# Array Databases:
# Agile Analytics for Spatio-Temporal Big Data

**5th Workshop on the use of GIS/OGC standards in meteorology**

**2014-oct-29, DWD, Offenbach, Germany**

**Peter Baumann**

Jacobs University | rasdaman GmbH

baumann@rasdaman.com

[gamingfeeds.com]

# 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.

# Arrays in [Geo] Science & Engineering

- spatio-temporal sensor, image, simulation, statistics data(cubes)

**sensor feeds**

**simulation data**

**Big Data server**
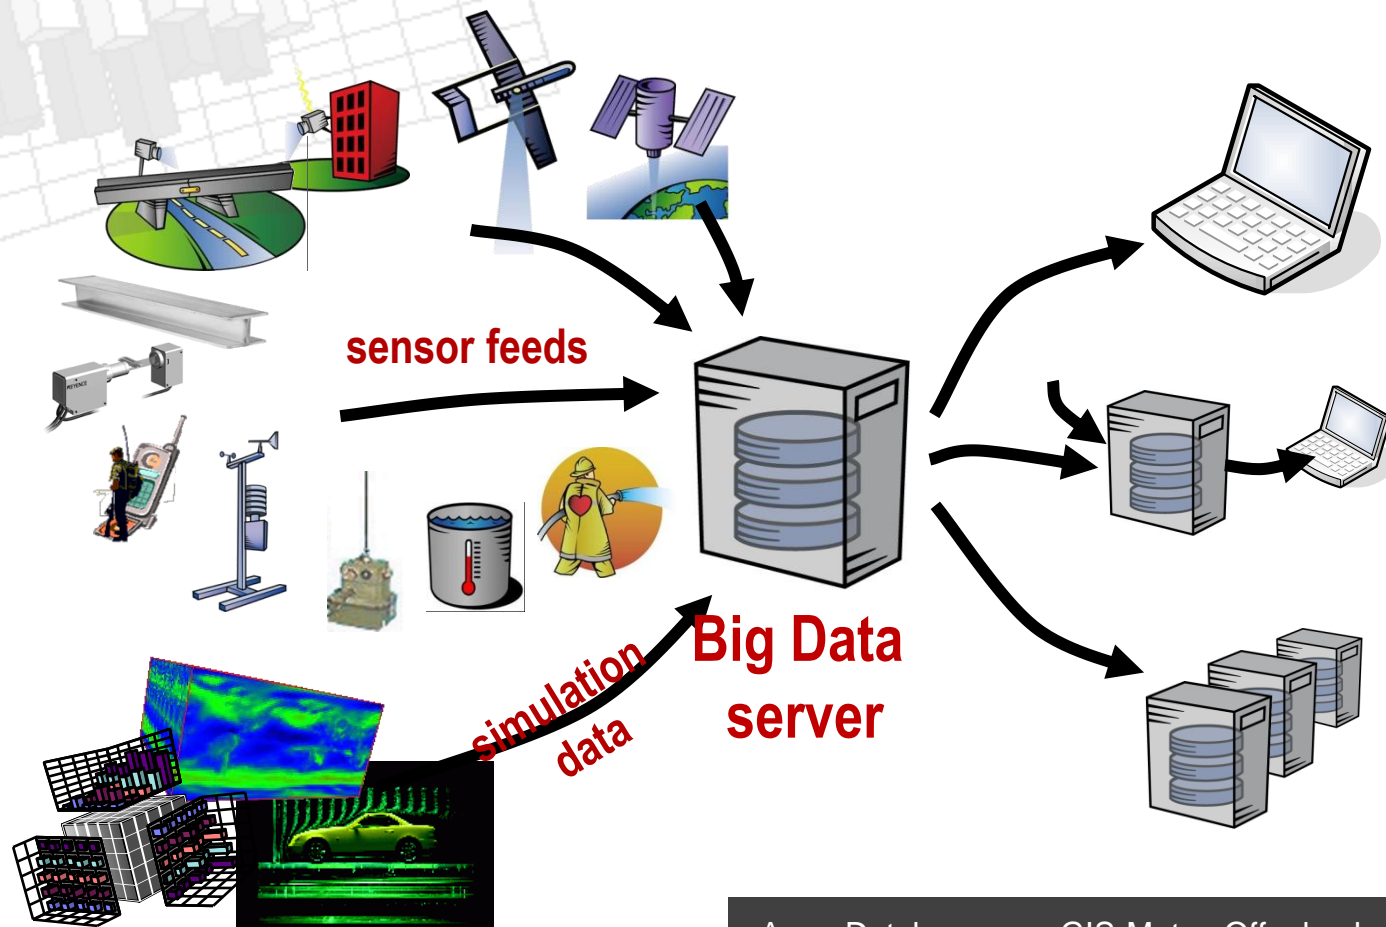
