# Computation of Large Sparse Aggregated Areas for Analytic Database Queries

Steffen Wittmer • Tobias Lauer

Jedox AG

Collaborators: Zurab Khadikov • Alexander Haberstroh • Peter Strohm

# Business Intelligence and Corporate Planning

# Jedox BI Workflow

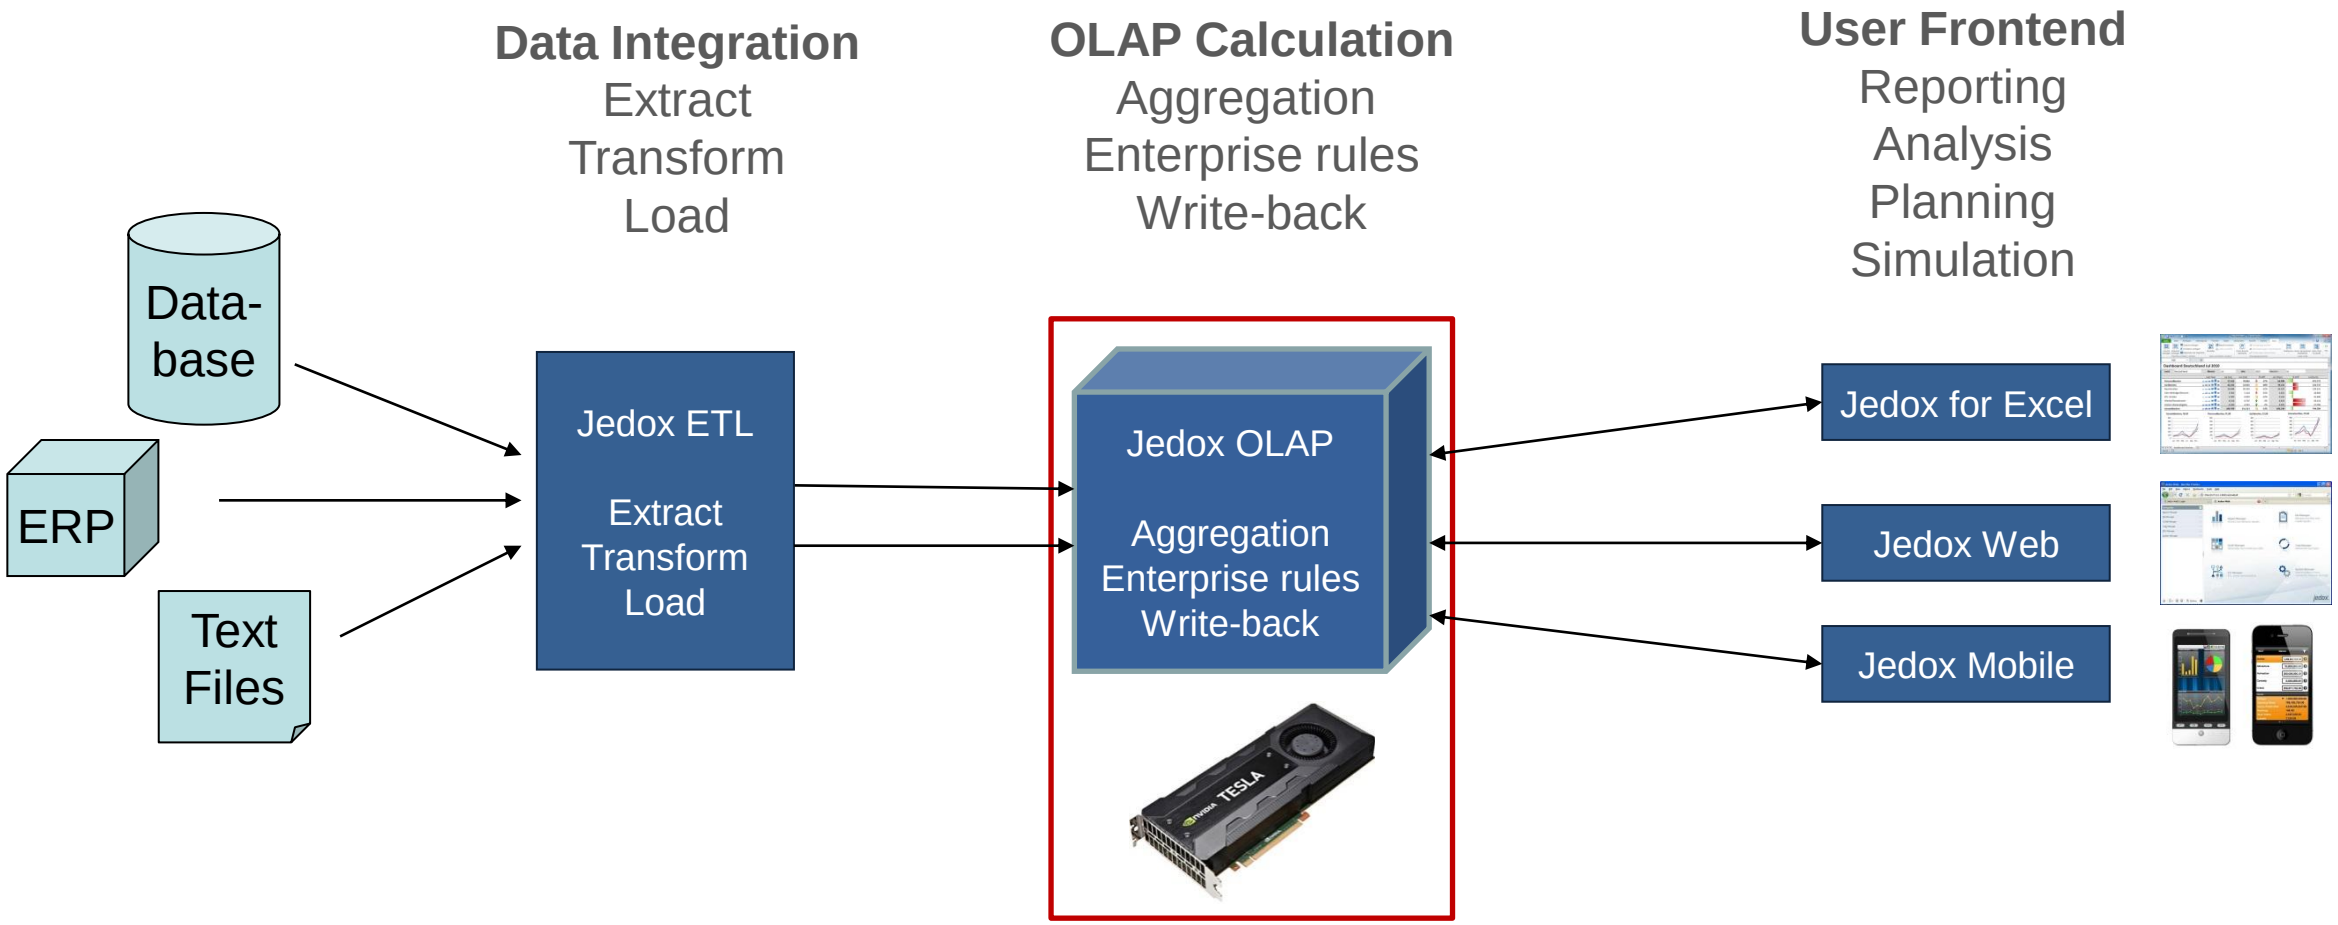


jedox.

