RDA PROPOSAL FOR

Array Database Assessment Working Group (ADA:WG)

Peter Baumann, Jacobs University[1]

Co-Chairs

Initial chairs will be: **Peter Baumann** (Jacobs University, Germany) and **Kwo-Sen Kuo** (Bayesics, US). Replacement of chairs, or extending the chairs list, is possible anytime through democratic voting.

Background

Multi-dimensional data arrays play a core role in many, if not most science and engineering domains where they typically represent spatio-temporal sensor, image, simulation output, or statistics data. The research field of Array Databases has emerged in the attempt to augment the traditional set-driven paradigm with modeling and query support for large, n-D arrays. Such systems attempt to combine the best functionality and performance of different worlds: the long-standing experience of array handling in the sciences, the flexibility of database query languages, and the parallelization and scalability methods developed in HPC, HPD and Cloud, as well as using new hardware. Various implementations are known [1][4][2][5][6][10][14], deployments of Array Databases today are in the hundreds of Terabytes [13], and single queries reportedly have been parallelized over more than 1,000 nodes [3]. Therefore, Array Databases have to be considered as a serious option for Big Data management in science, engineering and beyond.

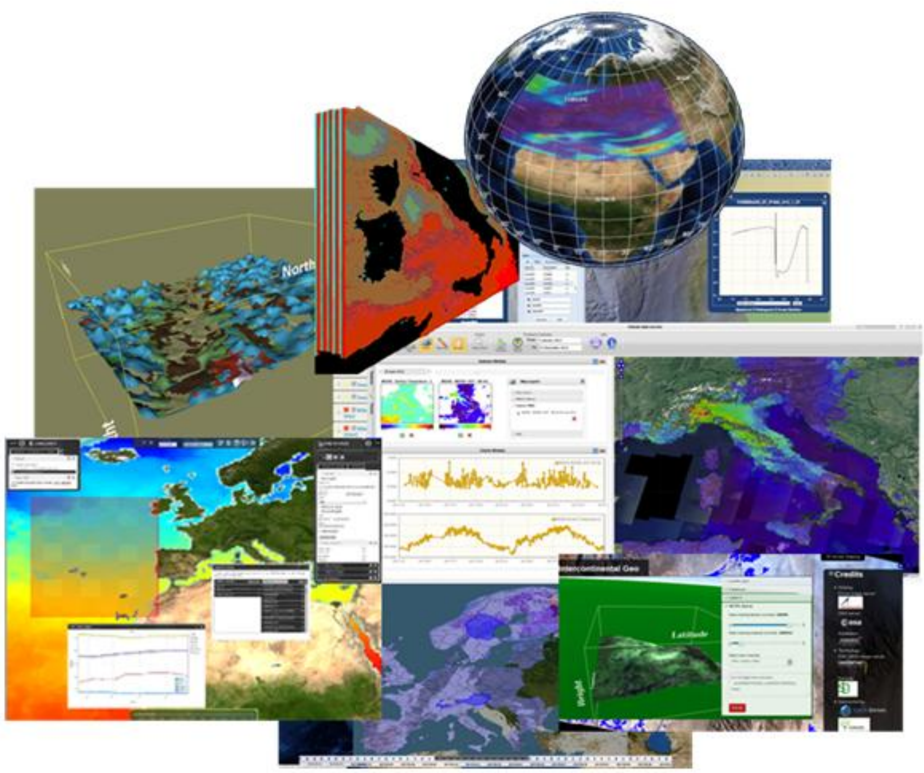

Fig. 1: Kaleidoscope of portals utilizing an Array Database today

[1] Several colleagues contributed important facets to this proposal; aside from the TAB reviewers, Peter Wittenburg, Ben Evans, and several more.

1. WG Charter

The Array Database WG will inspect the emerging technology of Array Databases to provide **support for technologists and decision makers** considering Big Data services in academic and industrial environments (such as in large-scale data centers) by establishing best-practice guidelines on how to optimally **serve multi-dimensional gridded Big Data through Array Databases**. This will be accomplished through a neutral, thorough **hands-on evaluation** assessing available Array Database systems and comparable technology …

- based on relevant **standards**, such as the NIST Big Data Reference Architecture [2], ISO "Array SQL" [8], and OGC Web Coverage Processing Service (WCPS) [9] for the geo domain;
- comparing **technical criteria** like functionality, thereby eliciting the state of the art;
- establishing and running a combination of domain-driven and domain-neutral **benchmarks** that will be run on each platform;
- as well as real-life, **publicly accessible deployments** at scale.

The result, consisting of the AD-WG **report** together with the **open-source benchmarking software** and the **services** established, will establish a **hitherto non-existing overview** on the state of the art and best use of Array Databases in science, engineering, and beyond.

