A strategy for semantic integrity checking in distributed databases

ABSTRACT

Integrity constraints represent knowledge about data with which a database must be consistent. The process of checking constraints to ensure that the update operations or transactions which alter the database will preserve its consistency has proved to be extremely difficult to implement efficiently, particularly in a distributed environment. In the literature, most of the approaches/methods proposed for finding/deriving a good set of integrity constraints concentrate on deriving simplified forms of the constraints by analyzing both the syntax of the constraints and their appropriate update operations. These methods are based on syntactic criteria and are limited to simple types of integrity constraints. Also, these methods are only able to produce one integrity test for each integrity constraint. In Ibrahim, Gray, and Fiddian (1997), we introduced an integrity constraint subsystem for a relational distributed database. The subsystem consists of several techniques necessary for efficient constraint checking, particularly in a distributed environment where data distribution is transparent to application domain. However, the technique proposed for generating integrity tests is limited to several types of integrity constraints, namely: domain, key, referential and simple general semantic constraint and only produced two integrity tests (global and local) for a given integrity constraint. In this paper, we present a technique for deriving several integrity tests for a given integrity constraint where the following types of integrity constraints are considered: static and transition constraints.

Keyword: Integrity constraints; Distributed database; Integrity tests