


Accelerating Vector Similarity Search Using FPGAs

Background Vector similarity search is increasingly important in vector databases and embedding-based applications. Given a query vector, the objective is to identify the vectors in a database that are most similar according to a distance.



Suggested Literature

- Fast Graph Vector Search via Hardware Acceleration and Delayed-Synchronization Traversal
- Approximate Nearest Neighbor Search on FPGA
- Co-design Hardware and Algorithm for Vector Search

Prerequisites: Basic knowledge in C++ and VHDL or Verilog

Type of Work: Theory (20%), Conception (40%), Implementation (40%) (BA, MA)

Supervisor: Chen Lei (lei.chen@tu-ilmenau.de)