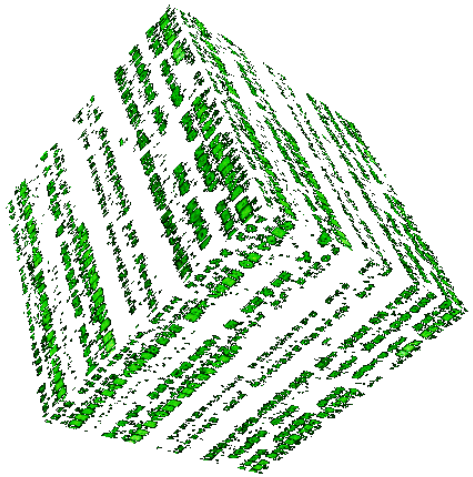

Datacubes on Steroids with ISO Array SQL

FOSDEM 2017, Brussels
Dimitar Misev, Peter Baumann
Jacobs University | rasdaman GmbH
baumann@rasdaman.com

[gamingfeeds.com]

[co-funded through EU EarthServer-2]

Jacobs MSc in Data Engineering

Machine Learning – Big Data – Cloud Computing – Visualization – and more...

all-English, international campus – research involvement – strong industry connections

Structural Variety in Big Data

- Stock trading: 1-D sequences (i.e., arrays)
- Social networks: large, homogeneous graphs
- Ontologies: small, heterogeneous graphs
- Climate modelling: 4D/5D arrays
- Satellite imagery: 2D/3D arrays (+irregularity)
- Genome: long string arrays
- Particle physics: sets of events
- Bio taxonomies: hierarchies (such as XML)
- Documents: key/value stores = sets of unique identifiers + whatever
- etc.

Structural Variety in Big Data

- arrays
- graphs
- graphs
- arrays
- arrays
- arrays
- sets
- hierarchies
- sets
-

Structural Variety in Big Data

sensor, image [timeseries], simulation, statistics data

sets + hierarchies + graphs + arrays

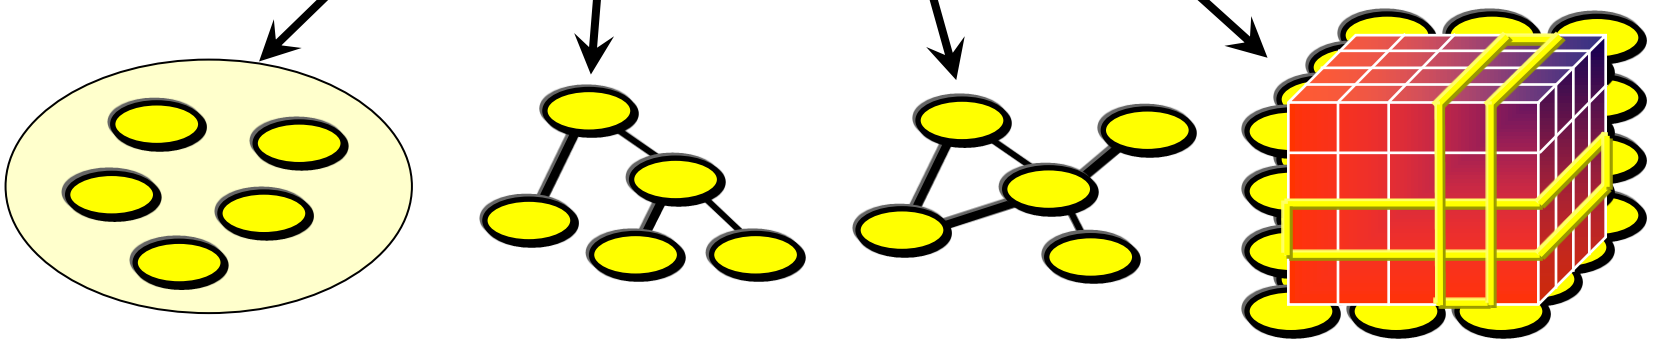

rasdaman raster data management

JACOBS UNIVERSITY

Array SQL

ISO/IEC JTC 1/SC 32

Date: 2014-06-04

WD 9075-15:2014(E)

Information technology — Database languages — SQL —

Part 15:
Multi-Dimensional Arrays (SQL/MDA)

ISO/IEC JTC 1/SC 32/WG 3

The United States of America (ANSI)