2. Value Proposition

Open-source and proprietary Array Databases are now becoming more readily available to the research community, and they promise to provide a new way of managing and analyzing regularly and irregularly gridded data in Earth, Space, and Life sciences, and beyond. For example, in the Earth sciences we find 1-D sensor time-series, 2-D satellite imagery, 3-D x/y/t image time-series and x/y/z geophysical voxel cubes, as well as 4-D x/y/z/t climate simulations. Similar data are used within industry, such as oil and gas exploration, insurance, pharmaceutical, automotive, shipping, aerospace industries, and may provide increased uptake of such technology.

The analysis of existing technology will focus on one question: **to what extent can data scientists and engineers benefit from Array Database technology?** To this end, the Working Group will address at least the following questions (greater detail to be determined during the first phase, i.e., until Milestone 1):

- **What systems are out there?** (currently considered are: rasdaman [1][4], SciQL [10], SciDB [5], PostGIS Raster [6], Ophidia[14], possibly Extascid [2]; further, a comparison with alternative technology is foreseen, such as SciHadoop [11]), HDF5 on filesystems [12]) and Python scientific data formats.
- **What are their features?** Currently recognized aspects include: query language power; scientific tool integration; adaptive data partitioning and distribution support, including alignment for data fusion; parallelization in clouds and data center federations; support for modern hardware; etc. – list to be completed by Milestone 1.
- **What is their performance, measured through objective, open tests?** To this end, a combination of domain-specific and domain-neutral ("kernel") benchmarks will be developed, agreed in the WG, and applied to the systems.
- **What are relevant tuning parameters, and what are the recommended settings for these systems?**
- **How can these systems be used in large-scale deployments?** This will be answered by setting up publicly accessible installations[2] (utilizing pre-existing installations where possible, to minimize

[2] This will depend on the WG member resources. It is expected that each entity participating will pick and run at least one system. Therefore, further members will be chosen so as to achieve a broad range of such installations. (Note that this applies only to the large-scale deployment – all other investigations will be done on several systems in any case.)

effort) and evaluating them from the perspective of IT, service operator, and user. A list of known installations will be compiled.

- **Generally, what are the strengths and weaknesses of Array Databases for Big Data in science & engineering?**

In the final report, the findings will be consolidated into a state-of-the-art report, accompanied by the benchmarks, as well as guidance to service operators on using an Array Database system. Any code developed will made available as open-source along with documentation and samples. It is hoped that this will be an important decision aid for science and industry alike.

3. Engagement with existing work in the area

Engaging and collaborating with other related existing works are important to our mission. This will include interaction with consortiums, alliances, and standards bodies. The knowledge gained is expected to enhance our development by aligning with high quality best practices. These include at a minimum:

- ISO: SC32/WG3 "SQL" [8]; ISO TC211/WG6 "Geo Imagery"; WG9 "Big Data" [7]
- OGC [9]
- EU INSPIRE

The list will be revisited and likely amended until Milestone 1 (see below).

4. Work Plan

- **Milestones**

milestone	month	outcome
M1	3	Evaluation criteria established, candidate systems chosen, external engagements defined
M2	6	Testbeds established
M3	9	Evaluations done
M4	12	Final report: results and recommendations

- **Mode and Frequency of Operation**

  It is anticipated that all discussions will be conducted via emails plus bi-weekly teleconferencing through a Web-based conferencing tool. The date/time of the bi-weekly telecon will be determined by Doodle Poll.

- **Membership**

  Participation in the WG is open to all interested parties. There are no membership fees.

- **Coordination/Interaction**

  The WG will function in close coordination with other Big Data related standards and best practices from industry, academia and government at the international level.

- **Standing Rules**

  All information exchanged within the WG will be freely available (using CC-BY license).
  All information exchanged within the WG will contain non-IPR materials.
  WG members should assume that all materials exchanged will be made public.
  Documents will be publicly accessible from the WG portal.

5. Adoption Plan

Following completion of the report (and likely at intermediate milestones, too) the AD-WG will perform outreach activities to promote insights into stakeholder communities. This includes standardization meetings (such as OGC, ISO, INSPIRE) as well as conferences (such as database and geospatial conferences). The outreach portfolio will be elaborated in the WG and will evolve over its lifetime.

