

Array Database Assessment Working Group

Peter Baumann, Kwo-Sen Kuo (chairs)

[gamingfeeds.com]



ADA:WG :: RDA P8 :: ©2016 rasdaman

About ADA:WG

- Array Database Assessment WG
- Creature of Big Data Interest Group
- Chairs: Peter Baumann, Kwo-Sen Kuo
- Info:
 - <https://www.rd-alliance.org/groups/array-database-working-group.html> (WG page)
 - <https://rd-alliance.org/group/case-statement/array-database-working-group.html> (case stmt)
 - <https://www.rd-alliance.org/group/array-database-assessment-wg/wiki/array-database-assessment-working-group> (wiki - DO NOT USE file repo!)

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

- 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

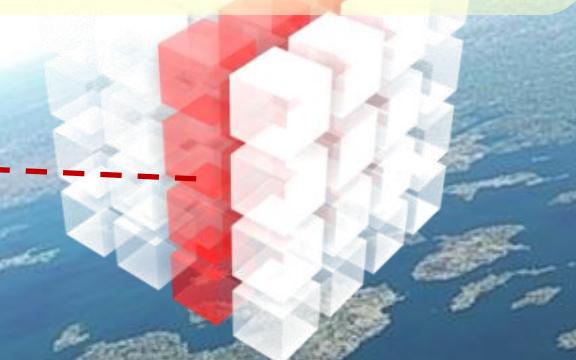
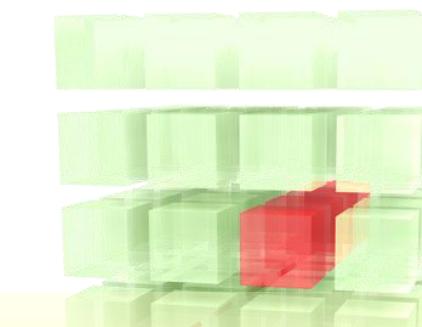
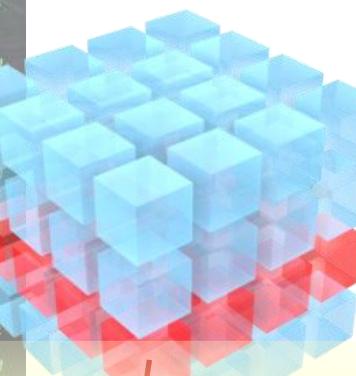
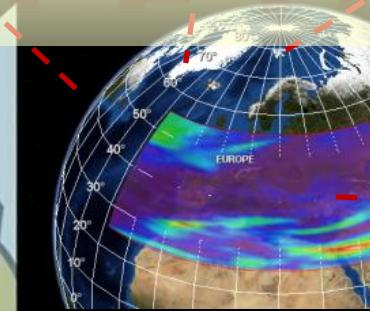
sets + hierarchies + graphs + arrays)

SERVICE QUALITY

SERVICE QUALITY

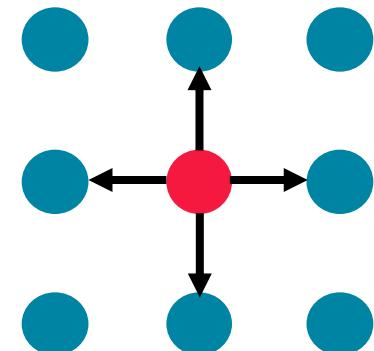
```
-rwx--x--- 1 rasdata users 1485 Oct 13 2004 4251NW.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251NWGR.tif
-rwx--x--- 1 rasdata users 640432 Oct 13 2004 4251NWGR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251NWGW.tif
-rwx--x--- 1 rasdata users 779368 Oct 13 2004 4251NWGW.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251NWRL.tif
-rwx--x--- 1 rasdata users 712492 Oct 13 2004 4251NWRL.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251NMWL.tif
-rwx--x--- 1 rasdata users 62830 Oct 13 2004 4251NMWL.tif
-rwx--x--- 1 rasdata users 1498 Oct 13 2004 4251SO.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SOGR.tif
-rwx--x--- 1 rasdata users 1076750 Oct 13 2004 4251SOGR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SOGW.tif
-rwx--x--- 1 rasdata users 197142 Oct 13 2004 4251SOGW.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SORL.tif
-rwx--x--- 1 rasdata users 936348 Oct 13 2004 4251SORL.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SOWL.tif
-rwx--x--- 1 rasdata users 119990 Oct 13 2004 4251SOWL.tif
-rwx--x--- 1 rasdata users 1485 Oct 13 2004 4251SW.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SWGR.tif
-rwx--x--- 1 rasdata users 577868 Oct 13 2004 4251SWGR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4251SWGW.tif
-rwx--x--- 1 rasdata users 252198 Oct 13 2004 4251SWGW.tif
-rwx--x--- 1 rasdata users 170000 Oct 13 2004 4251SWGW.tif
-rwx--x--- 1 rasdata users 180000 Oct 13 2004 4251SWGW.tif
-rwx--x--- 1 rasdata users 412000 Oct 13 2004 4251SWGW.tif
-rwx--x--- 1 rasdata users 1485 Oct 13 2004 4252N0.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252N0GR.tif
-rwx--x--- 1 rasdata users 355774 Oct 13 2004 4252N0GR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252N0GW.tif
-rwx--x--- 1 rasdata users 49046 Oct 13 2004 4252N0GW.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252N0RL.tif
-rwx--x--- 1 rasdata users 600964 Oct 13 2004 4252N0RL.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252N0WL.tif
-rwx--x--- 1 rasdata users 46714 Oct 13 2004 4252N0WL.tif
-rwx--x--- 1 rasdata users 1485 Oct 13 2004 4252NW.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252NWGR.tif
-rwx--x--- 1 rasdata users 1445064 Oct 13 2004 4252NWGR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252NWGW.tif
-rwx--x--- 1 rasdata users 410426 Oct 13 2004 4252NWGW.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252NWRL.tif
-rwx--x--- 1 rasdata users 655374 Oct 13 2004 4252NWRL.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252NMWL.tif
-rwx--x--- 1 rasdata users 108612 Oct 13 2004 4252NMWL.tif
-rwx--x--- 1 rasdata users 1485 Oct 13 2004 4252SO.ASC
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252SOGR.tif
-rwx--x--- 1 rasdata users 607646 Oct 13 2004 4252SOGR.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252SOGW.tif
-rwx--x--- 1 rasdata users 685092 Oct 13 2004 4252SOGW.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252SORL.tif
-rwx--x--- 1 rasdata users 632172 Oct 13 2004 4252SORL.tif
-rwx--x--- 1 rasdata users 216 Oct 13 2004 4252SOWL.tif
```

