Responsible Land Governance:
Towards an Evidence Based Approach
ANNUAL WORLD BANK CONFERENCE ON LAND AND POVERTY
WASHINGTON DC, MARCH 20-24, 2017

Insights from Space and Time:
Big Data for Land Governance

Peter Baumann
Jacobs University | rasdaman GmbH
baumann@rasdaman.com

Paper prepared for presentation at the
"2017 WORLD BANK CONFERENCE ON LAND AND POVERTY"
The World Bank - Washington DC, March 20-24, 2017

Responsible Land Governance:
Towards an Evidence Based Approach
ANNUAL WORLD BANK CONFERENCE ON LAND AND POVERTY
WASHINGTON DC, MARCH 20-24, 2017

Abstract

With the continuous stream of satellite imagery weather data, etc. we are obtaining sufficient material for doing large-scale timeseries analytics on demand, useful for land management, land use and environmental monitoring, disaster mitigation, and many more purposes. However, the massive data involved call for novel techniques for achieving high flexibility and short response times while keeping service provider costs reasonable.

The datacube paradigm is a suitable abstraction for “mix and match” services on massive spatio-temporal data which do not bother users with zillions of files anymore. Array Databases constitute technology specifically tuned towards serving such datacubes through high-level query languages.

In our paper we present Array Databases as an innovative service technology for massive multi-dimensional grids (“arrays”). We exemplify this using the example of rasdaman which is in practical use on databases exceeding 250 TB. Finally, we discuss why Array Databases are superior to MapReduce-based systems like Spark.

Key Words:

rasdaman, array database, datacube, Big Data, EarthServer

1. INTRODUCTION

With the unprecedented increase of orbital sensor, in-situ measurement, and simulation data as well as their derived products there is an immense potential for getting new and timely insights - yet, the value is not fully leveraged as of today. Incidentally, such spatio-temporal sensor, image, simulation, and statistics data in practice typically constitute prime Big Data contributors. In view of such data "too big to transport" the quest is up for high-level service interfaces for dissecting datasets and rejoining them with other datasets - ultimately, to allow users to ask "any question, anytime, on any size" enabling them to "build their own product on the go".

OGC/ISO coverage standards have proven instrumental in unifying regular and irregular grids, point clouds, and meshes so that such data can be accessed and processed through a simple, yet flexible and interoperable service paradigm. The OGC Web Coverage Service (WCS) comprises a modular suite for accessing large coverage assets, up to OGC's Big Earth Data query language, Web Coverage Processing Service (WCPS).

In this contribution, we inspect the OGC and EU-INSPIRE reference implementation for WCS, rasdaman ("raster data manager"), which falls into the category of Array Databases meaning that its conceptual model and query language supports massive multi-dimensional arrays. Goal is to demonstrate how selected architecture concepts can particularly contribute to flexibility and scalability of large-scale raster services.

Through WCS/WCPS rasdaman is serving out 250 Terabyte as of today, heading towards the Petabyte frontier, and WCPS queries have been distributed across more than 1,000 cloud nodes.

2. RASDAMAN

Being an Array DBMS rasdaman ("raster data manager") offers an array query language with a formal algebraic semantics underneath[1]. The query language offers high-level, declarative operations which roughly speaking add signal/image processing and statistics operations capabilities to SQL.

Figure 1 shows the overall system architecture of rasdaman. The central workhorse, rasserver, operates in a multi-parallel fashion without a single point of failure (contrary to Hadoop, for example). Array partitions can sit in some database or directly in the file system. As an additional option, external pre-existing archives can be included into queries. In this case, arrays are built from the files found there, whereby conventions like file and directory naming are considered during incremental registration.

As a domain-agnostic array engine, rasdaman operates on bare arrays. Geo semantics is added through an additional layer which implements coverages and offers them through OGC WMS, WCS, and WCPS interfaces.

[1] Actually, its language has served as blueprint for the forthcoming ISO SQL/MDA (Multi-Dimensional Arrays) standard.