# Online Analytical Processing (OLAP)

- Data modeled as multidimensional “cube”
  - *Operations:*
    - Analysis
    - Reporting
    - Planning
    - Simulation
  - *Dimensions are structured hierarchically:*
    - Consolidated elements
    - Base elements

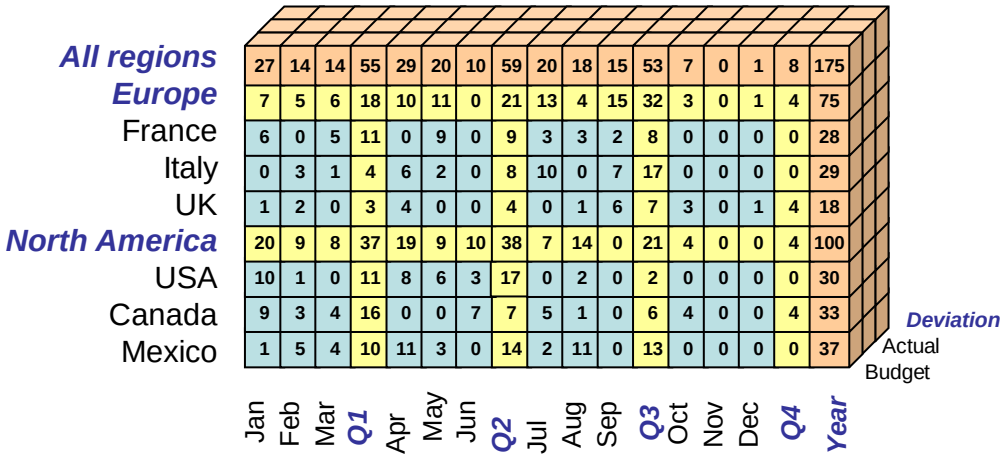


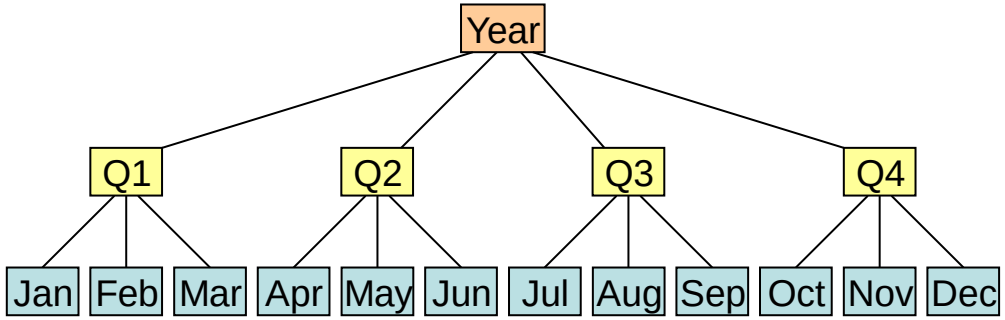

# In-memory OLAP storage model

- All data stored in main memory

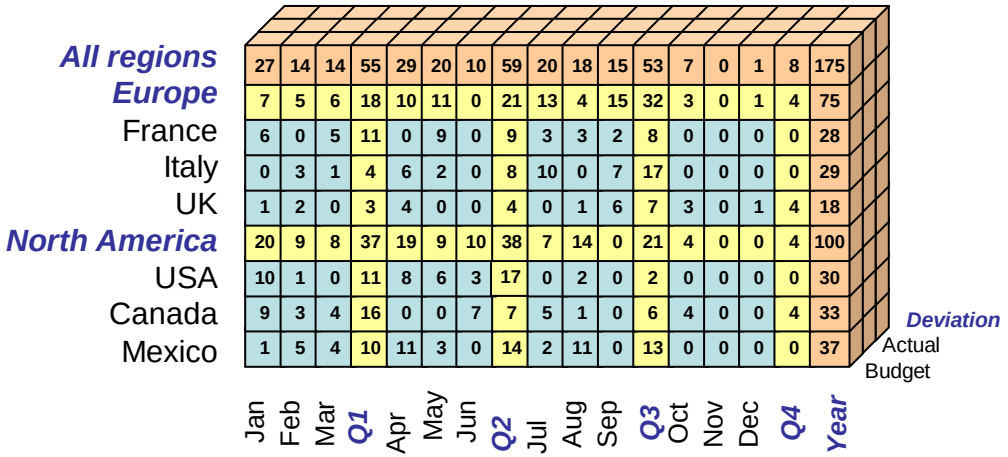

# In-memory OLAP storage model

- All data stored in main memory
- Only store base cells

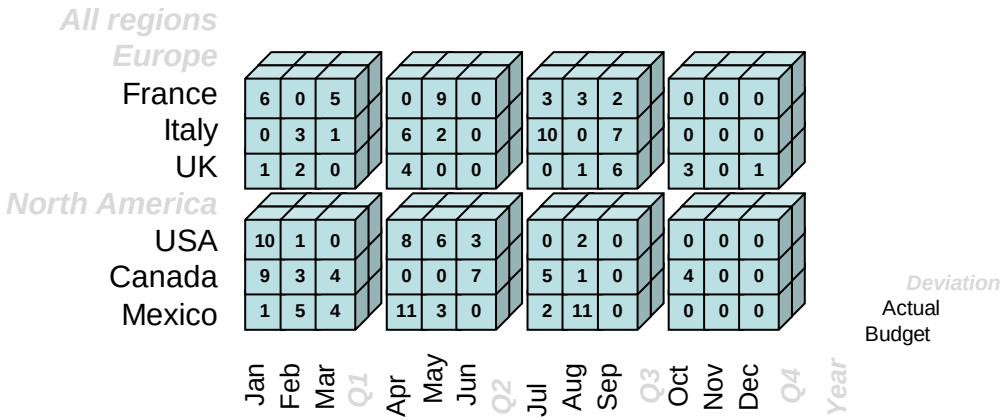

# In-memory OLAP storage model

- All data stored in main memory
- Only store base cells
- Do not store zero-value cells

