THE ALDEVCON 2018



TILEDB: ARRAY DATA NEEDS AN ARRAY STORAGE Manager

Tim Mattson May 24, 2018



ERAS OF DATABASE TECHNOLOGY

SQL Era

Common interface









SQL = Structured Query Language

Third party names are the property of their owners

Source: The BigDAWG Polystore System and Architecture, HPEC'2016, Vijay Gadepally



DATA IN THE REAL WORLD MESSY, HETEROGENEOUS, COMPLEX, STREAMING ...

• Consider patient data in an Intensive Care Unit (e.g. MIMIC II data set*)



Time series and tabular data are stored in a DBMS. Other data? Flat files

* MIMIC: Multiparameter Intelligent Monitoring in Intensive Care, http://www.physionet.org/mimic2/

Analysis of published MIMICII papers



We must Bring the power of a DBMS to all data

DBMS: database management systems

*Based on PhysioNet MIMIC2 ICU data

Source: Vijay Gadepally of MIT Lincoln labs, 2015

SHOULD WE CRAM ALL THE DATA INTO ONE DBMS? No!!! One size does not fit all*

Typical DB Operations



*Stonebraker, Michael, and Ugur Cetintemel. "" One size fits all": an idea whose time has come and gone." *21st International Conference on Data Engineering (ICDE'05).* IEEE, 2005.



ERAS OF DATABASE TECHNOLOGY



THE BIGDAWG POLYSTORE SYSTEM

- BigDAWG
 - Polystore: match data to the storage engine
- BigDAWG Islands
 - A data model + query operations
 - One or more storage engines
 - "Shim" connects a BigDAWG island to a data engine
 - "Cast" migrates data from one storage engine to another



W Polystores in the Cloud @Myria

Myria: a Polystore system from the University of Washington ... for on-premise clusters or in the cloud!

- A python-like, high level query language: MyriaL
- RACO: Relational Algebra COmpiler. Query optimizer/executor that targets multiple back-ends.
- A data-flow, execution engine for relational-algebras with iteration: MyriaX



Source: Daniel Halperin, Victor Teixeira de Almeida, Lee Lee Choo, Shumo Chu, Paraschos Koutris, Dominik Moritz, Jennifer Ortiz, Vaspol Ruamviboonsuk, Jingjing Wang, Andrew Whitaker, Shengliang Xu, Magdalena Balazinska, Bill Howe, and Dan Suciu. Database Group & eScience Institute, University of Washington

http://myria.cs.washington.edu

Third Party Names are the Property of their Owners.

Analytics to predict "fatal" cardiac events



- Goal: Find patients with similar ECG time-series*
 - Perform Discrete Wavelet Transform of ECG
 - Generate wavelet coefficient histogram
 - TF-IDF waveform coefficients (weight rare changes higher)
 - Cluster and correlate against other ECGs

TF-IDF=Term Frequency-Inverse Document Frequency ECG = Electrocardiogram



POLYSTORE ANALYTICS PERFORMANCE Medical informatics ECG Waveform/Clustering/Cardiac-event study



POLYSTORE ANALYTICS PERFORMANCE

Medical informatics ECG Waveform/Clustering/Cardiac-event study



CONCLUSION: THERE IS GREAT VALUE IN SPECIALIZING THE STORAGE ENGINE TO THE DATA





TYPES OF DATA AND THEIR STORAGE ENGINES

[Student ID]	[First Name]	[Last Name]	[State]
[1534]	[Robert]	[Paulson]	[Dead]
[1324]	[Honey]	[Bunny]	[Mushy]
[1134]	[Beatrix]	[Kiddo]	[Angry]



Key value pairs

Relation tables

[<id001> performer: Nancy Sinatra] [<id002> song: Woo Hoo] [<id002> performer: The 5,6,7,8's] [<id001> song: Bang Bang]



Arrays





ARRAYS IN BIG DATA PROBLEMS: Geospatial data with (latitude, longitude) coordinates

Automatic ID system (AIS) data set showing ship locations as a function of time in and around U.S. waters



Arrays in Big Data problems: Graphs can be represented as a 2D sparse array



from 7 to 7

AT

A = the adjacency matrix ... Elements denote edges between vertices



ARRAY DATA SETS

What do these two examples have in common?



The arrays are sparse. Existing array storage engines do OK with dense arrays, but sparse data is challenging for them.