# What Users Really Want

- "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 taken by the SEVIRI satellite on August 25, 2007 which contain fire hotspots in areas which have been classified as forests 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

# rasdaman: Agile Array Analytics

- „raster data manager": SQL + n-D raster objects

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

- Scalable parallel "tile streaming" architecture

- In operational use
  - OGC Web Coverage Service
    Core Reference Implementation
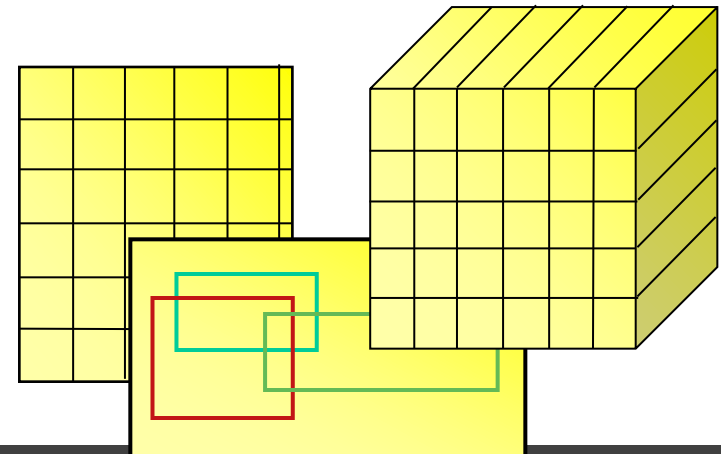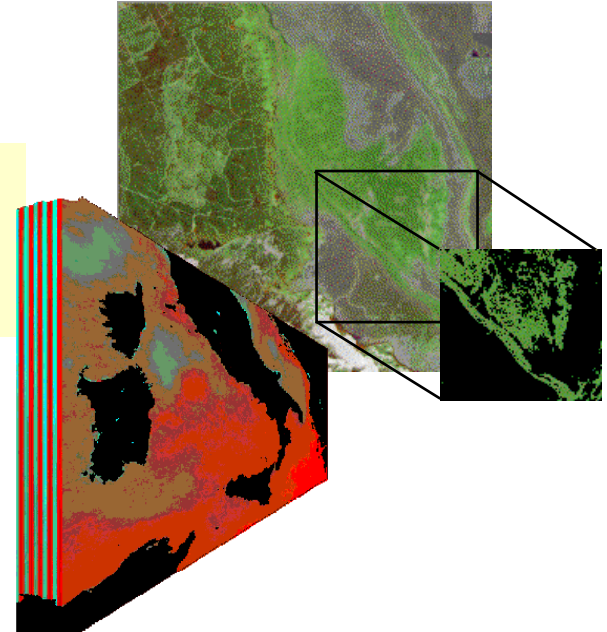
# Inset: Hadoop is not the Answer to All