→ Memory saving, data consistency

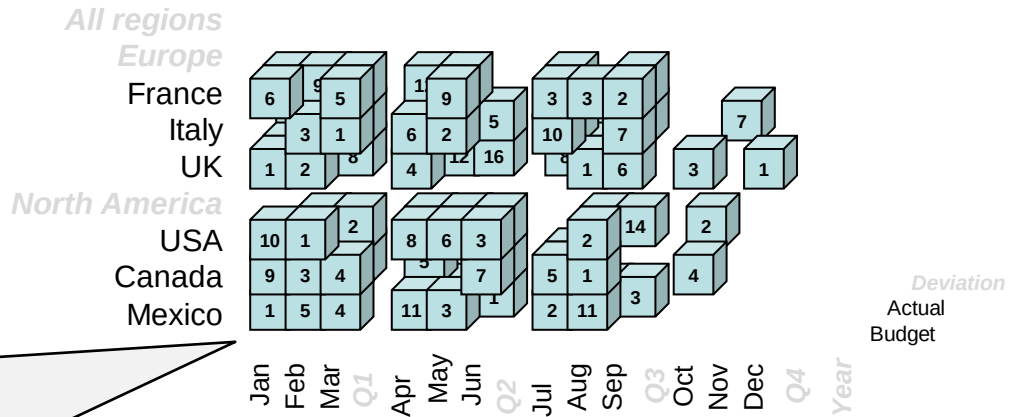


Represent cells as (key, value) pairs,
e.g.
( (2, 1, 0), 4.0 )
Note: Values are double precision!

# In-memory OLAP storage model

- All data stored in **GPU** memory
- Only store base cells
- Do not store zero-value cells

→ Memory saving, data consistency

- Compute other cells *on the fly* when needed

→ Use GPU to accelerate

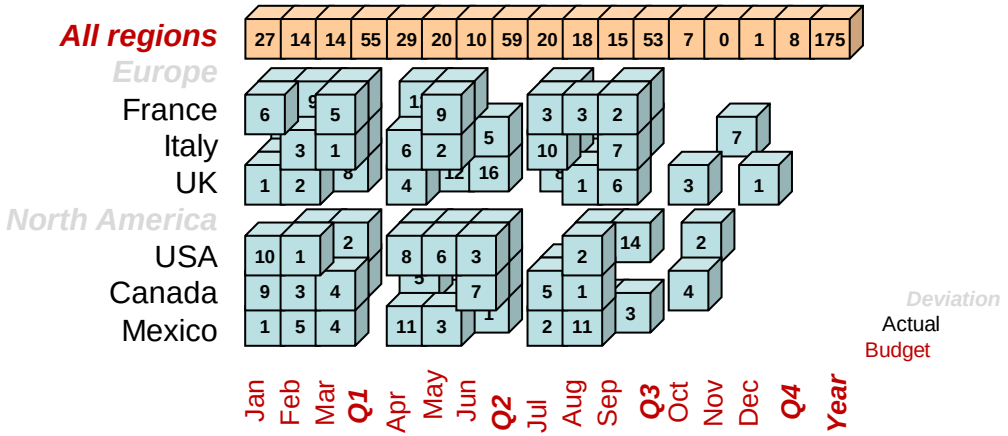


jedox.

# GPU aggregation solutions

## Target driven aggregation

- *Parallel reductions*
- *Very fast for small/dense target areas*

| | Last Year | Jul (Ist) | Jun (Ist) |
|---|---|---|---|
| **Personalkosten** | | **57.626** | **78.861** |
| **Sachkosten** | | **43.142** | **52.651** |
| Raumkosten | | 16.048 | 20.502 |
| Vers-Beiträge/Steuern | | 3.504 | 5.224 |
| Kfz.-Kosten | | 5.309 | 6.859 |
| Werbe/Reisekosten | | 8.518 | 8.797 |
| Kosten Warenabgabe | | 4.265 | 4.341 |
| **Gesamtkosten** | | **100.768** | **131.513** |

## Source driven aggregation

- *Exploits OLAP characteristic of sparsity*
- *Well suited for large aggregated areas*

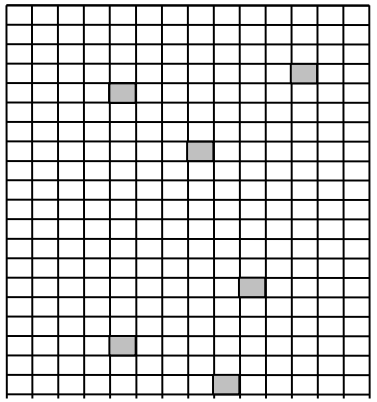

jedox.

# Target driven aggregation

- Fast parallel aggregation step
- Utilizes shared, global and constant memory
- Coalesced memory access
- Almost no thread divergence

