

A Look at Data Management Systems for Climate Research

James J. Pan (jamesjpan@outlook.com)

Database Group @ Tsinghua University



September 2, 2020

Rise of Machine Learning for Spatial...



NATURE | NEWS



How machine learning could help to improve climate forecasts

Mixing artificial intelligence with climate science helps researchers to identify previously unknown atmospheric processes and rank climate models.

Nicola Jones

23 August 2017

“If you go to major modeling centers and ask them how they work, the answer won’t be machine learning,” says Collins. **“But it will get there.”**¹

¹Nicola Jones. “How machine learning could help to improve climate forecasts”. In: *Nature* 548.7668 (2017), pp. 379–380.

ACM Special Interest Group on Knowledge Discovery and Data Mining (**SIGKDD**)

≡



APPLIED DATA SCIENCE INVITED TALKS



Professor Vipin Kumar

Professor
University of Minnesota

Big Data in Climate: Opportunities and Challenges for Machine Learning

Wednesday 10:00am - 12:00pm, Room 200D

[More Information](#)

*“We discuss the challenges involved in analyzing these **massive data sets** as well as opportunities they present for both advancing machine learning as well as the science of climate change...”²*

²Anuj Karpatne and Vipin Kumar. “Big Data in Climate: Opportunities and Challenges for Machine Learning”. In: *KDD (2017)*, pp. 21–22.

Very Large Databases (VLDB)

VLDB 2017 TUM



43rd International Conference on Very Large Data Bases

Tutorials

- The Era of Big Spatial Data

Ahmed Eldawy (UC Riverside) (eldawy@cs.ucr.edu)

Mohamed Mokbel (University of Minnesota) (mokbel@cs.umn.edu)

[Slides](#)

*"In this tutorial, we present the recent work in the database community for handling **Big Spatial Data**..."*³

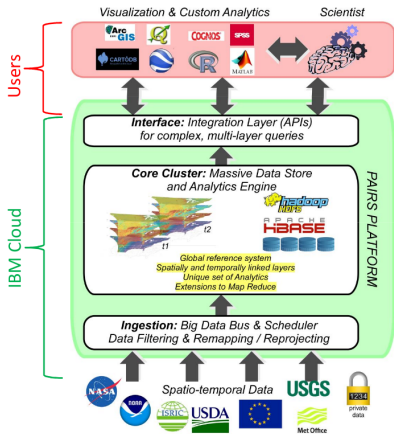
³Ahmed Eldawy and Mohamed Mokbel. "The Era of Big Spatial Data". In: *VLDB (2017)*.

Why is everyone so excited?

Meteorology in the eyes of data scientists:

- ▶ Terabytes of Earth imagery data get generated *per day*
 - ▶ Static analysis “macroscale” (e.g. deforestation, land-use, urban expansion, ...)
 - ▶ Real-time analysis “microscale” (e.g. flood monitoring, wildfires, landslides, tornados, ...)
 - ▶ Somewhere in between “mesoscale” (e.g. soil moisture, crop production, water availability, ...)
- ▶ Discovery and Monitoring: mine data to discover new climate relationships (e.g. teleconnections, tripoles), monitor real-time interactions (e.g. wildfires, floods, effects on rainfall...)
- ▶ Model Improvement: refine climate models to improve predictive and reproductive power
- ▶ **Data Management**: ingest, store, query, ... to efficiently support data applications

Example: IBM Physical Analytics Integrated Repository and Services (PAIRS)⁴



Examples



- ▶ E.g. “Give me two years of temperature in Kyoto, plus 2-week forecast”

⁴Siyuan Lu. “IBM PAIRS - A Big Physical Data Service to Accelerate Analytics and Discovery”. In: (2017).

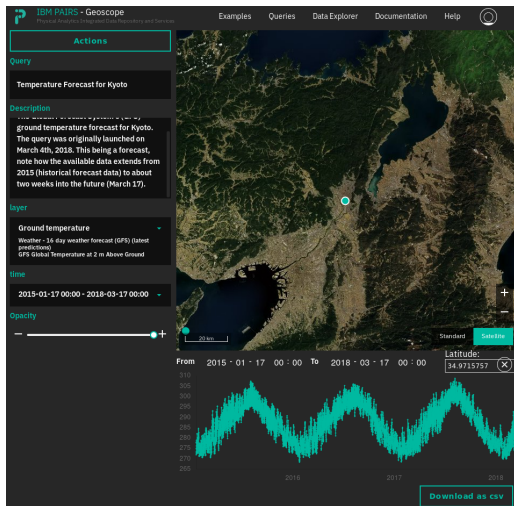


Figure: The IBM Pairs web interface.

