

Science SQL

OGC TC Barcelona, 2015-Mar 10

Peter Baumann

Jacobs University | rasdaman GmbH

baumann@rasdaman.com

[gamingfeeds.com]



EarthServer

www.earthserver.eu

- **Agile** Big Data Analytics for Earth & Planetary Sciences
- 20 ... 132 TB spatio-temporal databases
 - running rasdaman


Cryospheric Science

landcover mapping




EOX

Airborne Science



Atmospheric Science

climate variables



MEEO
Meteorological Environmental
Earth Observation

Geology


geological models



BGS
British Geological Survey
NATURAL ENVIRONMENT RESEARCH COUNCIL

Oceanography


marine model runs
+ in-situ data



PML PLYMOUTH MARINE
LABORATORY

Planetary Science

Mars geology



JACOBS
UNIVERSITY

- EarthServer-2 starting May 2015

What Users Really Want


- "Given me all of the images in this geographic area in this this time span that are at least 80% cloud free have been radiometrically corrected and are from these satellites and then pass those images into a workflow to perform functions x,y,z"
 - Carl Reed
- "Find images that contain fire hotspots according to CORINE Land Cover, and are located within 2km from an archaeological site in the Peloponnese."
 - INSPIRE related
- „propability that the combined wave and swell height in area X, for a daylight time window of 12 hous, will be less than 1.5m, is 90%"
 - WMS BP draft

complexity + flexibility
...aka „Variety“

Sample Application: Database Visualization

```

select
  encode (
    struct {
      red:      (char) s.image.b7[x0:x1,x0:x1],
      green:    (char) s.image.b5[x0:x1,x0:x1],
      blue:     (char) s.image.b0[x0:x1,x0:x1],
      alpha:    (char) scale( d.elev, 20 )
    },
    "image/png"
  )
from SatImage as s, DEM as d
  
```





Human Brain Imaging

- Research goal: to understand structural-functional relations in human brain
- Experiments capture activity patterns (PET, fMRI)
 - Temperature, electrical, oxygen consumption, ...
 - → lots of computations → „activation maps“
- Example: *“a parasagittal view of all scans containing critical Hippocampus activations, TIFF-coded.”*

```
select tiff( ht[ $1, ** , ** ] )
  from  HeadTomograms as ht,
        Hippocampus as mask
 where count_cells( ht > $2 and mask )
        / count_cells( mask )
        > $3
```



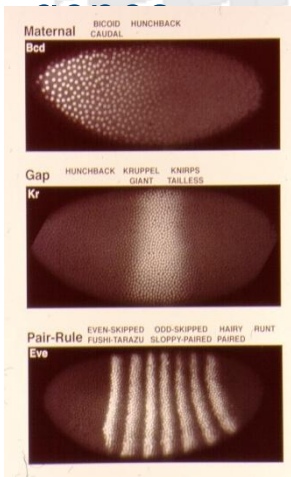
\$1 = slicing position, \$2 = intensity threshold value, \$3 = confidence



Gene Expression Analysis

<http://urchin.spbcas.ru/Mooshka/> [Samsonova et al]

- Gene expression = reading out genes for reproduction
- Research goal: capture spatio-temporal expression patterns in *Drosophila*



```
select jpeg( scale( {1c,0c,0c}*e[0,*,*,*:*]
                    +{0c,1c,0c}*e[1,*,*,*:*]
                    +{0c,0c,1c}*e[2,*,*,*:*], 0.2 ) )
from EmbryoImages as e
where oid(e)=193537
```

„Science SQL“

- **ISO 9075 Part 15: SQL/MDA**
 - resolved by ISO SQL WG in June 2014

- **n-D arrays as attributes**

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

- **declarative array operations**

```
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  
  avg( scene.band3-scene.band4)/(scene.band3+scene.band4) > 0
```

Information technology — Database languages — SQL —

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


Technologies de l'information — Langages de base de données — SQL —
Partie 15: Tableaux multi-dimensionnels (SQL/MDA)

Document type: Technical Report
Document subtype: Technical Report (TR)
Document stage: (3) CD under Consideration
Document language: English


Edited by: Jim Melton (Ed.) and Peter Baumann (Associate Ed.)