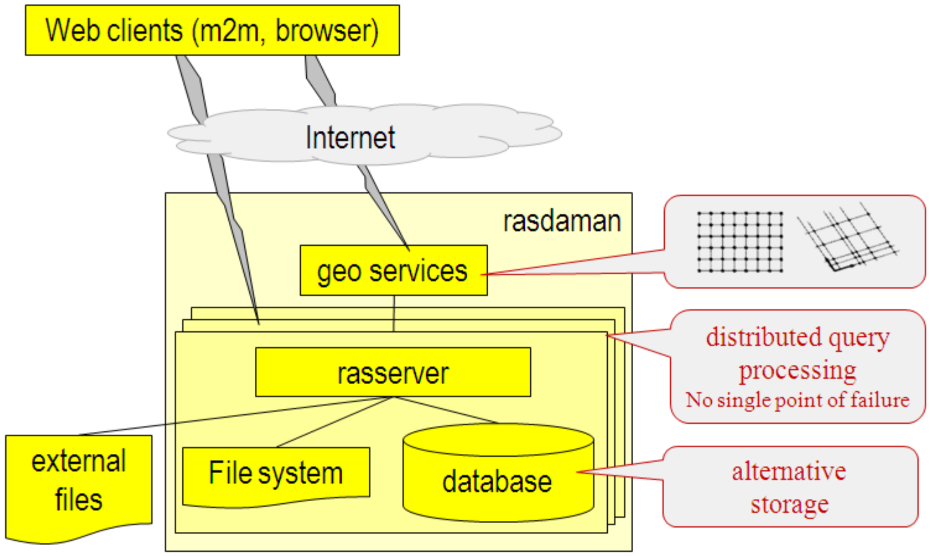

Figure 1. rasdaman federation architecture

In the server, arrays are partitioned into so-called tiles which form the unit of access. Partitioning can be streamlined to optimally support any given workload of access patterns (Figure 2). Under work is dynamic load-aware distribution of tiles.

Incoming queries undergo a series of optimizations and can be processed in distributed environments federating laptops, clouds, and data centers through query splitting techniques. Macro-splitting allows sending sub-queries to different nodes, thereby enabling not only cloud processing, but also data center federations (Figure 3). Single queries have successfully been split across more than 1,000 cloud nodes. In 2017, one rasdaman federation instance will get installed on board an ESA satellite. Micro-splitting utilizes heterogeneous hardware locals. Currently a new query engine is under development which compiles queries directly into machine code with query fragments sent to CPU cores, GPU, or other hardware based on a capacity-aware splitting.

Finally, caching query results helps to speed up requests which are completely or partially repetitive.