- ▶ Good for simple retrievals; limited processing capability
- ▶ No real-time monitoring
- ▶ Limited modeling capability (limited forecasting ability)

Example: ESRI ArcGIS⁵

Policymaker: “How is land use changing in my district?”




Figure 4. Training Samples Manager and Sample Labels

Figure: Train a convolution neural network based on expert labels, then feed new images per pixel and collect the predicted land use.



Figure 5. Raw NAIP Image versus Classified Image

Class	Precision	Recall
Buildings	82.50	81.28
Roads	84.78	85.13
Waters	86.14	85.55
Harvested	90.38	91.88
Planted	89.05	88.19
Forest	91.46	92.65

Table 1. Accuracy Assessment of U-Net Model (Precision and Recall in %)

⁵Amin Tayyebi. “High-Resolution Land Cover Mapping using Deep Learning.” *ISPRS International Journal of Geo-Information* (2016)

Overview

Introduction

Background: Filesystems and Data Management

Management Systems for Spatial Data

- Array-Based: SciDB and friends

- Spark-Based: GeoTrellis and ClimateSpark

Scalable Machine Learning

Summary

Background: Filesystems and Data Management

Why use data management software (databases) anyway? Isn't the filesystem good enough?




Figure: Magnetic hard drives store bits (0's and 1's) inside magnetic fields. A chunk of data (typically 512 bytes) is called a *sector* and is the minimum size the drive can write.

- ▶ Reading a particular “file” involves spinning the disk to visit sectors containing the file contents.
- ▶ Sometimes sectors belonging to one file are not physically near each other on the disk, known as “fragmentation”.

Comparison of Read Latencies

- ▶ **Magnetic disk** at 15,000 RPM: up to 4 milliseconds (4,000,000 ns)
- ▶ **Solid-state drive (NAND)**: 50,000–100,000 ns
- ▶ **RAM**: about 60 ns
- ▶ CPU L2 cache: about 10 ns
- ▶ CPU access: about 1 ns

Reading from disk is about 10^6 times slower than reading from RAM. In other words, relative difference between **1 second** and **11 days**.

Filesystems and Data Management

Both filesystems and databases deal with storage and I/O:

- ▶ Filesystems: ext4 (Unix), NTFS (Windows), APFS (Apple)
- ▶ Fairly uniform use case: fast file access, security, data integrity, compression, paging, volume resizing, locking
- ▶ E.g. journaling systems such as NTFS keep a master log file to support data restore in case of crashes
- ▶ Databases: Oracle, SQL Server, Postgres, Teradata, Vertica, MongoDB, MonetDB, kdb+ ...
- ▶ **Data contents are related**, e.g. “find total revenue of the Rome store over all Mondays of last year”
- ▶ **“No one size fits all”**: different approaches for different use cases

No One Size Fits All: Row-Store vs. Column-Store

Consider the I/O cost of reading/writing this table:

<i>Item</i>	<i>Category</i>	<i>Revenue</i>
<i>Glove</i>	<i>Sport</i>	500
<i>Cap</i>	<i>Sport</i>	200
<i>Chair</i>	<i>Housing</i>	450
<i>Table</i>	<i>Housing</i>	100
<i>Shoe</i>	<i>Sport</i>	600

Row-store (e.g. H-Store⁶): {Glove, Sport, 500, Cap, Sport, 200, Chair, Housing, ...}

Column-store (e.g. MonetDB⁷): {Glove, Cap, Chair, ..., Sport, Sport, Housing, ..., 500, 200, 450, ...}

⁶Robert Kallman et al. "H-store: a high-performance, distributed main memory transaction processing system". In: *VLDB 1.2 (2008)*, pp. 1496–1499.

⁷Stratos Idreos et al. "MonetDB: Two Decades of Research in Column-oriented Database Architectures". In: *IEEE Data Eng. Bull.* 35.1 (2012), pp. 40–45.

