What's next for the Berkeley Data Analytics Stack?

Michael Franklin
June 30th 2014

Spark Summit
San Francisco

AMPLab: Collaborative Big Data Research

60+ Students, Postdocs, Faculty and Staff

Systems, Networking, Databases, and Machine Learning

amplab UC BERKELEY

Franklin Jordan Stoica Goldberg Joseph Katz Patterson Recht Shenker

Industry Sponsors + White House Big Data Program (NSF and Darpa)

amazon.com web services, Google, SAP, Siebel The Thomas and Stacey Siebel Foundation, Adobe, Apple Inc., Cisco, Ericsson, facebook, NSF, DARPA, Berkeley Lab, hp, Huawei, intel, Microsoft, Energy, cloudera, guavus, NetApp, EMC² Pivotal, vmware, Yahoo!, splunk>, virdata, wanDISCO

In-house Apps:

Cancer Genomics, Mobile Sensing, Collaborative Energy Saving

AMPLab: Integrating 3 Resources

Algorithms

- Machine Learning, Statistical Methods
- Prediction, Business Intelligence

Machines

- Clusters and Clouds
- Warehouse Scale Computing

People

- Crowdsourcing, Human Computation
- Data Scientists, Analysts

Extreme Elasticity

Algorithms

- Approximate Answers: Trading time for error
- ML Libraries and Ensemble Methods
- Active Learning

Machines

- Cloud Computing; Multi-tenancy
- Deep and dynamic storage hierarchies
- Relaxed (eventual) consistency/ Multi-version methods

People

- Dynamic Task and Microtask Marketplaces
- Visual analytics
- Manipulative interfaces and mixed mode operation

Big BDAS Bets

Leverage commodity trends: in-memory processing, elastic compute and scalable storage (HDFS)

Optimize for common patterns: MR, graphs, SQL, iterative machine learning, matrix algebra

Tradeoff accuracy and runtime with sampling and relaxed consistency

Integration: expose a consistent API to support batch, stream and interactive computation

Build and support an Open Source community around it on an academic research budget

Berkeley Data Analytics Stack
(Spark is just part of what we do)

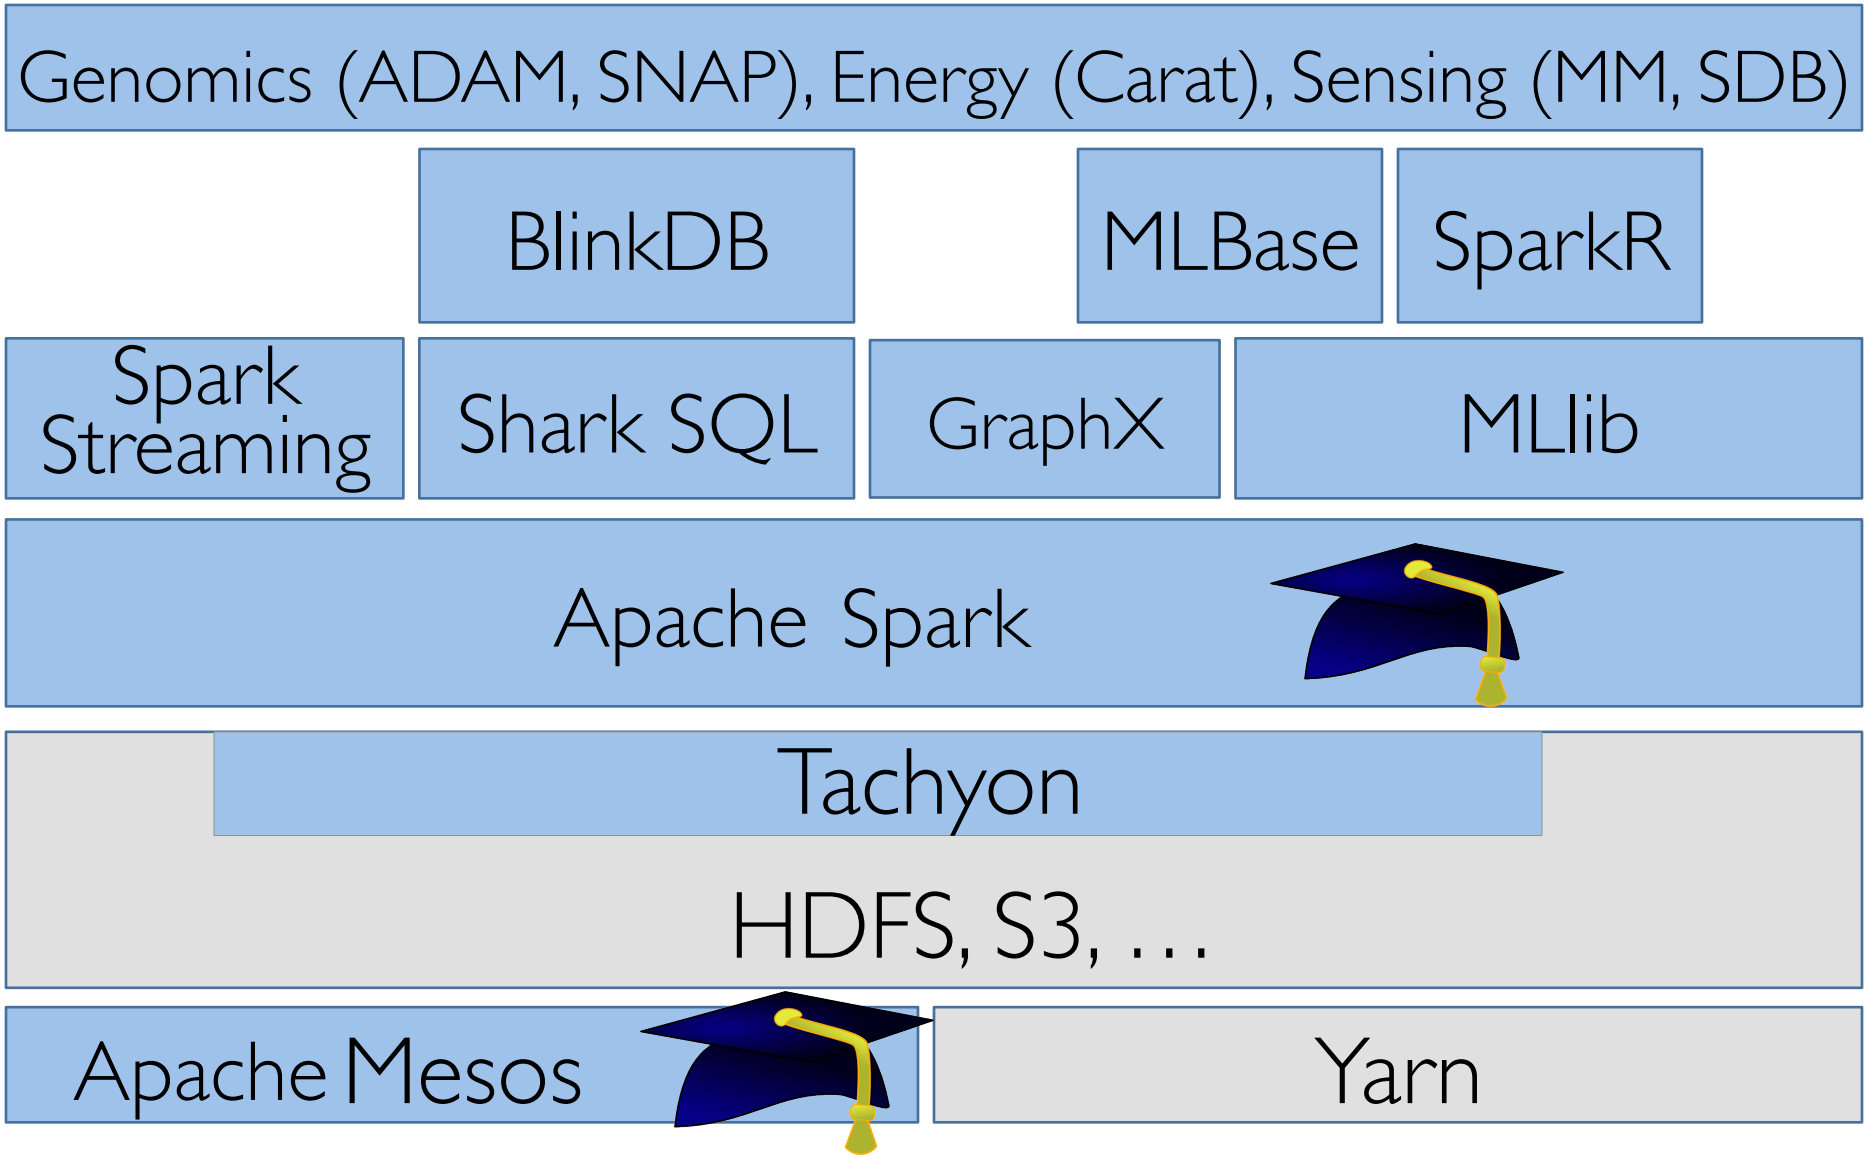

WHAT'S NEW IN BDAS?