OGC WCPS

- OGC **Web Coverage Processing Service (WCPS)**
= high-level geo raster query language; adopted 2008



- WCPS 2: all grid types:



- "From MODIS scenes M1, M2, M3: **difference between red & nir**, as TIFF"
 - ...but only those where nir exceeds 127 somewhere

```
for $c in ( M1, M2, M3 )
where some( $c.nir > 127 )
return encode( $c.red - $c.nir, "image/tiff" )
```


(tiff_A,
tiff_C)

rasdaman: Agile Array Analytics

- „raster data manager“: SQL + n-D arrays

```
select img.green[x0:x1,y0:y1] > 130
from   LandsatArchive as img
where  avg_cells( img.nir ) < 17
```


- Scalable parallel “tile streaming” architecture
- Implemented & in use: ESA, NASA, PML, DWD, ...
 - OGC WCS Core Reference Implementation
 - Copernicus Masters Big Data Challenge 2014 winner




Adaptive Tiling

- Sample tiling strategies [Furtado]:



- regular



directional



area of interest

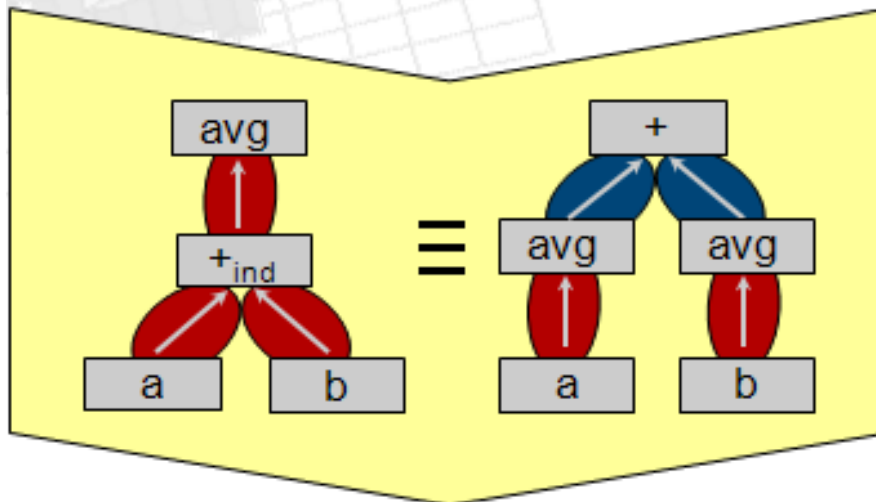


- rasdaman storage layout language

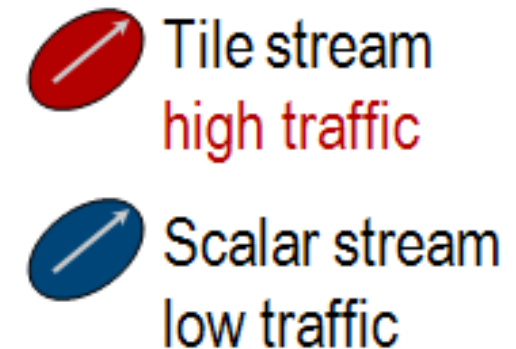
```
insert into MyCollection
values ...
tiling area of interest [0:20,0:40], [45:80,80:85]
tile size 1000000
index d_index storage array compression zlib
```

Optimization 1: Query Rewriting

```
select avg_cells( a + b )
from a, b
```



```
select avg_cells( a )
       + avg_cells( b )
from a, b
```




- *understood:*
heuristic optimization
– 150 rules in rasdaman [Ritsch 2002]
- *partially understood:*
cost-based optimization

Optimization 2: Just-In-Time Compilation

[Jucovschi, Stancu-Mara 2008]

- Observation: interpreted mode slows down
- Approach:
 - cluster suitable operations
 - compile & dynamically bind
- Benefit:
 - Speed up complex, repeated operations
- Variation:
 - compile code for GPU (later)

```
select x*x*...*x
from float_matrix as x
```