Multi-GPU solution

Performance optimized

*bulk aggregations*

*prefiltering*

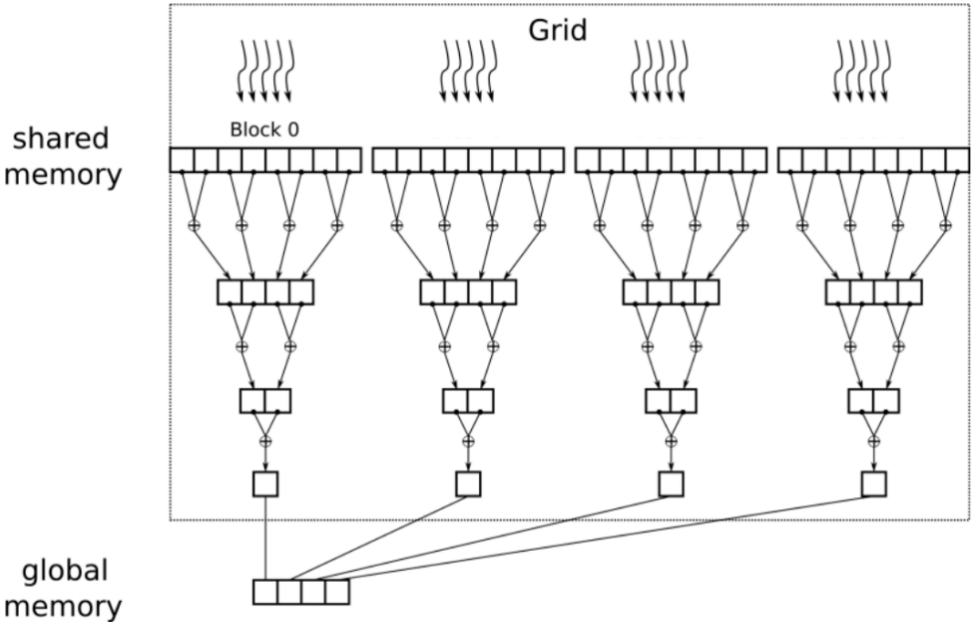

# Large sparse aggregated areas

- Large: up to millions of target cells

  *Top 10 products for each customer*

- Sparse: most target values are zero

  - □ *= zero*
  - ■ *= non-zero*

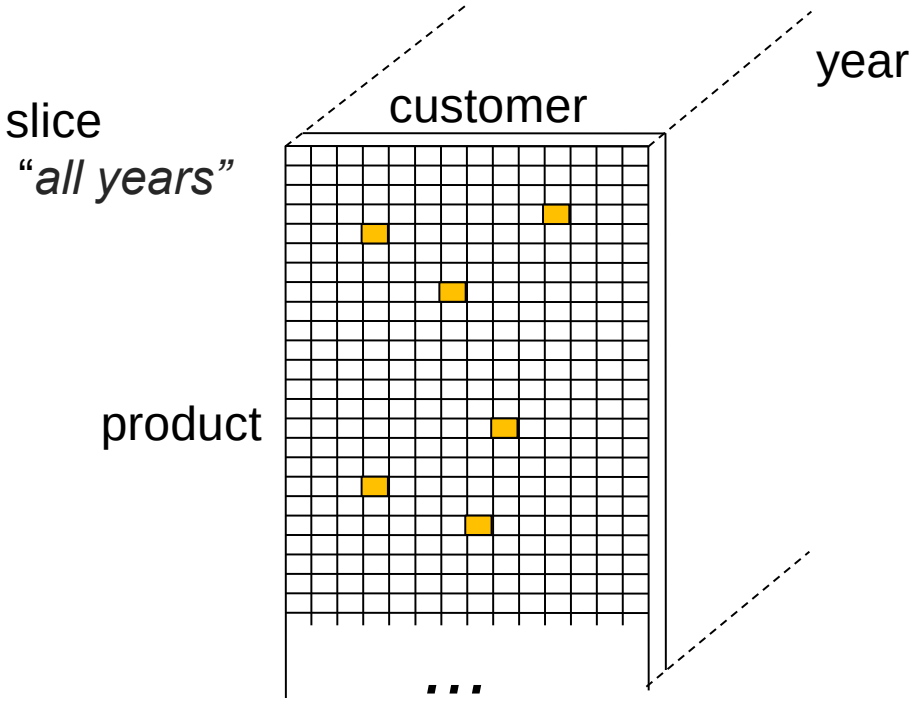

# Handling large sparse areas

- Requirements

  *Performance and memory consumption that correlate with number of **non-zero** target cells*