Tachyon

In-memory, fault-tolerant storage system

Easily and efficiently share data across frameworks

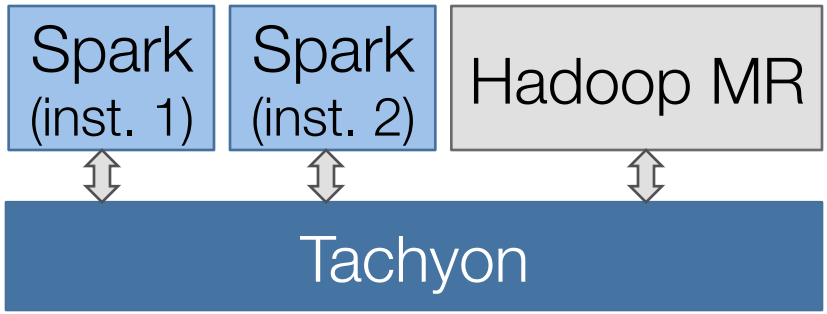

Off-heap allocation; HDFS compatible API

See Haoyuan Li's talk at 3pm this afternoon

Fast, approximate answers with error bars by executing queries on small, pre-collected samples of data

Recent focus on diagnostics for reliable error reporting

```
SELECT avg(sessionTime)
FROM Table
WHERE city='San Francisco'
WITHIN 2 SECONDS
```

Queries with Time Bounds

```
SELECT avg(sessionTime)
FROM Table
WHERE city='San Francisco'
ERROR 0.1 CONFIDENCE 95.0%
```

Queries with Error Bounds

S. Agarwal et al., "Knowing When You're Wrong: Building Fast and Reliable Approximate Query Processing Systems" *ACM SIGMOD Conf.*, June 2014.

Sampling + Data Cleaning using less data for better answers

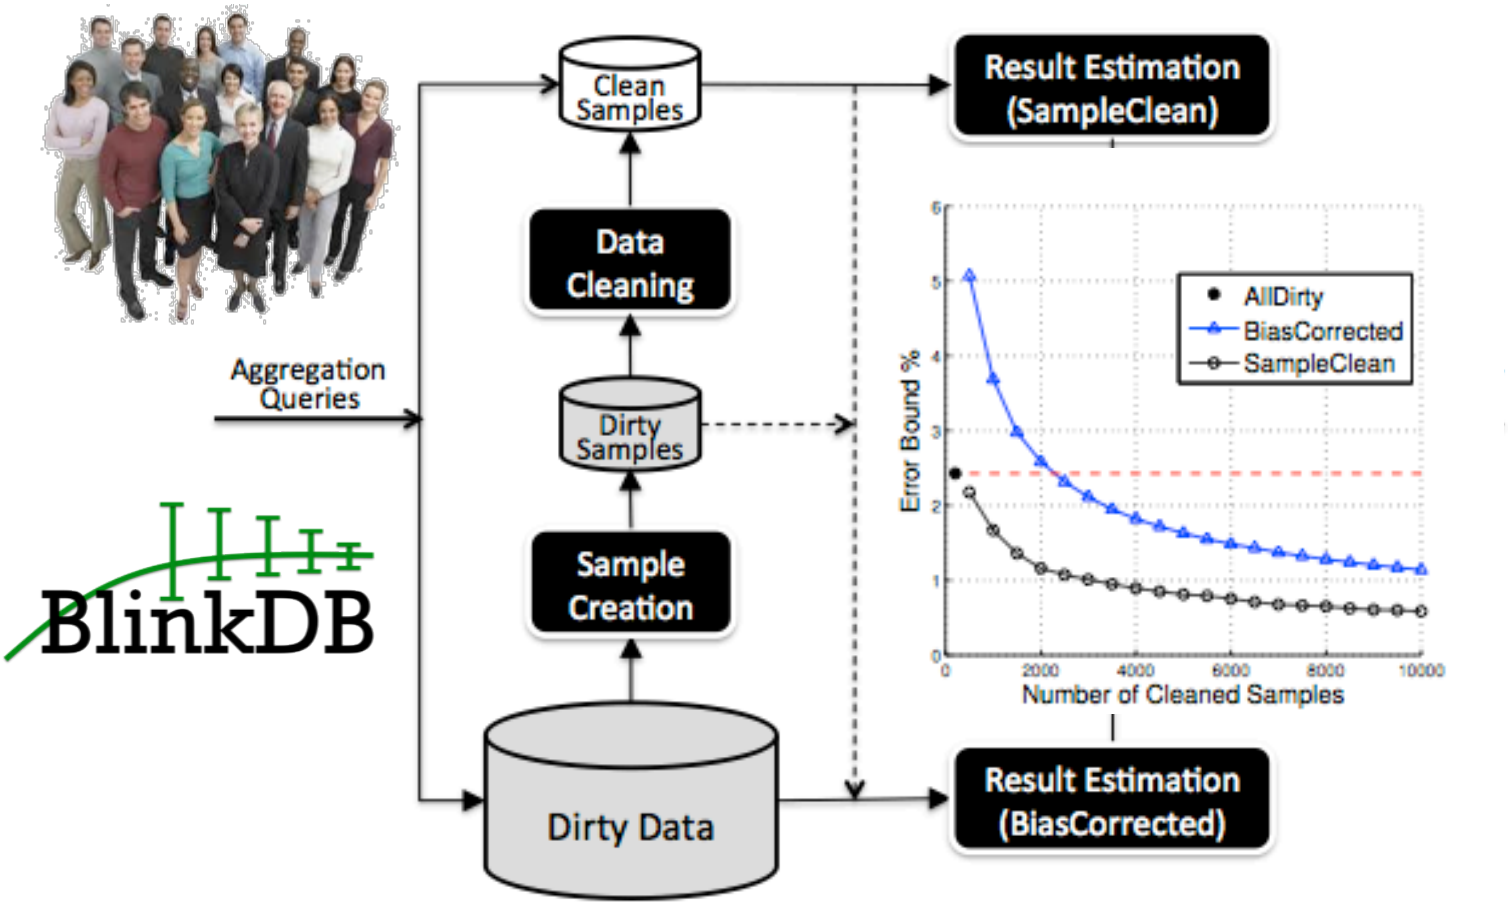

S. Krishnan, J. Wang et al., *"A Sample-and-Clean Framework for Fast and Accurate Query Processing over Dirty Data", ACM SIGMOD Conf.,* June 2014.

