

Topic 16:

An Introduction to 'NoSQL' DBMSes

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 1/12

Relational DBMSes: A Review

Hopefully mostly old news!

- Schemas describe highly-structured data
- Originally monolithic (centralized storage and control)
- ACID properties help to safeguard the data
- Designed for business data (numbers, strings)
- Performs well on transactional and analytical workloads
- Uses SQL as DDL, DML, DCL, and QL.

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 2/12

Relational DBMSes: Some Lingering Problems

- Challenging to Scale Up to Modern Data & Storage
- Rigid Schemas
- ACID properties limit DBMS performance

NoSQL DBMSes – CSc 460 v1.1 (McCann) – p. 3/12

The ‘Veas’ of Big Data

The original three issues/characteristics of Big Data:

- 1.
- 2.
- 3.

Then others were added, including:

4. Veracity (‘dirty’ data must be cleaned)
5. Value (of what use is the data?)
6. Validity (how relevant is the data for our needs?)
7. Volatility (for how long is the data useful?)

NoSQL DBMSes – CSc 460 v1.1 (McCann) – p. 4/12

DBMS Support for the 'Vees'

Big-data DBMSes usually need to:

- Accept/Store/Process lots of data in real time
- Scale to larger workloads with commodity hardware
- Allow DB schemas to be flexible
- Relax enforcement of ACID properties

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 5/12

What is a 'NoSQL' DBMS?

First: 'NoSQL' is a terrible name for such systems!

- Why? The ideas have little to do with SQL directly
- Thus, some people now say it means 'Not (only) SQL'; still not useful
- Maybe 'Beyond Relational'? Branding needs work!

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 6/12

Four Varieties of NoSQL DBMSes (1 / 5)

1. Attribute–Value Systems (a.k.a., Key–Value)

- In programming, also known as dictionaries / hash tables
- Are *opaque* — structure of values isn't a concern
-
-

NoSQL DBMSes – CSc 460 v1.1 (McCann) – p. 7/12

Four Varieties of NoSQL DBMSes (2 / 5)

2. Wide–Column Stores

- A *wide–column* is a group of attribute–value pairs describing one entity
- Unlike a relation, each group can have different collections of pairs
- Like a relation, each group can have a row id (a.k.a., a key)
- By contrast, a *column store* stores data by columns (often partitioned) instead of by rows

NoSQL DBMSes – CSc 460 v1.1 (McCann) – p. 8/12

Four Varieties of NoSQL DBMSes (3 / 5)

3. Document Systems

- Are based on attribute–value systems, as are wide–column stores
- Typically use XML (eXtensible Markup Language), JSON (Javascript Object Notation), or BSON (Binary JSON) for data storage
- Motivating example: Your resume has your name, address, list of schools attended, current employer, past employers, etc.
 - An RDBMS would store these items in separate tables:
 -

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 9/12

Four Varieties of NoSQL DBMSes (4 / 5)

3. Document Systems (continued)

Example(s): A JSON representation of resume info:

```
{ "name" : "Amy",
  "address" : "123 Main Street",
  "degree" : [ "M.S.", "B.S." ],
  "experience" : [ { "position" : "tutor",
                    "employer" : "TUSD" },
                  { "position" : "volunteer",
                    "employer" : "Community Food Bank" }
                ]
}
```

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 10/12

Four Varieties of NoSQL DBMSes (5 / 5)

4. Graph Systems

- Use graph data structures to store data and relationships
 - Vertices: Hold the data
 - Edges: Are the relationships
- Representative Applications:
 -
 -
 -

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 11/12

Closing Remarks

- Clearly, NoSQL is a large umbrella
- Many NoSQL systems still offer:
 -
 -
- Beware the marketing hype!
 -
 -

'NoSQL' DBMSes – CSc 460 v1.1 (McCann) – p. 12/12