6. Initial Membership[3]

Currently confirmed members
(note the balanced participation of academia, supercomputing centres, and industry SMEs):

First Name	Last Name	Organization	Country
Peter	Baumann	Jacobs University	Germany
Dimitar	Misev	rasdaman GmbH	Germany
Morris	Riedel	Juelich Supercomputing Centre	Germany
Oliver	Clements	Plymouth Marine Laboratory	UK
Mike	Grant	Plymouth Marine Laboratory	UK
Stephan	Siemen	European Centre for Medium-Range Weather Forecasts	UK
Julia	Wagemann	European Centre for Medium-Range Weather Forecasts	UK
Simone	Mantovani	MEEO s.r.l.	Italy
George	Kakaletris	CITE s.a.	Greece
Panagiota	Koltsida	ATHENA Research Centre	Greece
Patrick	Hogan	NASA Ames	US
Ben	Evans	National Computational Infrastructure	Australia
Joseph	Antony	National Computational Infrastructure	Australia
Sandro	Fiore	Euro Mediterranean Center on Climate Change (CMCC)	Italy
Kwo-Sen	Kuo	Bayesics. Inc.	US

7. References

[1] Peter Baumann: On the Management of Multidimensional Discrete Data. VLDB Journal 4(3)1994, Special Issue on Spatial Database Systems

[2] Y. Cheng, F. Rusu. *EXTASCID: An extensible system for the analysis of scientific data*. Poster XLDB, Standford, California, USA, September 10 – 13, 2012

[3] A. Dumitru, V. Merticariu, P. Baumann: Exploring Cloud Opportunities from an Array Database Perspective. Proc ACM SIGMOD Workshop on Data analytics in the Cloud (DanaC'2014), June 22 - 27, 2014, Snowbird, USA

[4] rasdaman. www.rasdaman.org

[5] SciDB. www.scidb.org

[6] R. O. Obe, L. S. Hsu: *PostGIS in Action*. Manning Publications Co., 2011

[7] NIST Big Data Reference Architecture, M0266

[8] 9075 SQL Part 15: MDA - Multi-Dimensional Arrays (under preparation)

[3] will invite further members once WG has initial agreement from RDA

[9] P. Baumann (ed.): OGC Web Coverage Processing Service (WCPS) Implementation Specification. OGC document 08-068

[10] SciQL. www.sciql.org

[11] J.B. Buck, N. Watkins, J. LeFevre, K. Ioannidou, C. Maltzahn, N. Polyzotis, S. Brandt. *SciHadoop: Array-based Query Processing in Hadoop*. Proc. 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC), pp. 66:1–66:11

[12] HDF5. www.hdfgroup.org

[13] P. Baumann, P. Mazzetti, J. Ungar, R. Barbera, D. Barboni, A. Beccati, L. Bigagli, E. Boldrini, R. Bruno, A. Calanducci, P. Campalani, O. Clement, A. Dumitru, M. Grant, P. Herzig, G. Kakaletris, J. Laxton, P. Koltsida, K. Lipskoch, A.R. Mahdiraji, S. Mantovani, V. Merticariu, A. Messina, D. Misev, S. Natali, S. Nativi, J. Oosthoek, J. Passmore, M. Pappalardo, A.P. Rossi, F. Rundo, M. Sen, V. Sorbera, D. Sullivan, M. Torrisi, L. Trovato, M.G. Veratelli, S. Wagner: Big Data Analytics for Earth Sciences: the EarthServer Approach. International Journal of Digital Earth, 0(0)2015

[14] Ophidia. http://ophidia.cmcc.it/

Annex: Responses to public comment period

The following comments below have been posted. Those excerpts containing a “do message” have been extracted, and are commented on.

Suggestion	Comment
deepseadawn: • I wonder if this WG will be considering unstructured (big) data as well?	• The WG, by its definition, will be looking at data that can be represented as arrays (i.e., gridded rasters). Therefore, general unstructured data (such as text) is out of scope. That said, many people in the past have considered array data (when coming as images or data) have been (wrongly) considered as “unstructured”. In this sense, such data are under consideration, too.
Peter Wittenburg: • benefits (performance gain, management easiness, etc.) should be compared with the costs (investments in time and money) • Comparisons are essential, the question is comparison between what? If I just compare between different database concepts not so much is gained, what people in general using is some form of sliding window across files. A comparison against optimal procedures of traditional type would be excellent. No idea how this can be done, but ... Here you need to be a bit more specific I guess and of course examples from different communities would be great. • adoption plan is not yet satisfying. It is important to know who is going to test things etc. I think that it is important to mention scientific communities that have an interest and will participate in this. In the membership list you have quite a number of experts partly engaged in communities, but it is not clear what their role in adoptions, testing etc is. • should start with two co-chairs which can be temporary and be replaced by new ones later.	• For sure there will be experience provided on ease/difficulty of installation and operation. Price lists etc. cannot be provided (interested parties should contact vendors for conditions). • See above – non-database concepts are included, such as SciHadoop. Generally, the more activists we find the more systems and facets can be addressed. Examples from different communities are planned, but again: we need stakeholders from domains in order to investigate them. Of course, the more specific you suggest items the better we can try planning them in. For example, we will be most happy if you can provide more domain stake-holders to us in addition to those on board! • Who is going to test: this will be the WG members. Detailed work distribution to be done as a first workpackage of the group (see M1). Of course, it critically depends on the level of commitment of partners. • Good point, will take that.