- Solution

  *Source driven approach*

  *- Serialized aggregation with atomics*

  *- Utilize hash tables on GPU*

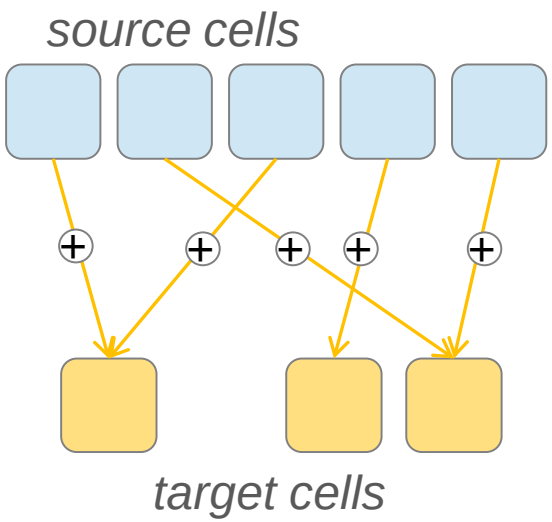

# Source driven aggregation

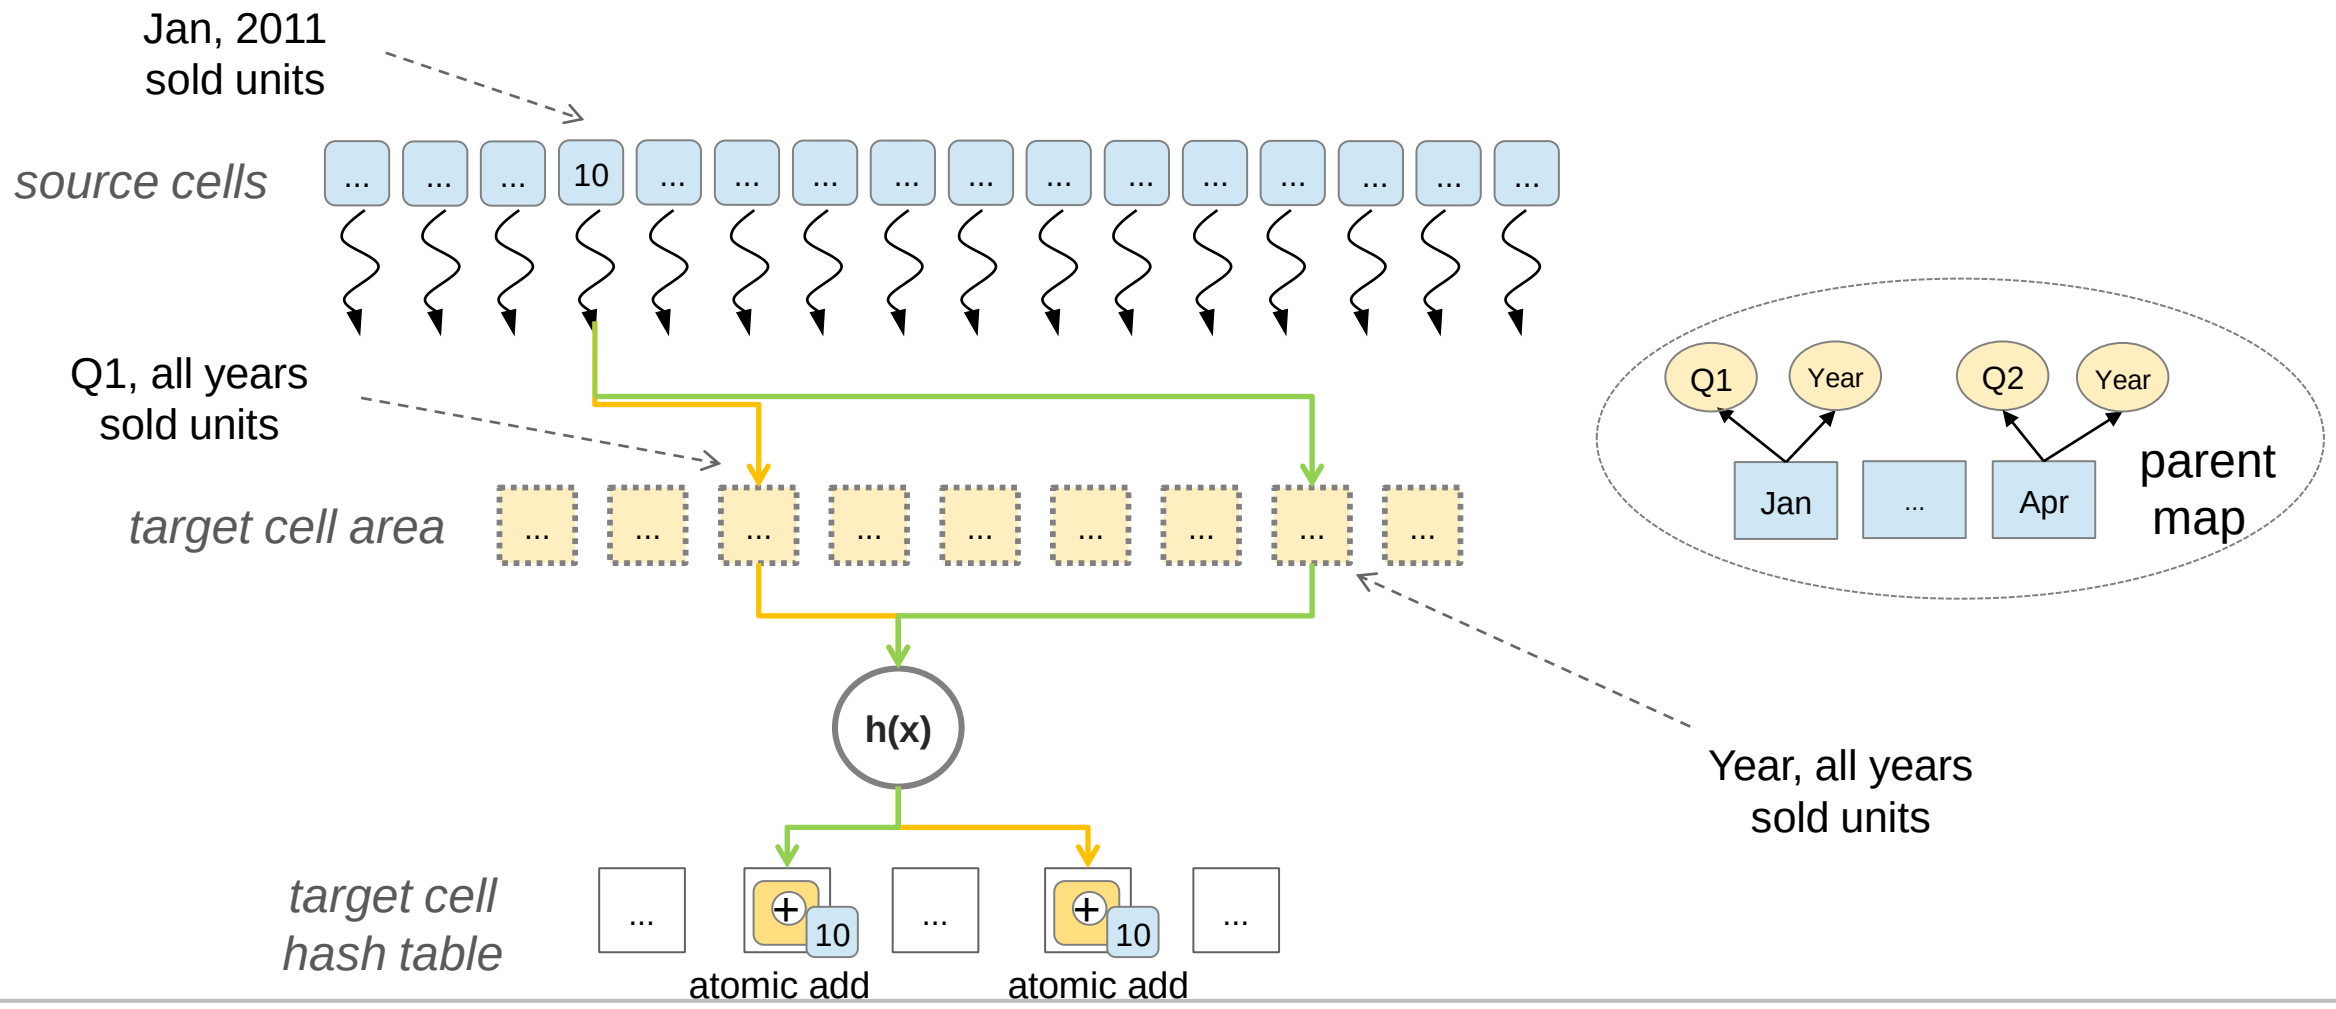

