

CC5212-1

PROCESAMIENTO MASIVO DE DATOS

OTOÑO 2016

Lecture 11: NoSQL II

Aidan Hogan

aidhog@gmail.com

RECAP: NOSQL

NoSQL



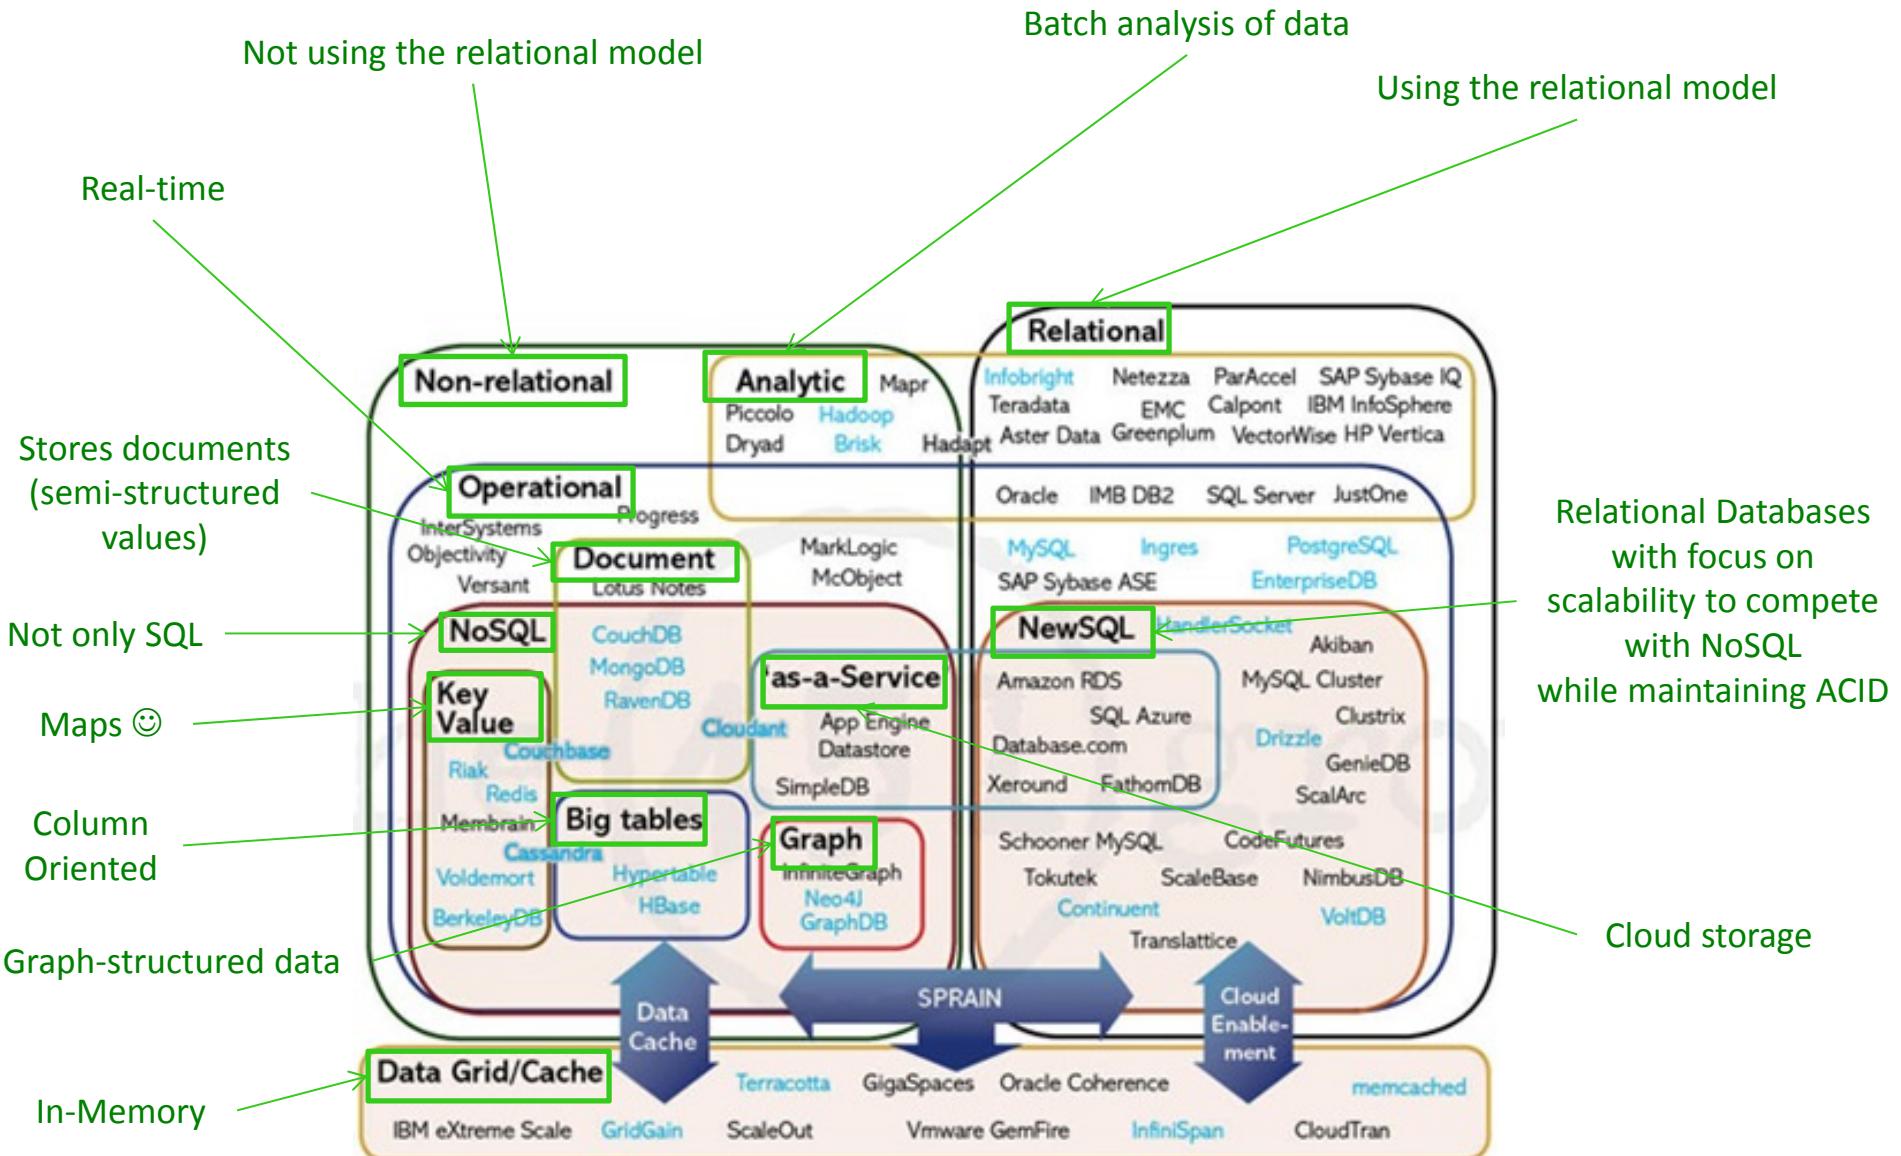
NoSQL

NoSQL vs. Relational Databases

What are the big differences between relational databases and NoSQL systems?

What are the trade-offs?

The Database Landscape



RECAP: KEY–VALUE

Key–Value = a Distributed Map

Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201}
...	...