- 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"
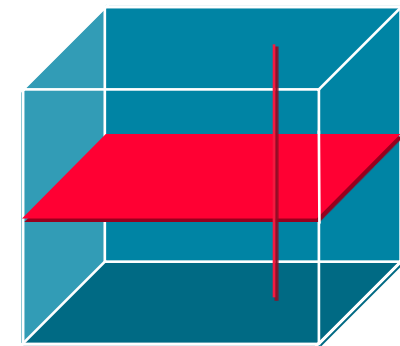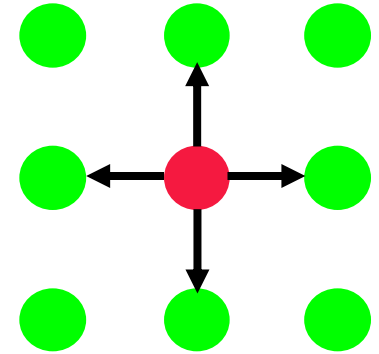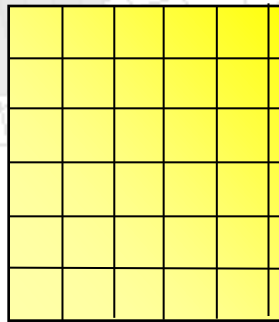


**COMMON SENSE**
Just because you can, doesn't mean you should.
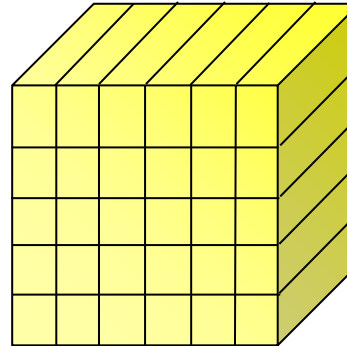
# Adaptive Tiling

- Sample tiling strategies [Furtado]:
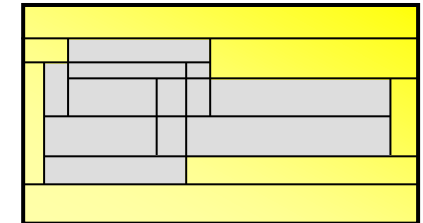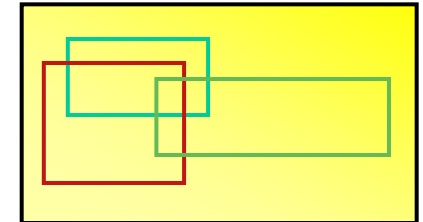
  - regular             directional            area of interest



- 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
```

# 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
```

[Jacobs U, Fraunhofer: data courtesy
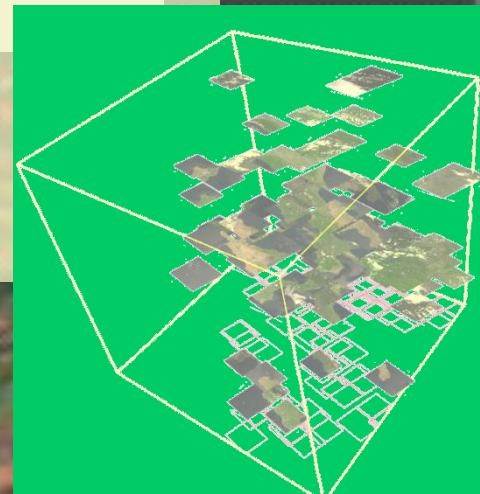
# Use Case: Plymouth Marine Laboratory

- "Avg chlorophyll concentration for given area & time period, from x/y/t cube"
  - 10, 60,120, 240 days

- Conclusions:
  - „we must minimise data transfer as well as [client] processing"
  - "standards such as WCPS provide the greatest benefit"





rasdaman

## North Sea

Response Time (S) vs North Sea (subset, wps, wcps)