GraphX – Adding Graphs to the Mix

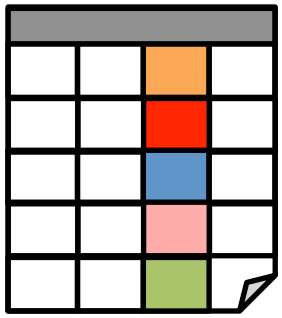

Table View

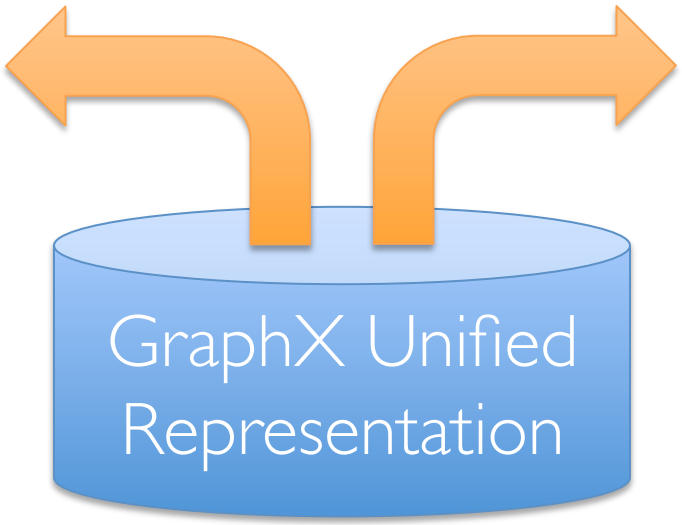

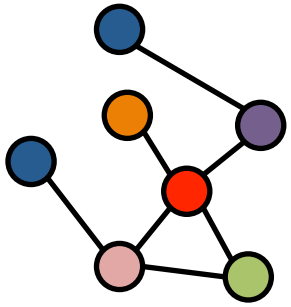

Graph View

Tables and Graphs are composable *views* of the *same physical* data

Each view has its own operators that exploit the semantics of the view to achieve efficient execution

R. Xin, J. Gonzalez, M. Franklin, I. Stoica, "GraphX: In-situ Graph Computation Made Easy" GRADES Workshop at SIGMOD, June 2013.