# Atomics: Contention

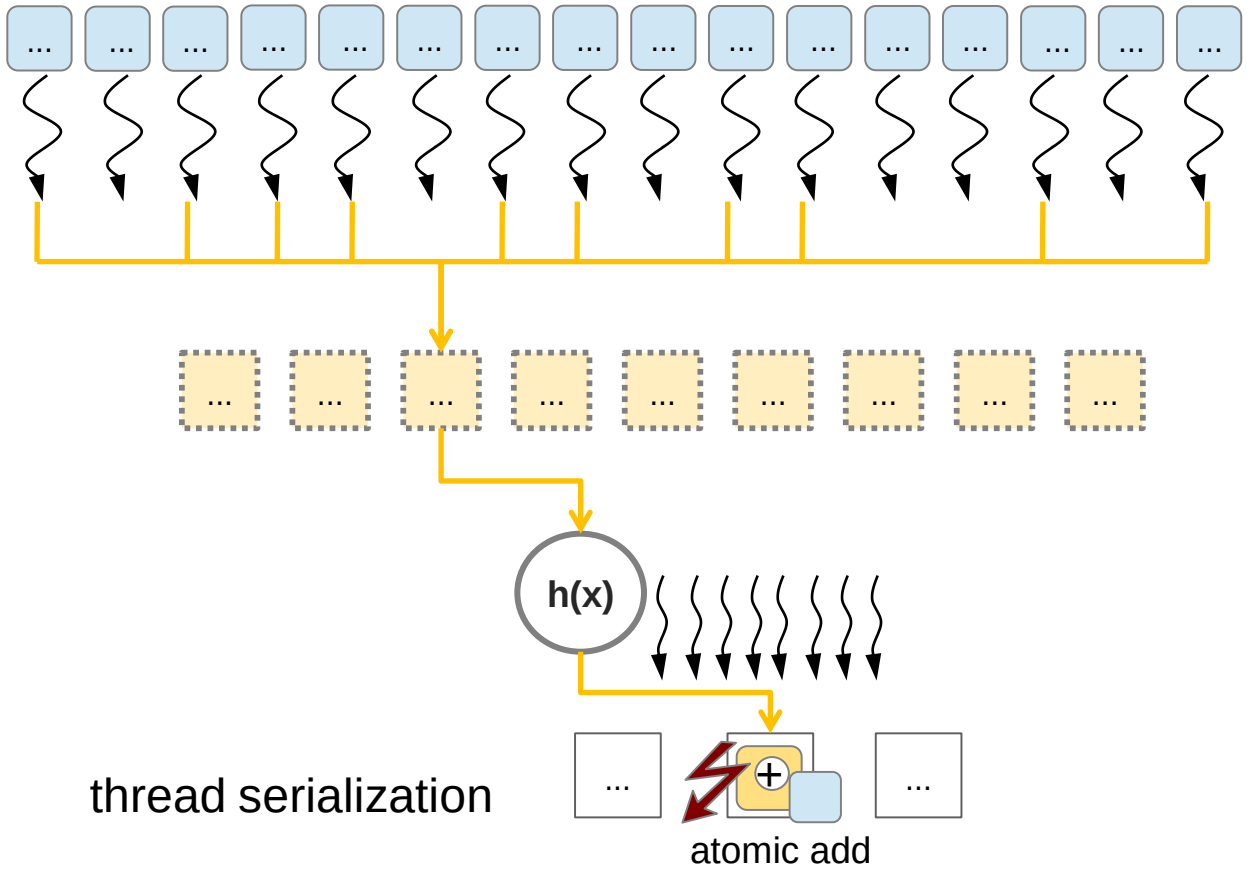


*Great improvement in Fermi and Kepler over CC 1.X*

# Reducing contention

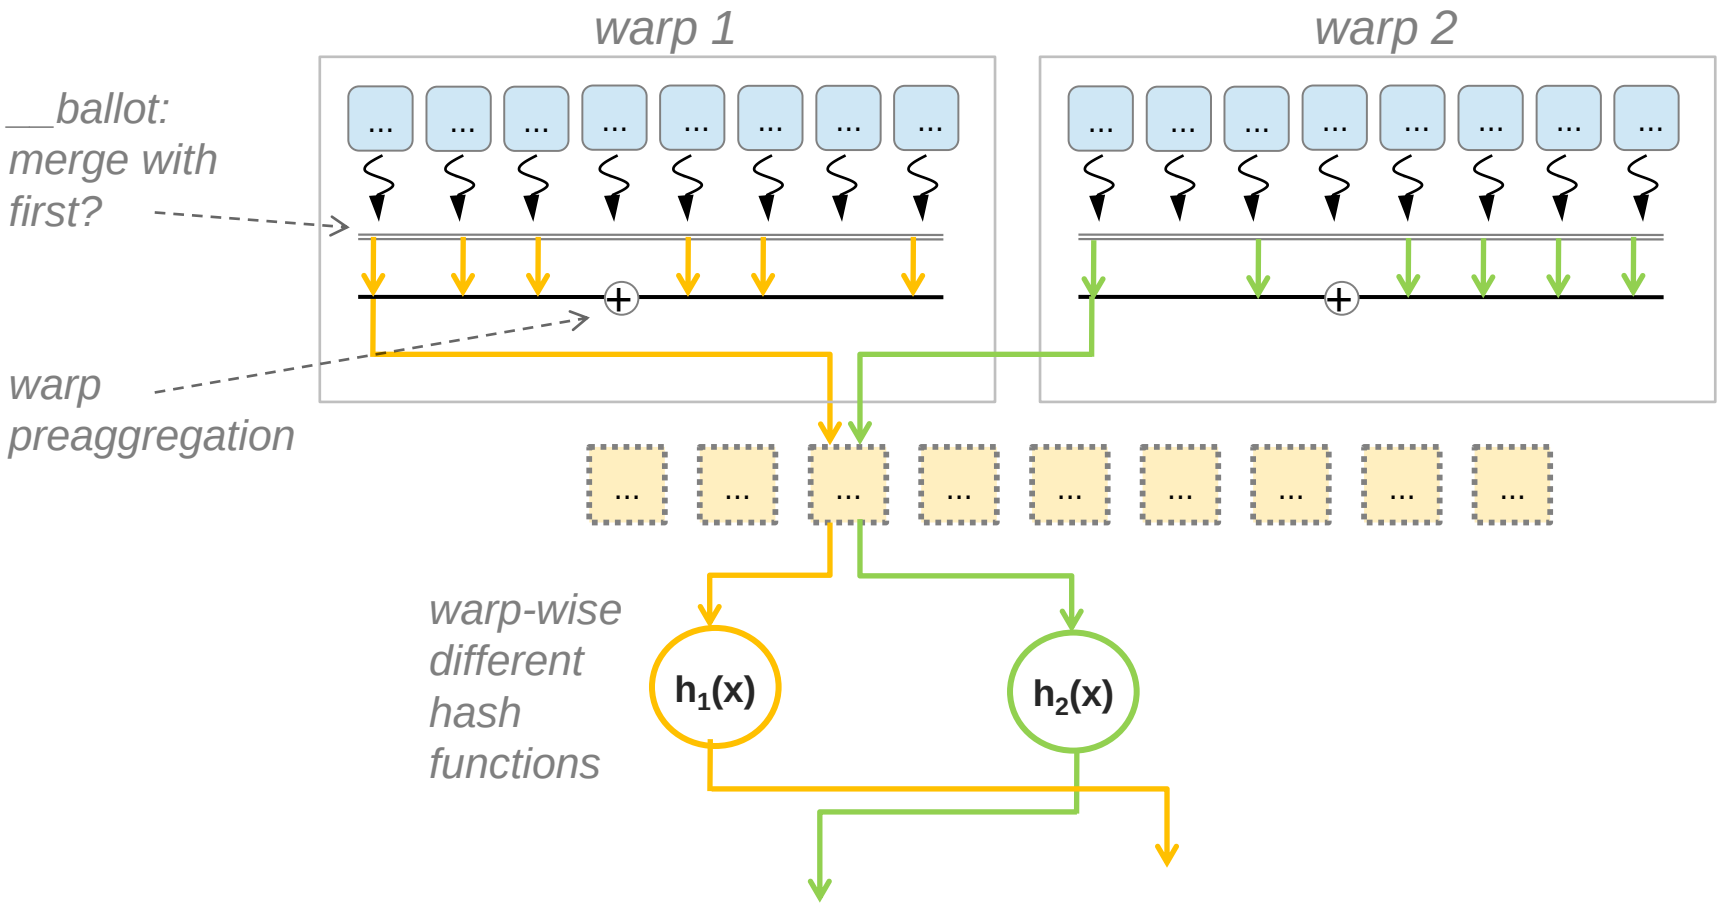

