Chapter 3
A Survey of Parallel Indexing Techniques for Large-Scale Moving Object Databases

Eleazar Leal
University of Minnesota – Duluth, USA

Le Gruenwald
University of Oklahoma, USA

Jianting Zhang
City College of New York, USA

ABSTRACT

A moving object database is a database that tracks the movements of objects. As such, these databases have business intelligence applications in areas like trajectory-based advertising, disease control and prediction, hurricane path prediction, and drunk-driver detection. However, in order to extract knowledge from these objects, it is necessary to efficiently query these databases. To this end, databases incorporate special data structures called indexes. Multiple indexing techniques for moving object databases have been proposed. Nonetheless, indexing large sets of objects poses significant computational challenges. To cope with these challenges, some moving object indexes are designed to work with parallel architectures, such as multicore CPUs and GPUs (graphics processing units), which can execute multiple instructions simultaneously. This chapter discusses business intelligence applications of parallel moving object indexes, identifies issues and features of these techniques, surveys existing parallel indexes, and concludes with possible future research directions.

DOI: 10.4018/978-1-5225-4963-5.ch003

INTRODUCTION

Through the use of location-sensing devices, very large moving object datasets can be collected. These datasets make it possible to issue spatio-temporal queries with which users can gather real-time information about the characteristics of the movements of objects involved in these datasets, derive patterns from that information, and then make decisions based on these patterns. Examples of these large datasets are Geolife (Zheng, Xie, & Ma, 2010) and T-drive (Yuan et al., 2010).

Geolife consists of trajectories (i.e., the time-ordered sequence of positions that an object occupies in time) collected with the use of GPS phones by researchers of Microsoft Research Asia as they went through their daily lives. T-drive contains the GPS logs of the positions occupied by taxis in Beijing. Both datasets are large: Geolife contains 17,000+ trajectories whose lengths add up to 1,200,000+ kilometers, and span an interval of 48,000+ hours. T-drive, on the other hand, contains the trajectories of 10,000+ taxis, whose lengths span 9,000,000 Km. Datasets like these can be used to support decisions in the transportation domain, such as: helping taxis find the fastest routes by mining taxi trajectories (Yuan et al., 2010), finding driving directions, and for urban planning (Wang, Zheng, & Xue, 2014).

Other uses of moving object datasets are the following: in epidemiology, to help centers for disease control and prevention make decisions on how to avoid the spread of the avian influenza, by tracking and studying the movements of mallards (Hill et al., 2017); in meteorology, to help predict the path of a developing hurricane, which can be done by exploiting the tendency of hurricanes to follow similar trajectories, thereby aiding meteorologists issue more accurate recommendations on which areas must be evacuated (Li et al., 2010); in law enforcement, to automatically detect drunk drivers, and then help police departments make decisions on how to better allocate their police force by area (Ge et al., 2010); in trajectory-based mobile advertising (Ammar, Elsayed, Sabri, & Terry, 2015), where shopping malls, by tracking the positions of shoppers using the mall's WiFi, can increase their revenue by sending online advertising that has been tailored to the shoppers based on their movement patterns around the mall (Ghose, 2017); for city planning in places like Shanghai, to help planners decide where to build new bike lanes, while taking into account Shanghai's budget limitations, and the way existing bike lanes are utilized (Bao, He, Ruan, Li, & Zheng, 2017); in online trajectory-sharing applications, to suggest attractive travel destinations based on the trajectories that others have enjoyed (Zheng et al., 2010); and in sports, to deduce the common plays of a given sports team (Buchin, Dodge, & Speckmann, 2014) from video footage, and then help coaches make decisions about their team's next play.