Chapter 32 HERMES: A Trajectory DB Engine for Mobility-Centric Applications

Nikos Pelekis University of Piraeus, Greece

Elias Frentzos University of Piraeus, Greece Nikos Giatrakos University of Piraeus, Greece

Yannis Theodoridis University of Piraeus, Greece

ABSTRACT

This chapter presents HERMES, a prototype DB engine that defines a powerful query language for trajectory databases, which enables the support of mobility-centric applications, such as Location-Based Services (LBS). HERMES extends the data definition and manipulation language of Object-Relational DBMS (ORDBMS) with spatio-temporal semantics and functionality based on advanced spatio-temporal indexing and query processing techniques. Its implementation over two ORDBMS and its utilization in various domains proves the expressive power and applicability of HERMES in different application domains where knowledge regarding mobility data is essential. As a proof-of-concept, in this chapter HERMES is applied to a case study related with vehicle traffic analysis, demonstrating its flexibility and usefulness for delivering custom-defined LBS.

INTRODUCTION

Moving Object Databases (MOD) (Güting et al., 2000) and more specifically Trajectory Databases (TD) being at the core of spatio-temporal database research, have emerged due to the explosion of mobile devices and positioning technologies. A MOD is the basic component of any mobility-centric application (Kargin, Basoglu, & Daim, 2009). However, although such LBS applications are already in the air for some years, the services currently provided are rather naive, not exploiting the current software capabilities and the recent advances in MOD research field. We argue that one of the reasons for this is due to the common practice in existing approaches, which provides services to mobile users by just taking into account the current location-time and velocity information, arriving at the MOD server as a

DOI: 10.4018/978-1-4666-9845-1.ch032

sequence of updates. Given this model and the fact that LBS applications need to handle huge volumes of data, it rationally arises that performance is a significant problem; therefore, efficient query processing and indexing techniques should be applied. Moreover, this model has limited applicability in real-world applications, since safe estimations about future positions should involve past positions as well.

The key observation that motivates Hermes is that the more the knowledge in hand about the trajectory of a mobile user, the better the exploitation of the advances in spatio-temporal query processing for providing intelligent LBS. Based on this motivation, the aim of this paper is to describe a robust framework capable of aiding either an analyst working with mobility data, or more technically, a developer who models, queries a TD and builds a mobility-centric application on top of the TD.

Moreover, given the ubiquitousness of location-aware devices, databases handling moving objects will, sooner or later, face enormous volumes of data. It consequently arises that performance in the presence of vast data sizes, is a significant problem for moving object databases and the only way to deal with such enormous sizes is the exploitation of specialized access methods used for spatio-temporal indexing purposes. The domain of spatio-temporal indexing, as well as other related domains, such as multimedia (Chatterjee, & Chen, 2010) and spatial indexing, is dominated by the presence of the R-tree, along with its variations and extensions. Among others, 3D R-trees (Theodoridis, Vazirgiannis, & Sellis, 1996), TB-trees and STR-trees (Pfoser, Jensen, & Theodoridis 2000), and PA-trees (Ni, & Ravishankar, 2007) are considered as extensions of the R-tree in the spatio-temporal domain. As in the case of appropriate moving object data types and methods for extending the type system of ORDBMS, except the well-known R-trees, which are suitable only for static spatial data, none of the above proposals have been incorporated into existing ORDBMS. Among them, the Trajectory Bundle tree (TB-tree) (Pfoser et al., 2000), is adopted in this work and appropriately designed and implemented inside Hermes taking advantage of the indexing extensibility interface of ORDBMS. Being a member of the R-tree family, TB-tree is able to support traditional queries such as range and distance-based queries. At the same time, it supports objects moving on the unconstrained space, and is the only one that fulfills the need for trajectory preservation so as to efficiently support trajectory-based operations.

Furthermore, apart from simple query operators (e.g. range queries) natively supported by R-trees, there is a variety of spatio-temporal operators which require more sophisticated query processing techniques in order to be efficiently processed. Among them, an important class of queries is the so-called k nearest neighbor (k-NN) search, where one is interested in finding the k closest trajectories to a predefined query object Q (stationary or moving). Thus, one of the challenges being present in the domain of trajectory databases is to develop mechanisms to perform k-NN search on MODs exploiting spatio-temporal indexes storing historical information. Among the solutions proposed in the literature we adopt the one proposed by (Frentzos, Gratsias, Pelekis, & Theodoridis 2007) which efficiently supports Nearest Neighbor (NN) queries over historical trajectory data.

Finally, as we aim at providing a powerful toolkit for analysts, Hermes provides qualitatively different techniques for trajectory similarity search, which is exploited to support trajectory clustering and classification mining tasks that imply a way to quantify the distance between two trajectories. More specifically, we adopt a novel set of trajectory distance functions (Pelekis, Kopanakis, Ntoutsi, Marketos, Andrienko & Theodoridis, 2007; Pelekis, Andrienko, Andrienko, Kopanakis, Marketos, & Theodoridis, 2010) based on primitive (space and time) as well as derived parameters of moving objects (speed, acceleration, and direction), which are also capable to support sub-trajectory similarity matching. The overall framework advances the contribution of our approach by two inter-related facts: firstly, the combination of the similarity operators in the extended with MOD semantics SQL-like query language (using AND/ 24 more pages are available in the full version of this document, which may

be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/hermes/149520

Related Content

Mapping the Hopi Landscape for Cultural Preservation

Saul L. Hedquist, Stewart B. Koyiyumptewa, Wesley Bernardini, T. J. Ferguson, Peter M. Whiteleyand Leigh J. Kuwanwisiwma (2015). *International Journal of Applied Geospatial Research (pp. 39-58).* www.irma-international.org/article/mapping-the-hopi-landscape-for-cultural-preservation/121570

Framework for Graphical User Interfaces of Geospatial Early Warning Systems

Martin Hammitzsch (2013). Geographic Information Systems: Concepts, Methodologies, Tools, and Applications (pp. 449-464).

www.irma-international.org/chapter/framework-graphical-user-interfaces-geospatial/70455

Valid Space Description in BIM for 3D Indoor Navigation

Abdoulaye Abou Diakitéand Sisi Zlatanova (2016). International Journal of 3-D Information Modeling (pp. 1-17).

www.irma-international.org/article/valid-space-description-in-bim-for-3d-indoor-navigation/177237

Developing an E-Planning System Compatible With Smart City Design Principles: Case of Bayrakli

Tugce Altinkilitand Muhammed Aydogan (2022). International Journal of Digital Innovation in the Built Environment (pp. 1-18).

www.irma-international.org/article/developing-an-e-planning-system-compatible-with-smart-city-design-principles/306254

GeoCache: A Cache for GML Geographical Data

Lionel Savary, Georges Gardarinand Karine Zeitouni (2009). *Handbook of Research on Geoinformatics* (pp. 350-368).

www.irma-international.org/chapter/geocache-cache-gml-geographical-data/20422