Technologies de l'information — Langages de base de données — SQL —

Partie 15: Tableaux multi-dimensionnels (SQL/MDA)

```
create table LandsatScenes(
    id: integer not null, acquired: date,
    scene: row( band1: integer, ..., band7: integer ) mdarray [ 0:4999,0:4999] )
```

```
select  id, encode(scene.band1-scene.band2)/(scene.nband1+scene.band2)), „image/tiff“ )
from    LandsatScenes
where acquired between „1990-06-01“ and „1990-06-30“ and
          mdavg( scene.band3-scene.band4)/(scene.band3+scene.band4)) > 0
```

Linear Algebra

- Matrix multiplication

$$(\mathbf{AB})_{ij} = \sum_{k=1}^{m} A_{ik} B_{kj}$$

```
select mdarray  i in [0:m], j in [0:p]
       elements condense +
                over     k in [0:n]
                using    a [ i, k ] * b [ k, j ]
from   matrix as a, matrix as b
```

- Histogram

```
select mdarray  bucket in [0:255]
       elements mdcount( img = bucket )
from   img
```

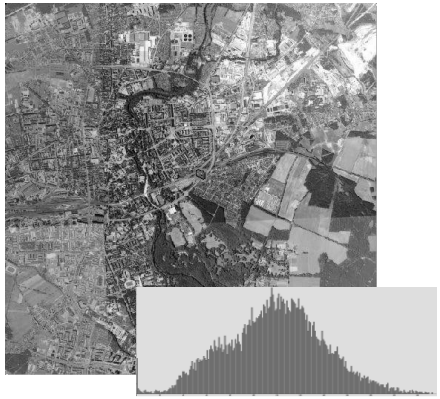

rasdaman

www.rasdaman.org

= „raster data manager“: SQL+ n-D arrays
  - Scalable parallel “tile streaming” architecture

- Mature, in operational use, on OSGeo Live
  - Ex: www.planetserver.eu
- ESA 2017: “world leading environment”, “standard working horse for OGC standardisation on these innovative data access interfaces“

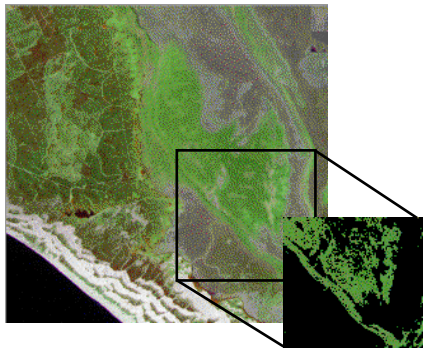

CIO 100 MOST PROMISING BIG DATA Review SOLUTION PROVIDERS 2016

GEOSPATIAL WORLD FORUM

The Earth Monitoring Competition
Copernicus masters
WINNER
T-SYSTEMS BIG DATA CHALLENGE
2014
T · · Systems ·

Adaptive Partitioning („Tiling“)

- Any tiling, canned into strategies [ICDE 1999]
  - *250+ TB datacubes*

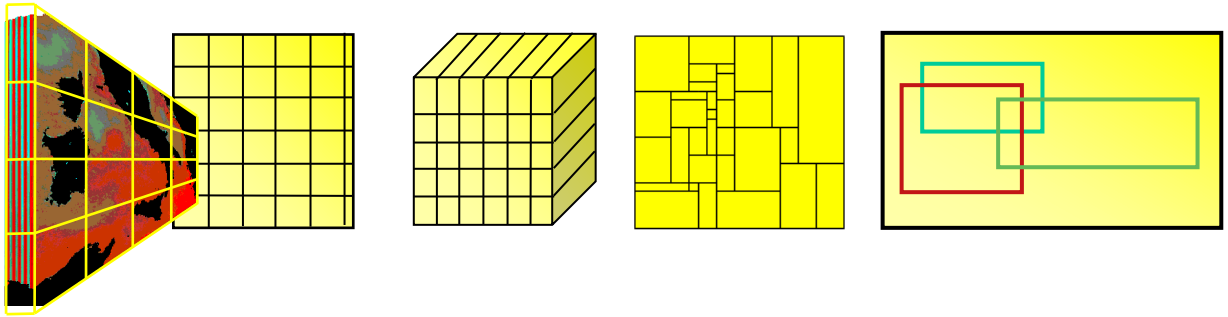

- rasdaman storage layout language [IEEE SSTDM 2010]

```
insert into MyCollection
  values ...
  tiling
    area of interest [0:20,0:40], [45:80,80:85]
    tile size 1000000
    index d_index storage array compression zlib
```

Adaptive Partitioning („Tiling“)