Management Systems for Spatial Data: Array Databases

Consider the cost of the blue subarray query:

Relational Database

<i>i</i>	<i>j</i>	<i>value</i>
0	0	32.5
1	0	90.9
2	0	42.1
3	0	96.7
0	1	46.3
1	1	35.4
2	1	35.7
3	1	41.3
0	2	81.7
1	2	35.9
2	2	35.3
3	2	89.9
0	3	53.6
1	3	86.3
2	3	45.9
3	3	27.6

Array Database

32.5	46.3	81.7	53.6
90.9	35.4	35.9	86.3
42.1	35.7	35.3	45.9
96.7	41.3	89.9	27.6

Traditional RDBMS:

- ▶ Basic types (integer, float, string, ...)
- ▶ Relational operators (select, join, group by, ...)
- ▶ Good for queries such as “What are the phone numbers for all employees in Delaware?”

<i>ID</i>	<i>Name</i>	<i>Location</i>	<i>PhoneNumber</i>
1	<i>JohnSmith</i>	<i>Delaware</i>	123 – 456 – 7890
2	<i>JaneDoe</i>	<i>Delaware</i>	123 – 456 – 7891
...

Array-Based DBMS:

- ▶ **Array** data model
- ▶ Array operators (structural operators, content operators)
- ▶ Good for multidimensional queries such as point time-range query, spatial aggregation query

2	5	4
2	1	8
...

⁸The SciDB Development Team. “Overview of SciDB”. In: *SIGMOD* (2010).

SciDB: Chunk Storage

$J \setminus I$	[0]	[1]	[2]	[3]	[4]
[0]	(2, 0.7)	(5, 0.5)	(4, 0.9)	(2, 0.8)	(1, 0.2)
[1]	(5, 0.5)	(3, 0.5)	(5, 0.9)	(5, 0.5)	(5, 0.5)
[2]	(4, 0.3)	(6, 0.1)	(6, 0.5)	(2, 0.1)	(7, 0.4)
[3]	(4, 0.25)	(6, 0.45)	(6, 0.3)	(1, 0.1)	(0, 0.3)
[4]	(6, 0.5)	(1, 0.6)	(5, 0.5)	(2, 0.15)	(2, 0.4)

Step 1: Vertically partition *attributes* in the *logical array*.

$J \setminus I$	{A}					{B}				
[0]	2	5	4	2	1	0.7	0.5	0.9	0.8	0.2
[1]	5	3	5	5	5	0.5	0.5	0.9	0.5	0.5
[2]	4	6	6	2	7	0.3	0.1	0.5	0.1	0.4
[3]	4	6	6	1	0	0.25	0.45	0.3	0.1	0.3
[4]	6	1	5	2	2	0.5	0.6	0.5	0.15	0.4

Step 2: Decompose each attribute array into equal sized, and potentially overlapping, *chunks*.

$J \setminus I$	{A ₁ }			{A ₂ }			{A ₃ }			{A ₄ }		
[0]	2	5	4	4	2	1	4	6	6	6	2	7
[1]	5	3	5	5	5	5	4	6	6	6	1	0
[2]	4	6	6	6	2	7	6	1	5	5	2	2

Figure 5: SciDB Storage Manager

SciDB

- ▶ Designed to support spatial operators (e.g. Gaussian smoothing)
- ▶ Chunk partitioning motivated by typical access patterns

MySQL compared with SciDB; Cluster size 10 nodes, 2GB Ram + 3.2 GHz CPU per node

DBMS	Dataset	Loading/Cooking [min]				Query Runtimes [min]									
		Load	Obsv	Group	Total	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Total
MySQL	<i>small</i>	760	110	2	872	123	21	393	0.4	0.36	0.6	0.6	49	50	638
	<i>normal</i> (scaleup)	770 (1.0)	200 (1.8)	90 (45)	1060 (1.2)	54 (0.4)	44 (2.1)	161 (0.4)	50 (125)	32 (89)	51 (85)	52 (87)	395 (8.1)	395 (7.9)	1234 (1.93)
SciDB	<i>small</i>	34	1.6	0.6	36	8.2	0.2	3.7	0.007	0.01	0.01	0.01	1.8	1.9	16
	<i>normal</i> (scaleup)	67 (2.0)	1.9 (1.2)	15 (25)	84 (2.3)	3.6 (0.4)	0.07 (0.4)	1.7 (0.4)	0.015 (2.1)	0.017 (1.7)	0.02 (2)	0.11 (11)	2.2 (1.2)	2.3 (1.2)	10 (0.63)
(MySQL /SciDB)	<i>small</i>	(22)	(69)	(3.3)	(24)	(15)	(105)	(106)	(57)	(36)	(60)	(60)	(27)	(26)	(40)
	<i>normal</i>	(12)	(105)	(6)	(13)	(15)	(630)	(95)	(3330)	(1880)	(2550)	(470)	(180)	(170)	(120)

TABLE I

BENCHMARK RESULTS. (*num*) IS A RATIO OF RUNTIMES, EITHER *normal* VS. *small* (SCALEUP) OR MYSQL VS. SCIDB.