# Speedup: Small target areas (1-11 cells)

## GPU Target Driven

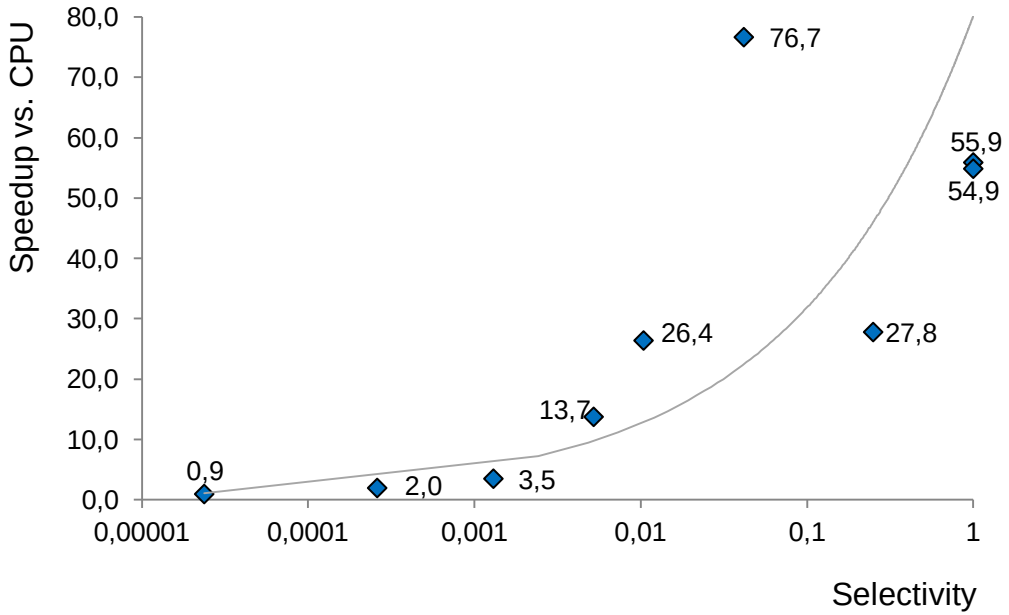


## GPU Source Driven

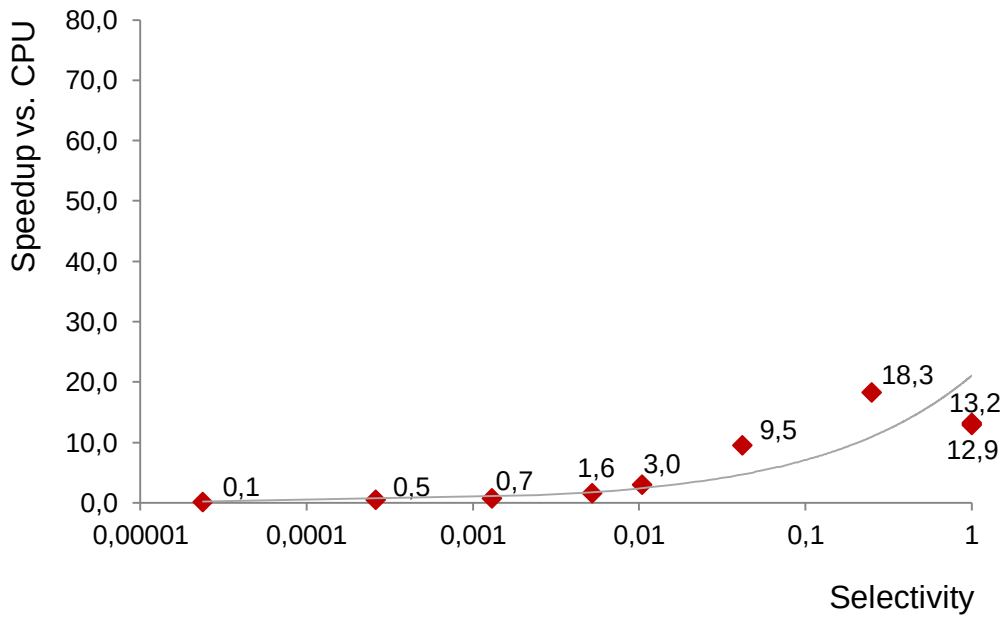


Database: 1B records (filled cube cells)
GPU: 3x Tesla C2070 (18 GB RAM)

jedox.