Distributed Graphs as Distributed Tables

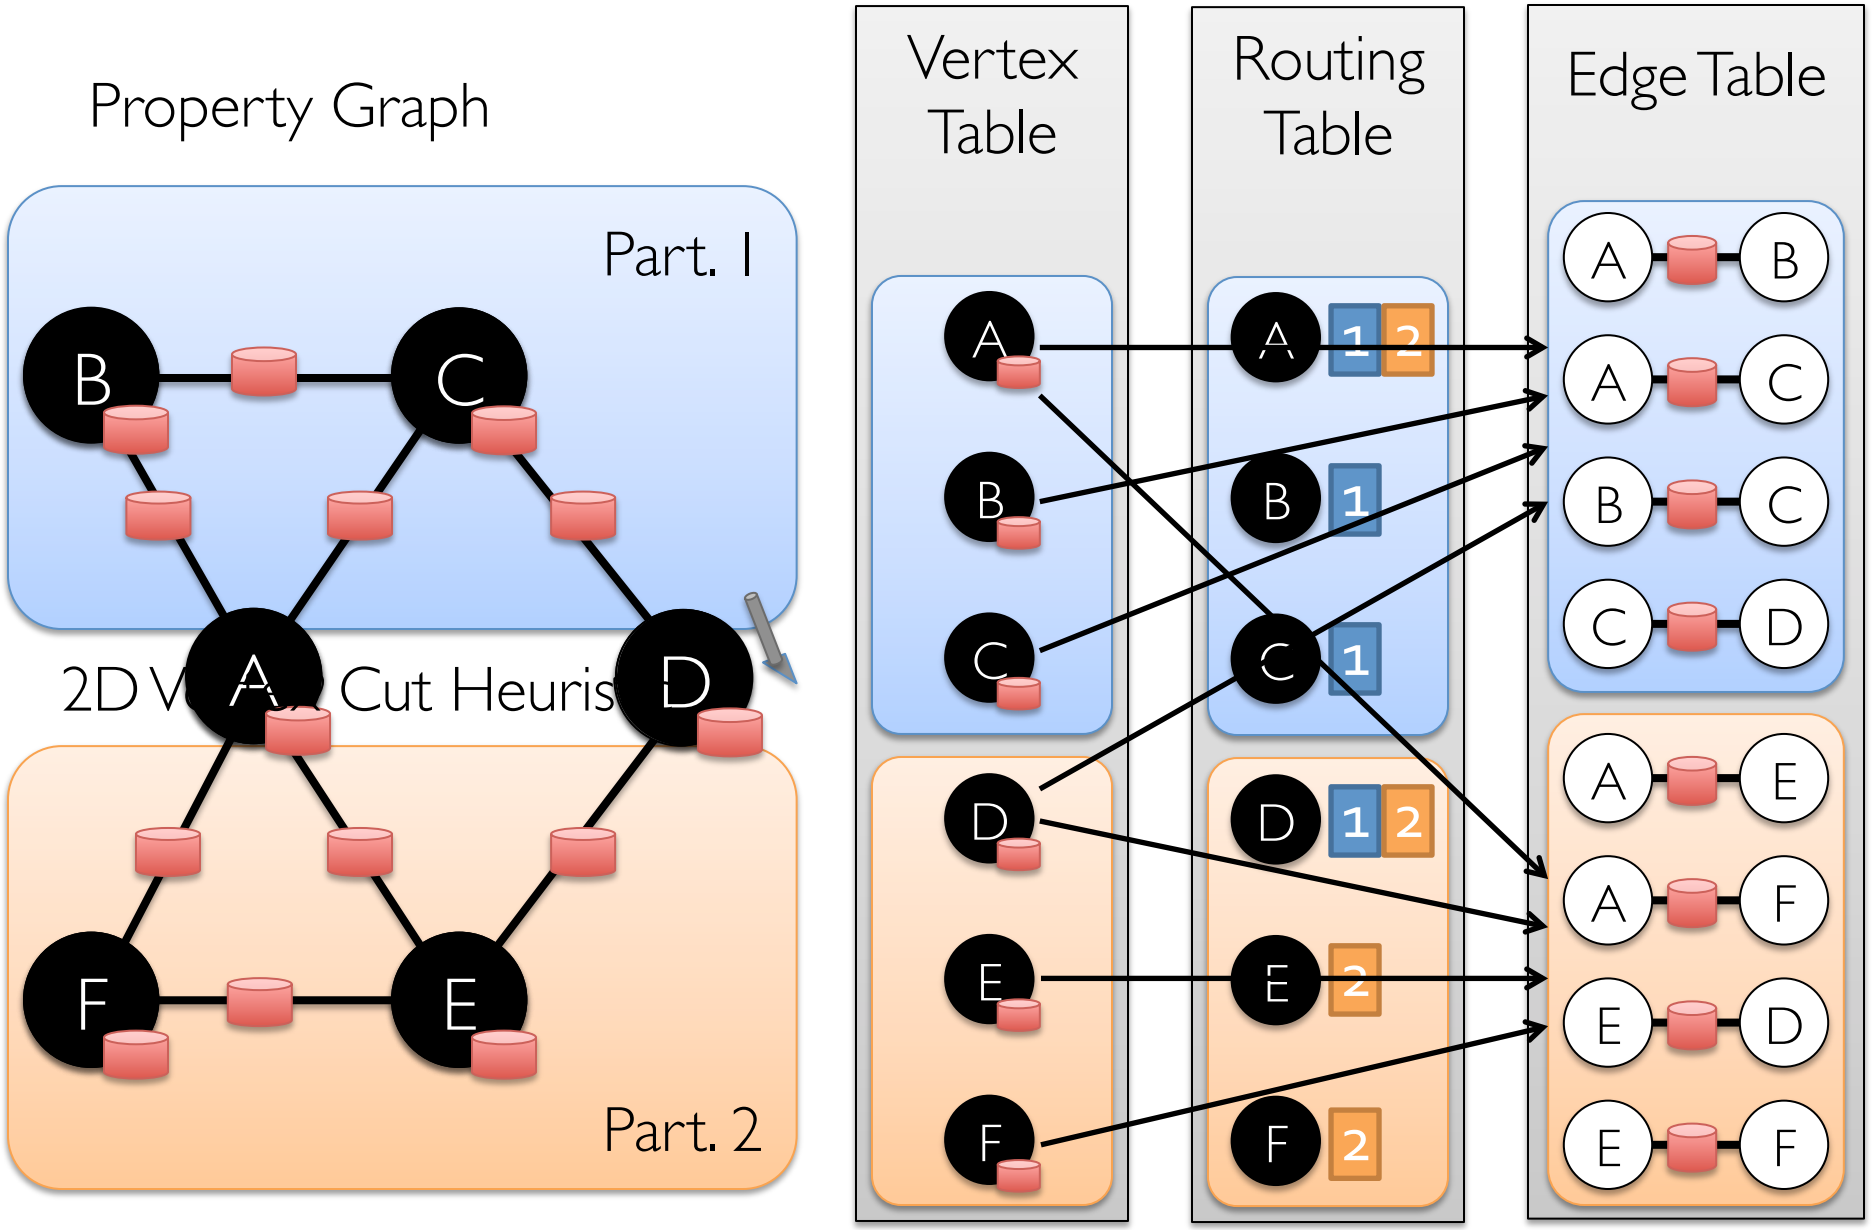

GraphX: Graphs → Dataflow

1. Encode graphs as distributed tables
2. Express graph computation in relational ops.
3. Recast graph systems optimizations as:
   1. Distributed join optimization
   2. Incremental materialized maintenance

Integrate Graph and Table data processing systems.

Achieve performance parity with specialized systems.

WHAT'S NEXT IN BDAS?

MLBase: Distributed ML Made Easy

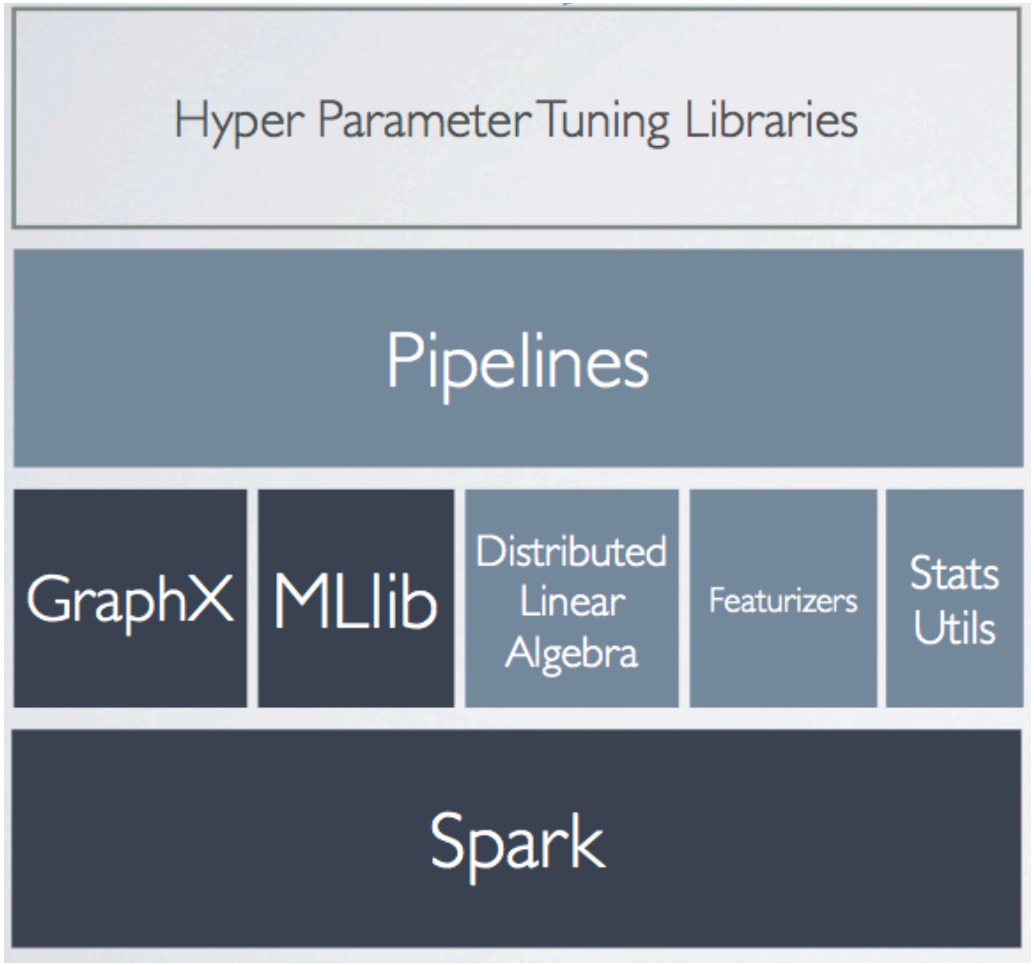

DB Query Language Analogy:

Specify **What** not **How**

MLBase chooses:

- Algorithms/Operators
- Ordering and Physical Placement
- Parameter and Hyperparameter Settings
- Featurization