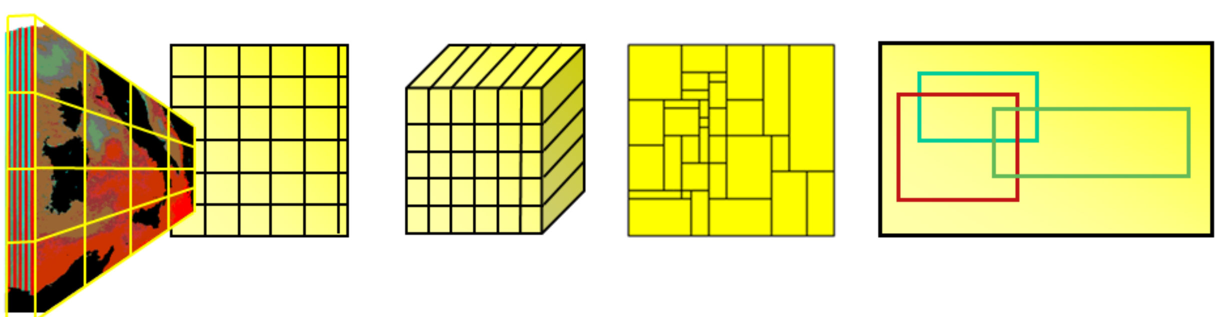

Figure 2. Adaptive partitioning of arrays in rasdaman

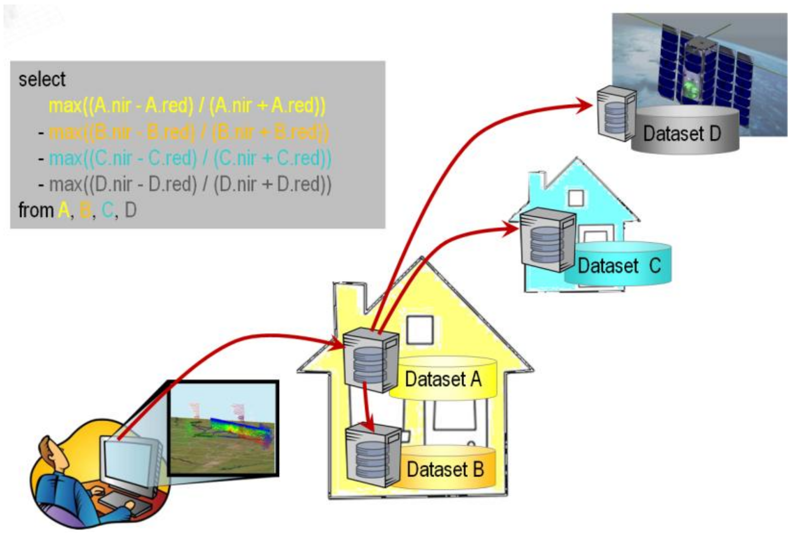

Figure 3. rasdaman query splitting and macro-distribution

3. DEPLOYED SERVICES

In the EarthServer initiative (www.earthserver.eu, Figure 4), large-scale data centers use rasdaman for establishing a planetary-scale datacube federation. The European Centre for Medium-Range Weather Forecast (ECMWF) maintains a climate data archive amounting to 120 PB. European Space Agency (ESA) Sentinel imagery stored at currently MEEO amounts to 250 TB. Plymouth Marine Laboratory (PML) has entered the federation with ocean colour analysis based on Landsat, MODIS, SEAWIFS, etc., and increasingly Sentinel. National Computational Infrastructure (NCI) Australia contributes reprocessed Landsat data, among others. Currently, datasets are being extended to cross the Petabyte barrier within the next months. In 2016, datacube fusion (i.e., array joins) have been shown involving a 4-D climate datacube hosted by ECMWF close to London, UK and a Landsat image datacube hosted by NCI at Canberra, Australia (Figure 9). Additionally, the PlanetServer data hub offers data on Mars, Moon, and soon Vesta and more planetary bodies.

According to independent reviewers, EarthServer will "significantly transform the way that scientists in different areas of Earth Science will be able to access and use data in a way that hitherto was not possible". They attested "proven evidence" that rasdaman will “significantly transform [how to] access and use data“.

Importantly, datacube languages like WCPS and Array SQL are not envisioned as user interfaces. Similar to WCS, they rather constitute a convenient machine-to-machine interface between servers and client frontends like visual globes or desktop GISs (Figure 5and 6), but in particular in the middle of automated processing chains where there is no human interaction and control. In particular the latter use case is an important scenario for the “Big Data” services under discussion here. Generally speaking, array interfaces should at best be visible to experts as “power tools” – but even then well-known environments like python and R offer, once coming with a tightly integrated rasdaman backend, a more convenient experience as experts do not need to leave their comfort zone of tool skills.

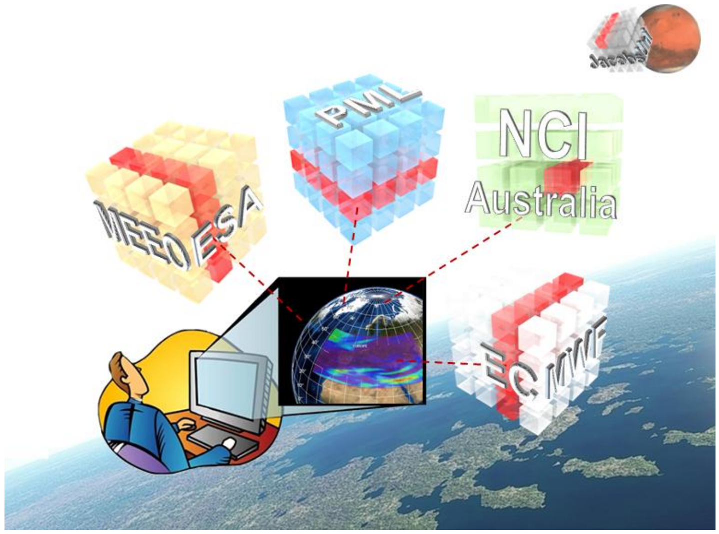

Figure 4. Intercontinental datacube mix and match in the EarthServer initiative, plus the PlanetServer service on Mars and Moon (source: EarthServer)

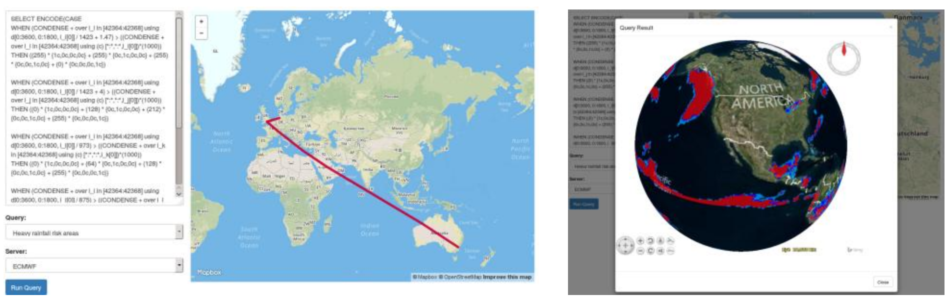

Figure 5. Datacube fusion between ECMWF / UK and NCI Australia: query distribution path (left) and result visualization (right) (source: EarthServer)