SERVICE QUALITY



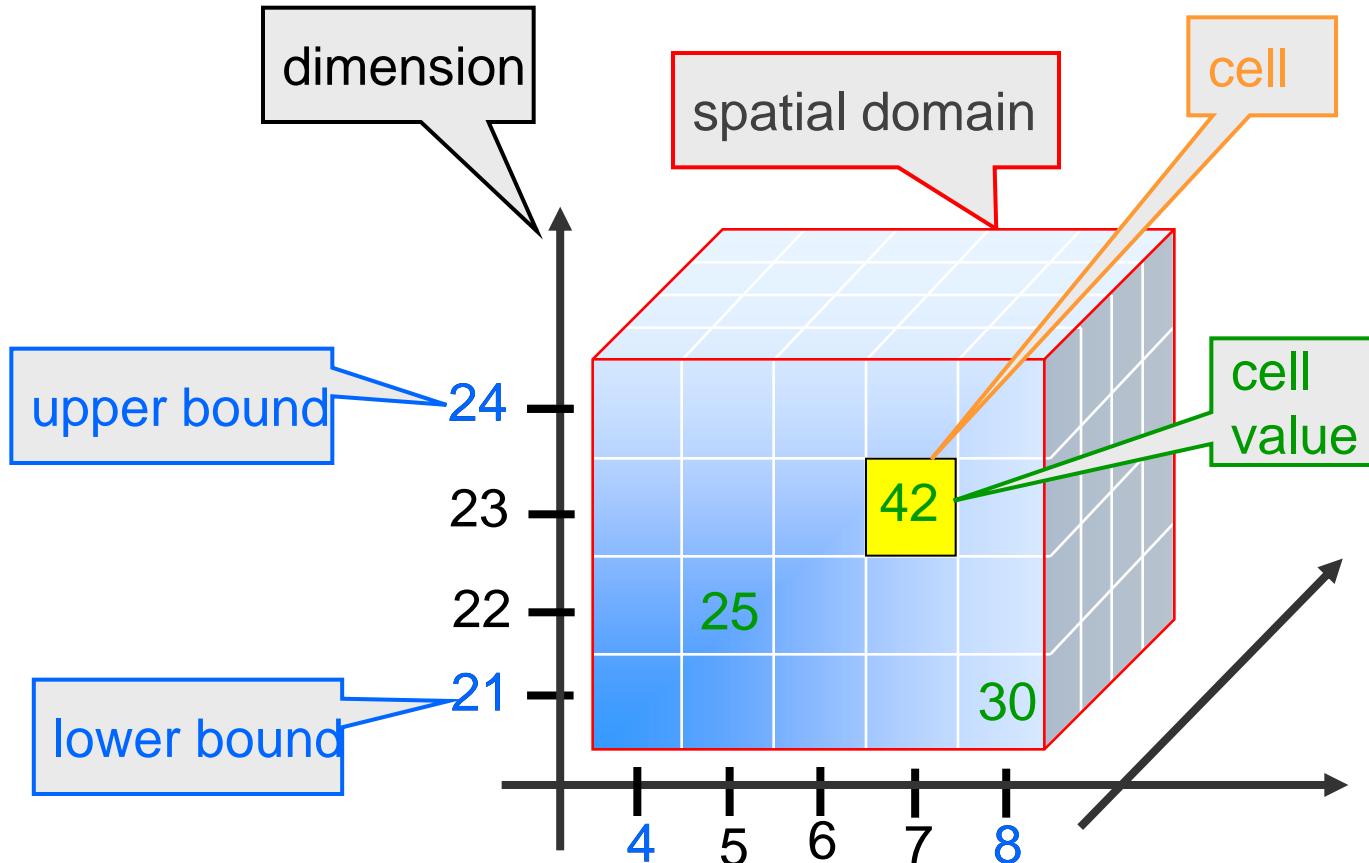
Array Analytics

- Array Analytics :=
Efficient analysis on multi-dimensional arrays of a size several orders of magnitude above evaluation engine's main memory
- Essential **data** property: n-dimensional Euclidean neighborhood
 - Secondary: #dimensions, density, ...
- **Operations**: statistics, image/signal processing, Linear Algebra++



[EDBT/ICDT Array Databases Workshop, 2011]

The Multidimensional Data Model



Hadoop – Not the Answer to All

- “Since it was not originally designed to leverage the **structure** its **performance is suboptimal**” [Daniel Abadi]
- Similar for SPARK