Leverages Spark for Speed and Scale

ML Pipeline Generation & Optimization

Ex: Image Classification Pipeline

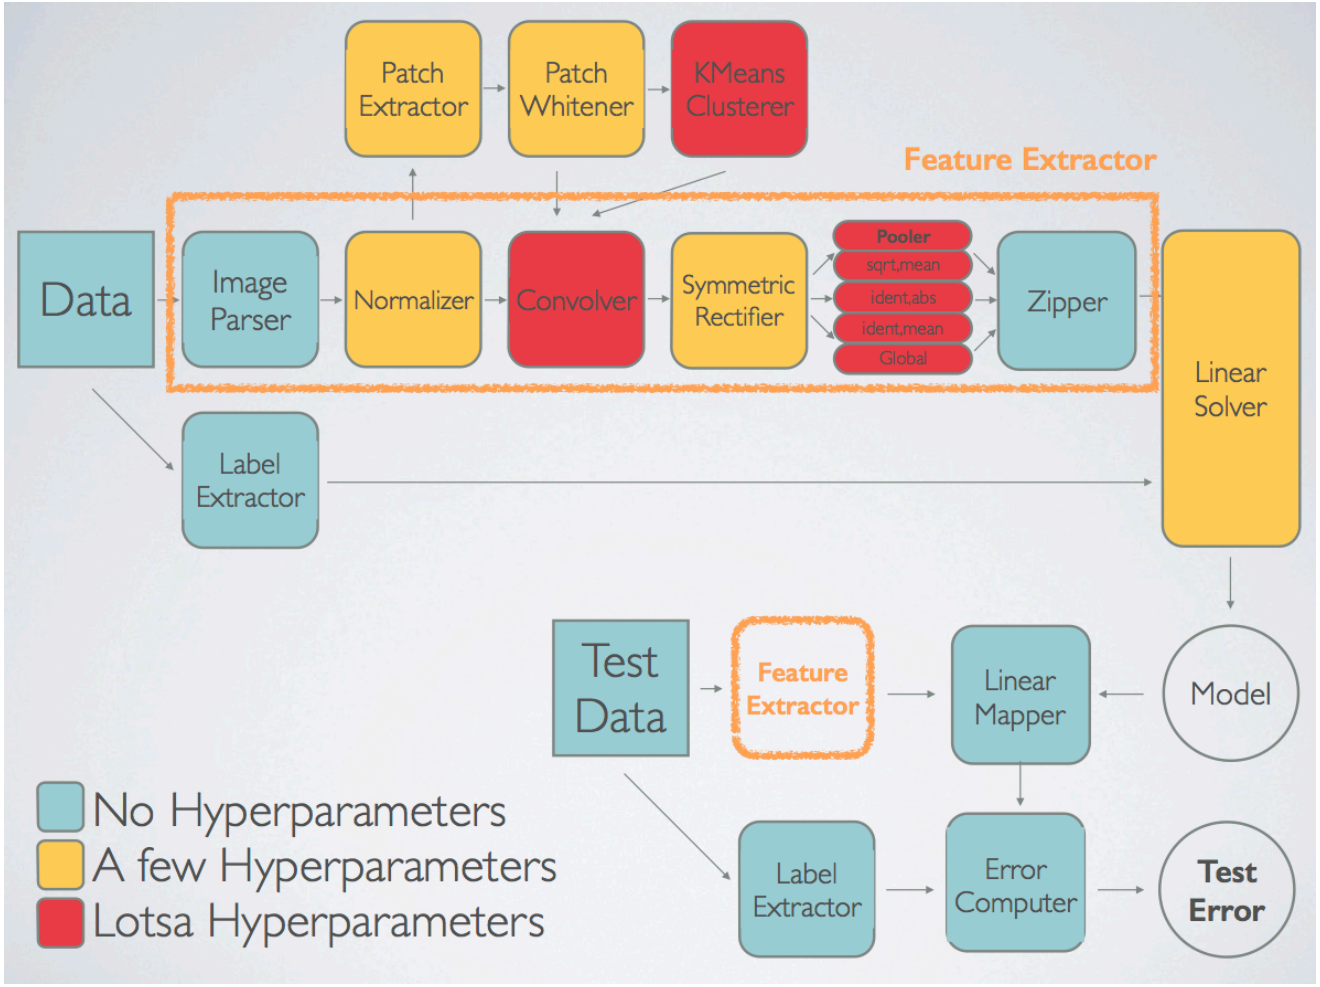

MLBase Optimizer Focus:

- Better Resource Utilization
- Algorithmic Speedups
- Reduced Model Search Time

Work in Progress: See Evan Spark's talk at 1pm tomorrow afternoon

What about Updates?

Big Data Analytics assume **Append-mostly** data

Many applications require "Point Updates"

Real-time ML and Analytics require fast Model Updates

We have several projects exploring this space

- **MDCC**: Mutli Data Center Consistency
- **HAT/RAMP**: Highly-Available Transactions/Read Atomicity
- **PBS**: Probabilistically Bounded Staleness
- **PLANET**: Predictive Latency-Aware Networked Transactions

(see SIGMOD 2014 for details on RAMP and PLANET)

