summer school

An Explosion of Databases

The XML Summer School
Oxford 12/09/2018

Adam Retter


adam@evolvedbinary.com


@adamretter

EVOLVED BINARY

Adam Retter

- **eXist-db - Core Developer (13 years!)**
  - Native XML Database
  - Implemented in Java
  - Open Source: LGPL v2.1
- **RocksDB - Developer (4 years)**
  - Key / Value Database
  - Implemented in C++ (and Java API)
  - Open Source: GPL v2 / Apache 2.0
- **Granite - Developer (4 years)**
  - Polystore: XML, Key/Value (JSON, MarkDown... DOM)
  - Implemented in C++ and Java
  - Will be Open Source: likely AGPL v3

In the beginning... ~1960s

- **General Electric - IDS (Integrated Data Store)**
  - Possibly the first DBMS
  - Network Model
  - Schema
  - CODASYL
  - Tuple-at-a-time queries
- **IBM - IMS (Information Management System)**
  - Developed for the Apollo moon mission (purchasing)
  - Hierarchical Model
  - Programmer defined physical storage (Hash / Tree / etc.)
    - Determines the API you can use to query
  - Tuple-at-a-time queries

Things Start Improving... ~1970s

- **Ted Codd (IBM)**
  - Avoid rewritting applications for every schema change
  - Need more abstraction
    - Logical vs. Physical
    - Let the database engine worry about physical storage
    - Let the user query their logical model
    - Query through a high-level language
  - The Relational Data Model is born

- **Implementations**
  - System R (Jm Gray - IBM)
  - INGRES (Michael Stonebraker - U.C. Berkeley)
  - Oracle (Larry Ellison)

Things Stabilise... ~1980s

- **Mostly Improvements to the Relational Model**
- **SQL is The standard - 1986 ANSI**
- **Further notable implementations:**
  - Informix (1981 / SQL 1985)
  - IBM DB2 (1983)
  - Sybase (1987)
    - Partnership created Microsoft SQL Server (1989)
    - Later SAP!
  - Postgres (1989)
    - Stonebraker - Post-Ingres
- **Oracle dominates!**
- **New: Object Database (1985) / Object-Relational hybrids**

Kinda dull, until... ~1990s

- **Postgres95 (1995)**
  - 1994 - Berkley shutters Postgres
  - Released as Open Source under MIT
  - Forked as Postgres95 (later PostgreSQL)
- **MySQL (1995)**
  - Open Source rewrite of mSQL
- **The Web Takes off!**
  - The rise of the LAMP stack!
  - 1995 - 16 million users (0.4% world pop.)
  - 1999 - 248 million users (4.1% world pop.)
- **Cap Theorem (Eric Brewer - 1999)**

Scaling... ~2000s

- **SQLLite (2000)**
- **The Web (2009)**
  - Reaches 1,802 million users (26.6% world pop.)
- **Big Web Companies:**
  - Commercial databases are too expensive and don't scale
  - Open Source databases lack features
  - Each building middle-ware to distribute load, e.g.:
    - eBay and Amazon - Oracle
    - Facebook - MySQL
  - Start building their own DBMS:
    - Google - BigTable/LevelDB (2004), Spanner (2012)
    - Amazon - DynamoDB (2012)

The explosion... ~2010s

- **Much data!**
  - Facebook - 6 billion photos a month / 100 petabytes (2012)
  - Google - 40,000 searches per second (2014)
- **The NoSQL "Movement"**
  - Not SQL => Not (only) SQL
  - Rejects classic DMBS in favour of lighter faster storage
  - Compromises - Consistency, Availability, Durability vs. Performance
  - Full-circle. The SQL vendors fought back! - NewSQL
- **New Hardware**
  - RAM is cheap
  - SSD / NVMe / RDMA
  - GPU and FPGA

Interesting databases today

- **RocksDB**
  - Facebook's Open Source LevelDB fork... for SSD/NVMe etc.
  - Key/Value
  - Powers almost everything at Facebook (and others)
  - Used in: ArangoDB, Cassandra, CockRoach DB, MongoRocks (MongoDB), MyRocks (MySQL), many more...
- **MapD**
  - Database core is in-memory and GPU optimized
  - SQL
  - Optimized for data analytics
- **CockrochDB**
  - Open Source. Distributed database.
  - SQL

- **ScyllaDB**
  - Cassandra compatible implementation in C++
  - Column Store
  - 2x - 10x faster than Cassandra
  - Optimised for multi-threaded machines. Clustering also.
- **FoundationDB**
  - Previously closed source, now Open Source (under Apple)
  - Key/Value Store
  - Designed for performance (after durability)
  - Compromises - Transaction lifetime
- **FasterDB**
  - Microsoft Open Source
  - Embedded key/value store
  - Impressive performance "claims"

Where are we heading?