Legend: ☐10 ■60 ☐120 ■240
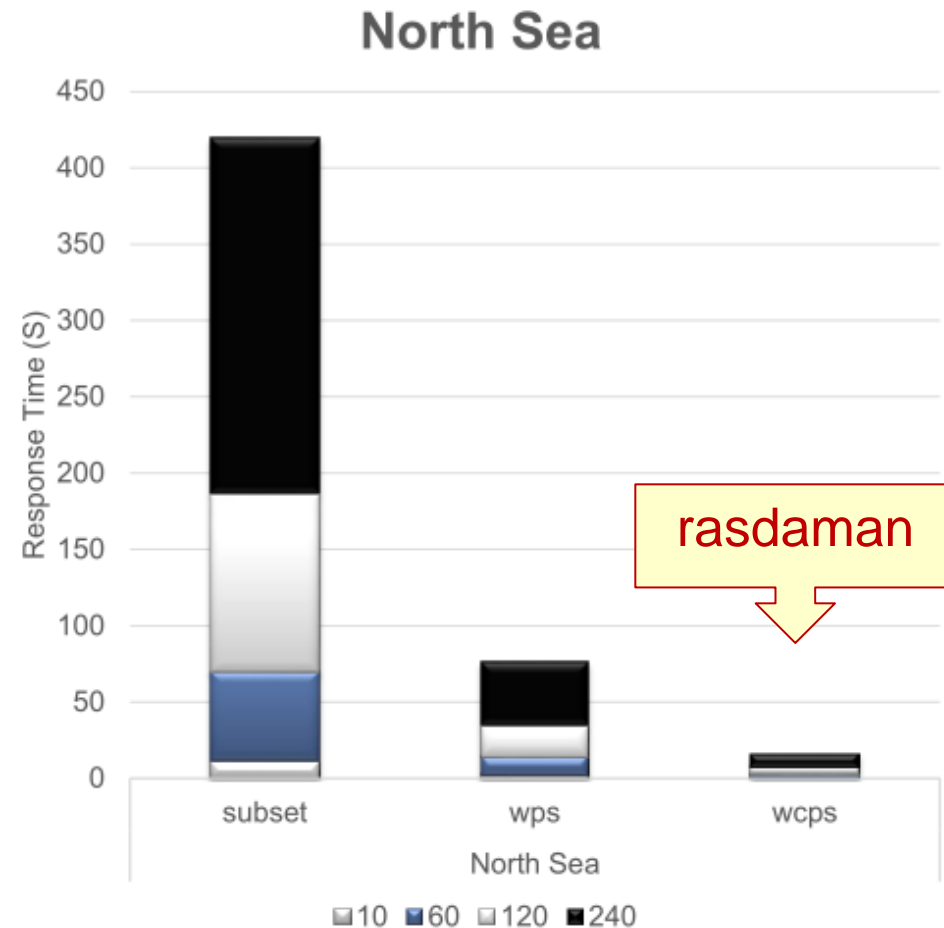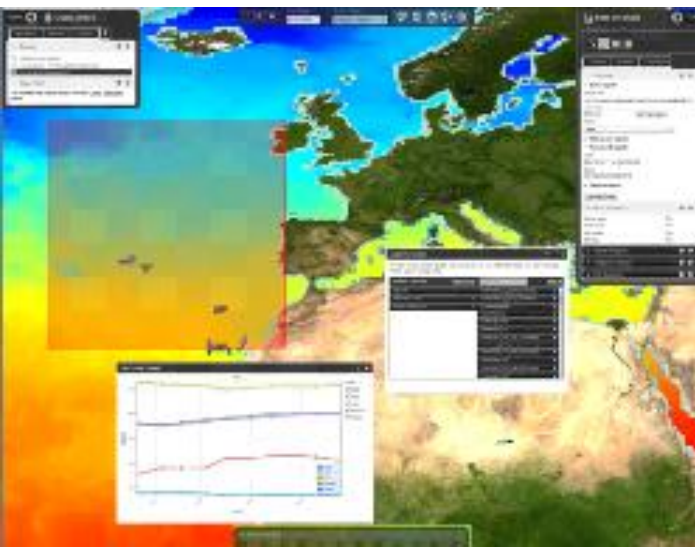
# Parallel / Distributed Query Processing