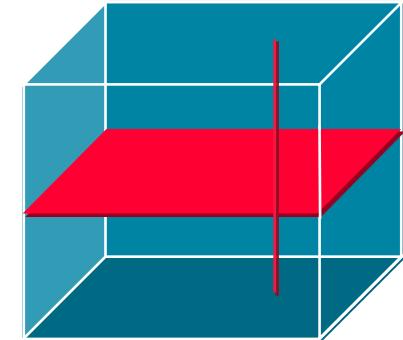
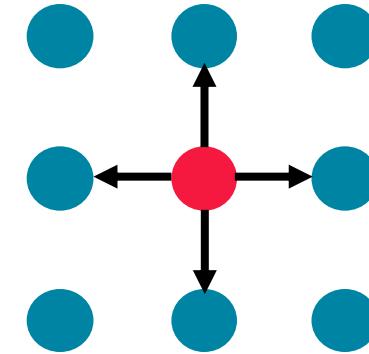


COMMON SENSE

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



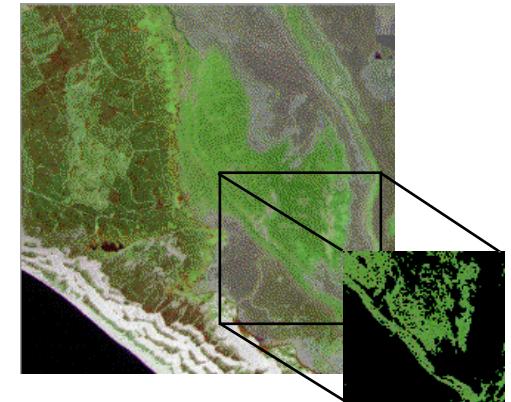
ADA:WG :: RDA P8 :: ©2016 rasdaman



rasdaman

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

- Scalable parallel “tile streaming” architecture
- Pioneer Array DBMS
- in operational use, 145+ TB databases
- Blueprint for standards in ISO, OGC, INSPIRE



Direct Data Visualization



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

Linear Algebra Ops

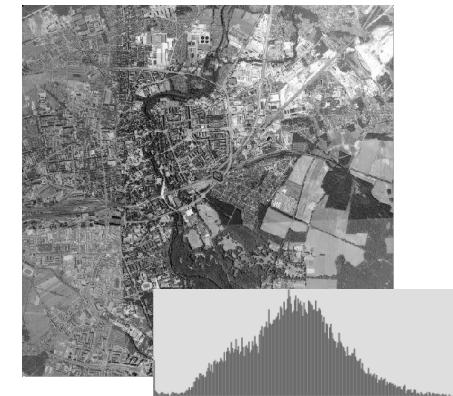
- Matrix multiplication

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

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

- Histogram

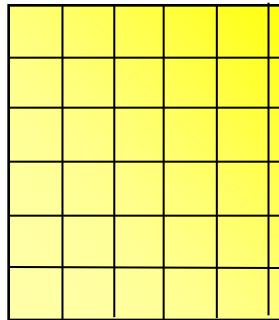
```
select marray bucket in [0:255]
      values count_cells( img = bucket )
from    img
```



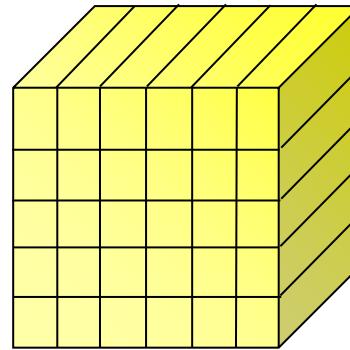
Adaptive Tiling

- Sample tiling strategies [ICDE 1999]:

