A secondary partitioning technique for indexing multidimensional ranges

Speaker: Professor Nikos Mamoulis Department of Computer Science and Engineering, University of Ioannina

Date: February 24, 2022 (Thursday) Time: 2:00pm (HK Time) Format: Online via Zoom (Registration is required)

Link for Registration: https://hku.zoom.us/meeting/register/tJUud uioqzMsH9Bsf11CKLAxPh2s8n09OOTy





While research has primary focused on the effective management of scalars or vectors, i.e., (multidimensional) points, in many applications there is also a great need for indexing (multidimensional) ranges, i.e., intervals or multidimensional boxes. This problem is also a classic one and it has been supported by effective access methods, such as the interval tree and the R-tree. Recently, given the fact that main memories have become bigger and cheaper, there is an increasing interest for efficient main-memory indices, where the objective is to minimize the computational cost of queries and not the number of I/O accesses. Based on this, we revisit the design and implementation of indexing (multidimensional) ranges. Driven by the fact that the most efficient methods for points are those which define and use disjoint space partitions, we focus on adaptations of such methods for ranges. An inherent problem of indexing ranges using disjoint partitions is that a data unit may have to be assigned to multiple partitions, which causes data redundancy and duplicate results in queries. We propose a simple but effective secondary data partitioning, which (i) reduces redundancy, (ii) avoids the generation of duplicate results, and (iii) minimizes the number of necessary comparisons. For 1D intervals, we propose HINT, an efficient hierarchical index that has controlled space requirements and includes storage optimization techniques for the effective handling of data sparsity and skewness. Our indices target range overlap queries, which are a basic component of many search and analysis tasks. Experimental results on real and synthetic data show that our approaches are typically one order of magnitude faster than the state-of-the-art indexing methods.



XXXX			
Оххх		1xxx	
00xx	01xx	10xx	11xx

Presenter's short biography:

Nikos Mamoulis received a diploma in Computer Engineering and Informatics in 1995 from the University of Patras, Greece, and a PhD in Computer Science in 2000 from the Hong Kong University of Science and Technology. He is currently a Professor at the Department of Computer Science and Engineering at the University of Ioannina. He was a faculty member at the Department of Computer Science, University of Hong Kong from 2001 to 2017. His main research interests are in the management, search and analysis of complex data, with a focus on spatially enriched data. He served as PC co-chair of SSTD 2009, COMAD 2013, and HDMS 2014, as a general chair of SSDBM 2008, and as a PC member in more than 100 conferences. He has been an Associate Editor for the VLDB Journal, IEEE Transactions on Knowledge and Data Engineering, Knowledge and Information Systems, Geoinformatica Journal. He is a Senior ACM member. He has published over 200 papers in top international conferences and journals of data management and mining.