The Big Picture: Serving Models & Data

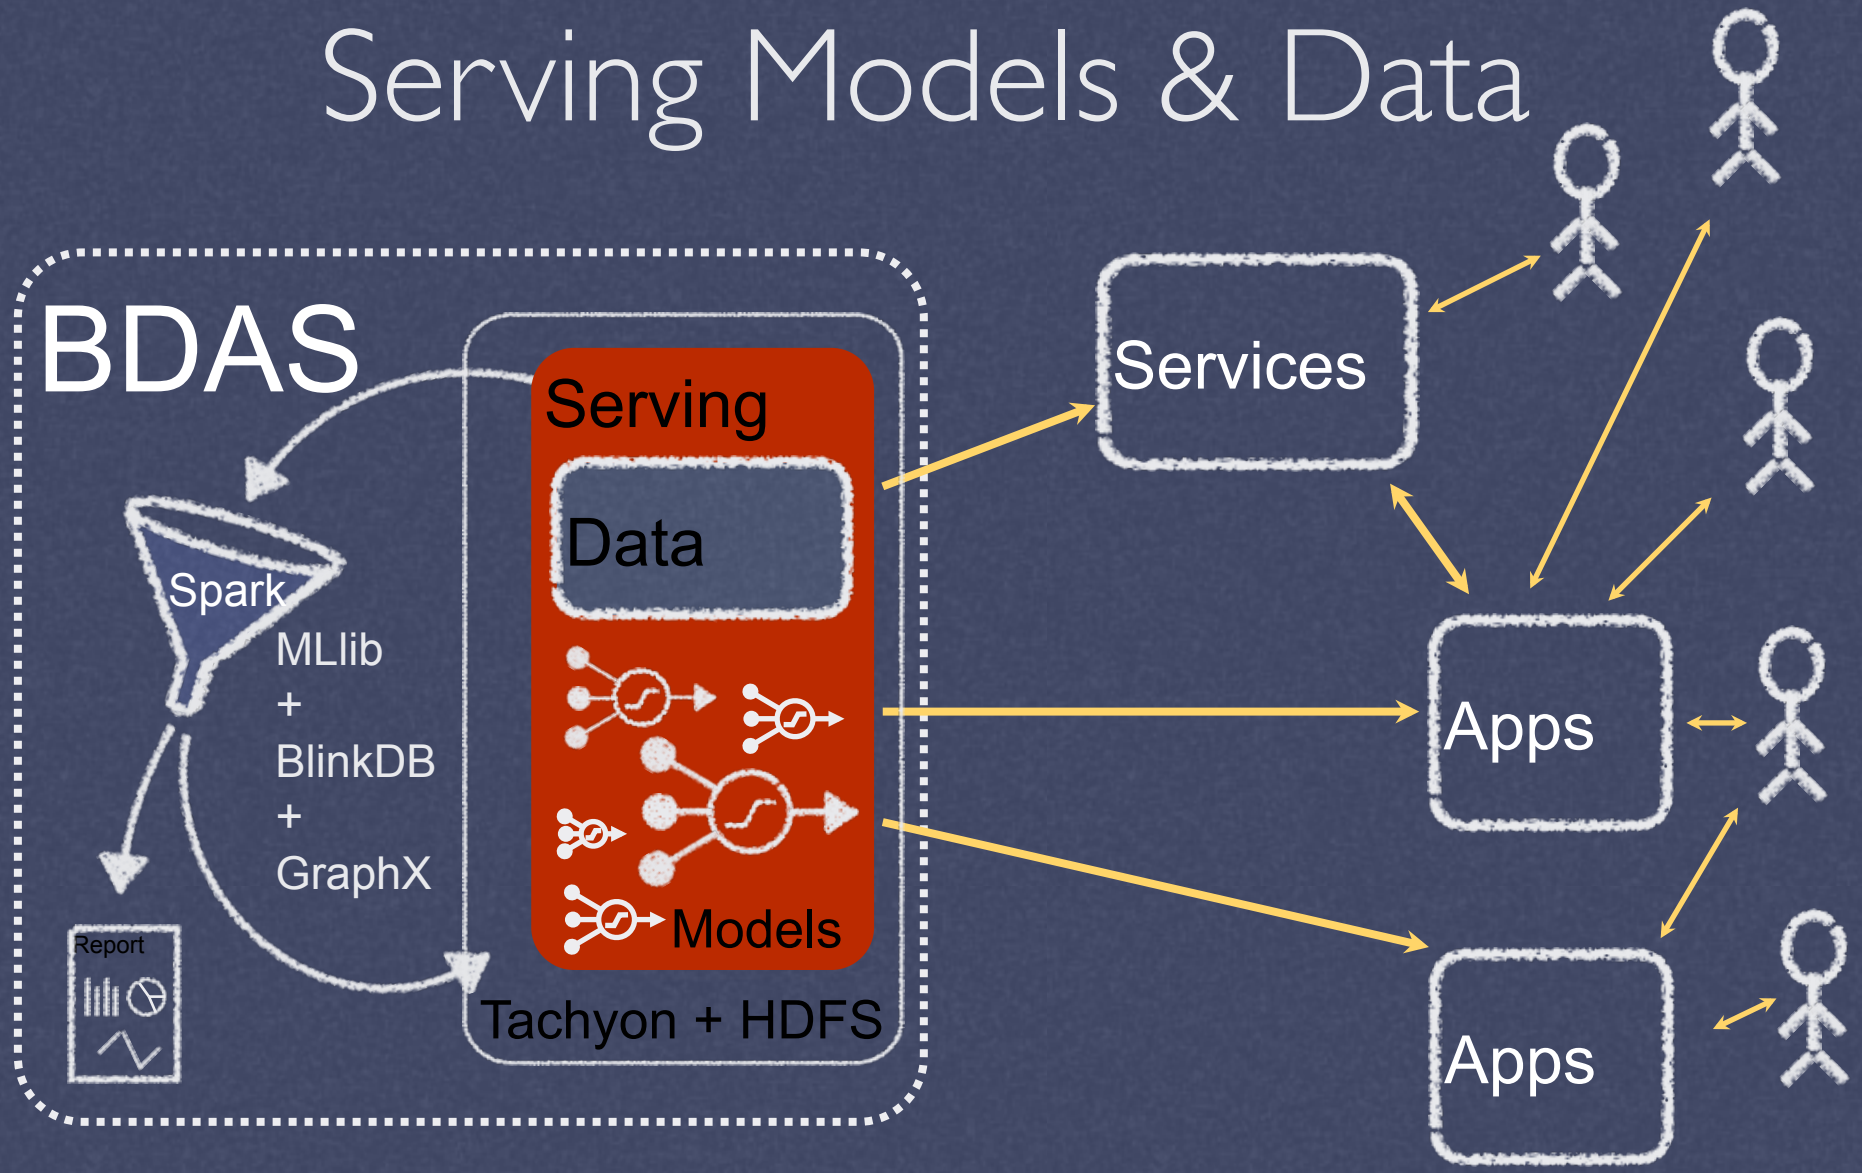

Enables real-time Analytics and Model Updates

Looking Further…

The big breakthrough of the "Big Data" ecosystem isn't scalability

It's not really speed either

It's flexibility!