- In-Memory?
  - Memory / Persistent disk is now blurred (NVMe etc).
  - Custom Hardware - ASIC, FPGA, RDMA Network etc.
- Consistency is back in vogue.
- Likely SQL (or similar) for the user.
- Distributed. Sharding. Clustered.
  - Node failure happens! Data centre failure happens!
- Common Core, e.g.: RocksDB.
- Polystore vs. Multiple databases

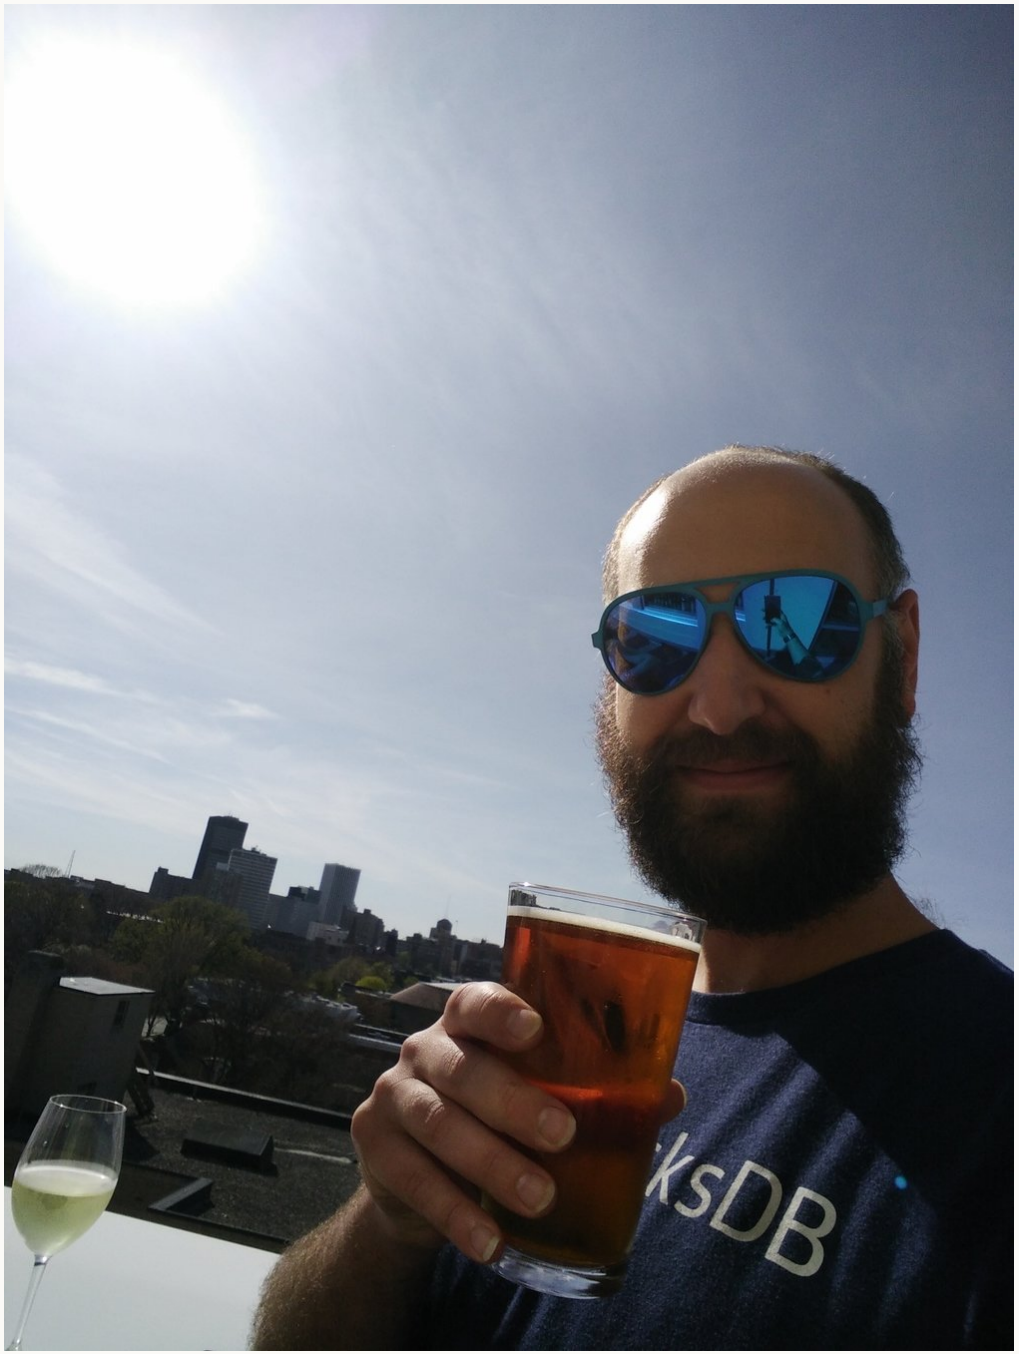

Questions?

Learn More:

CMU Advanced Database Systems

https://15721.courses.cs.cmu.edu/spring2018/

...The YouTube Videos are excellent!

EVOLVED BINARY