# Larger areas

## Calculation times

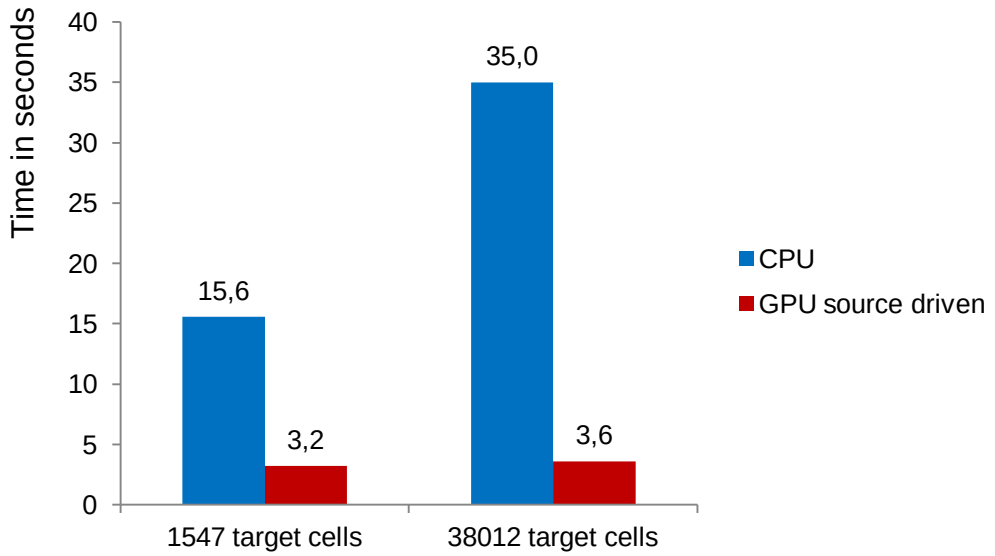


## Speedup

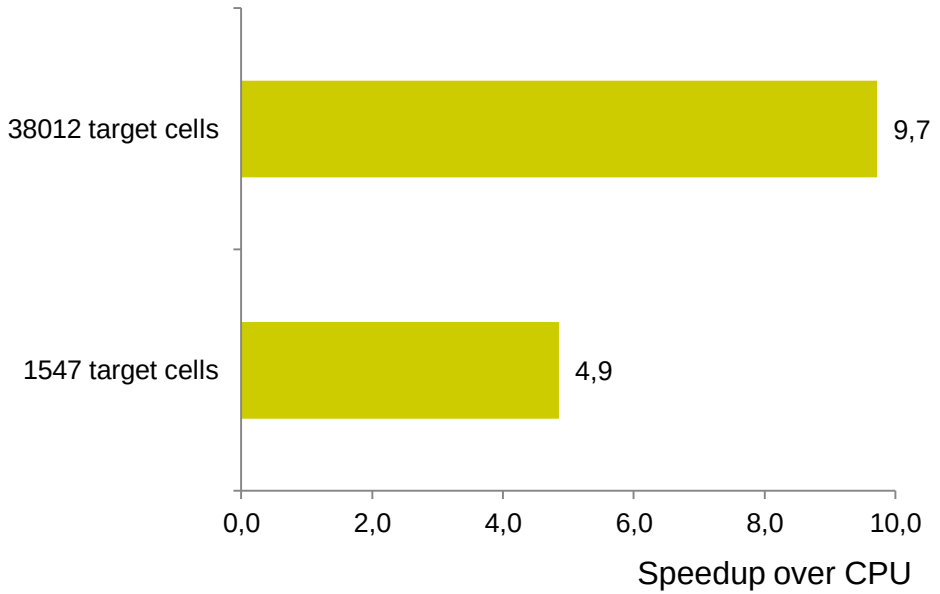


Database: 40M records (filled cube cells)
GPU: 2x Tesla K20 (10 GB RAM)

jedox.

# Comparison: aggregation algorithms

- CPU algorithm good for low aggregation
- Target driven algorithm: good for small and/or dense target areas
- Source driven algorithm: optimized for large and sparse areas

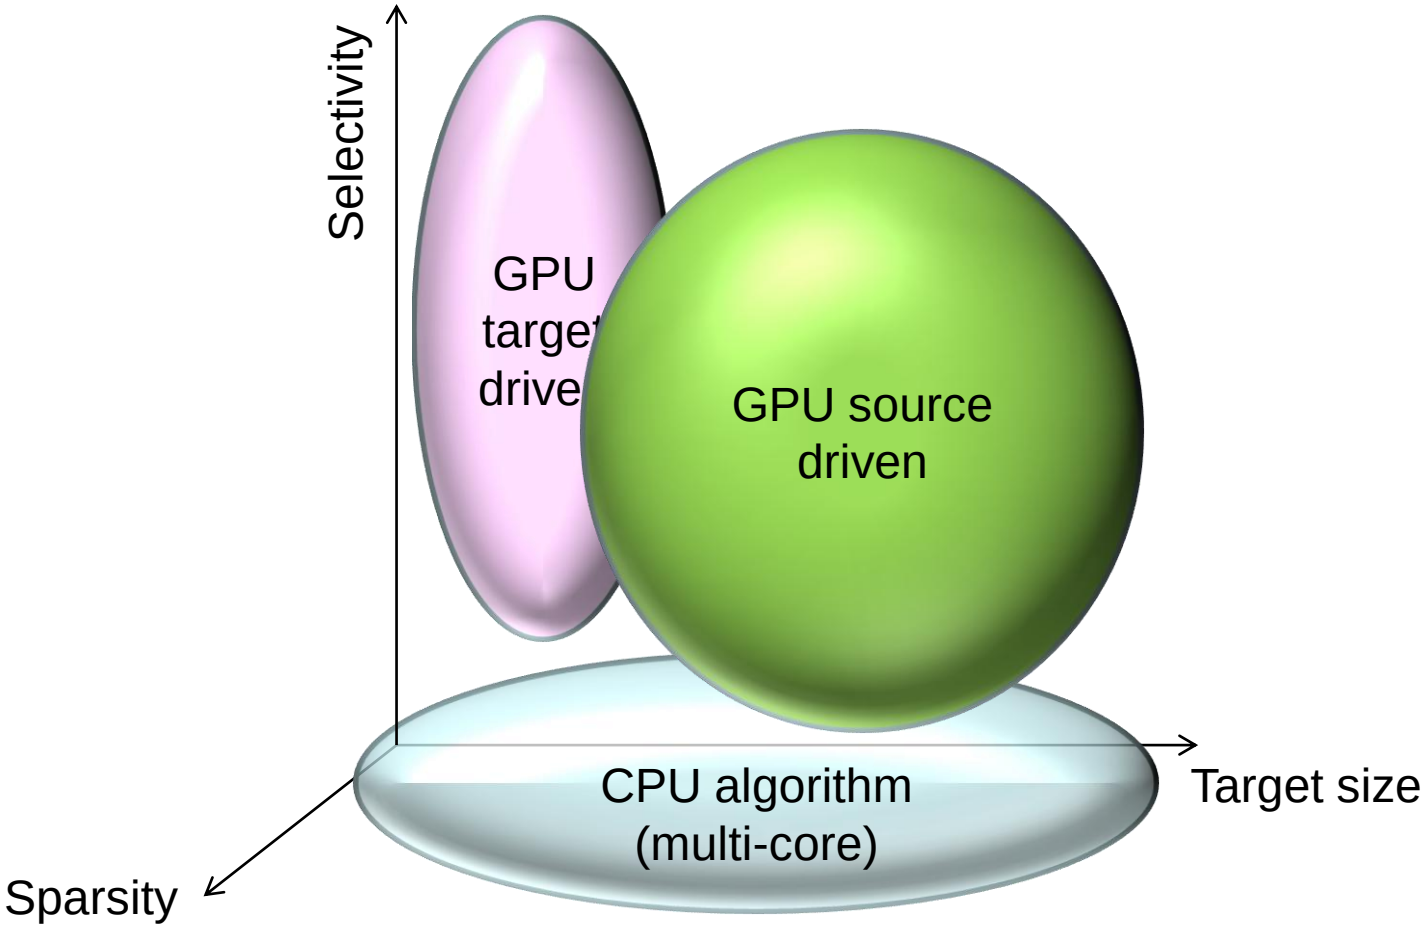

# Future challenges

- Decision mechanism
  - *When GPU, when CPU?*
  - *Which GPU algorithm in which situation?*

- Find suitable thresholds

- What about large and *dense* target areas?
  → GPU memory problem

# Come visit us!

Exhibit Hall: Booth 718

Online: www.jedox.com

tobias.lauer@jedox.com
steffen.wittmer@jedox.com


jedox.