TILEDB A NEW ARRAY DATA STORAGE MANAGER: Optimized for sparse arrays

Logical representation



Physical representation



Manage array storage as tiles of different shape/size in the index space, but with ~equal number of non-empty cells



ARRAYS IN BIG DATA PROBLEMS: Geospatial data with (latitude, longitude) coordinates

Automatic ID system (AIS) data set showing ship locations as a function of time in and around U.S. waters





TileDB Concepts & Features **CELL ORDER**



Although data are **multi-dimensional** ...



data storage is **single-dimensional**

 \checkmark We need to define a **cell order**



Tile DB Concepts & Features **CELL ORDER & TILING**



<u>space tile extents:</u> 2x2 <u>tile order:</u> column-major <u>cell order:</u> row-major





TileDB Concepts & Features **CELL ORDER & TILING**



space tile extents: 2x2 tile order: column-major cell order: row-major capacity: 2





SPARSE DATA: LOAD PERFORMANCE



Third Party Names are the Property of their Owners.

Source: Papadopoulos, Datta, Madden, Mattson, VLDB 2017

SPARSE DATA: SUBARRAY QUERIES

- Select a sparse region and a sparse sub-array of size "Result size".
- Randomly shift the subarray 50 time.
- Return cell coordinates that fall inside the subarray





Third Party Names are the Property of their Owners.

TileDB Concepts & Features **UPDATES**



3

212

х

4

213

yy





DENSE ARRAY: RANDOM UPDATES TO 4 GB ARRAY



Third Party Names are the Property of their Owners.

Source: Papadopoulos, Datta, Madden, Mattson, VLDB 2017

Tile DB Concepts & Features **PHYSICAL ORGANIZATION** Logical View





File organization





Tile DB Concepts & Features **PHYSICAL ORGANIZATION** Logical View

 $\mathbf{2}$

3

4



File organization





TILEDB VIRTUAL FILE SYSTEM

The virtual filesystem exposes basic filesystem operations through an API embedded in the TileDB C interface

Lets users write portable code that moves between different file systems





Applications **GENOMICS:** THE DATA



... ATTCGGAGCATCGGAAATT ... Reference

··· AGTCGGATCATCGGAAATA ···

- ... TTTGAGATCATCGGAACTT ... Samples
- $\cdots \quad \mathsf{TTTCAGATCATCGGAAAGT} \cdots$



Applications **GENOMICS:** THE DATA







We only store locations where samples differ from the reference genome



Applications **GENOMICS:** MAPPING ONTO TILEDB

Z

Genome Positions (~3 Billion)



col-major order



Source: The TileDB Array Data Storage Manager, Stavros Papadopoulos, Kushal Datta, Samuel Madden, Tim Mattson, VLDB 2017

3

Applications **GENOMICS: GATK 4***



- GenomicsDB: A data management system built on top of TileDB
- GenomicsDB was incorporated into GATK 4 from the Broad Genomics Institute.
 Combine and Merge whole genomes for processing (units of 1K genomes)



APPLICATION: VISUAL ANALYTICS: TDB, A LOSSLESS IMAGE FORMAT FOR VISUAL ANALYTICS Christina Strong, Ian Adams, Vishakha Gupta-Cledat and Luis Remis of Intel Labs and the Visual Cloud Istc

Store image pixels as a dense array in TileDB



- Analytics integrated with storage. Example, load & crop:
 - TileDB loads only the tiles you need to crop.
 - Other formats must load the full image before the cropping can be done.



Raw: Work from the OpenCV array of pixels ... no compression or metadata so this is not practical. Just a performance reference point

The Future of TileDB

- TileDB: just a library so easy integration with your code.
- Uses a non-viral open source license (MIT License).



• With TileDB, you get the best of both worlds:

- A growing open source community plus commercial support when you need it. TileDB Inc website: https://tiledb.io



CONCLUSION

- AI is of little value without lots of data.
- Think about how you should manage your data ... flat files will hurt you sooner or later.
 - -You need a real data management system.
- Data storage engines should fit the data:
 One size does not fit all. ← Polystores are the next big thing!!
- If you have Array data (and you probably do) give TileDB a try



GET THE CODE: WE LOVE OPEN SOURCE CODE AT INTEL

- Polystore Data management systems
 - -BigDAWG (for DBMS researchers):
 - -Myria (ready-to-use tool for data scientists): https://myria.cs.washington.edu/
- TileDB storage engine for array data
 - -Intel Health/Life-sci repo (with genomicsDB): https://github.com/Intel-HLS/TileDB

https://bigdawg.mit.edu/

https://github.com/IntelLabs/vdms

https://github.com/IntelLabs/pmgd

- TileDB Inc. repo (latest TileDB developments): https:// github.com/TileDB-Inc/TileDB
- Visual Data Management System and associated tools
 - -VDMS (Manage visual data):
 - -PMGD (graph DBMS for metadata):
 - -VCL(visual data storage engine with TileDB): https://github.com/IntelLabs/vcl



WIN DEDENS AWS DEEPLENS AIDEVCONAPP.INTEL.COM

aws

intel

Completing a session evaluation in the mobile app by 10:00 a.m. tomorrow automatically enters you in a drawing to win.

Copies of the complete sweepstakes rules are available at the Concierge Desks.