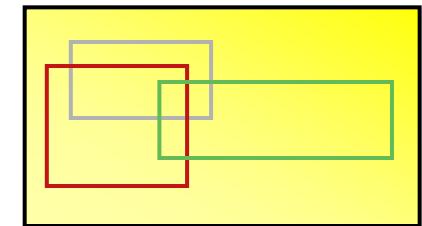
regular



directional



area of interest



- rasdaman storage layout language [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
```

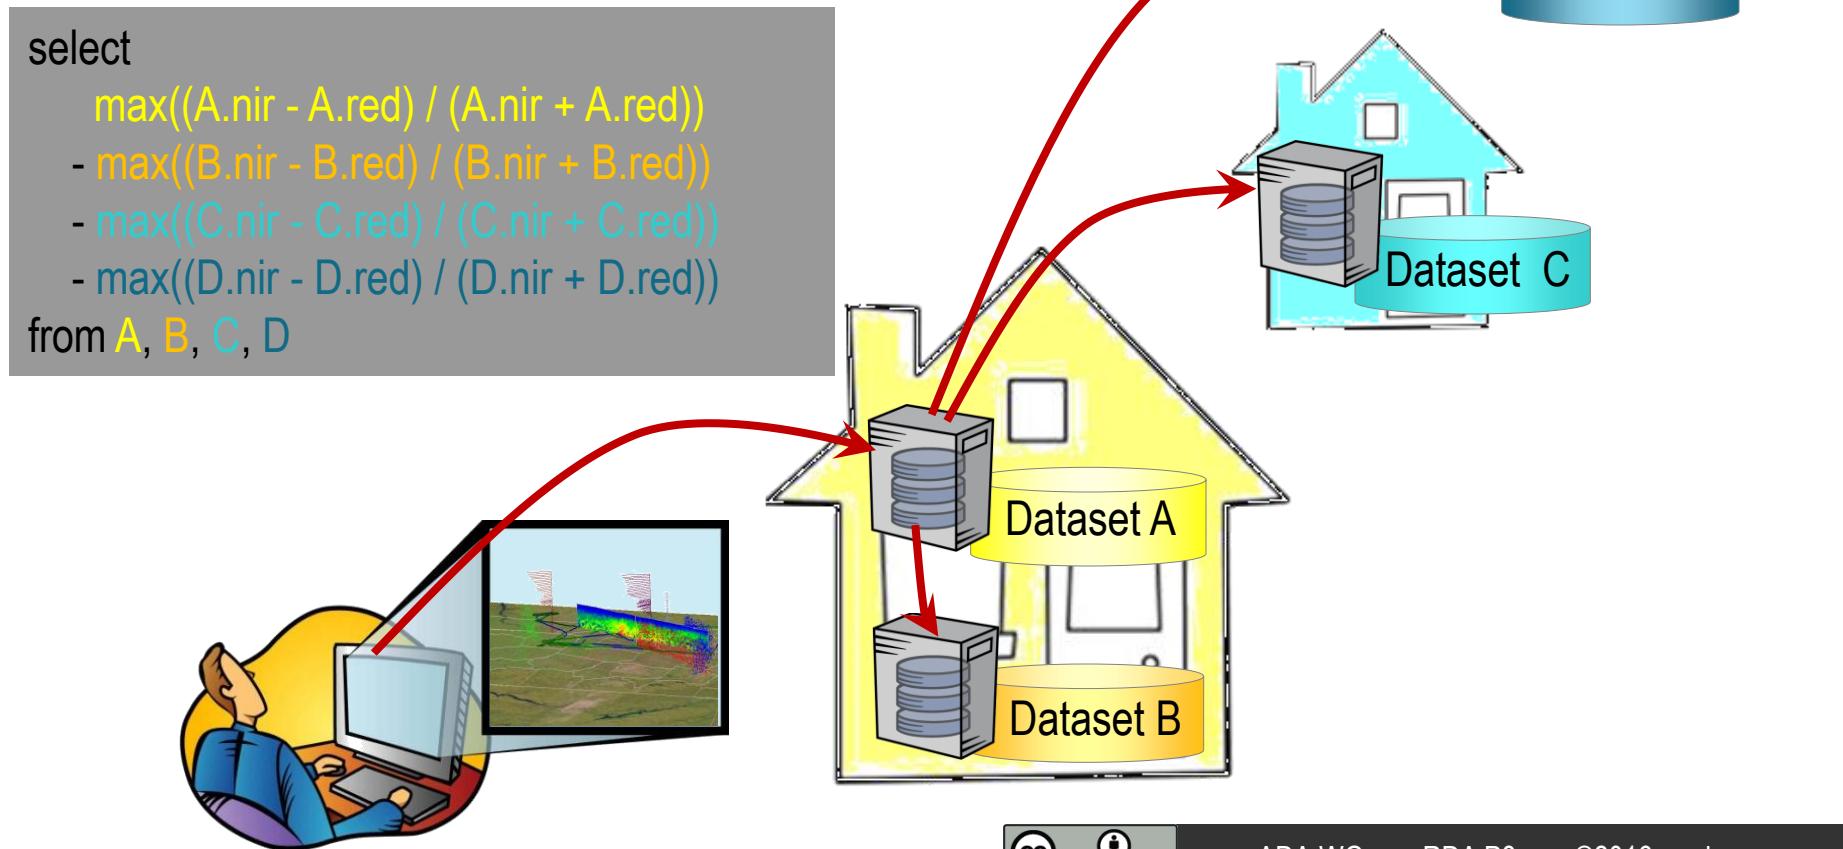
Parallel / Distributed Query Processing

1 query → 1,000+ cloud nodes

[SIGMOD DANAC 2014]

select

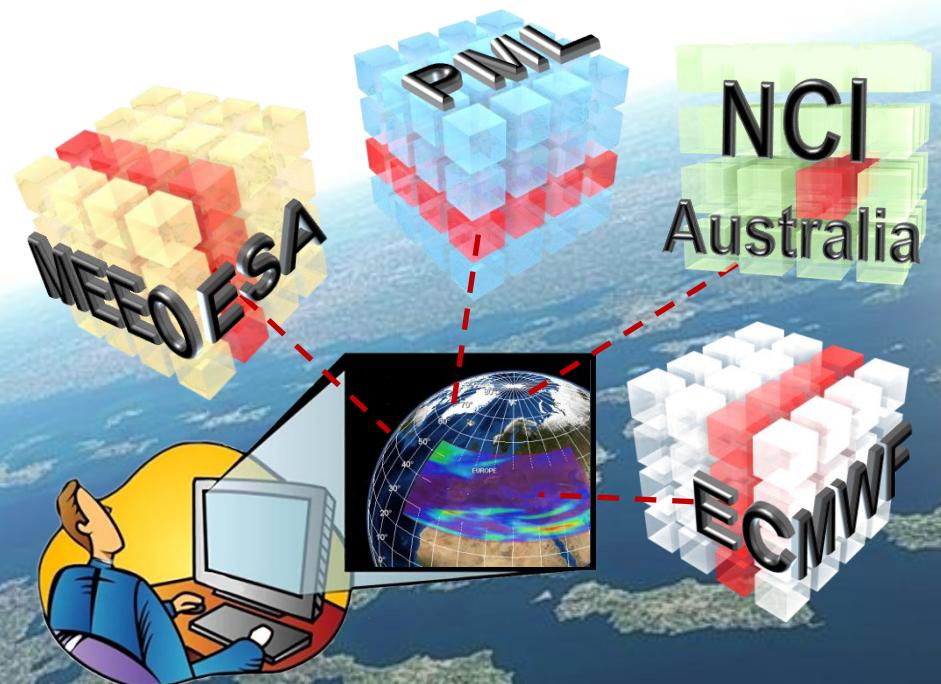
```
max((A.nir - A.red) / (A.nir + A.red))  
- max((B.nir - B.red) / (B.nir + B.red))  
- max((C.nir - C.red) / (C.nir + C.red))  
- max((D.nir - D.red) / (D.nir + D.red))  
from A, B, C, D
```





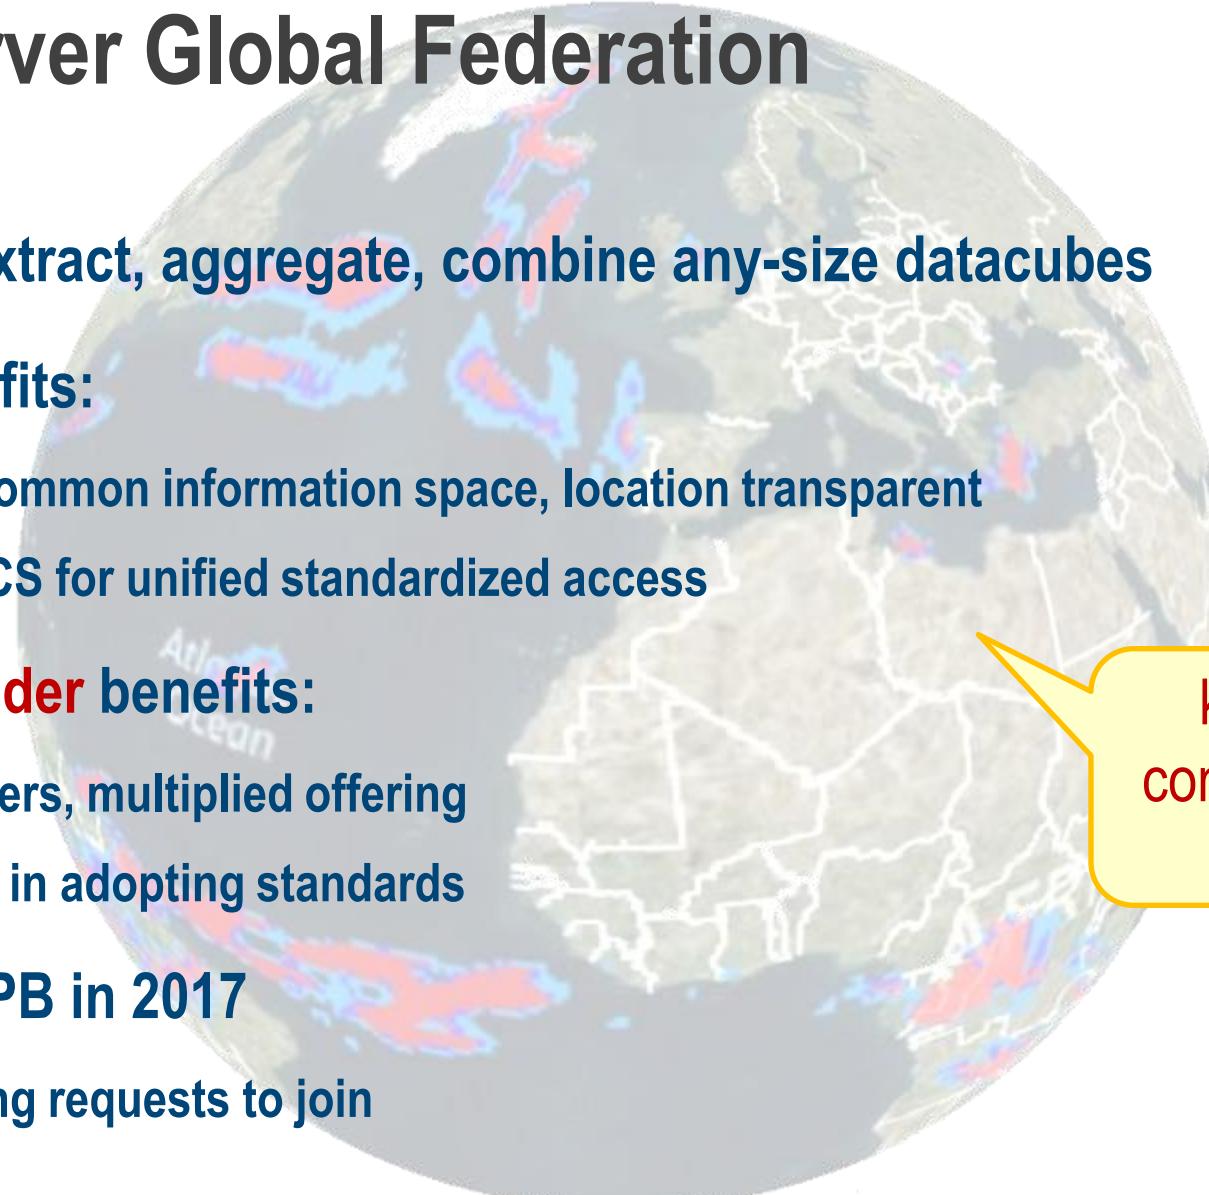
EarthServer: Datacubes At Your Fingertips

- Intercontinental initiative, 3+3 years: EU + US + AUS
- Agile Analytics on x/y/t + x/y/z/t Earth & Planetary datacubes
 - Rigorously standards: OGC WMS + WCS + WCPS
 - EU rasdaman + NASA WorldWind
 - 100s of TB sites now, next: 1+ PB