- 1 query → 1,000+ cloud nodes

```
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
```
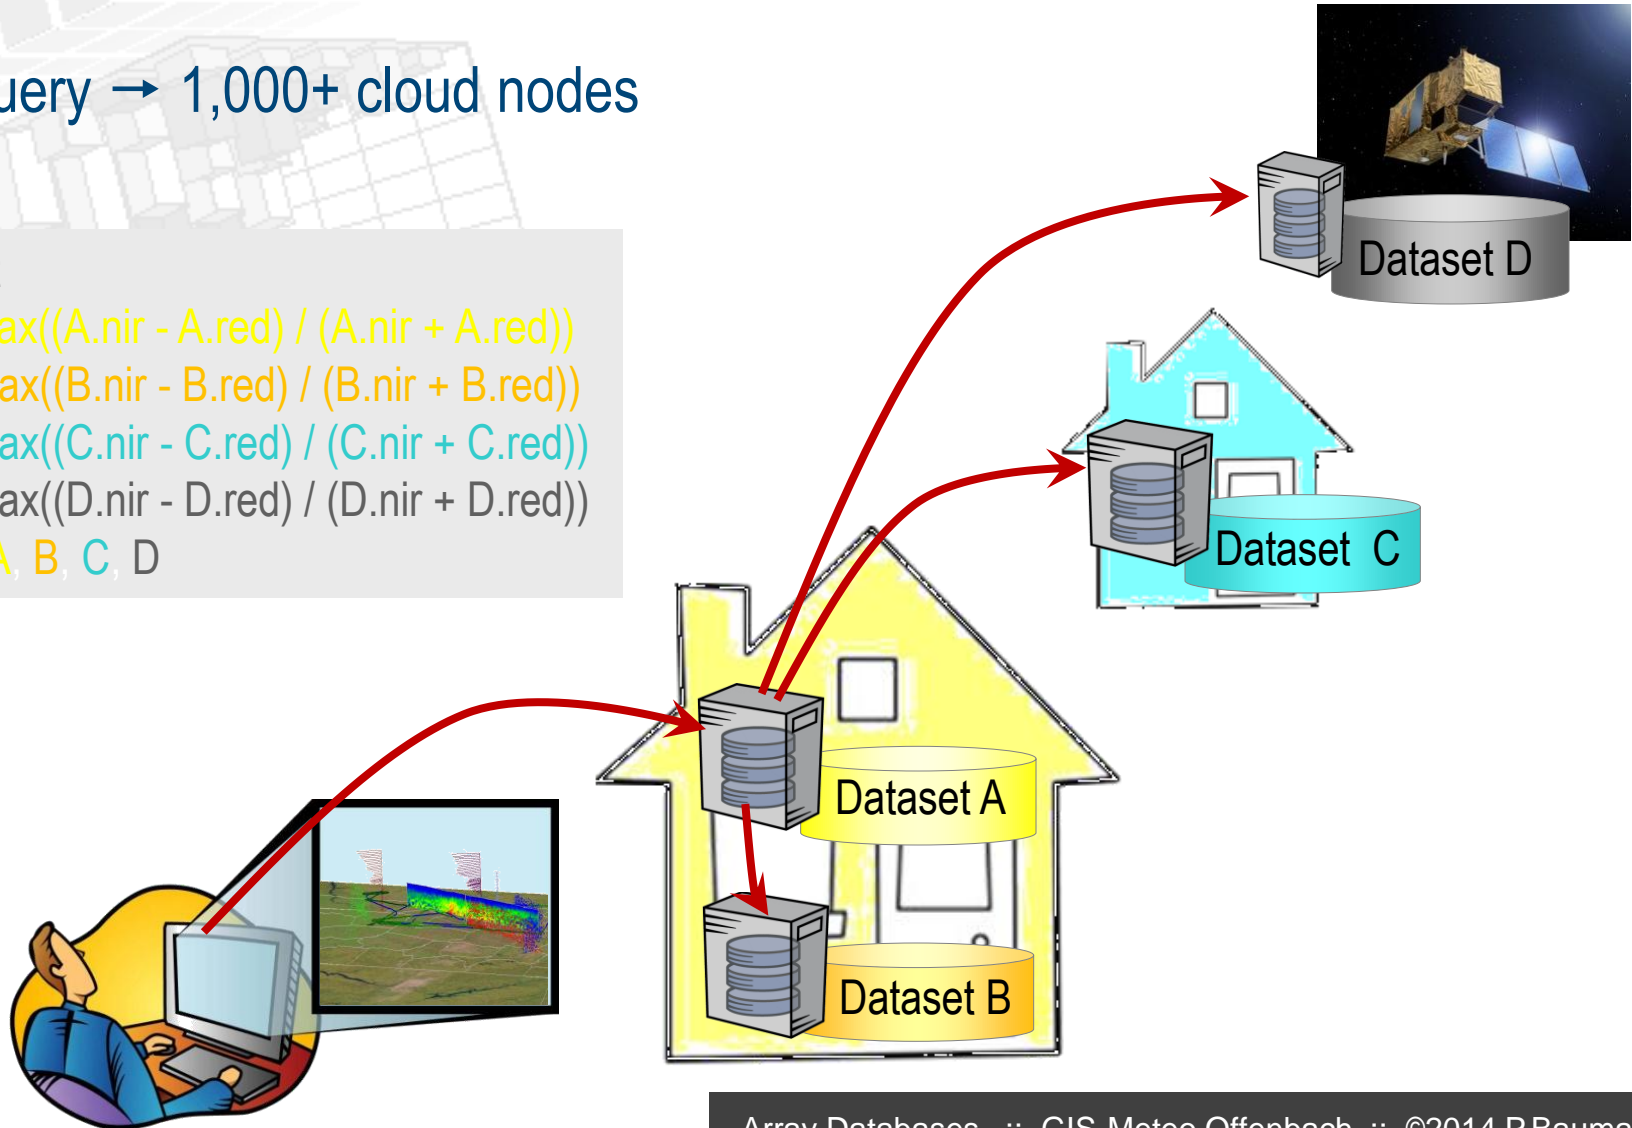