- Any tiling, canned into strategies [ICDE 1999]
  - *250+ TB datacubes*

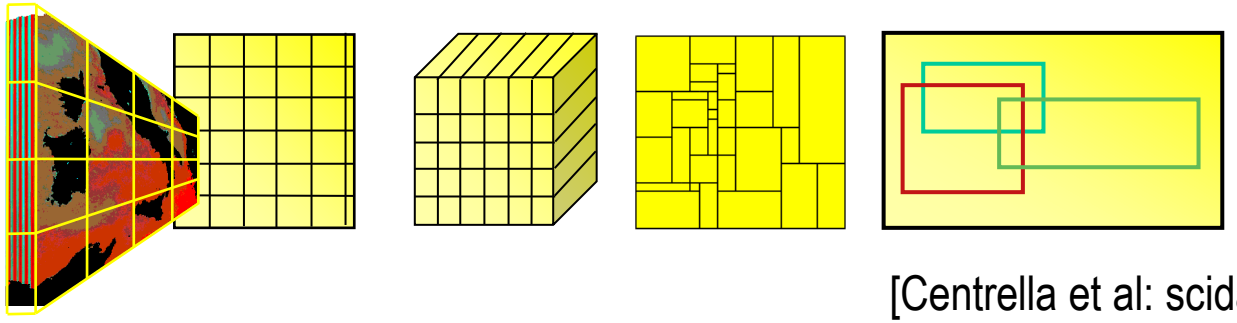

[Centrella et al: scidacreviews.org]

- Why irregular tiling?

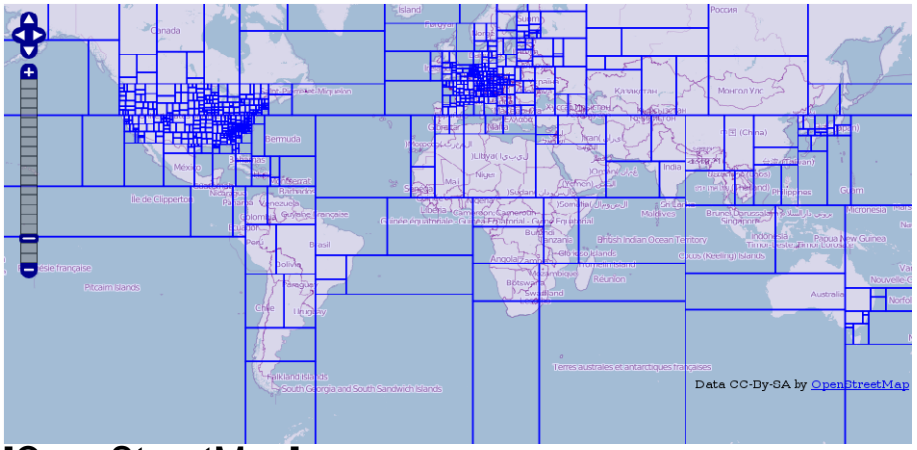

[OpenStreetMap]

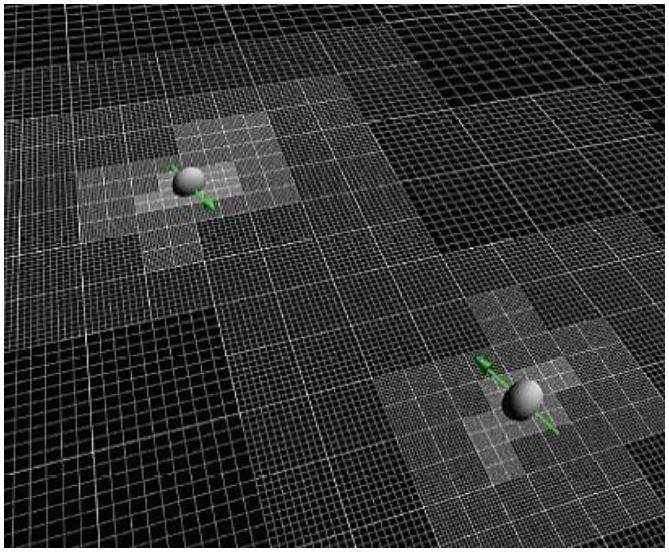

Parallel, Distributed Processing

1 query → 1,000+ cloud nodes
[ACM SIGMOD DanaC 2014]
[VLDB BOSS 2016]

```
select
      max((A.nir - A.red) / (A.nir + A.red))
    - max((B.nir - B.red) / (B.nir + B.red))
    - max((C.nir - C.red) / (C.nir + C.red))
    - max((D.nir - D.red) / (D.nir + D.red))
from A@A, B@B, C@C, D@D
```

