INF3100: Databasesystemer

Oppgavesett 12

Oppgave 20.5.2: In this exercise, we need a notation for describing sequences of messages that can take place during a two-phase commit. Let *(i,j, M*) mean that site *i* sends the message *M* to site *j,* where the value of *M* and its meaning can be *P* (prepare), *R* (ready), *D* (don't commit), *C* (commit), or *A* (abort). We shall discuss a simple situation in which site 0 is the coordinator, but not otherwise part of the transaction, and sites 1 and 2 are the components. For instance, the following is one possible sequence of messages that could take place during a successful commit of the transaction:
(0,1,P), (0,2,P), (2,0,R), (1,0,R), (0,2,C), (0,1, C)

a) Give an example of a sequence of messages that could occur if site 1 wants to commit and site 2 wants to abort.

Oppgave 20.7.1: Given the circle of nodes of Fig. 20.14, where do key-value pairs reside if the key hashes to: *(a) 35 (b) 20 (c) 60?*

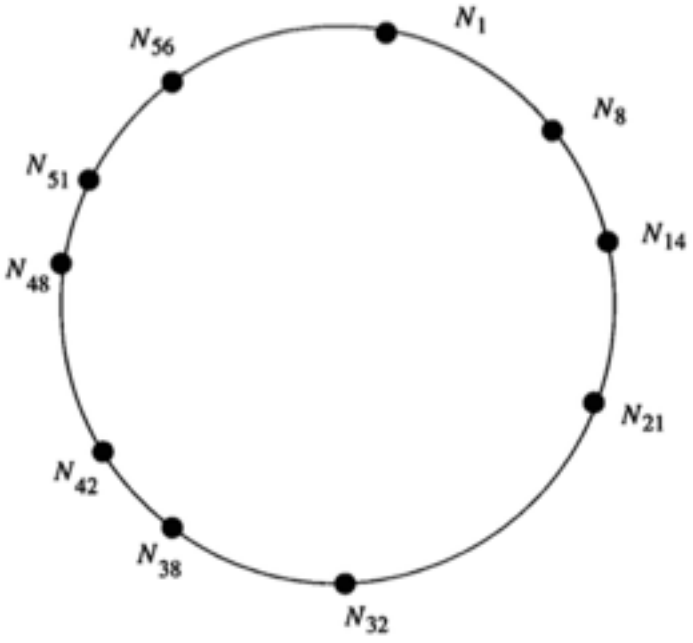

Figure 20.14: A chord circle

To place a node in the circle, we hash its ID i, and place it at position h(i). We shall henceforth refer to this node as N(hi). Thus, for example, in Fig. 20.14, N21 is a node whose ID i has h(i) = 21. The successor of each node is the next higher one clockwise around the circle. For example, the successor of N21 is N32 , and N1 is the successor of N56. Likewise, N21 is the predecessor of N32 , and N56 is the predecessor of N1.

The nodes are located around the circle using a hash function h that is capable of mapping both keys and node ID's (e.g., IP-addresses) to m-bit numbers, for some m. In Fig. 20.14, we suppose that m = ***6****, so there are 64 different possible locations for nodes around the circle. In a real application, m would be much larger.*

Key-value pairs are also distributed around the circle using the hash function h. If (K, V) is a key-value pair, then we compute h(K) and place (K, V) at the lowest numbered node Nj such that h(K) < j. As a special case, if h(K) is above the highest-numbered node, then it is assigned to the lowest-numbered node. That is, key K goes to the first node at or clockwise of the position h(K) in the circle.

Oppgave 20.7.2: Given the circle of nodes of Fig. 20.14, construct the finger tables for: (a) N14 (b) N51

Oppgave 20.7.3: Given the circle of nodes of Fig. 20.14, what is the sequence of messages sent if:
a) N14 searches for a key that hashes to 27
b) N8 searches for a key that hashes to 5
c) N56 searches for a key that hashes to 54

Oppgave 16.6.4: Consider the join of relations R(a,b), S(b,c), T(c,d), and U(a,d), where R and U each have 1000 tuples, while S and T each have 200 tuples. Further, there are 200 values of all attributes of all relations, except for attribute c, where
V(S,c) = V(T,c)=20.

a) What is the order selected by the greedy algorithm? What is its cost?
b) What is the optimum join ordering and its cost?