- ▶ “small”: single-machine, 160 3750x3750 images, 99 GB
- ▶ “normal”: distributed, 400 7500x7500 images, 990 GB

Takeaway: SciDB achieves 2-orders speedup compared to MySQL across the 9 benchmark queries on the cluster

⁹Philippe Cudre-Mauroux et al. “SS-DB: a standard science DBMS benchmark”. In: *XLDB* (2012).

Application: EarthDB (MODIS)¹⁰

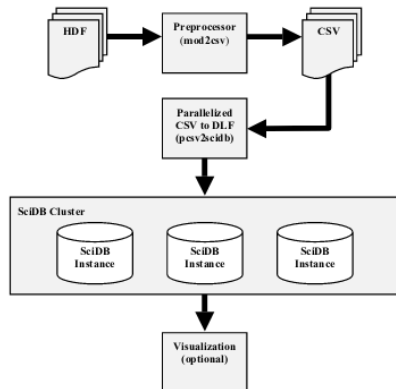


Figure 3: EarthDB High-Level Data Flow

- ▶ Extends SciDB to support MODIS data

¹⁰Gary Planthaber, Michael Stonebraker, and James Frew. "EarthDB: Scalable Analysis of MODIS Data using SciDB". In: *SIGSPATIAL (2012)*, pp. 11–19.

Application: AscotDB (Astronomy)¹¹




Figure 1: AscotDB architecture: SciDB as back-end, python middleware, Ascot and IPython as front-ends.

- ▶ Adds spherical support to SciDB's Cartesian operators
- ▶ Adds efficient spherical operators
- ▶ Adds Python bindings and graphical interface
- ▶ Interdisciplinary collaboration between astronomers and database experts

¹¹Jacob Vanderplas et al. "Squeezing a Big Orange into Little Boxes: The AscotDB System for Parallel Processing of Data on a Sphere". In: (2015).

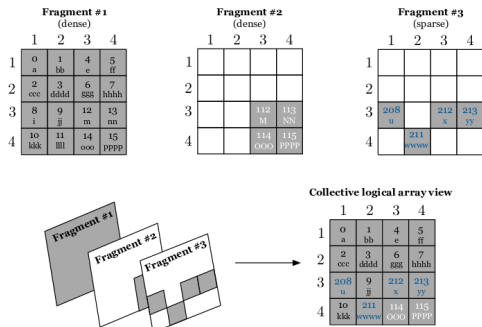


Figure 4: Fragment examples

- ▶ Better support for sparse arrays
- ▶ Avoids fixed-dimension chunking and full chunk reads
- ▶ Key idea: convert random-access writes into sequential appends by storing update **fragments**

¹²Stavros Papadopoulos et al. "The TileDB Array Data Storage Manager". In: *VLDB* (2016).

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

	1	2	3	4
1	0 a	1 bb	4 e	5 ff
2	2 ccc	3 dddd	6 ggg	7 hhhh
3	8 i	9 jj	12 m	13 nn
4	10 kkk	11 lll	14 ooo	15 pppp

Files
 (binary format)




Figure 6: Physical organization of dense fragments

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

	1	2	3	4
1	0 a	1 bb		2 ccc
2			3 dddd	
3	4 e		6 ggg	7 hhhh
4		5 ff		

Files
 (binary format)




Figure 7: Physical organization of sparse fragments




Figure 10: Random update performance of dense arrays

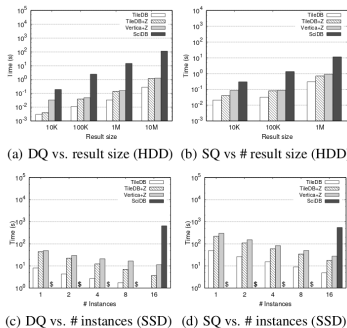


Figure 14: Subarray performance for sparse arrays

- ▶ Random-access update on dense arrays is 2-order faster than HDF5, 3-order faster than SciDB
- ▶ Subarray on sparse arrays is 2-order faster than SciDB
- ▶ Dense read comparable to HDF5

