IDEA GROUP PUBLISHING
701 E. Chocolate Avenue, Suite 200, Hershey PA 17033-1240, USA
Tel: 717/533-8845; Fax 717/533-8661; URL-http://www.idea-group.com
ITB10787

Chapter X

Management of Large Moving Objects Databases:
Indexing, Benchmarking and Uncertainty in Movement Representation

Talel Abdessalem
Ecole Nationale Supérieure des Télécommunications, France

Cédric du Mouza
Conservatoire National des Arts et Métiers, France

José Moreira, Universidade de Aveiro, Portugal

Philippe Rigaux, University of Paris Sud, France

Abstract

This chapter deals with several important issues pertaining to the management of moving objects datasets in databases. The design of representative benchmarks *is closely related to the formal characterization of the properties (that is, distribution, speed, nature of movement) of these datasets;* uncertainty *is another important aspect that conditions the accuracy of the representation and therefore the confidence in query results; finally, efficient* index structures, *along with their compatibility with existing softwares, is a crucial requirement for spatio-temporal databases, as it is for any other kind of data.*

Introduction

A lot of emerging applications (traffic control, mobile computing, vehicles tracking) rely on large datasets of dynamic objects. This proliferation is encouraged by mature technologies (for example, the Global Positioning System, or GPS) that provide online information on mobile devices and enable communication between a centralized system and mobile users. Most of the services that can be provided by the system to a user are based on the location of the latter at a given instant (for instance, the case of company searching for the taxi nearest to a customer calling on his mobile phone, or a tourist in his car looking for his next hotel). But, apart from these so-called location-based services that deal with the present or future (expected) positions of objects, we can also envisage applications that study the past movements within large moving objects datasets (for data-mining purposes, for instance).

These examples illustrate some new requirements that address the core functionalities of Database Management Systems (DBMS). Indeed, we must consider new data models (as any previously proposed model falls short in representing continuous movements), new query languages and new system-level support. In this chapter we focus on the latter aspect. More specifically, we propose a survey of the following issues: benchmarking of operations on large moving objects datasets, uncertainty in trajectories representation and database indexing. Let us be more specific by shortly developing each topic.

Benchmarking

In computing, a benchmark is the result of running a set of standard tests on one component or system to compare its performance and capacity to that of other components or systems. They are designed to simulate a particular type of workload, running actual real-world programs on the system "application benchmarks," or using specially created programs that impose the workload on the component "synthetic benchmarks." Application benchmarks are meant to be representative of real-world applications and potentially give a better measure of real-world performance. On the other hand, synthetic benchmarks offer a sizeable workload of data sets and operations, allowing testing individual components (such as indexing methods or hard disks) and stressing the strengths and weaknesses of each one individually. In the spatio-temporal database context, benchmarks help to experiment with new approaches (for example, new languages or new indexing structures); they can be used to assess the effectiveness of a new system; and finally, they contribute to characterizing the properties of datasets.