

CS 386: Quiz 6, 23 May 2006

Name: _____ KEY _____

No books or notes. Work individually.

Assume we have a table that describes auctions on Ebay:

Ebay(AuctionID, Item, Category, Date, FinalBid)

There are three indexes on this table:

I1: Clustered B+ tree on Item, Category

I2: Unclustered hash table on AuctionID

I3: Unclustered B+ tree on FinalBid

Some people missed that these indexes are defined on particular attributes, and tried to use them on other attributes.

For each of the WHERE conditions below, say which index you think is best to use and **why** you think it's the best.

Question 1 (3 points):

FinalBid <= 325 AND FinalBid >= 300 AND AuctionID = 1337

I2 is probably the best. Since AuctionID is a key, accessing the hash table will return at most one record, for which we can easily test the conditions on FinalBid. (I3 probably wouldn't work so well – it might return a substantial number of tuples, only one of which could satisfy the WHERE condition.)

Question 2 (3 points):

FinalBid <= 325 AND FinalBid >= 300 AND AuctionID <> 1337

I3 is really the only choice here, since a hash table can only be used with an equality condition. Note that all but one of the tuples retrieved using I3 will satisfy the WHERE condition.

Question 3 (4 points):

Item = 'iPod Nano' AND
(Category = 'Electronics' OR Category = 'Music')
AND FinalBid >= 10

I1 is likely the best choice. Using I3 on FinalBid > 10 is likely to return a large portion of the tuples. On the other hand, we can use I1 to get two groups of contiguous tuples: One group for (Item = 'iPod Nano' AND Category = 'Electronics'), and another group for (Item = 'iPod Nano' AND Category = 'Music').