Array DB vs HDF5

Array DB

- ▶ Is a database
- ▶ Can serve as warehouse of files from various sources
- ▶ Supports native operators, e.g. join, query
- ▶ Different flavors support different use cases, e.g. dense/sparse arrays, writes/reads, ...
- ▶ Easy scale-out on clusters
- ▶ Suitable for frequent, varied access

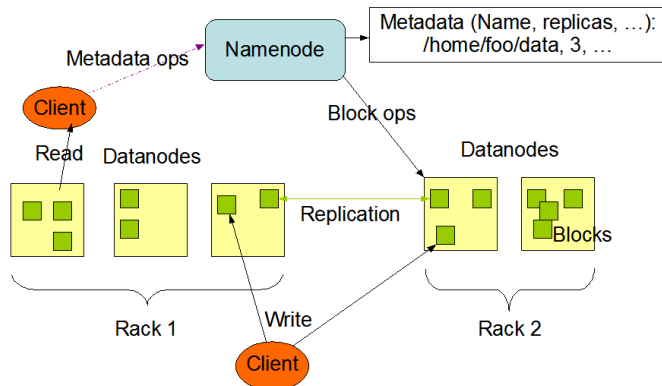
HDF5

- ▶ Is a file format and library
- ▶ Typically stores one “data product” per file
- ▶ Operators are typically at the application level, e.g. a Python program to join two datasets
- ▶ Not meant to be distributed in a cluster
- ▶ Suitable for one-shot processing

GeoTrellis (<http://geotrellis.io>)

- ▶ In-memory big raster analysis on top of Apache Spark
- ▶ Data is managed by Hadoop Filesystem¹³ (HDFS) for fault-tolerance and to support map-reduce style computation
- ▶ “Moving computation is cheaper than moving data”

HDFS Architecture



¹³ “HDFS Architecture Guide”. In: ().

GeoTrellis: NDVI Example

$NDVI = (NIR - Red) / (NIR + Red)$ In GeoTrellis:

```
rdd.mapValues { (red, nir) => (nir - red) /  
(nir + red) } .reduceByKey(_.localMax(-))
```

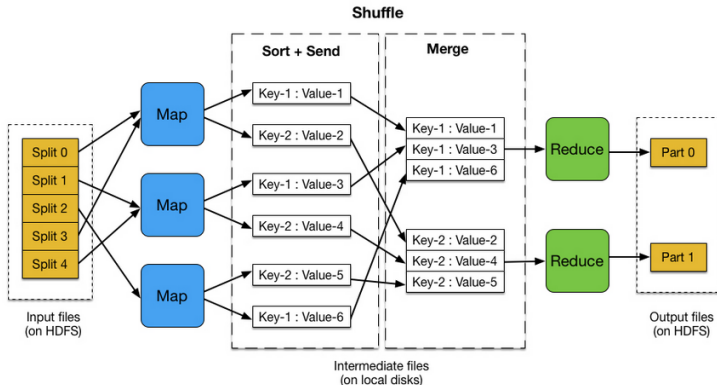
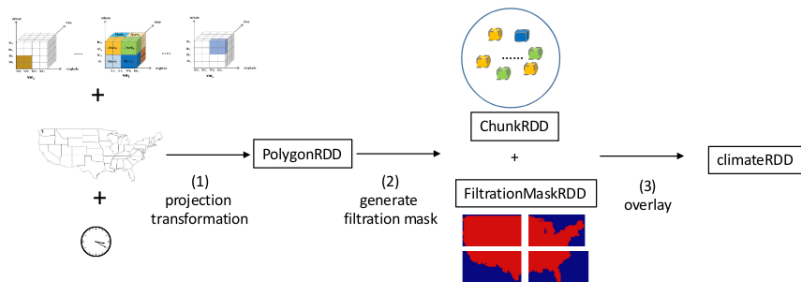


Figure: From <http://www.sunlab.org/teaching/cse6250/spring2017/lab/image/post/mapreduce-flow.jpg>

- ▶ Adds spatiotemporal index to improve data access
- ▶ HDF, netCDF files stored inside HDFS, then indexed, chunked, and finally converted to in-memory RDD's
- ▶ RDDs are processed in Spark via map reduce, same as GeoTrellis



¹⁴Fei Hu et al. "ClimateSpark: An in-memory distributed computing framework for big climate data and analytics". In: *Computers and Geosciences* (2018).