EarthServer Global Federation

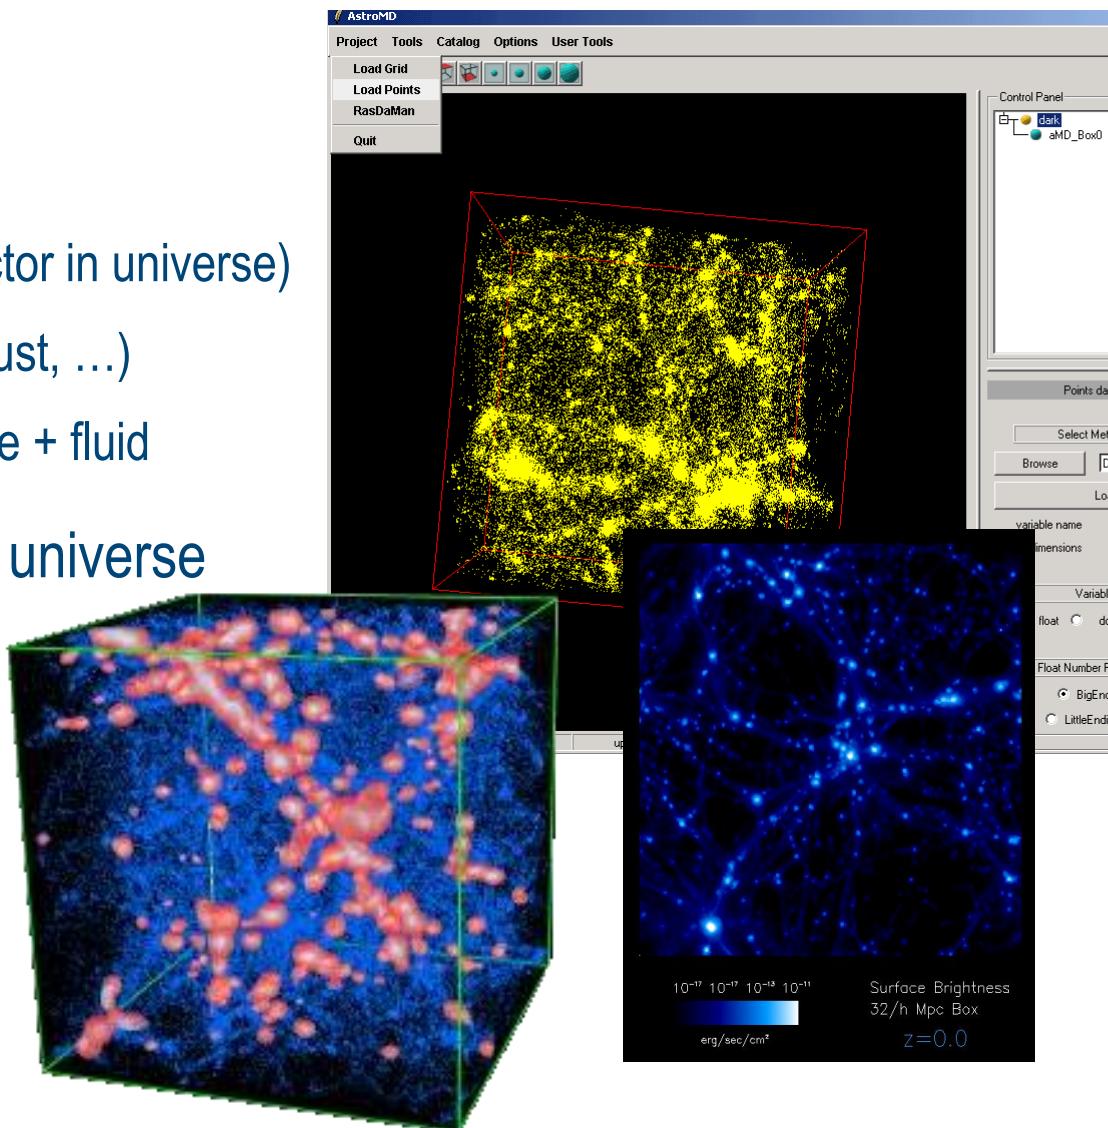
- **Access, extract, aggregate, combine any-size datacubes**
- **User benefits:**
 - single common information space, location transparent
 - OGC WCS for unified standardized access
- **Data provider benefits:**
 - More users, multiplied offering
 - Support in adopting standards
- **Goal: 3.5 PB in 2017**
 - Receiving requests to join



keeping contributions visible

Cosmological Simulation

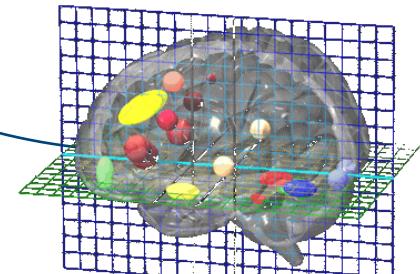
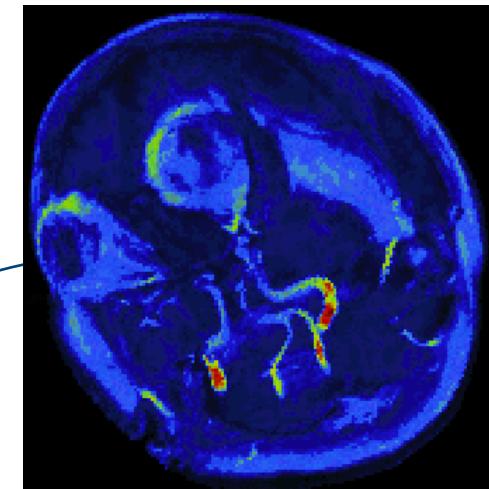
- Modelling domain: 4D
 - Dark matter (highest mass factor in universe)
 - Baryonic matter (stars, gas, dust, ...)
 - → Coupled simulation: particle + fluid
- Results: 3D/4D cutouts from universe
 - Eg, 64 Mpc³
(1 pc = 3.27 light years)
- Screenshots: AstroMD
[Gheller, Rossi 2001]



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
```

