

A relation is defined as a set of *tuples*. By definition, all elements of a set are distinct; hence all tuples in a relation must also be distinct. This means that no two tuples can have the same combination of values for all their attributes. Usually there are other **subsets of attributes** of a relation schema R with the property that no two tuples in any relation state r of R should have the same combination of values for these attributes.

Every relation has atleast one default superkey—the **set of all it's attributes**. A superkey can have redundant attributes, however, so a more useful concept is that of a *key*, which has no redundancy.

Hence a *key* is a *minimal superkey*—that is, a superkey from which we cannot remove any attributes and still have the uniqueness constraint hold.