Accepted paper with abstract, SoCG'09

Peyman Afshani, Chris Hamilton and Norbert Zeh

A Unified Approach for Cache-Oblivious Range Reporting and Approximate Range Counting

In this paper we present cache-oblivious solutions to two important variants of range searching: range reporting and approximate range counting. The main contribution of our paper is a unified approach for constructing cache-oblivious data structures that provide relative $(1 + \varepsilon)$-approximations for a general class of range counting queries. This class includes three-sided range counting in the plane, dominance counting in three dimensions, and halfspace range counting in three dimensions. Our technique allows us to obtain data structures that use linear space and answer queries in the optimal query bound of $O(\log_B (N / K))$ in the worst case, where K is the number of points in the query range.

Except for orthogonal range counting in the plane, no such structures were known in the cache-oblivious model before. Furthermore, our approach is powerful enough to provide the first approximate halfspace range counting data structure with a *worst-case* query time of $O(\log (N / K))$ in internal memory; previously, the same query bound was only known to hold in the expected case.

An easy but important consequence of our main result is the existence of $O(N \log N)$-space cache-oblivious data structures with an optimal query complexity of $O(\log_B N + K / B)$ for the reporting versions of the same problems. Through standard reductions, we also obtain the first cache-oblivious data structures that use near-linear space and achieve the optimal query bound for circular range reporting and k-nearest neighbor search in the plane, and for orthogonal range reporting in three dimensions.