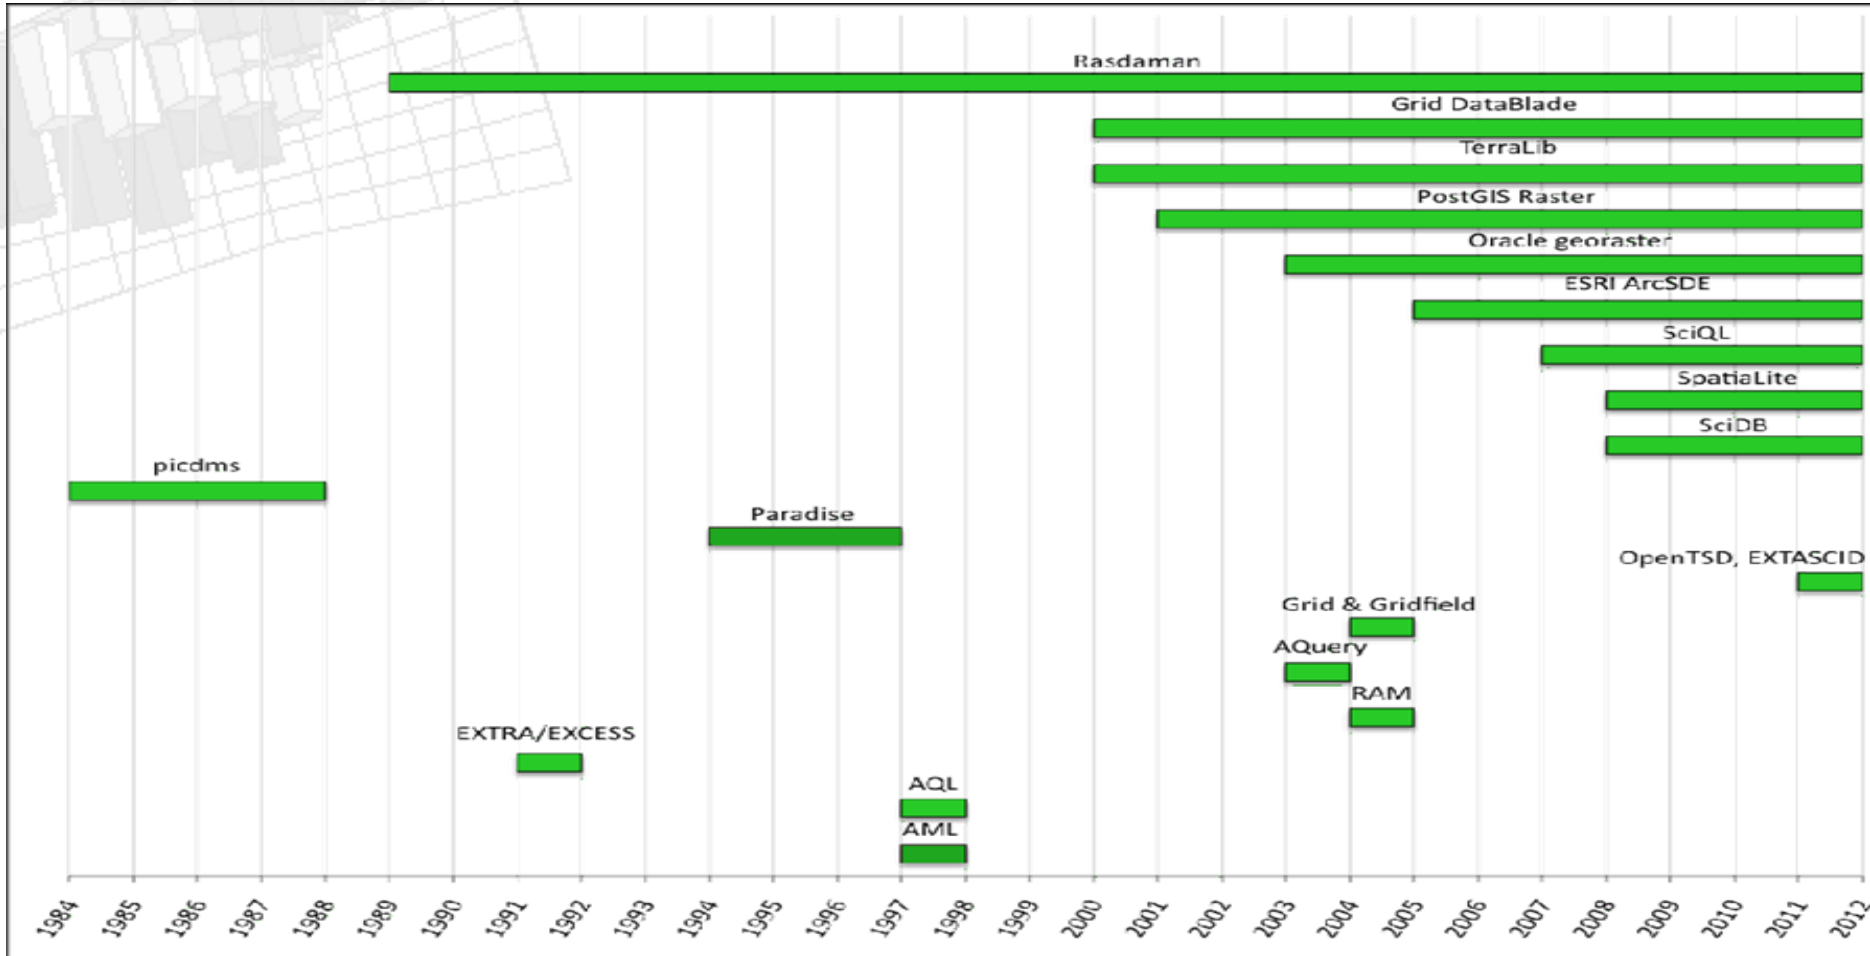


