ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ Η/Υ & ΠΛΗΡΟΦΟΡΙΚΗΣ • ΠΑΝΕΠΙΣΤΗΜΙΟ ΙΩΑΝΝΙΝΩΝ
DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING • UNIVERSITY OF IOANNINA

ΠΑΡΟΥΣΙΑΣΗ
ΔΙΔΑΚΤΟΡΙΚΗΣ ΔΙΑΤΡΙΒΗΣ

ΗΜΕΡΟΜΗΝΙΑ: Τρίτη, 17 Δεκεμβρίου 2024

ΩΡΑ: 11:30 – 12:30

ΑΙΘΟΥΣΑ: Αίθουσα Σεμιναρίων ΤΜΗΥΠ

ΟΜΙΛΗΤΗΣ: **Δημήτρης Τσιτσίγκος**

Θ έ μ α

«In-Memory Indexing for Parallel Processing of Single and Multi-Dimensional Queries»

Επταμελής Εξεταστική Επιτροπή:

1. **Nikolaos Mamoulis**, Professor, Department of Computer Science and Engineering, University of Ioannina
2. **Panagiotis Vassiliadis**, Professor, Department of Computer Science and Engineering, University of Ioannina
3. **Manolis Terrovitis**, Professor, Information Management Systems Institute (IMSI) of the Research and Innovation Centre in Information, Communication and Knowledge Technologies "Athena"
4. **Evaggelia Pitoura**, Professor, Department of Computer Science and Engineering, University of Ioannina
5. **Vassilios V. Dimakopoulos**, Professor, Department of Computer Science and Engineering, University of Ioannina
6. **Panagiotis Bouros**, Assistant Professor, Institute of Computer Science, Johannes Gutenberg University Mainz
7. **Christos Doulkeridis**, Professor, Department of Digital Systems, University of Piraeus

ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ Η/Υ & ΠΛΗΡΟΦΟΡΙΚΗΣ
ΠΑΝΕΠΙΣΤΗΜΙΟ ΙΩΑΝΝΙΝΩΝ
Τ.Θ. 1186, ΙΩΑΝΝΙΝΑ 45110
Τ: 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/

DEPT. OF COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF IOANNINA
P.O. BOX 1186, IOANNINA
GR - 45110 GREECE
T: +30 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/

Π ε ρ ί λ η ψ η:

Database Systems are essential for modern applications, providing structured, efficient, and reliable ways to manage data in fields like banking, healthcare, and scientific research. They are designed to handle large amounts of data and answer queries quickly. A key factor in making queries fast is the use of indices. Efficient indexing can significantly improve the speed, credibility, and overall performance of a database. A good index supports fast searches and updates, while having low space requirements. With the advancements in hardware, we can redesign index structures to be faster and more efficient. For example, larger memory capacities allow us to move computations from disk storage to faster in-memory processing. Additionally, the latest processors offer significantly more cores, enhancing parallel computing by distributing tasks across multiple cores to improve performance. These modern hard-ware components are affordable and found in commodity computers, making them accessible for a wide range of applications.

This dissertation explores how in-memory parallel indexing techniques can improve the performance of relational, temporal, and spatial databases. It addresses challenges like managing large-scale data and running complex queries on modern systems.

Interval joins are crucial for temporal databases and are also useful in many applications. However, a major challenge in parallel implementations of interval joins is dividing the workload effectively across processor cores. A simple approach is to split the data domain into disjoint partitions and assign the data in each partition to a different core. However, this creates a new problem: when an interval spans multiple stripes, it must be processed by multiple cores, possibly generating duplicate join results. One way to handle duplicates is by using a data structure like a set to store unique results. However, this increases memory usage and slows down query performance because of the extra checks for duplicates. This research introduces a novel domain-based partitioning approach to divide data into partitions that can be processed independently and in parallel, maximizing multi-core hardware capabilities. Three distinct strategies for partitioning, applicable to both hash-based and domain-based methods, are proposed. These strategies significantly accelerate interval join operations by minimizing computational overhead and enhancing scalability.

Indexing spatial data is challenging due to their shape and dimensionality. This complexity affects query performance, particularly for spatial intersection joins, which are the most resource-intensive. Like temporal databases, spatial databases also face duplication issues. For example, in a 2D grid, duplicates can occur when spatial objects overlap multiple grid tiles. The most common technique to avoid duplicates is the reference point approach, which reports a result in only one tile of the grid but requires additional computations. This dissertation primarily focuses on non-point data, where duplication issues arise.

ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ Η/Υ & ΠΛΗΡΟΦΟΡΙΚΗΣ
ΠΑΝΕΠΙΣΤΗΜΙΟ ΙΩΑΝΝΙΝΩΝ
Τ.Θ. 1186, ΙΩΑΝΝΙΝΑ 45110
Τ: 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/

DEPT. OF COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF IOANNINA
P.O. BOX 1186, IOANNINA
GR - 45110 GREECE
T: +30 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/

Our first study focuses on improving the Partition-Based Spatial Join (PBSM) algorithm for in-memory and parallel evaluation of spatial joins. We show how to choose the best partitioning settings based on data statistics to fine-tune the algorithm for specific join inputs. In our second study, we introduce a new secondary partitioning technique for space-oriented partitioning (SOP) indices, such as grids. This technique removes duplicate results during spatial queries by organizing objects within each partition and only accessing classes that do not produce duplicates. This innovation greatly improves the efficiency of grid-based spatial indices for range (disk and rectangle) and intersection joins. Finally, we propose a parallel method that boosts the performance of range queries.

Relational databases also face challenges in creating efficient indices. This dissertation focuses only on tree-like indices. One of the most well-known and efficient indexing structures in this domain is the B^+-tree, particularly suited for skewed workloads with dynamic data and for supporting range queries. With the rise of machine learning, many learned indices have been introduced. A learne index uses ML models to “learn” the data distribution and predict the location of the search key within a dataset aiming to reduce the space requirements and the memory accesses. The main difference between a traditional B^+-tree and a learned index is that learned indices replace inner nodes with machine learning models. We believe that B^+-tree are more efficient than learned indices because their query performance is stable and not affected by the data distribution. We also believe that advancements in hardware technology create new opportunities to further improve B-tree performance. Building on the foundational B^+-tree structure, this research introduces B^S-tree, a new indexing structure designed for main memory and modern hardware. The B^S-tree leverages data parallelism and integrates innovative optimizations, offering significant advancements over both traditional B^+-tree and emerging learned indices. Key features include a data-parallel branching mechanism implemented using SIMD instructions, a gap management strategy employing duplicate keys to delay splits and reduce data-shifting overhead, and a node compression scheme that minimizes memory usage while maintaining high throughput.

In summary, this dissertation provides a comprehensive framework for enhancing the performance of relational, temporal, and spatial databases. Through the integration of in-memory indexing, parallel processing techniques and duplicate avoidance techniques, this dissertation delivers robust, scalable, and efficient solutions for the evolving needs of modern data-intensive applications.

ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ Η/Υ & ΠΛΗΡΟΦΟΡΙΚΗΣ
ΠΑΝΕΠΙΣΤΗΜΙΟ ΙΩΑΝΝΙΝΩΝ
Τ.Θ. 1186, ΙΩΑΝΝΙΝΑ 45110
T: 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/

DEPT. OF COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF IOANNINA
P.O. BOX 1186, IOANNINA
GR - 45110 GREECE
T: +30 265100 8817 - 8813 - 7196
http:/www.cse.uoi.gr/