

Tobias Lauer
 University of Freiburg
 Georges-Köhler-Allee 051
 79110 Freiburg, Germany
 lauer@informatik.uni-freiburg.de

Amitava Datta
 University of Western Australia
 35 Stirling Highway
 Perth, WA 6009, Australia
 datta@csse.uwa.edu.au

Zurab Khadikov
 Jedox AG
 Bismarckallee 7a
 79098 Freiburg, Germany
 zurab.khadikov@jedox.com

Online Analytical Processing (OLAP)

OLAP is a core technology in *Business Intelligence* and *Corporate Performance Management*, allowing users to navigate and explore corporate data (usually extracted from a data warehouse) and to roll up or drill down along different hierarchical levels. Also, updates to the data must be supported for planning and forecasting.

Due to the highly interactive nature of OLAP analysis, query performance is a key issue.

Multidimensional Aggregation

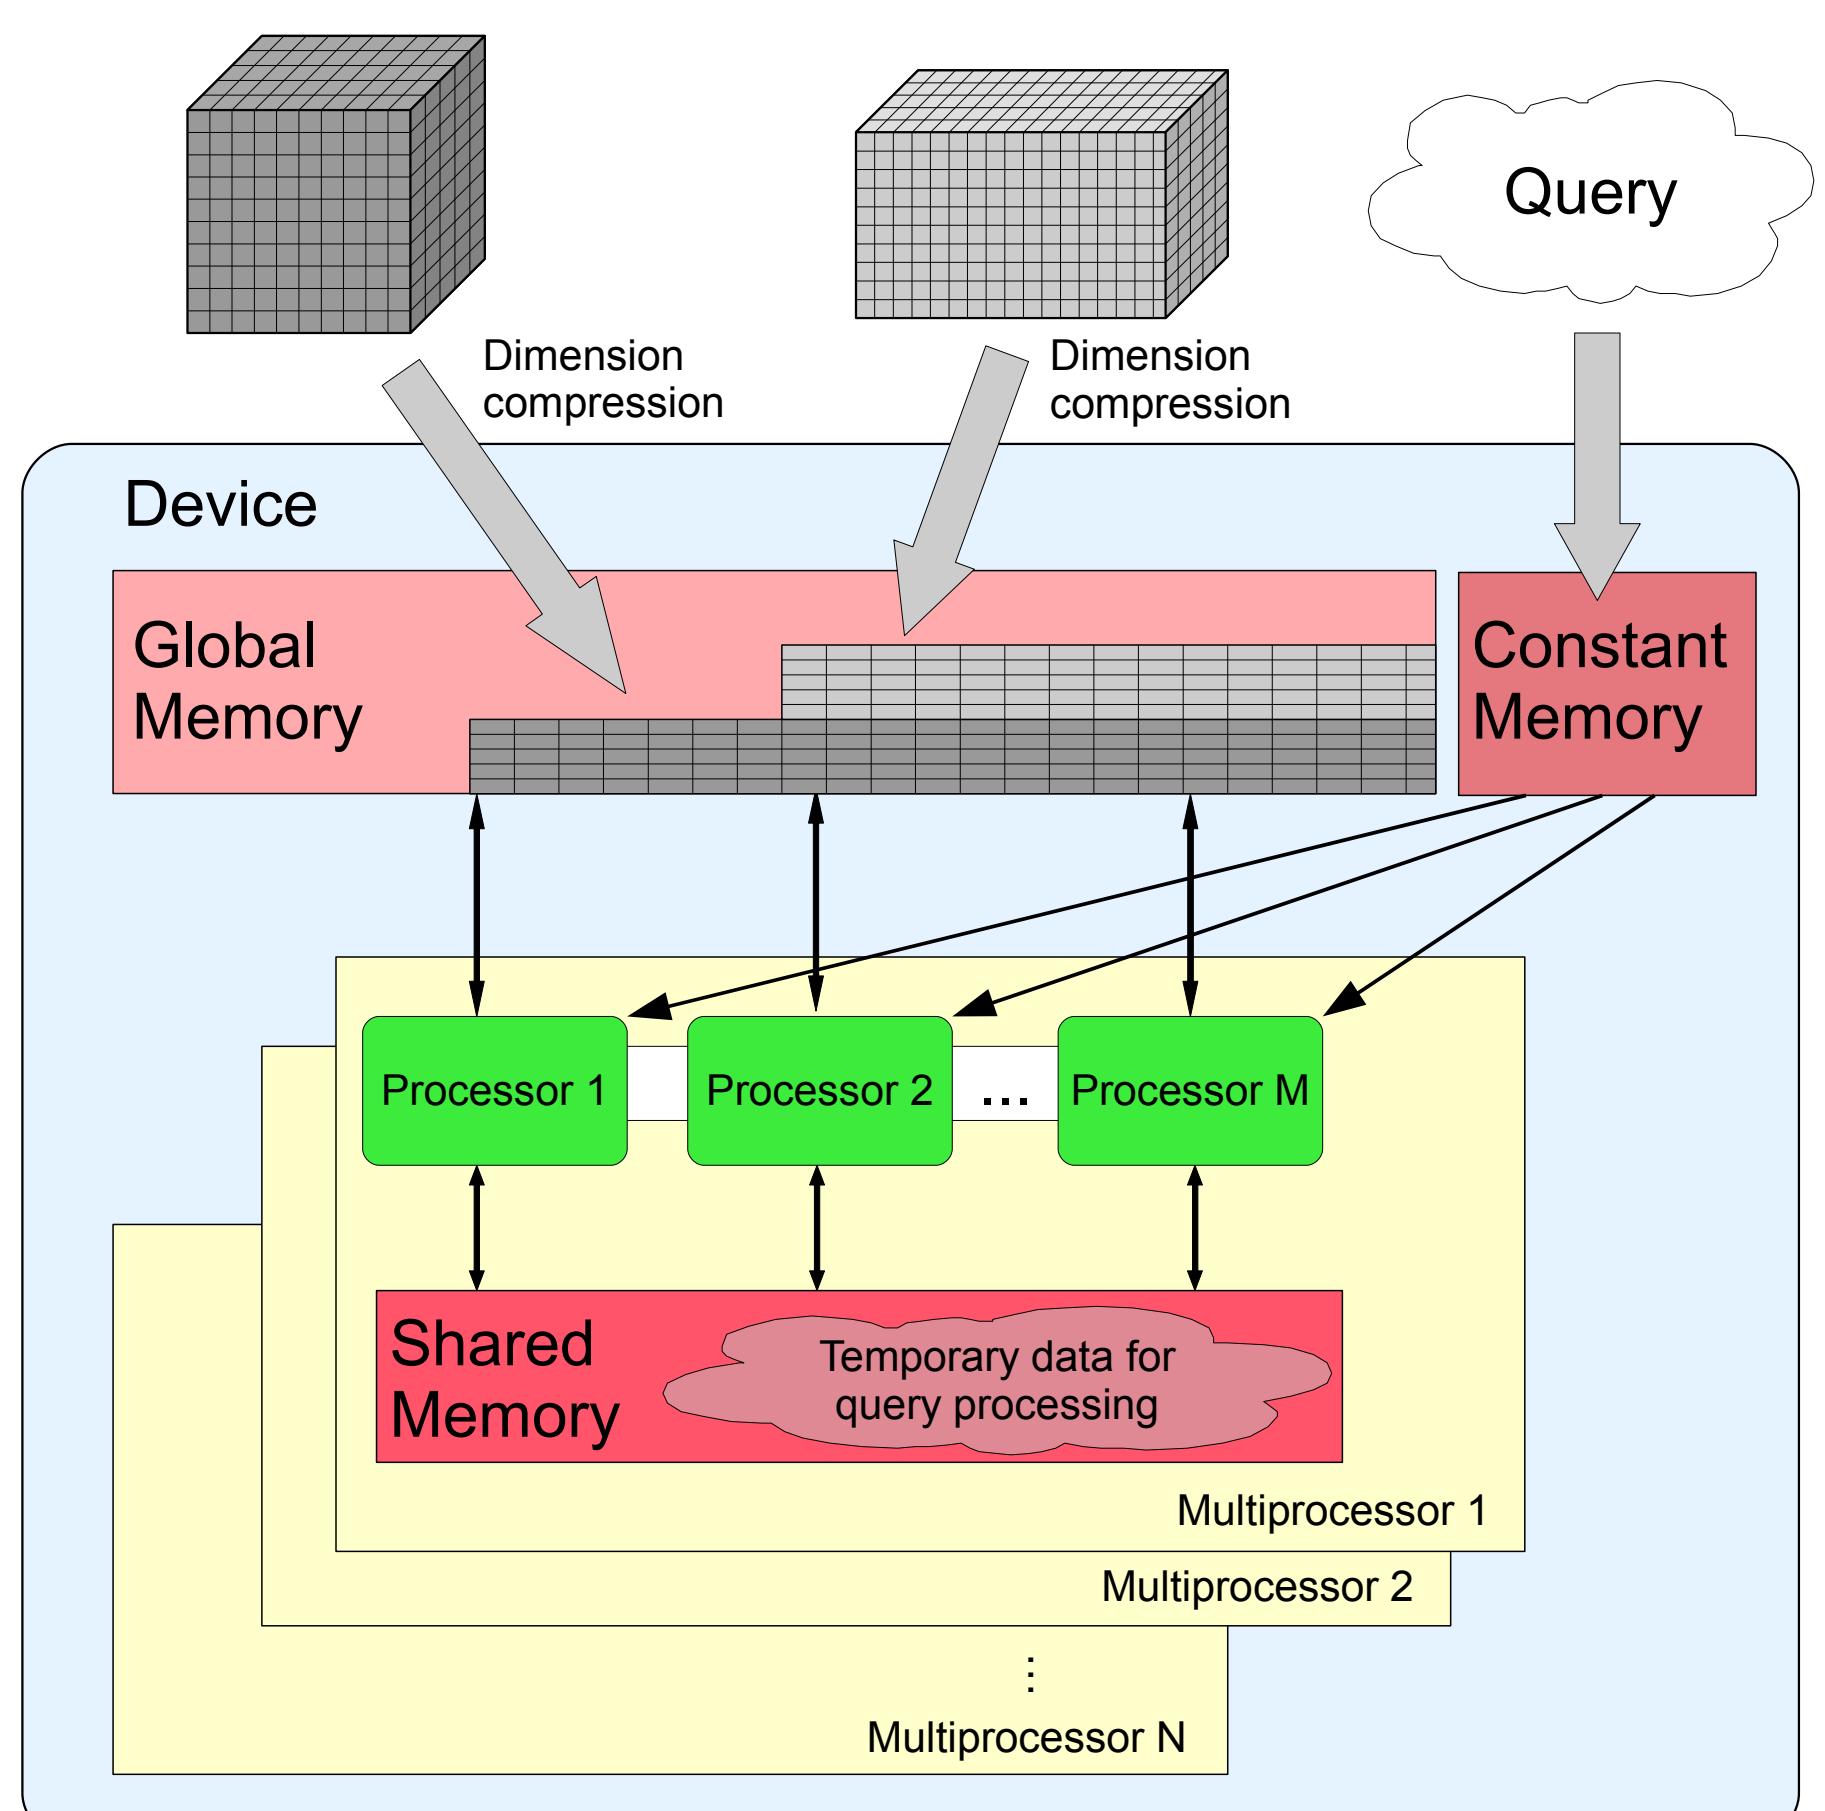
The conceptual model central to OLAP is the *Data Cube*, which is a view of the data as cells in a multidimensional table ("cube").