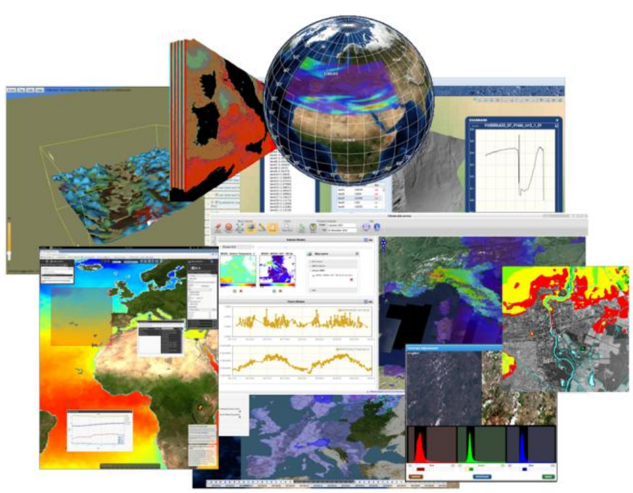

Figure 6. Screenshots of domain and task specific client frontends accessing rasdaman via WCPS (source: EarthServer)

4. ARRAY DATABASES VERSUS MAPREDUCE

The MapReduce principle has been developed mainly for data whose elements are unconnected ("un-structured") or where a regrouping independently from the original structure is to be developed – in short: MapReduce benefits from set-oriented structures. Arrays, however, are extremely structured with a well-defined, rigorous neighbourhood between cells. Figure 7 illustrates this relationship which is based on the Euclidean structuring of the grid space. This structure naturally induces basic operations of trimming and slicing a datacube.

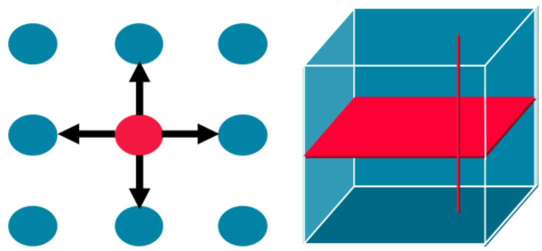

Figure 7. Euclidean neighbourhood of grid cells, naturally leading to trim&slice operations

Array databases are aware of this grid structure and support them through their complete architecture, from spatially clustered storage to array-aware processing. The result is higher efficiency in array query evaluation; in comparison between rasdaman, Spark, and Hive this has been observed by researchers not affiliated with any of the corresponding developer groups (Figure 8).

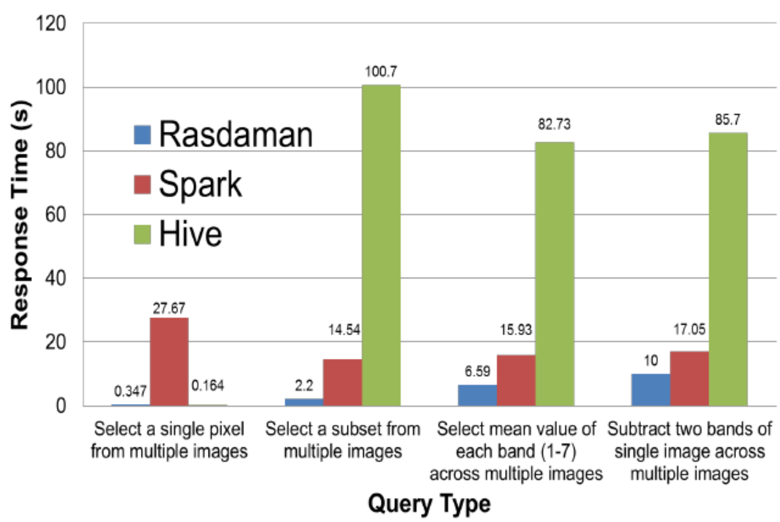

Figure 8. Independent performance comparison between rasdaman, Spark, and Hive
[C. Scheele, F. Hu, M. Yu, M. Xu, K. Liu, Q. Huang, C. Yang 2015]

5. CONCLUSION

With the OGC/ISO coverage standards, a suitable abstraction for all kind of gridded data (and beyond) has been established, such as satellite imagery, image timeseries, elevation data, and weather data. OGC

Coverage data and services have seen massive uptake – the major open source as well as proprietary tools today support it, including MapServer, GeoServer, QGIS, and ESRI. WCS-based coverage services are in use on large scale deployments of 250 TB and growing. The standards observing group of the US Federal Geographic Data Committee (FGDC) sees coverage processing a la WCS/WCPS as a future "mandatory standard".

In this contribution, we presented an implementation architecture based on Array Databases for flexible, scalable services on massive spatio-temporal grids. Based on a suitable modularization, effective optimizations can be found; we presented three of them: adaptive data partitioning, query splitting with distributed processing, and query result caching.

Altogether, coverage data and service models represent a major step forward towards a common information space for "Big Geo Data" enabling unified, interoperable access, processing, and fusion of spatio-temporal data.

6. ACKNOWLEDGEMENT

EarthServer is co-funded by EU H2020 EINFRA project no. 654367.