Times [ms] for 512² * n ops

Parallel / Distributed Query Processing

select

$\max((A.nir - A.red) / (A.nir + A.red))$

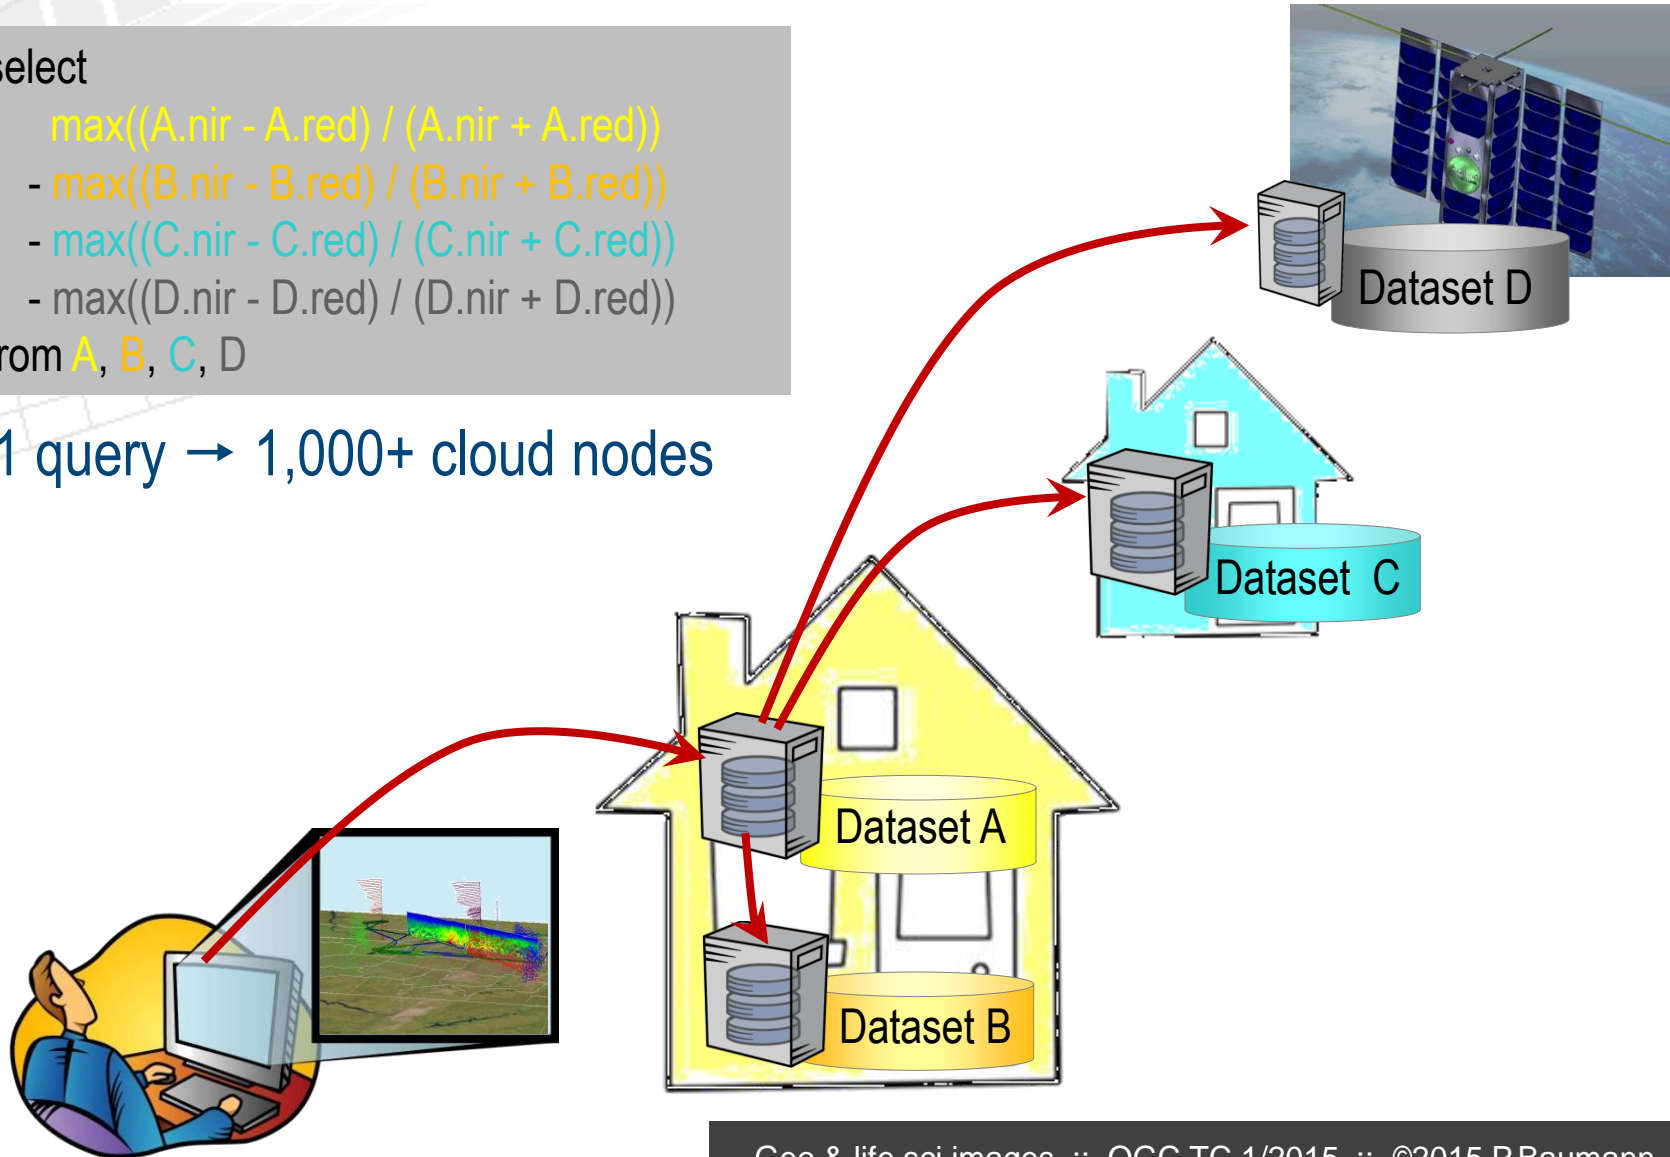
- $\max((B.nir - B.red) / (B.nir + B.red))$