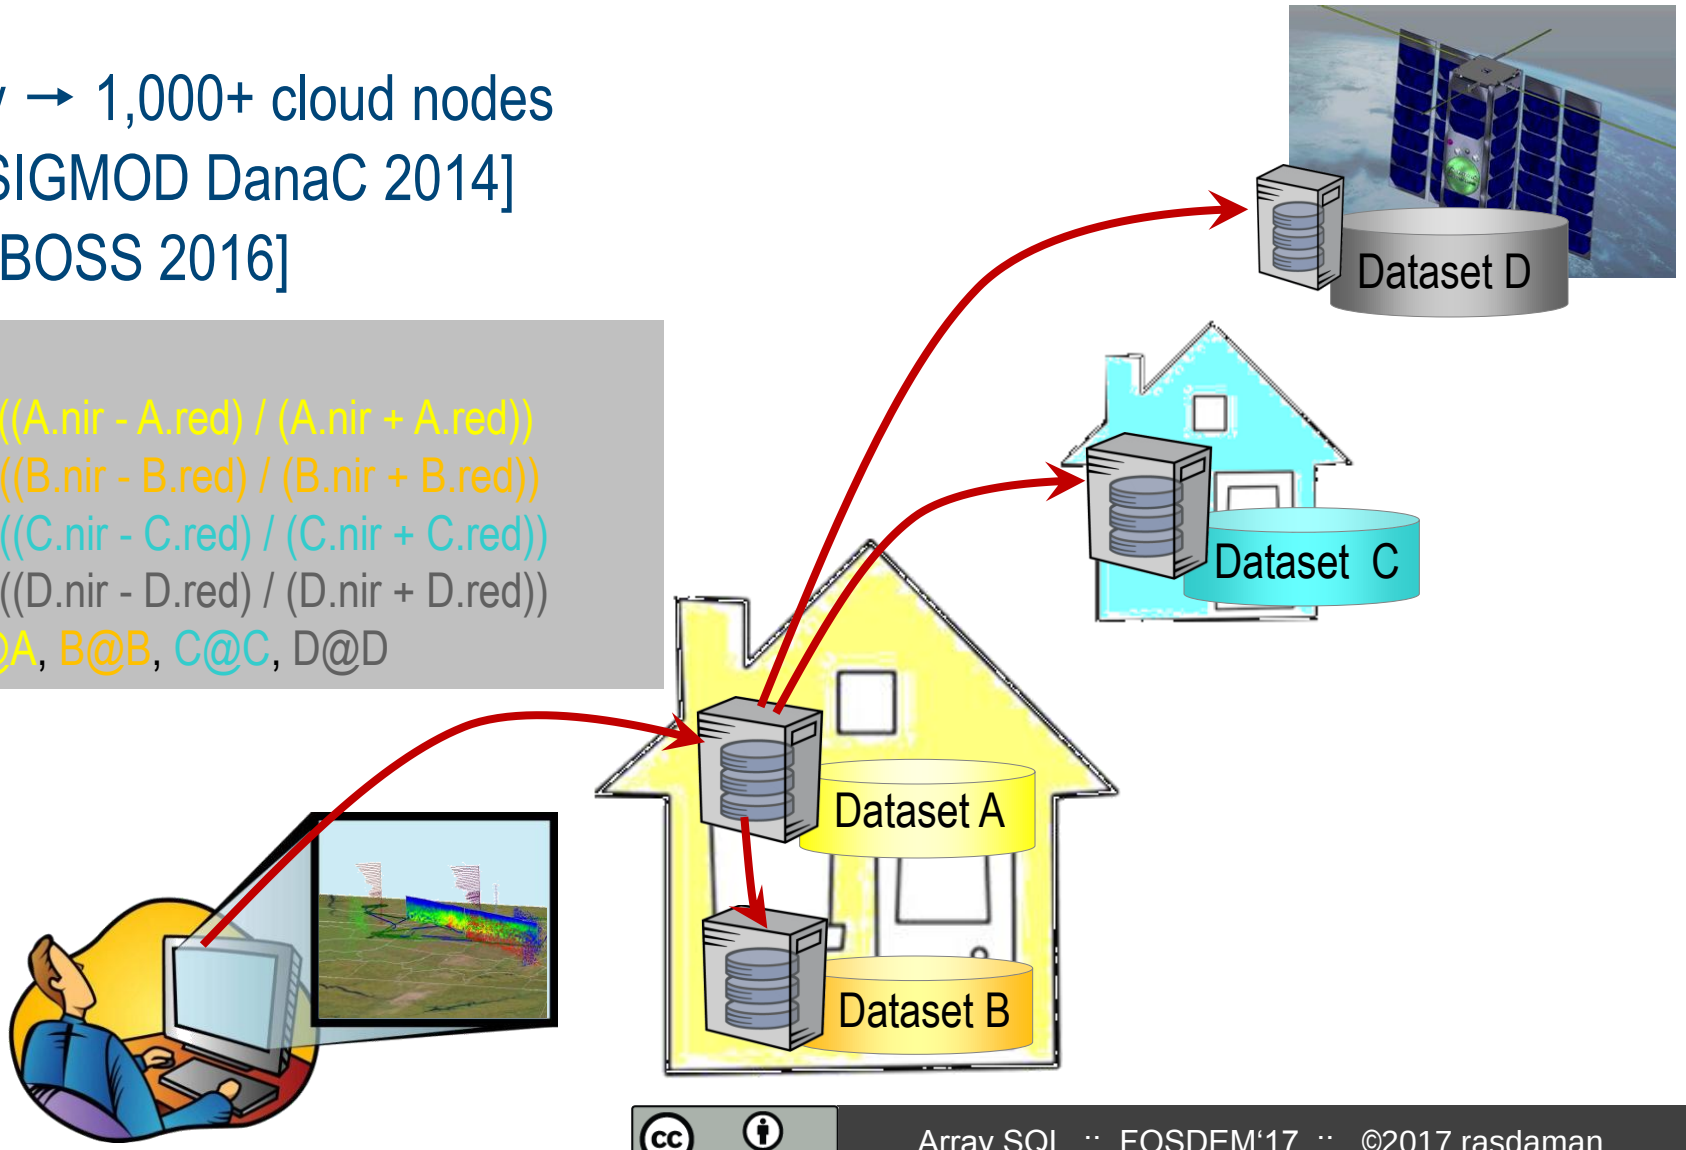

Architecture

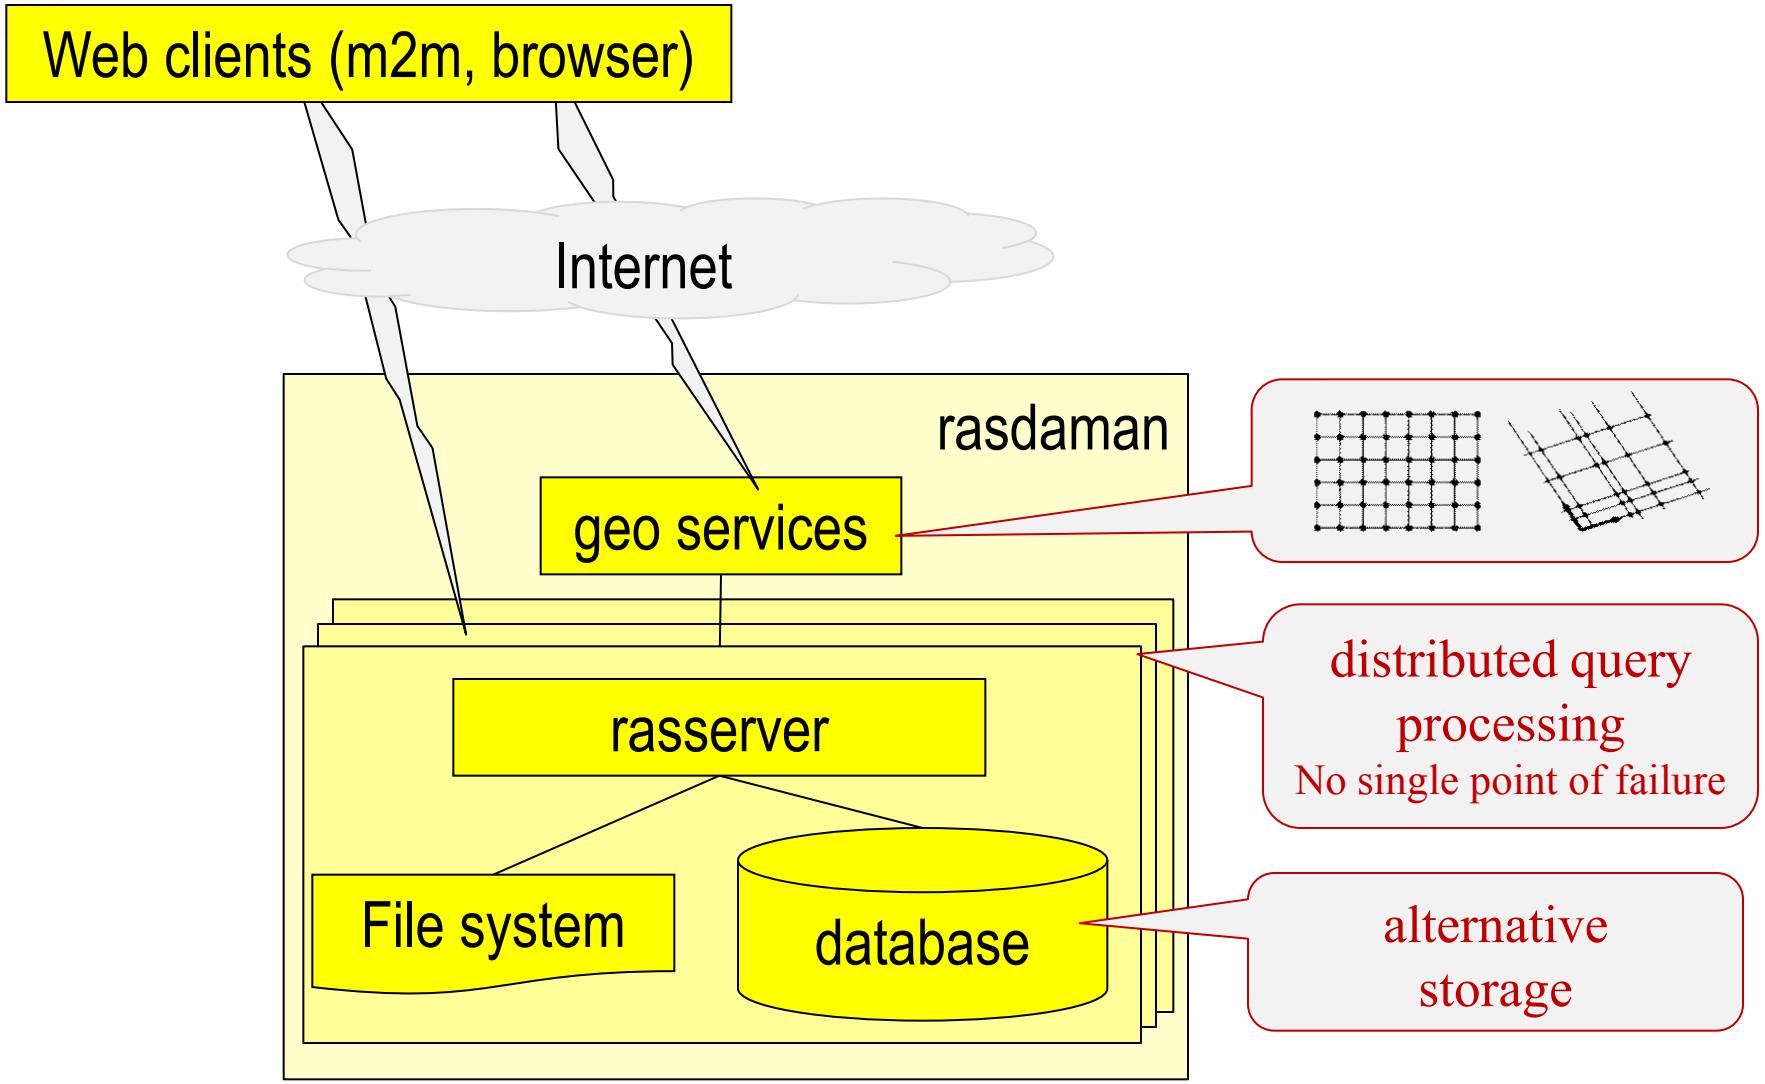

Hadoop/Spark – *one size does not fit all*

- "Since it was not originally designed to leverage the structure its performance is suboptimal" [Daniel Abadi]
- U Madison, GMU benchmark confirms [AGU 2015]