Aggregation along dimensional hierarchies is a basic building block involved in most OLAP operations. The main problems to solve in order to compute aggregates efficiently is the *sparsity* of data in high-dimensional spaces and the dimensional "explosion", since the number of possible aggregates grows exponentially.

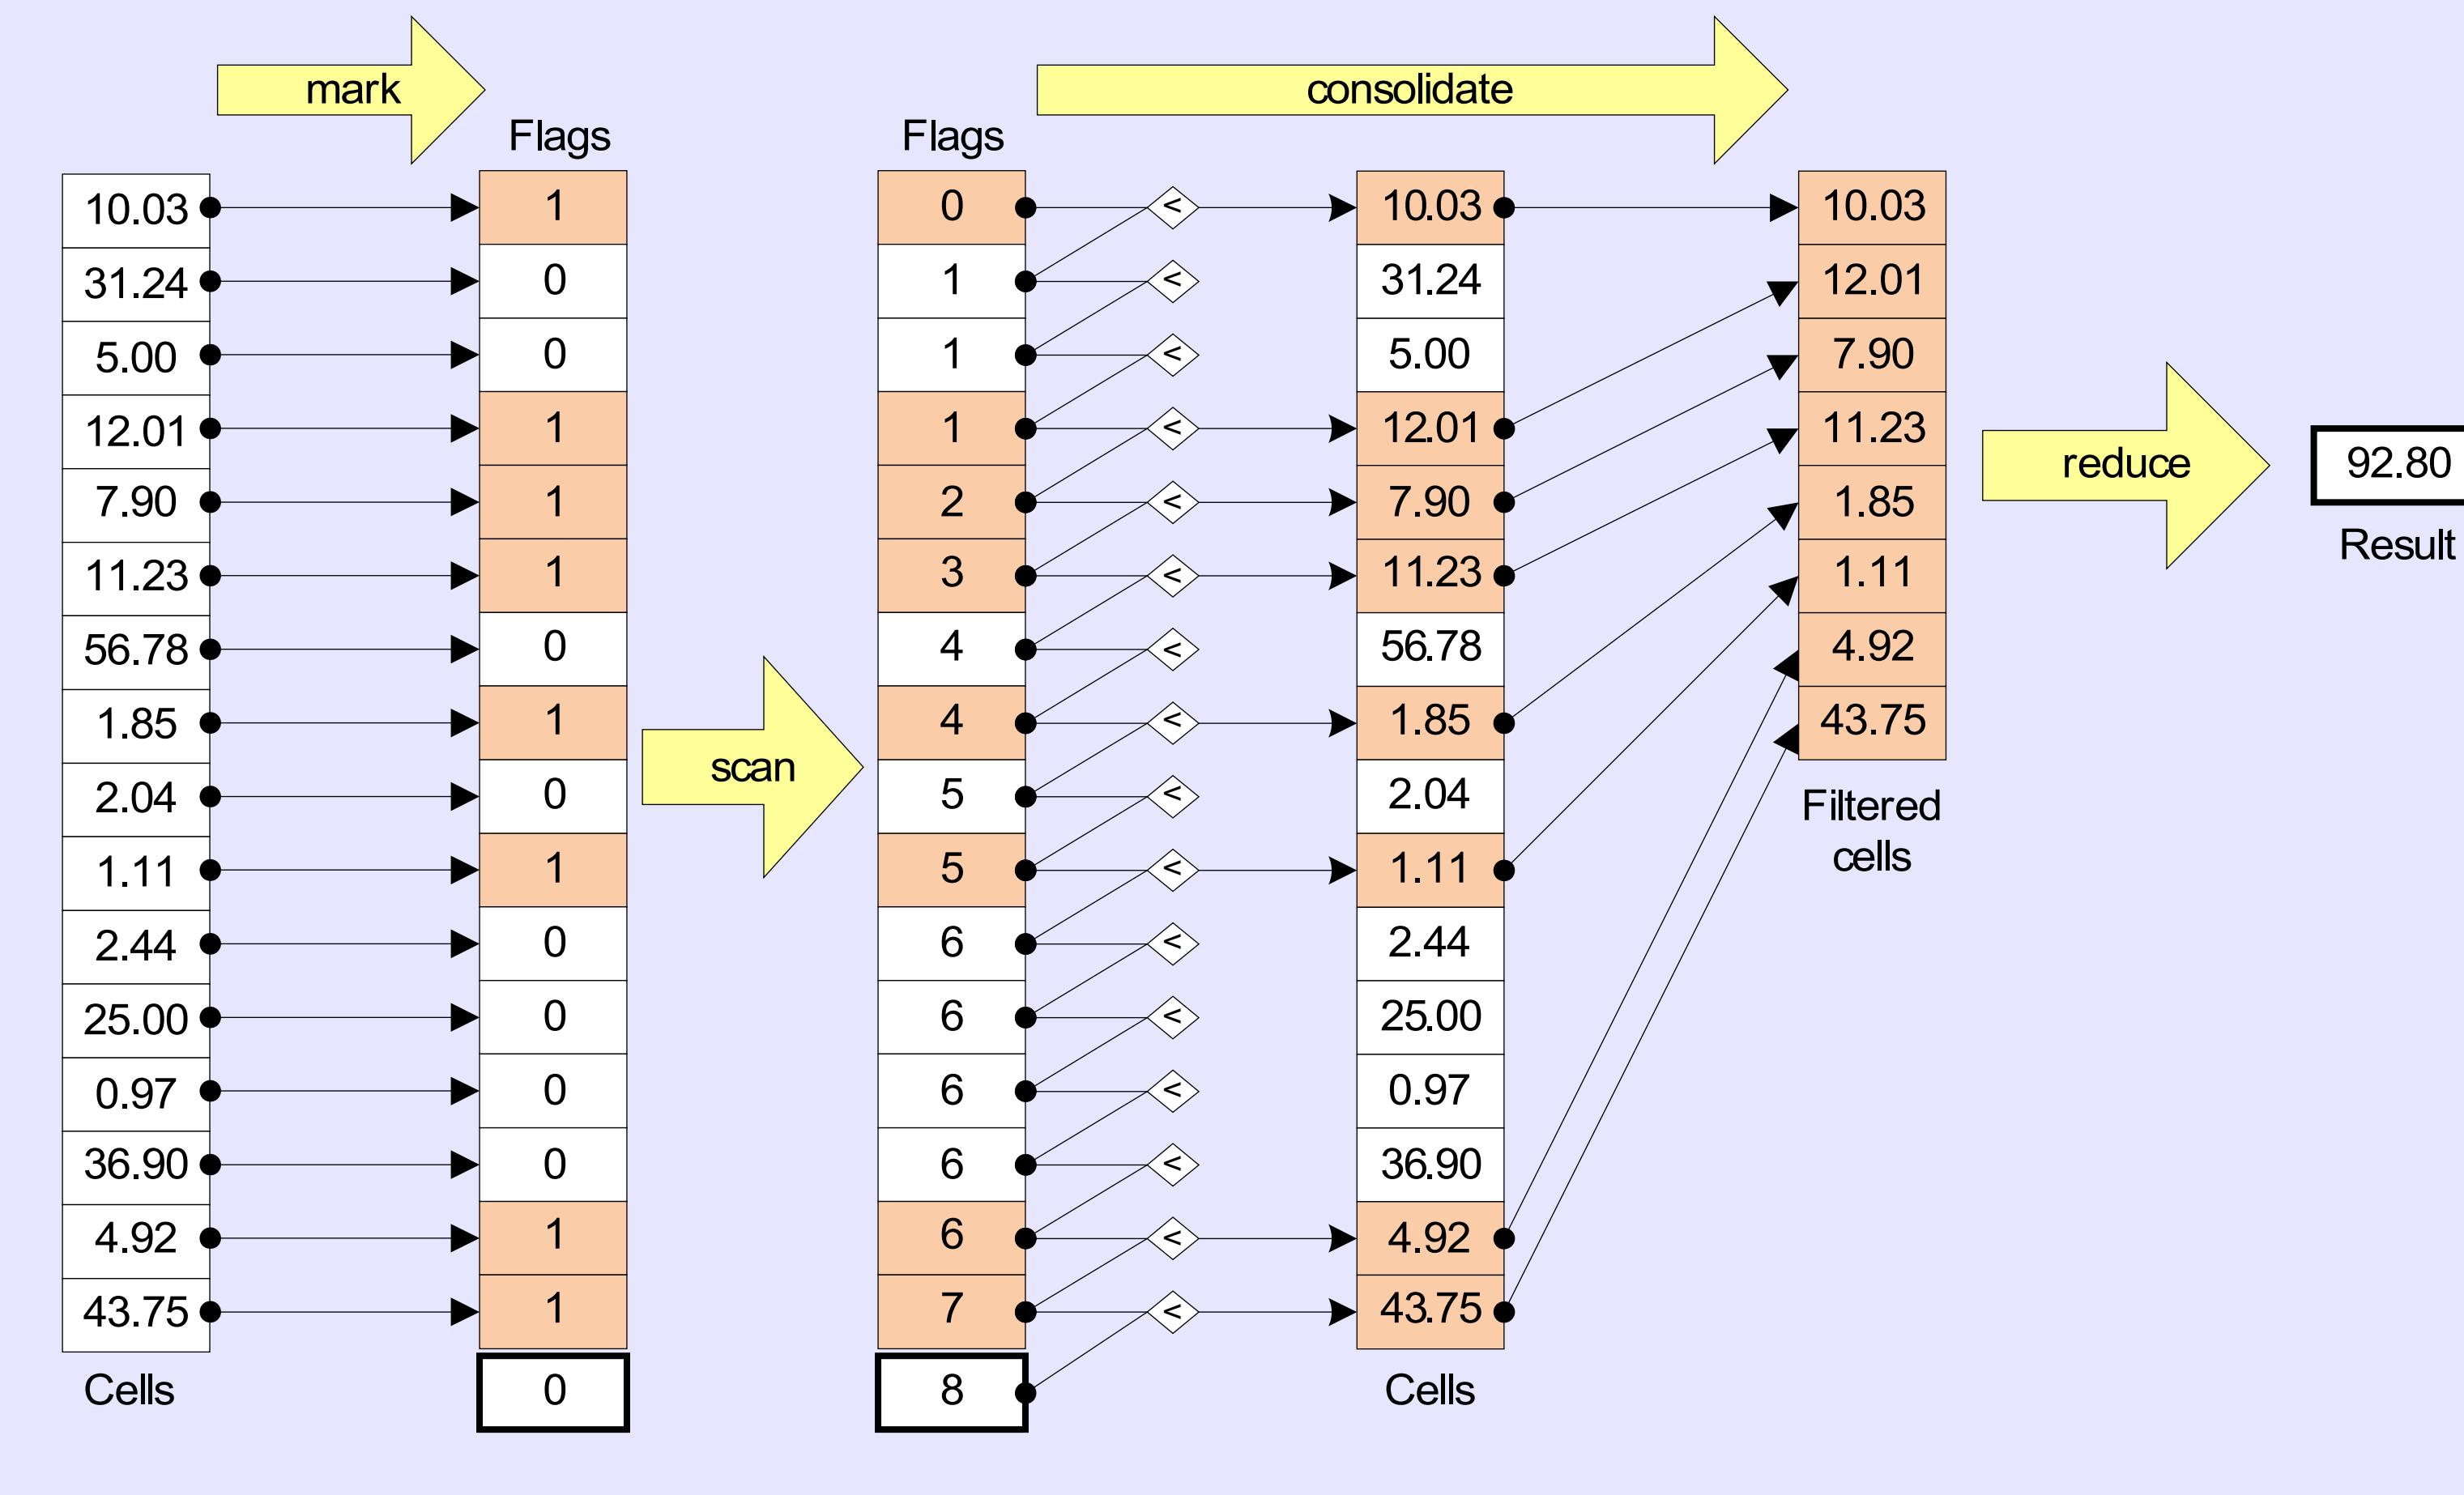
Our approach aims at speeding up aggregations by using the massively parallel architecture of GPUs.