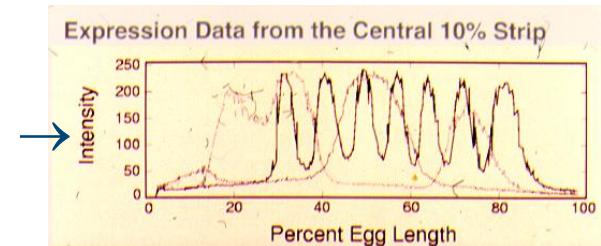
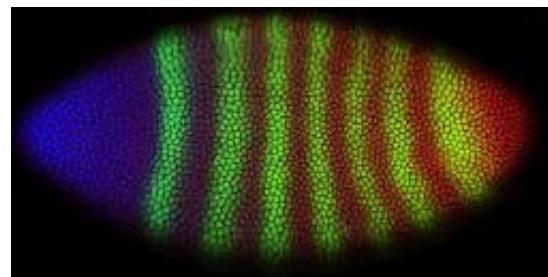
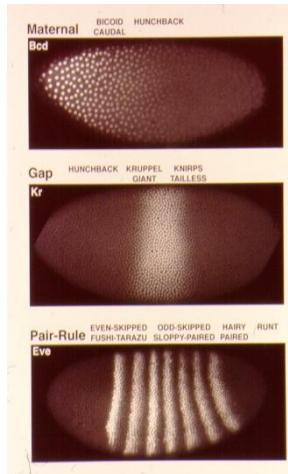