- $\max((C.nir - C.red) / (C.nir + C.red))$

- $\max((D.nir - D.red) / (D.nir + D.red))$



from A, B, C, D

1 query \rightarrow 1,000+ cloud nodes



Inset: Hadoop is not the Answer to All


- **no builtin knowledge** about structured data types
 - “Since it was not originally designed to leverage the structure [...] its **performance** [...] is therefore **suboptimal**” [Daniel Abadi]
 - M. Stonebraker (XLDB 2012): „will hit a **scalability wall**“



COMMON SENSE

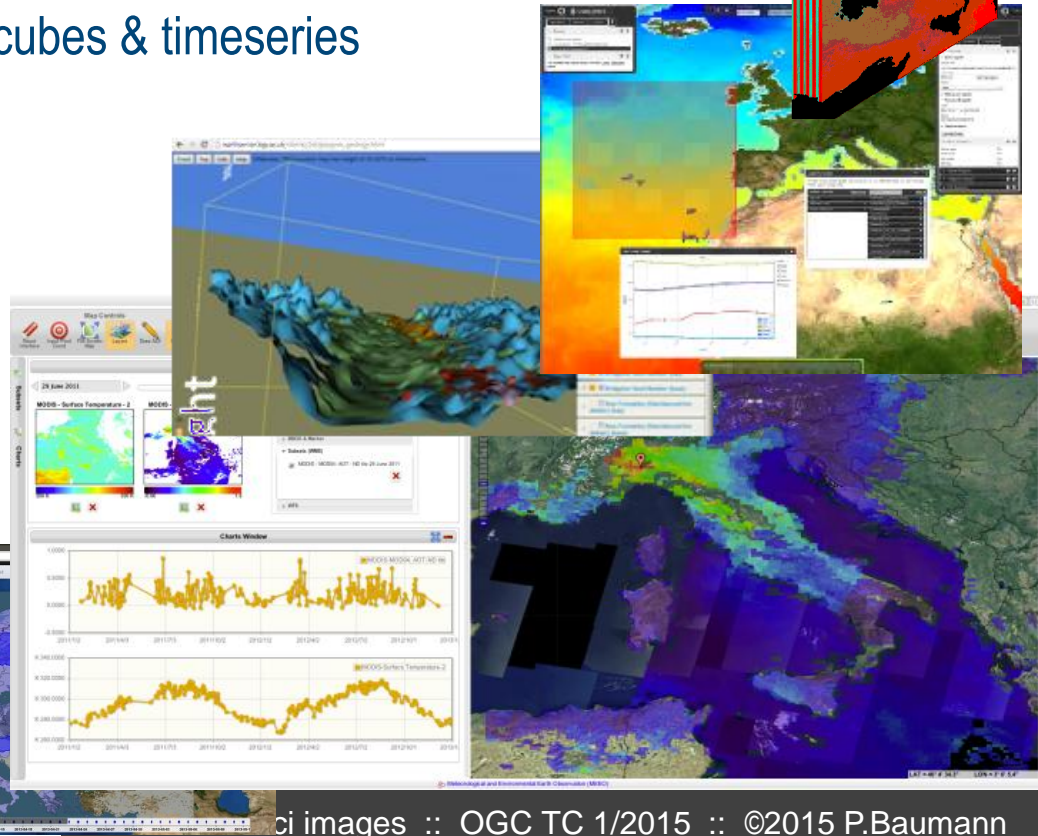

Just because you can, doesn't mean you should.

A Brief History of Array Databases



Conclusion

- **n-D Arrays** a major datatype, central to **science, engineering, business**
 - Massive spatio-temporal *sensor, image, simulation, statistics* data
- **Query language** = flexibility + scalability + information integration
 - 130+ TB databases, 2D...5D datacubes & timeseries
 - 1 query → 1,000+ cloud nodes
- ISO SQL/MDA a **game changer**
 - Any question, any time
 - No **data/metadata divide**
- Visit us:
 - www.rasdaman.org
 - www.earthserver.eu





JACOBS
UNIVERSITY

JADE

Jacobs Academy of Data Engineering