The (Traditional) Big Picture

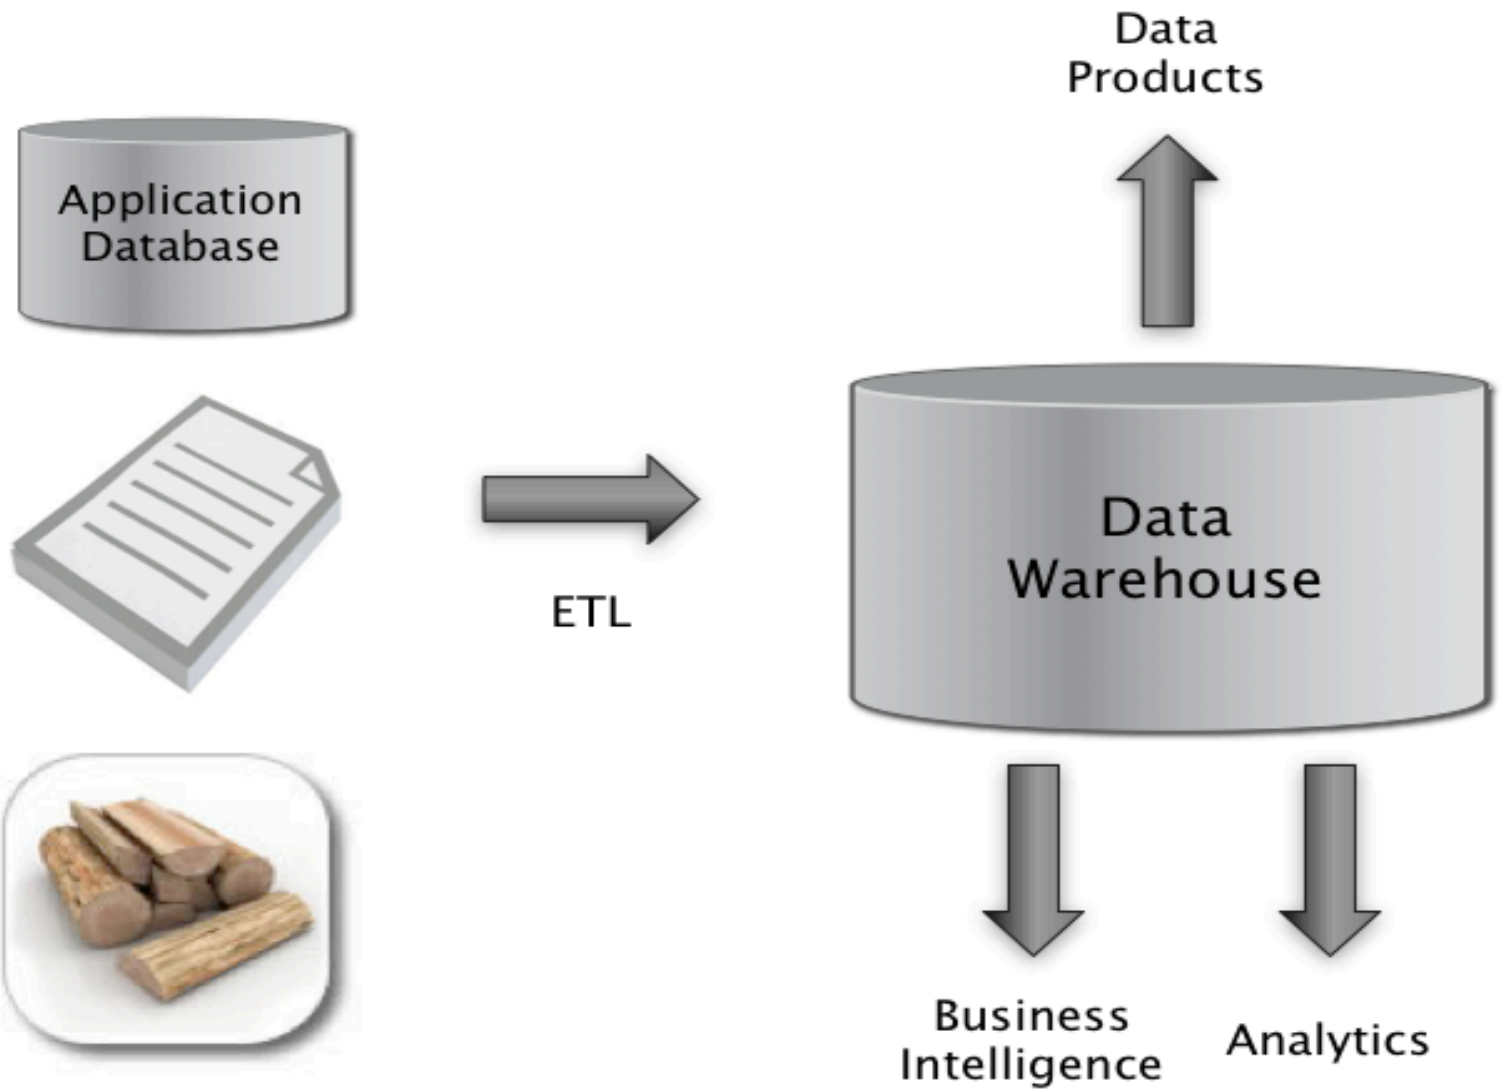

Database Philosophy

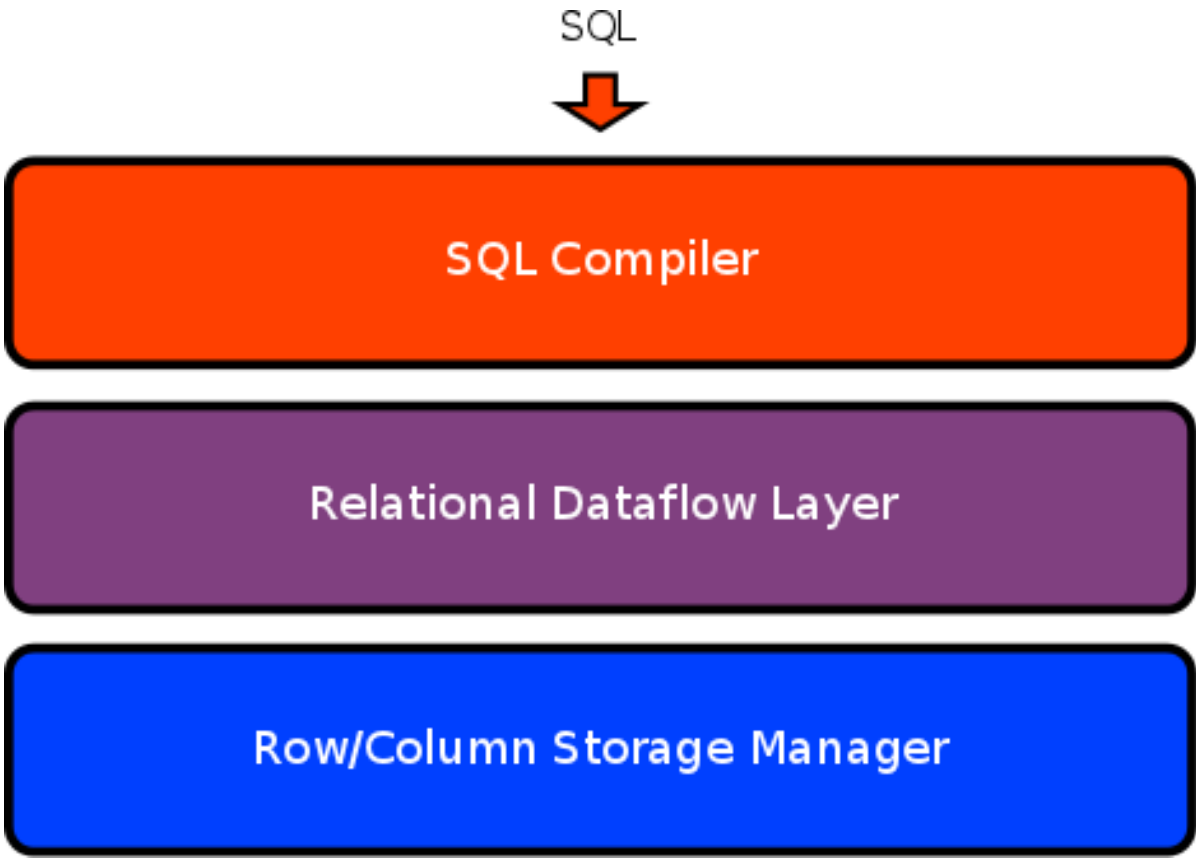

One way in/out: through the SQL door at the top

From Mike Carey, UCI

Open Source Big Data Stack

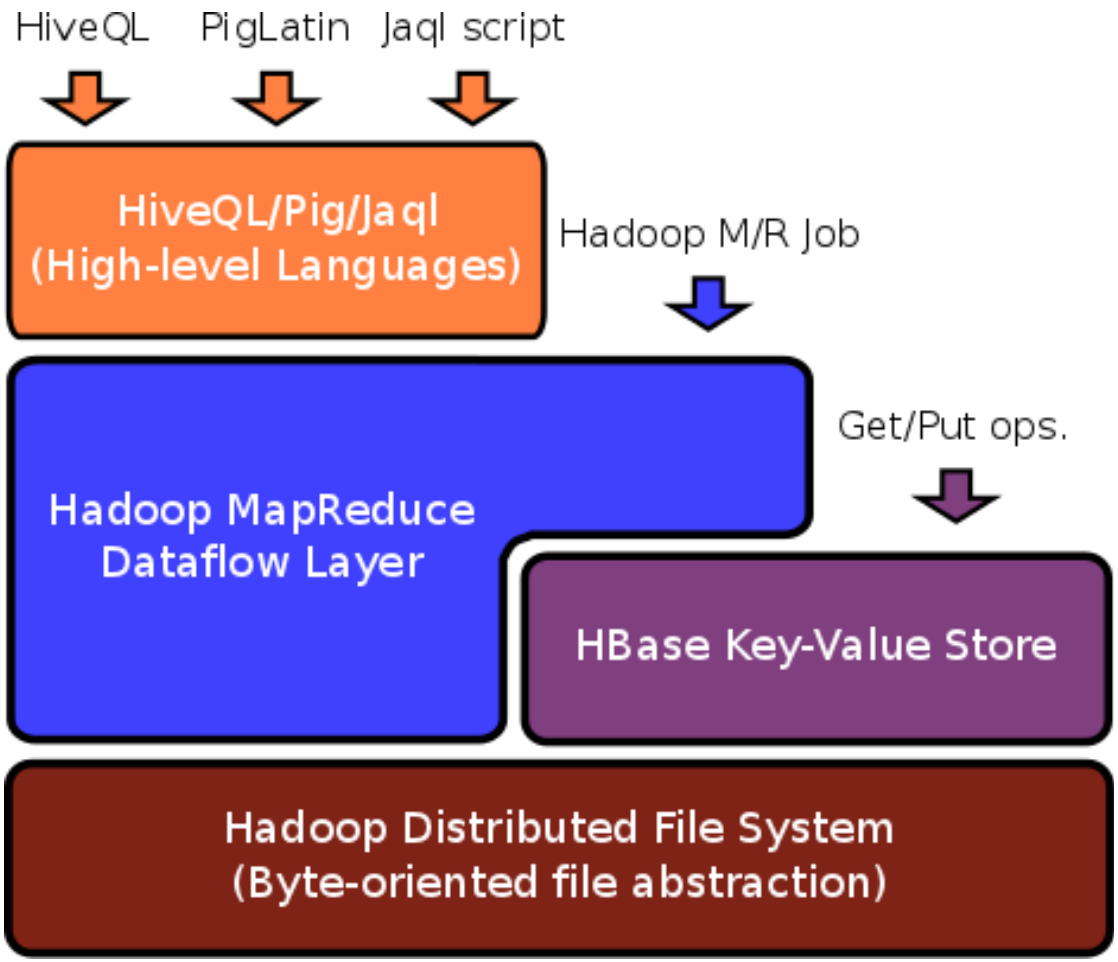

Notes:

- Giant byte sequence at the bottom
- Map, sort, shuffle, reduce layer in middle
- Possible storage layer in middle as well
- HLLs now at the top

