Checking integrity constraints - how it differs in centralized, distributed and parallel databases

ABSTRACT

An important aim of a database system is to guarantee database consistency, which means that the data contained in a database is both accurate and valid. Integrity constraints represent knowledge about data with which a database must be consistent. The process of checking constraints to ensure that update operations or transactions which alter the database will preserve its consistency has proved to be extremely difficult to implement, particularly in distributed and parallel databases. In distributed databases the aim of the constraint checking is to reduce the amount of data needing to be accessed, the number of sites involved and the amount of data transferred across the network. In parallel databases the focus is on the total execution time taken in checking the constraints. This paper highlights the differences between centralized, distributed and parallel databases with respect to constraint checking.