Could a B-tree help a vector lookup? Why or why not?

Describe the tradeoffs between number of clusters and number of probes?

Suppose a database contains 1,000,000 vectors, each with 384 dimensions. We construct an IVF index with 100 equally sized clusters.

A. How many dot-product coordinate multiplications are required for an exhaustive search?

B. If nprobe = 5, approximately how many vectors must be compared with the query?

C. Approximately how many coordinate multiplications does the IVF search require, including comparisons with all 100 centroids?

D. Approximately what percentage of the exhaustive-search multiplications have we eliminated?