Data Storage on GPU

Usage of different GPU memory types for storage and processing of data cubes:



GPU Aggregation Algorithm

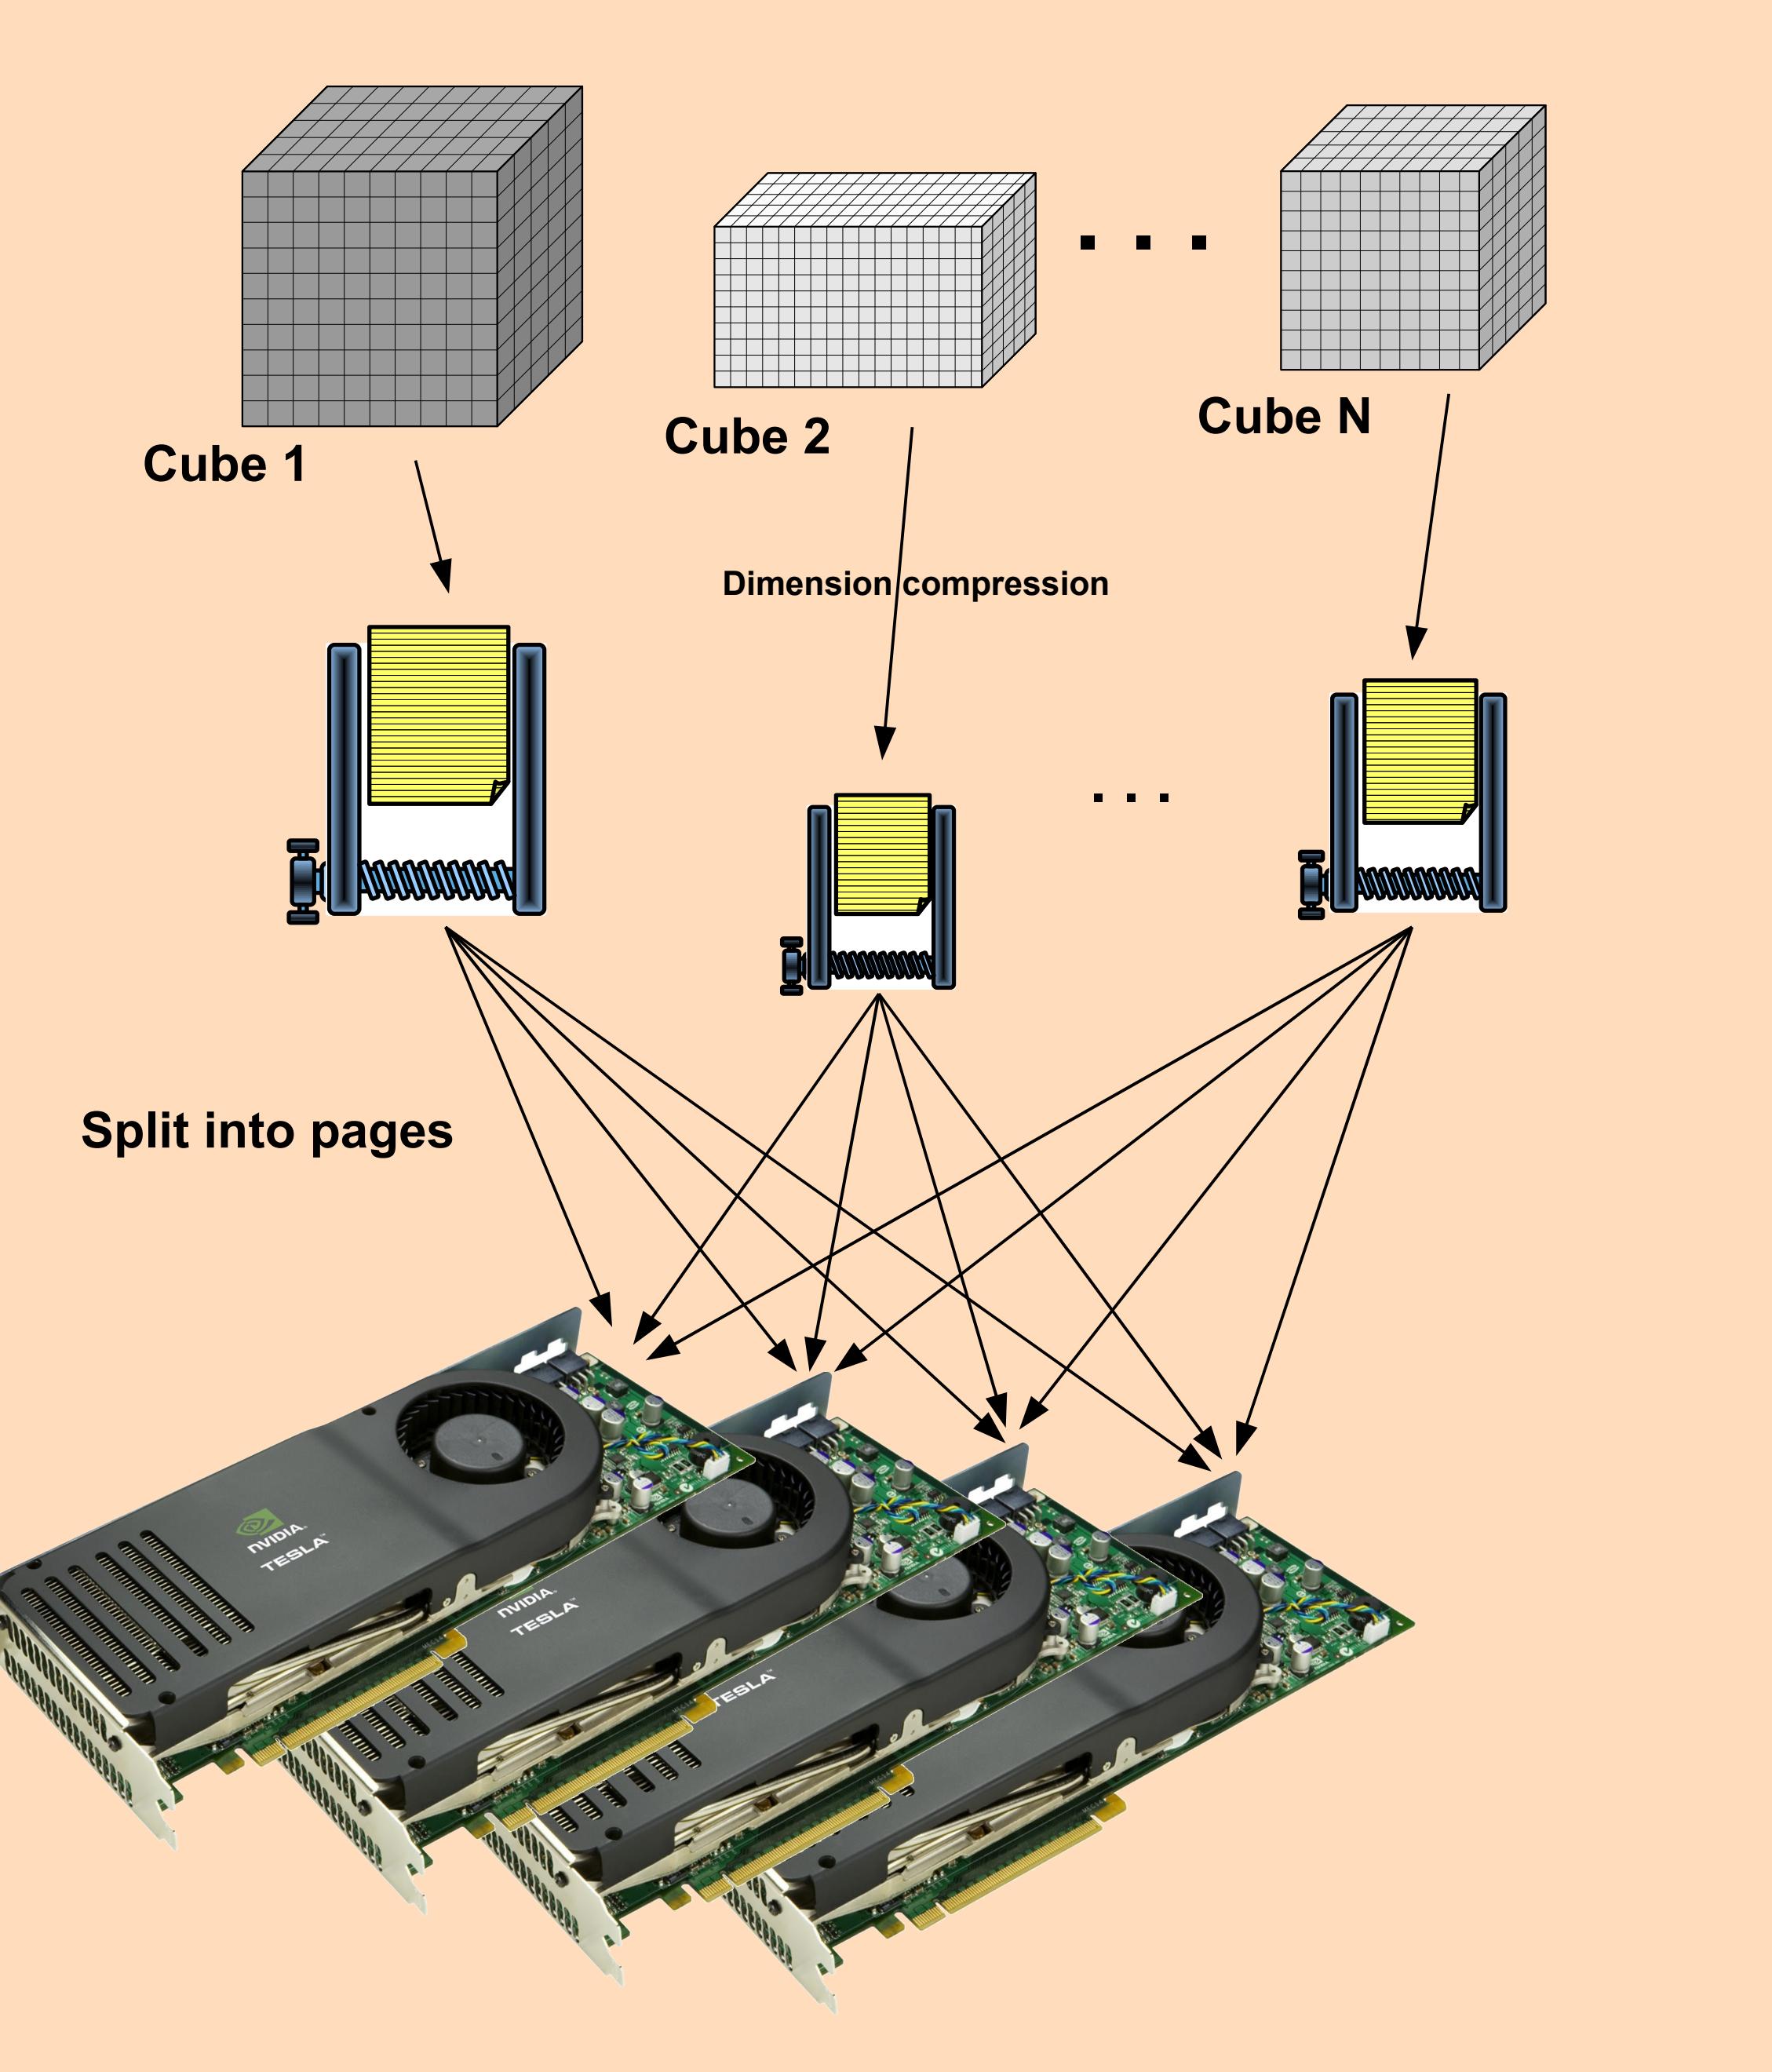


Scalability

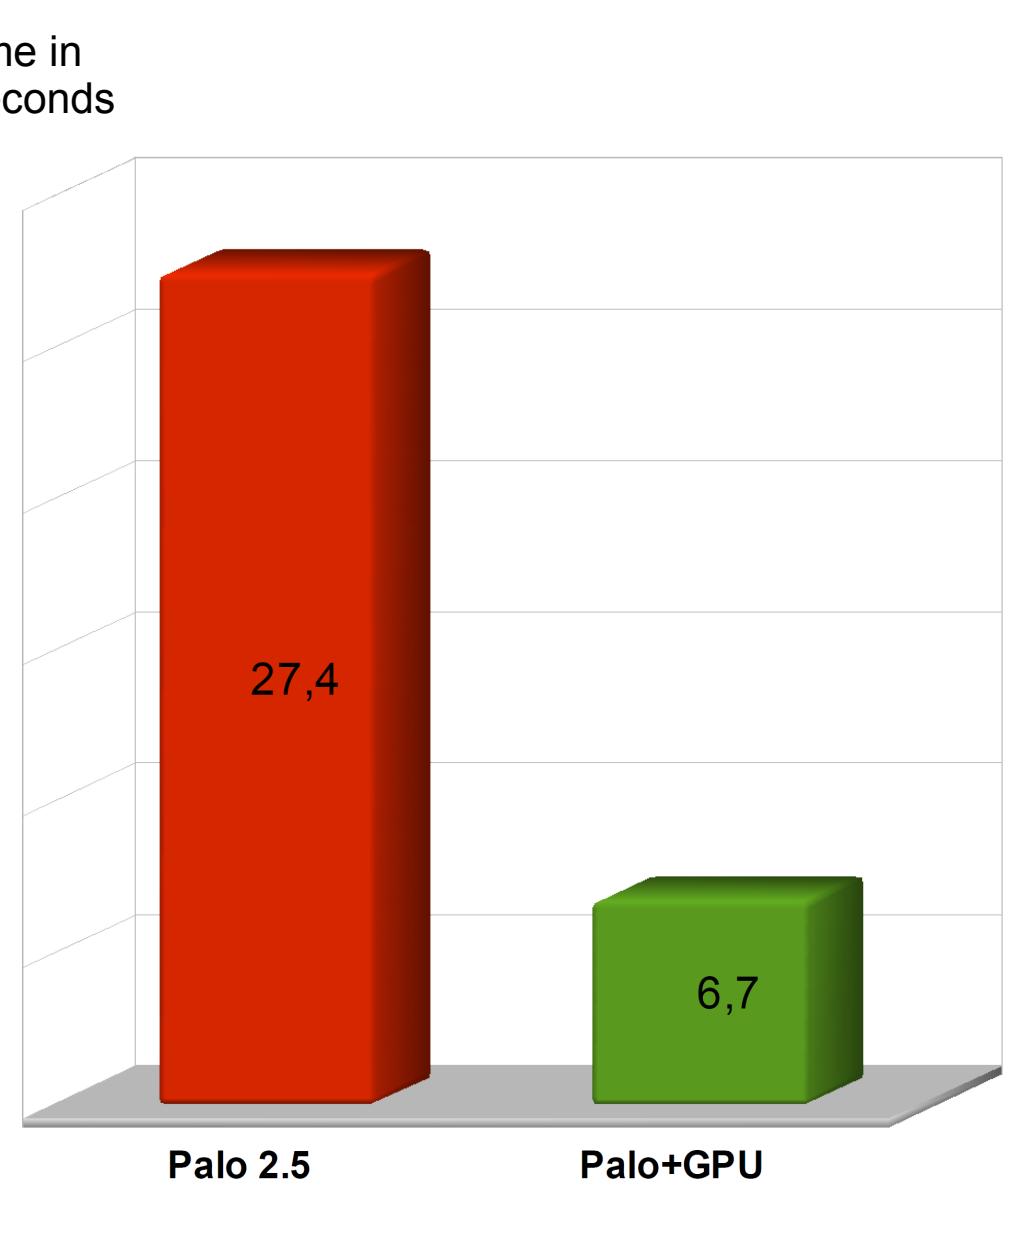
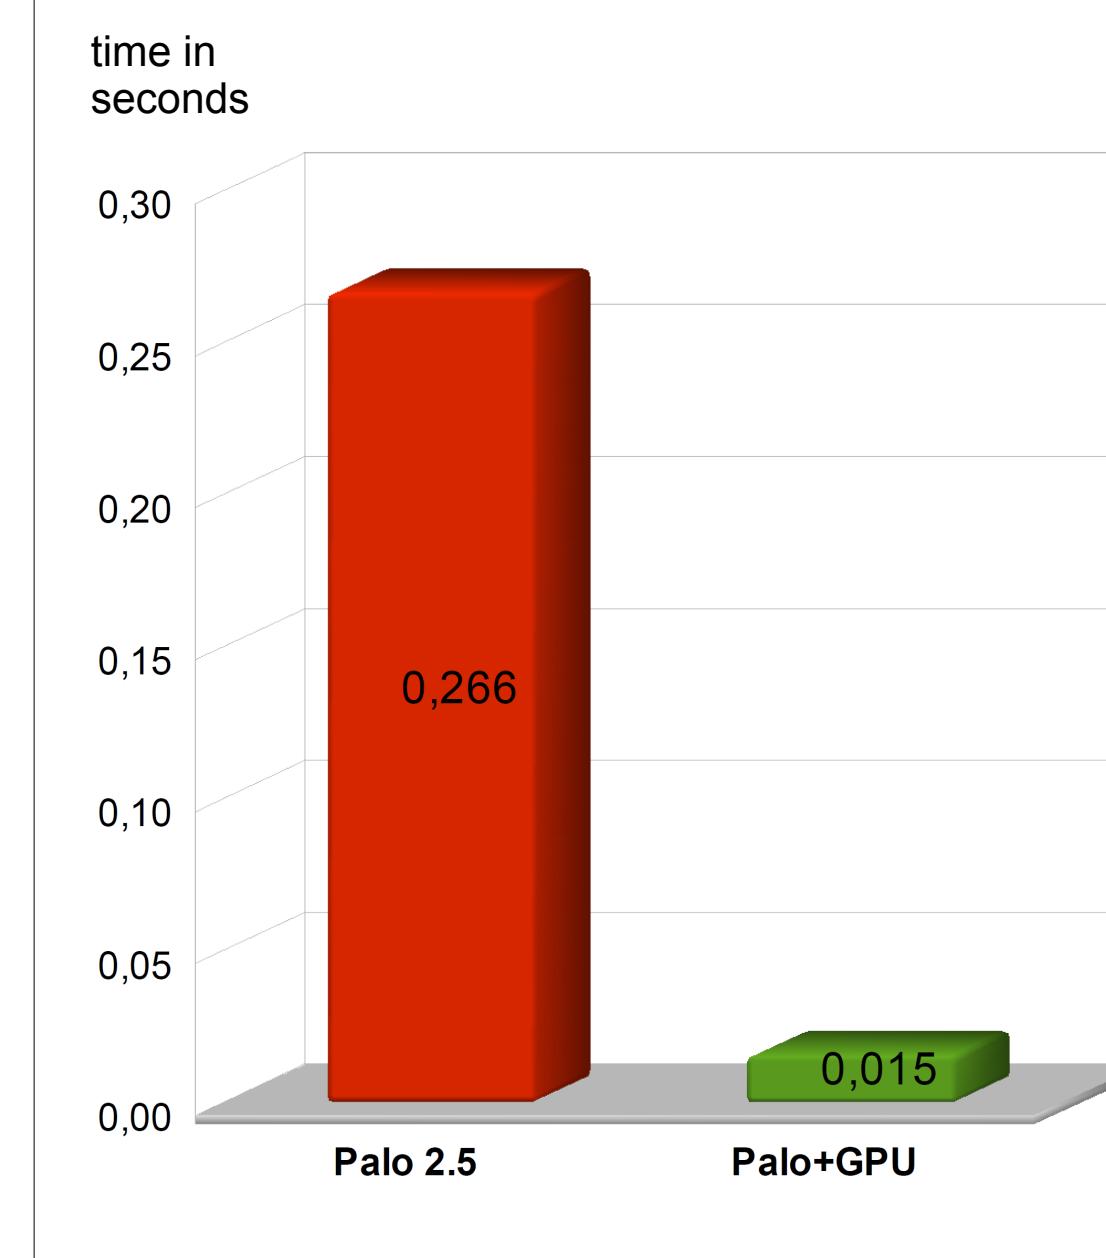
The proposed approach scales well to multiple GPUs.

All cube data are distributed among all available cards, ensuring that the workload is always divided evenly between the devices.

At the end, the individual results are aggregated by the main thread on the host.



Performance Evaluation



Current and future work

- Optimization of bulk queries
 - Two-stage algorithm
 - Pre-filtering
- Efficient updates to database
 - Spreading updated aggregate values to base facts
 - Allocation of new records in GPU memory
- Computation of advanced business rules
- Seamless integration of CUDA algorithms into *Palo*

About Palo

Jedox AG, headquartered in Freiburg (Germany) with offices in Great Britain and France, is one of the leading suppliers of Open-source Business Intelligence and Corporate Performance Management solutions in Europe. Jedox' core product, *Palo BI Suite*, accommodates the entire range of BI requirements including planning, reporting and analysis.

The multidimensional *Palo OLAP Server* at the core of the *Palo BI Suite* integrates simply and easily existing MS-Excel solutions and optimizes planning, reporting and analysis.