From Mike Carey, UCI

The (Traditional) Big Picture

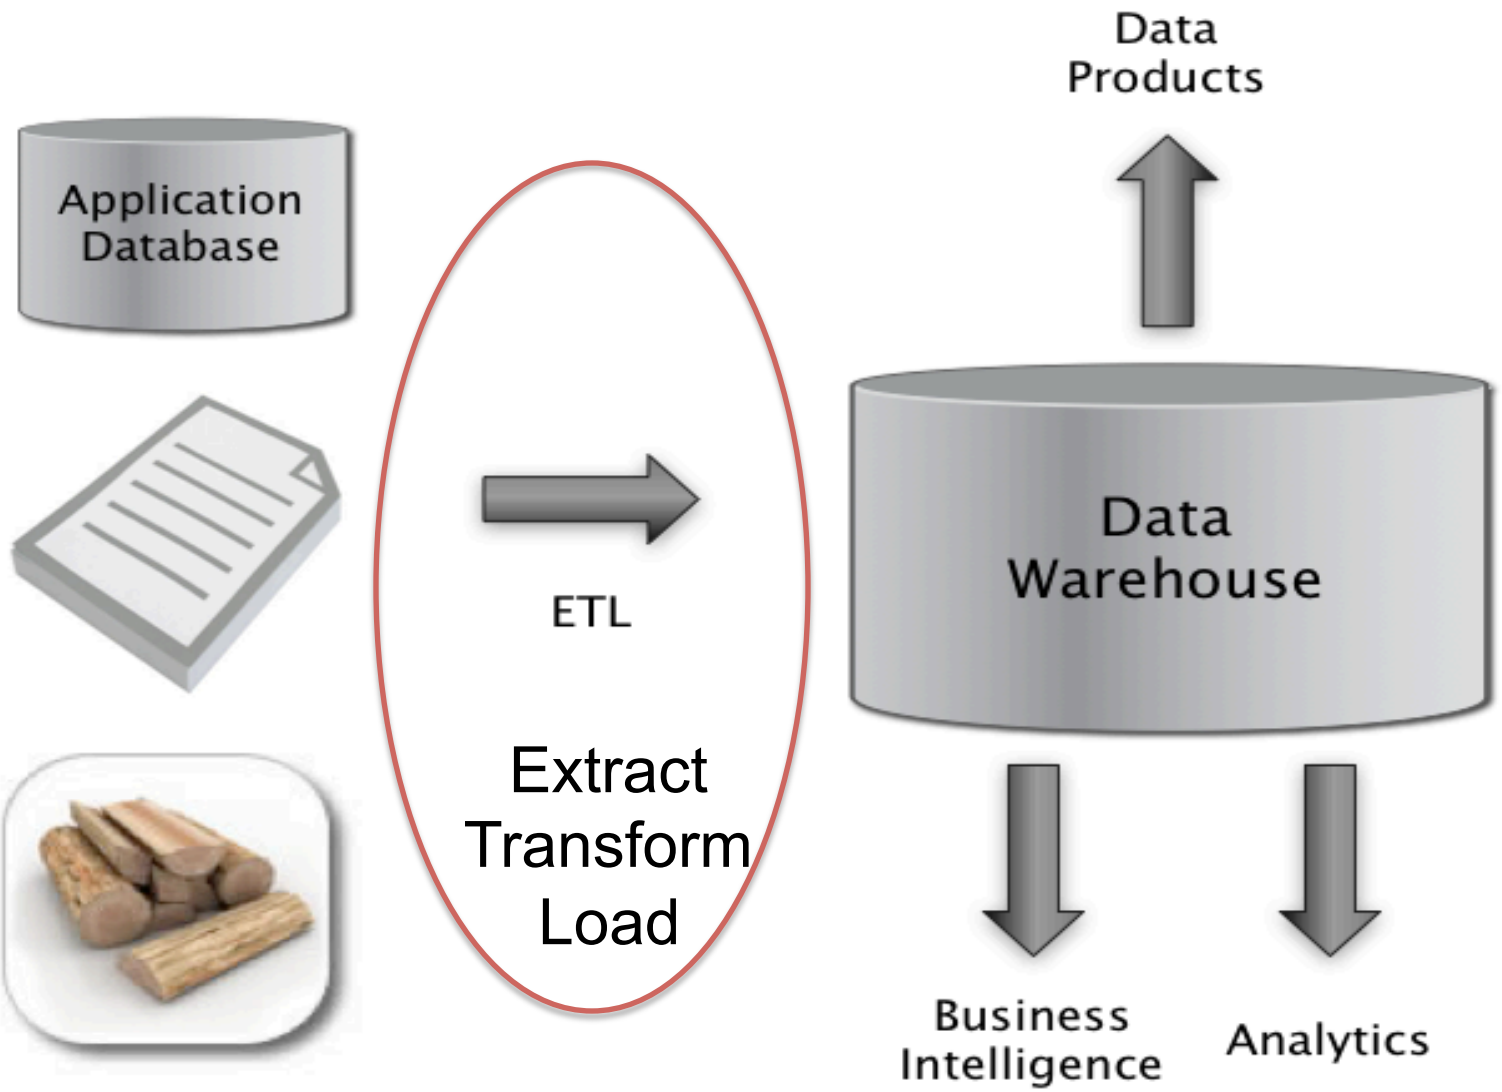

The Structure Spectrum

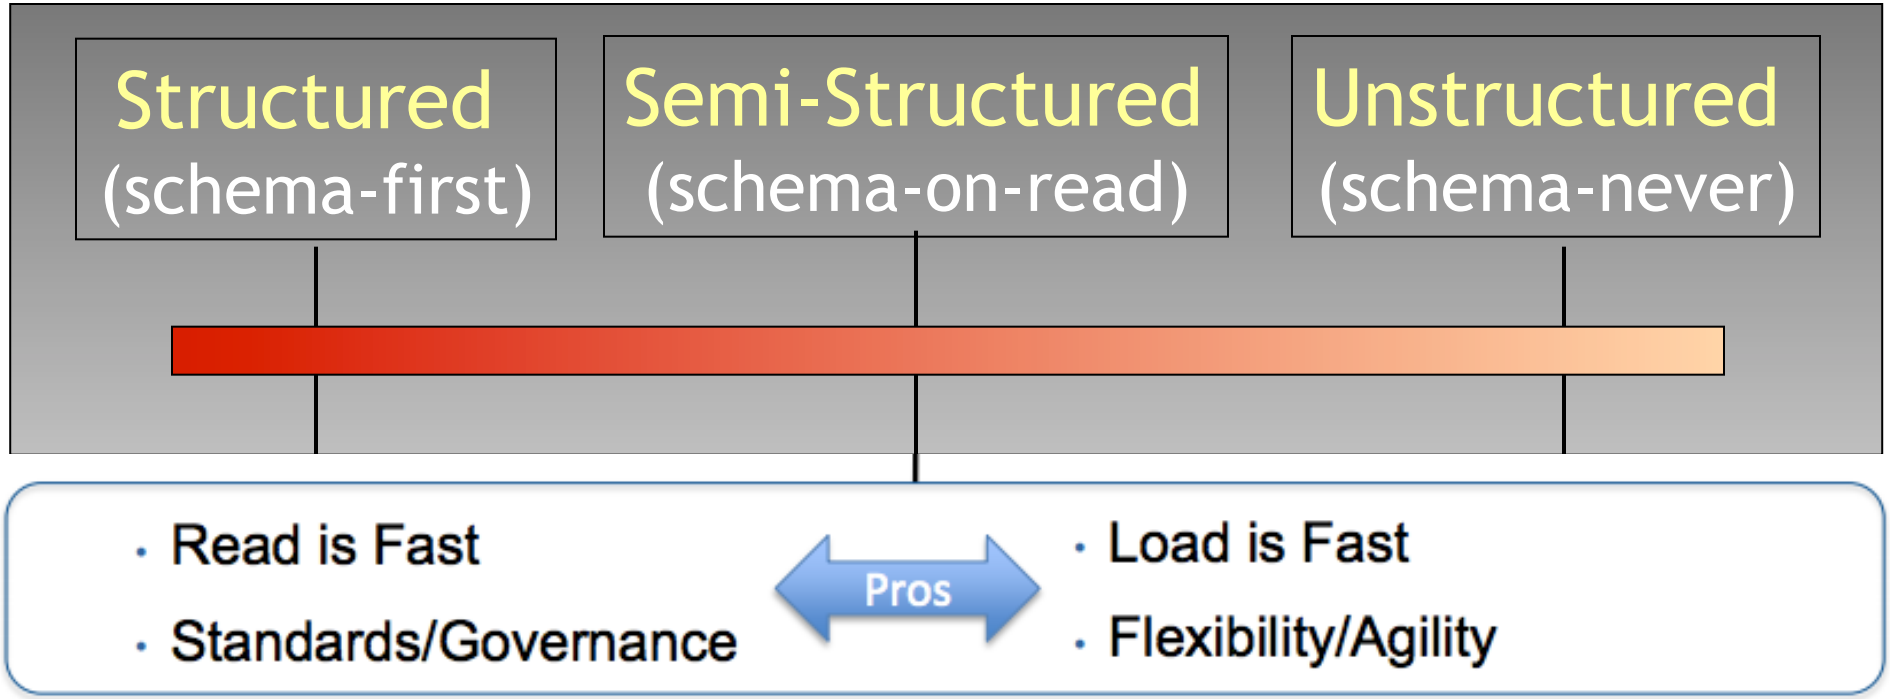

Schema-as-needed

- ETL Becomes a Continuous/On-Demand Process
- But, may lose access to some data in the meantime

Where will the next Spark Come From?

- Easy-to-use Machine Learning (a la MLBase)
- Integrating Serving and Analytics
  - more than the old OLTP + OLAP story
  - serving Data and Models; supporting Real-time
- Continuous Data Integration and ETL
- Renewed focus on Single-Node performance
- Pushing processing to the Edge – e.g., for IoT
- Tracking new server/data center technologies

To find out or (better, yet) get involved:

amplab.berkeley.edu

franklin@berkeley.edu

Thanks to NSF CISE Expeditions in Computing, DARPA XData,
Founding Sponsors: Amazon Web Services, Google, and SAP,
the Thomas and Stacy Siebel Foundation,
and all our industrial sponsors and partners.