Dataset D

Dataset C

Dataset A

Dataset B

# A Brief History of Array Databases

# Recent Progress: ISO Array 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

# OGC WCPS

- OGC Web Coverage Processing Service (WCPS)
  = high-level geo raster query language; adopted 2008
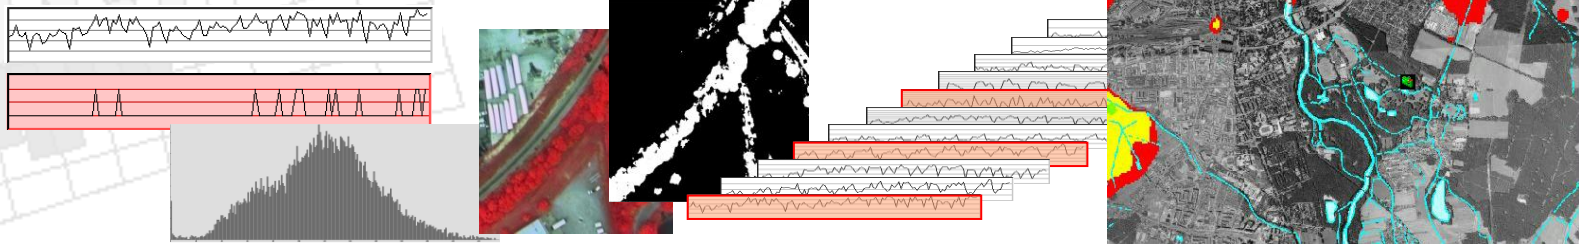
    - WCPS 2: all grid types:
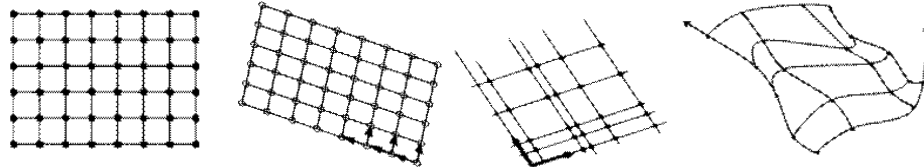
- "From MODIS scenes M1, M2, M3: difference between red & nir, as TIFF"

    - …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" )
```

$$( tiff_A, \; tiff_C )$$

# Outlook: WCPS 2

- Idea: merge WCPS with Xquery

    - Ex1: „difference of red, nir bands for all coverages on Austria"

    ```
    for $c in doc("http://acme.com")//coverage
    where
        some( $c.nir > 127 ) and metadata/@region = "Austria"
    return
        encode( $c.red - $c.nir, "image/tiff" )
    ```

    - Ex2: „name & location of coverages showing some phenomenon"

    ```
    for $c in doc("WCPS")//coverage/[ some( $c.nir > $c.red )]
    return
        <id> { $c/@id } </id>
        <area> { $c/boundedBy } </area>
    ```

- WCPS 2.0, in progress

    - Implementation: federation of eXist + rasdaman

[Jacobs University, Athena Research]

# 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 more data/metadata divide

- Visit us:

  - www.rasdaman.org

  - www.earthserver.eu