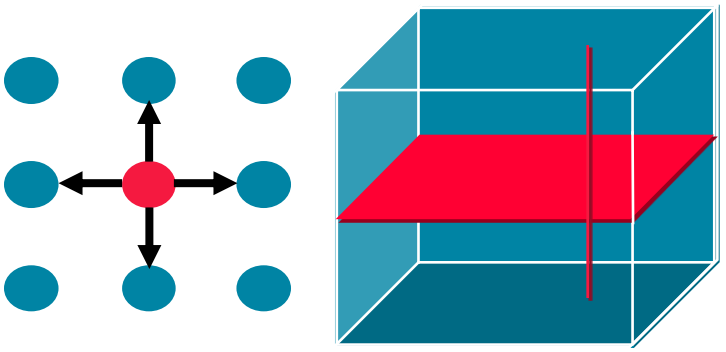

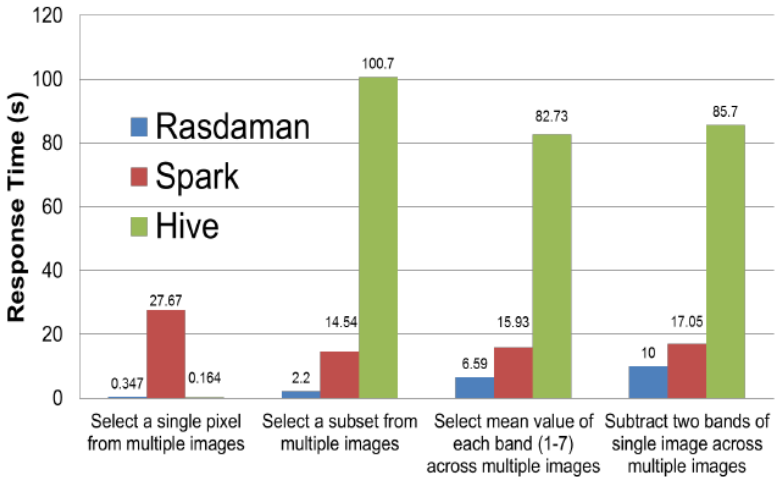

[C. Scheele, F. Hu, M. Yu, M. Xu, K. Liu, Q. Huang, C. Yang 2015]

rasdaman raster data management

JACOBS UNIVERSITY

EarthServer: Datacubes At Your Fingertips

- Agile Analytics on x/y/t + x/y/z/t Earth & Planetary datacubes
  - EU rasdaman + NASA WorldWind
- Global data federation
  - 250+ TB → 1+ PB
- Intercontinental initiative,
  3+3 years:
  EU + US + AUS

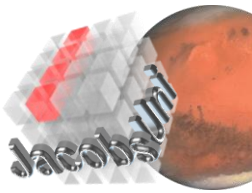

Co-funded by the European Union

www.earthserver.eu

Domains Investigated

- Geo
  - Environmental sensor data, 1-D [Sensors 2009]
  - Satellite / seafloor maps, 2-D [VLDB 1999, 2003, ...]
  - Geophysics (3-D x/y/z)
  - Climate modelling (4-D, x/y/z/t) [DB Spektrum 2012]
- Life science
  - Gene expression simulation (3-D) [InfSys 2003]
  - Human brain imaging (3-D / 4-D) [TiNS 2001]
- Other
  - Computational Fluid Dynamics (3-D)
  - Astrophysics (4-D)
  - Statistics (n-D)

