

The goal of this thesis is to implement an inverted-index database software that provides improvements in handling raw non-textual data, which is beneficial for several areas of research. The main internal structures of the library are designed to be cache-oblivious, also aiming to reduce the size of stored data. This thesis includes an overview of common inverted index implementation methods and describes related structures in a suitable cache-based model. This resulted in improvements of compression ratio, and performance similar to currently available highly optimized databases. The benchmark conducted on cheminformatic data has shown that the resulting software is applicable as an immediate, efficient replacement of the storage back-ends of specialized molecule databases.