

Research Directions in Moving Objects Databases

Goce Trajcevski, Ouri Wolfson, Bo Xu

University of Illinois - Chicago

Consider a database that represents information about moving objects and their location. For example, for a database representing the location of taxi-cabs a typical query may be: retrieve the free cabs that are currently within 1 mile of a customer at 33 N. Michigan Ave., Chicago; or for a trucking company database a typical query may be: retrieve the closest truck to truck ABT312 (which needs assistance); or for a database representing the current location of objects in a battlefield a typical query may be: retrieve the friendly helicopters that are expected to enter the region within the next 10 minutes. We will refer to applications with the above characteristics as moving-objects-database (MOD) applications, and to queries as the ones mentioned above as MOD queries. In the military, MOD applications arise in the context of the digital battlefield, and in the civilian industry they arise in transportation systems, mobile electronic commerce, and context aware mobile computing.

Currently, MOD applications are being developed in an ad hoc fashion. Database Management System (DBMS) technology provides a potential foundation upon which to develop MOD applications, however, DBMS's are currently not used for this purpose. The reason is that there is a critical set of capabilities that are needed by MOD applications and are lacking in existing DBMS's. The following is a discussion of the needed capabilities.

(A) Location Modeling.

Existing DBMS's are not well equipped to handle continuously changing data, such as the location of moving objects. The reason for this is that in databases, data is assumed to be constant unless it is explicitly modified. Thus, in order to represent moving objects (e.g. vehicles) in a database and answer queries about their location, the vehicle's location has to be continuously updated. This is unsatisfactory for performance and resource utilization considerations.

(B) Linguistic Issues.

Generally, a query in MOD applications involves spatial objects (e.g. points, lines, regions, polygons) and temporal constraints. Consider for example the query: "Retrieve the objects that will intersect the polygon P within the next 3 minutes". This is a spatial and temporal range query. The spatial range is the polygon P, and the temporal range is the time interval between now and 3 minutes from now. Similarly, there are spatio-temporal join queries such as: "Retrieve the pairs of friendly and enemy aircraft that will come within 10 miles of each other, and the time when this will happen." Traditional query languages such as SQL are inadequate for expressing such queries. Although spatial and temporal languages have been studied in the database research community, the languages are inappropriate for MOD applications.

(C) Indexing.

Observe that the number of moving objects in the database may be very large (e.g., in big cities with millions of inhabitants). Thus, for performance considerations, in answering MOD queries we would like to avoid examining the location of each moving object in the database. In other words, we would like to index the location attribute. The problem with a straight-forward use of spatial indexing for this purpose is that the continuous change of the locations implies that the spatial index has to be continuously updated. This is clearly an unacceptable solution.

(D) Uncertainty/Imprecision.

The location of a moving object is inherently imprecise because, regardless of the policy used to update the database location of the object (i.e. the object-location stored in the database), the database location cannot always be identical to the actual location of the object. This inherent uncertainty has various implications for database modeling, querying, and indexing.

In this work we will address the above problems, and present possible solutions.