ClimateSpark

- ▶ Experiment: 20 nodes, 24 CPU@2.35 GHz + 24GB RAM per node; $1/2^\circ * 5/8^\circ$ image resolution, roughly 9 TB MERRA2 data spanning 16 years

Varying the query time

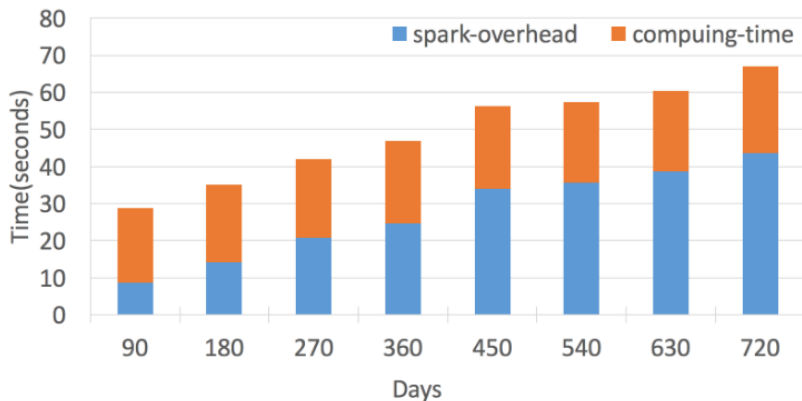


Figure 1. Run time for querying the data in Virginia and computing the monthly mean when varying the query time

Scalable Machine Learning: MLlib

Many ML operations are a good fit for map reduce. MLlib¹⁵ adds machine learning operators onto Apache Spark. Example, gradient descent:

$$w \leftarrow w - \alpha \cdot \sum_i g(w; x_i, y_i)$$

```
for (i <- 1 to n) {  
  val gradient =  
    points.map { p =>  
      (1/(1 + exp(-p.y*w.dot(p.x)) - 1)*p.y*p.x  
    ).reduce(_ + _)  
  w -= alpha*gradient  
}
```

¹⁵Xiangrui Meng. "MLlib: Scalable Machine Learning on Spark". In: ().

Scalable Machine Learning: MLlib

MLlib is a standard Spark component, thus is tightly integrated with Spark data operations such as Spark SQL:

```
val trainingTable = sql("""SELECT... FROM... JOIN...""")
val training = trainingTable.map { row => ... }
val model = SVMWithSGD.train(training)
```

Integrating machine learning with existing data management tools is an exciting area of research¹⁶.

¹⁶Arun Kumar, Matthias Boehm, and Jun Yang. “Data Management in Machine Learning: Challenges, Techniques, and Systems”. In: *SIGMOD* (2017).

Ridesharing: A different kind of spatial problem

17

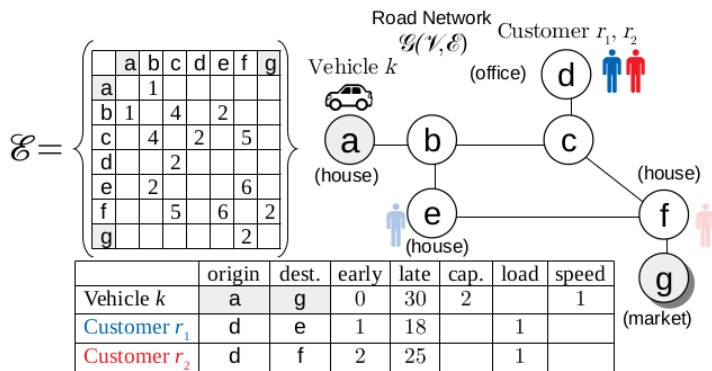


Figure 1: *Ridesharing example.*

Figure: How to optimally match vehicles to customers and design the service route?

¹⁷James Pan, Guoliang Li, and Juntao Hu. "Ridesharing: Benchmark, Simulator, and Evaluation". In: *VLDB (2019)*.

18




Figure 4: Jargo usage scenario.

Figure: How to evaluate ridesharing algorithms? Jargo can achieve real-time simulation of 1,000's vehicles on commodity machine by manipulating data of the system instead of physical simulation.

Summary

Key takeaways:

- ▶ New knowledge hiding in vast amounts of data
- ▶ Data management helps work around physical limitations of storage
- ▶ Array DB (SciDB, EarthDB, TileDB) for new science applications
- ▶ Integration of ML with data management, e.g. Spark MLlib
- ▶ **Very active space**

Thank you!

Q&A