is, one of its several advantages over other type-specific structures is that the same concept of its internal structure can be used to administer and solve different problems on arbitrary types of multidimensional data.

Thesis Description

In this Bachelorarbeit, we shall study an advanced variant of the R-Tree, known as the Priority R-Tree (PR-Tree) introduced by Arge *et al*. [3], which guarantees the worst-case range-query performance of $O((n/b)^{1-1/d} + k)$ time, where n refers a set of input d-dimensional rectangles stored in the R-Tree with a branching factor of b, and k is the number of outputs answering a given range-query. This is a stark improvement over the original R-Tree's non-worst-case efficient query performance. The construction of the PR-Tree is also an enhancement over the original R-Tree (and its early variants, such as the R+ and R* Trees) in terms of runtime and space complexities. It relies on a bulk-loading algorithm similar to that of the *kd*-tree's approach, where each input rectangle is treated like a point in d-space to first construct the *Pseudo*-PR-Tree, in order to obtain the final PR-Tree.

We shall be continuing on a current work of thoroughly analysing the PR-Tree, and verifying its performance on range-queries, particularly for d=2. On top of this, we shall also be reviewing the IP-lookup problem and its related issues [4, 5], particularly on conflict detection of 2D range-filters under the most-specific tie-breaking rule (MSTB). We shall look at the viability of extending several methodologies within the framework of the PR-Tree to conceptualise a technique to detect conflicts and analyse its performance. In retrospect, we will refer to the articles by Eppstein and Muthukrishnan [6], and Lu and Sahni [7] to base the work of this thesis.

References

1. A. Guttman. R-trees: a dynamic index structure for spatial searching. In *Proceedings of the International Conference on Management of Data (ACM SIGMOD)*, pages 47-57, 1984. ACM Press.
2. D. Comer. Ubiquitous B-tree. *ACM Computing Surveys*, 11(2):121-137, 1979. ACM Press.
3. L. Arge, M. de Berg, H. J. Haverkort, and K. Yi. The Priority R-tree: a practically efficient and worst-case optimal R-tree. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pages 347-358, 2004. ACM Press.
4. A. Datta and Th. Ottmann. A note on the IP table lookup problem. 2004.
5. C. Kupich and K. A. Mohamed. Conflict detection in internet router tables. *Institut für Informatik, Albert-Ludwigs-Universität Freiburg*, Technical Report 225, 2006.
6. D. Eppstein and S. Muthukrishnan. Internet packet filter management and rectangle geometry. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 827-835, 2001. ACM Press.
7. H. Lu and S. Sahni. Conflict detection and resolution in two dimensional prefix router tables. In *IEEE/ACM Transactions on Networking*, 13(6):1353-1363, 2005. ACM Press.