CS347 Homework Assignment

Due: Friday April 9, 2010 (5pm)

The following paper discusses an interesting fragmentation problem, where we use fragmentation to enhance the privacy of data:

Two Can Keep a Secret: A Distributed Architecture for Secure Database Services Available at: http://dbpubs.stanford.edu/pub/2004-42

(OK, I am biased regarding the "interesting" aspect since I am one of the authors :-)

(1) Read the paper, EXCEPT Sections 5.1, 5.2, 6.

(2) Consider a relation R with attributes X, A, B, C, D, E, F, G, H, I, J, K. Attribute X is the primary key of the relation. We know the following are privacy constraints:

{A, B, C, D}, {E, F}, {E, G}, {E, H}, {F, G, H}, {I, J}, {I, K}, {J, K}

We also know that in the affinity matrix all entries M[x, y] where x is not y have a value of 0, EXCEPT the entries

M[A, B], M[B, A], M[C, D], M[D, C], M[F, G], M[G, F]

which all have a value of 10. All entries M[x, x] have a value of 20 except for entries M[E, E] and M[I, I] which have a value of 1.

Describe a good, privacy-preserving fragmentation across 2 databases. You should encrypt as few attributes as possible. Briefly explain why your design is good. (Note there may be more than one good solution.)