\$1 = slicing position, \$2 = intensity threshold value, \$3 = confidence

Gene Expression Analysis

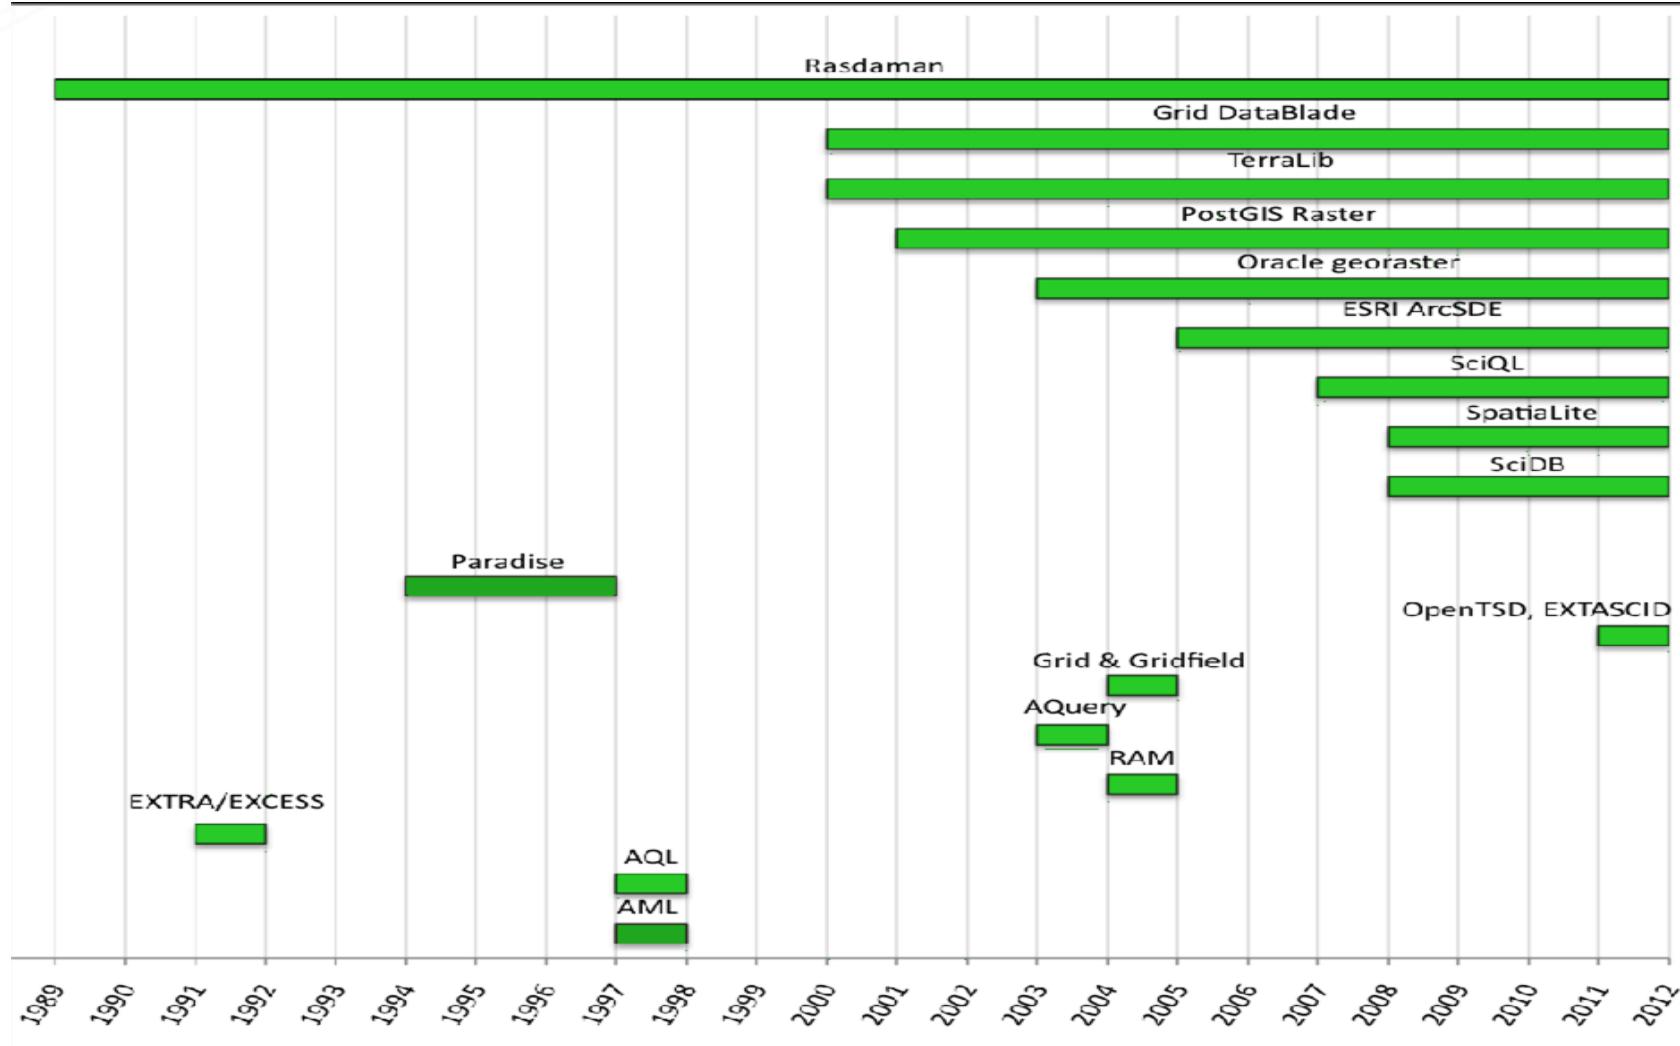
<http://urchin.spbcas.ru/Mooshka/> [Samsonova et al]

