

Accepted paper with abstract, SoCG'09

Peyman Afshani, Chris Hamilton and Norbert Zeh

Cache-Oblivious Range Reporting With Optimal Queries Requires Superlinear Space

We consider a number of range reporting problems in two and three dimensions and prove lower bounds on the amount of space used by any cache-oblivious data structure for these problems that achieves an optimal query bound of $O(\log_B N + K / B)$ memory transfers in the worst case, where K is the size of the query output.

The problems we study are three-sided range reporting in the plane, dominance reporting in three dimensions, and halfspace range reporting in three dimensions. We prove that, in order to achieve the above query bound or even a bound of $O((\log_B N)^c (1 + K / B))$ for any constant $c > 0$, the structure has to use $\Omega(N (\log \log N)^\varepsilon)$ space, where $\varepsilon > 0$ is a constant that depends on c and the constants hidden in the big-Oh notation of the query bound.

Our result has a number of interesting consequences. The first one is a new type of separation between the I/O-model and the cache-oblivious model, as I/O-efficient structures with the optimal query bound and using linear or $O(N \log^* N)$ space are known for the above problems. The second consequence is the non-existence of linear-space cache-oblivious persistent B-trees with worst-case optimal range queries.