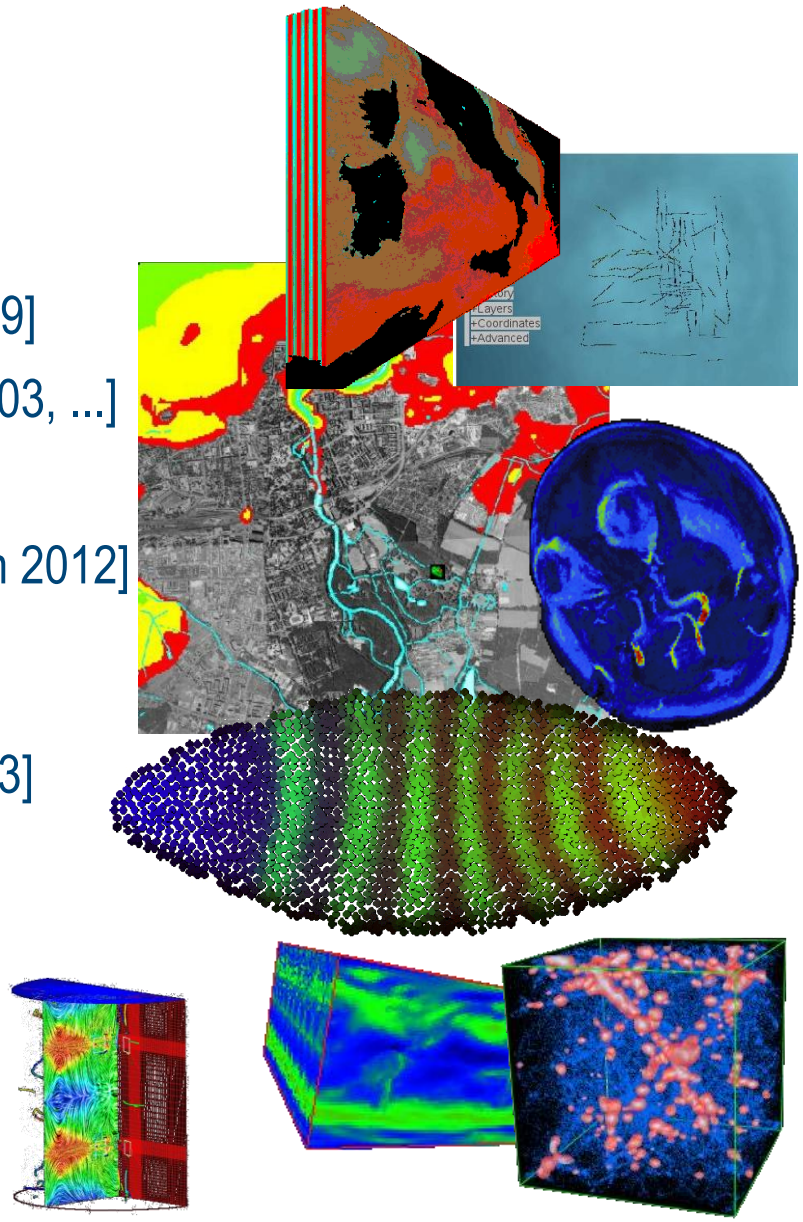

OSGeo Experiences & Thoughts

- Opportunities for improvement, based on:
  - OSGeo Charter Member
  - Executed project incubation procedure with rasdaman
  - Since 2008 engaged in open-source software

OSGeo Experiences & Thoughts

- Opportunities for improvement
- Organizational Maturity: Process definition & implementation, QM
  - No inventory of decisions taken
  - Incubation management: pyWPS 8 years; rasdaman 6.5 years
  - Apply incubation criteria to itself (first)
  - In elections, typically insiders recommend each other

OSGeo Experiences & Thoughts

- Opportunities for improvement
- Organizational Maturity: Process definition & implementation, QM
- Focus on Core Mission
  - Should brand „good software“, not conquer project
  - „design by committee“ over „expert leadership“

OSGeo Experiences & Thoughts

- Opportunities for improvement
- Organizational Maturity: Process definition & implementation, QM
- Focus on Core Mission
- Dogmatic „Software Communism“
  - „all software free“ - why?
  - Large companies don‘t care, small companies vulnerable

  → Need inclusive approach

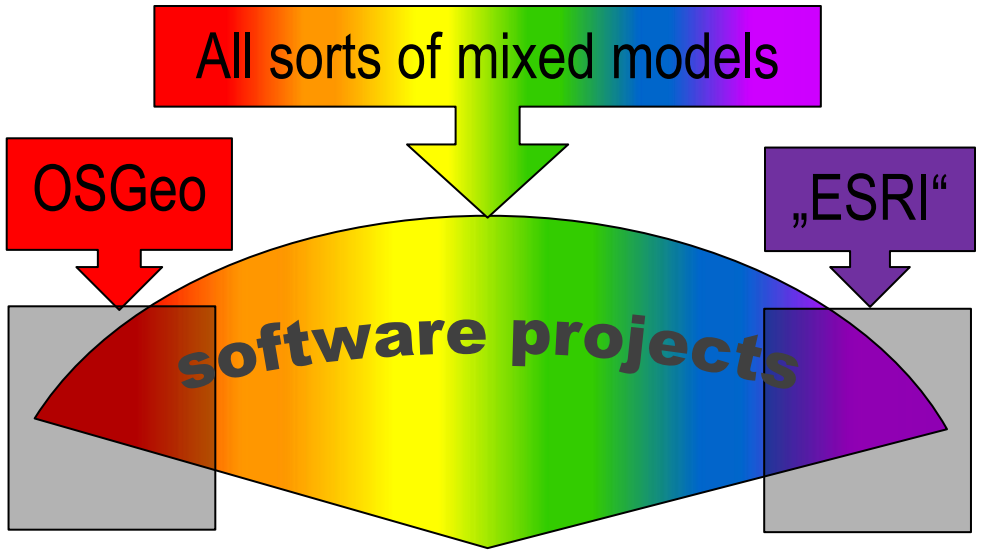

Wrap-Up

„One cube says more than a million images“

- ISO SQL/MDA candidate standard
- Array Databases bring flexibility + scalability
- rasdaman community
  - OGC, INSPIRE WCS reference implementation
  - blueprint for „Big Data“ standards
- FOSS needs to be inclusive