- 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
```

Array Databases: Historical Evolution





Array SQL

[SSDBM 2014]

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)

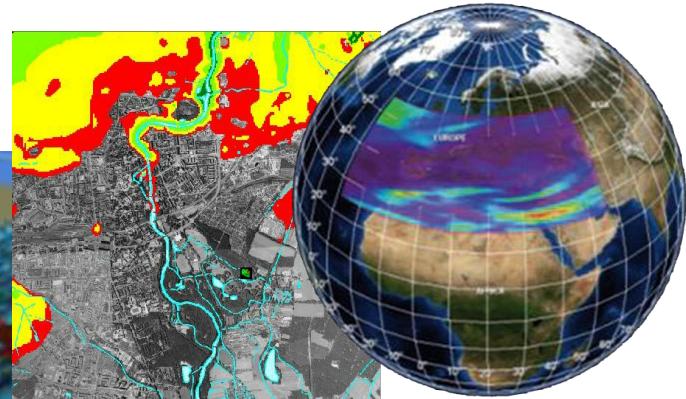
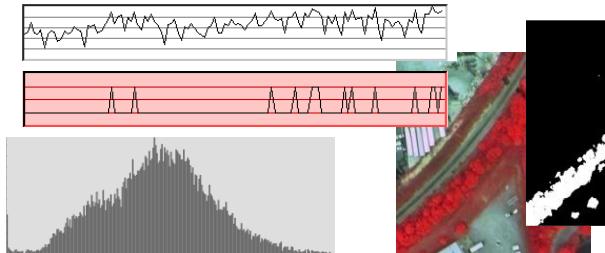
```
create table LandsatScenes(  
    id: integer not null, acquired: date,  
    scene: row( band1: integer, ..., band7: integer ) mdarray [ 0:4999,0:4999 ] )
```

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

OGC WCPS

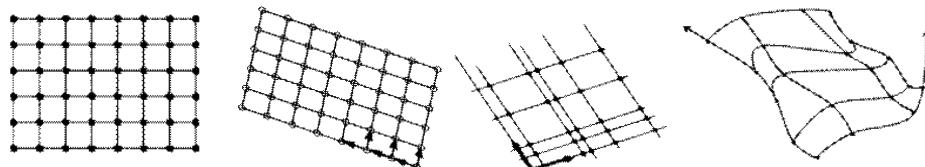
[Geoinformatica 2010]

- = OGC Web Coverage Processing Service
- geo raster query language



[JacobsU, FhG; NASA; data courtesy BGS, ESA]

- Part of OGC “Big Geo Data” suite, Web Coverage Service (WCS)
 - Any type of spatio-temporal grid



Array Database WG: Charter

- Spinoff from BigData IG
- Mission: to provide support for technologists and decision makers considering Big Data services using Array Databases in academic and industrial environments
- Approach: neutral, thorough hands-on evaluation of ADBMSs & comparable technology ...
 - based on relevant standards, such as ISO “Array SQL” [8], OGC WCPS
 - comparing technical criteria like functionality, thereby eliciting the state of the art;
 - establishing & running a combination of domain-driven & domain-neutral benchmarks;
 - plus real-life, publicly accessible deployments at scale.

Array Database WG: Value Proposition

- Determine to what extent can data scientists and engineers benefit from Array Database technology?
 - What **systems** are out there?
 - *currently considered ADBMSs: rasdaman, SciQL, SciDB, PostGIS Raster, ...*
 - *alternative technology: SciHadoop, HDF5 on filesystems, Python data formats*
 - What are their **features**?
 - What is their **performance**, measured through objective, open tests?
 - What are relevant **tuning** parameters & recommended settings?
 - How can these systems be used in large-scale **deployments**?
 - Generally: strengths & weaknesses of ADBMSs

Array Database WG: Collaboration Plans

- Standardization:
 - ISO SC32/WG3 “SQL”, TC211/WG6 “Geo Imagery”; WG9 “Big Data”
 - OGC
 - INSPIRE
- Related Activities:
 - EUDAT Array Database WG
 - EarthServer
- Open for more!
 - ...enough work on the plate

Let's Roll!

- Participant types (can overlap):
 - Use case providers: data & desirable queries
 - Service providers: storage & compute facilities, hosting datacubes
 - ADBMS software providers: software + support
- Outcome: N public „datacube“ services; sufficiently large

*Interested?
Contact us today!*



„One cube tells more than a million images“