Amazon Dynamo(DB): Model

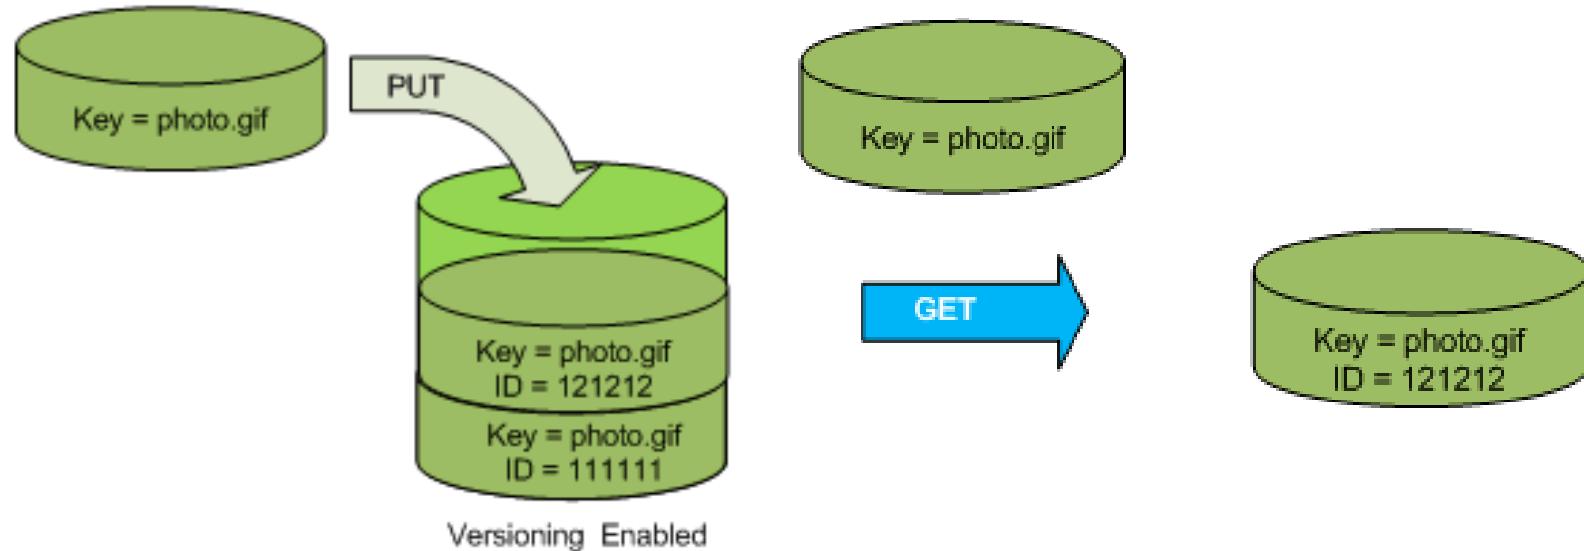
- Named table with primary key and a value

Countries	
Primary Key	Value
Afghanistan	capital:Kabul,continent:Asia,pop:31108077#2011
Albania	capital:Tirana,continent:Europe,pop:3011405#2013
...	...

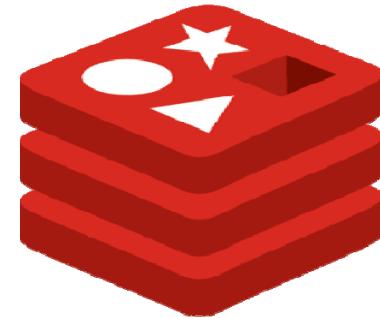
Cities	
Primary Key	Value
Kabul	country:Afghanistan,pop:3476000#2013
Tirana	country:Albania,pop:3011405#2013
...	...

Amazon Dynamo(DB): Object Versioning

- Object Versioning (per bucket)
 - PUT doesn't overwrite: pushes version
 - GET returns most recent version



Other Key–Value Stores



redis

RECAP: DOCUMENT STORES

Key–Value Stores: Values are Documents

Key	Value
country:Afghanistan	<country> <capital>city:Kabul</capital> <continent>Asia</continent> <population> <value>31108077</value> <year>2011</year> </population> </country>
...	...

- Document-type depends on store
 - XML, JSON, Blobs, Natural language
- Operators for documents
 - e.g., filtering, inv. indexing, XML/JSON querying, etc.

MongoDB: JSON Based

Key	Value (Document)
6ads786a5a9	{ “_id” : ObjectId(“6ads786a5a9”), “name” : “Afghanistan”, “capital”: “Kabul”, “continent” : “Asia”, “population” : { “value” : 31108077, “year” : 2011 } }
...	...



- Can invoke Javascript over the JSON objects
- Document fields can be indexed

```
db.inventory.find({ continent: { $in: [ 'Asia', 'Europe' ] } })
```

Document Stores



Couchbase



TABLULAR / COLUMN FAMILY

Key–Value = a Distributed Map

Countries	
Primary Key	Value
Afghanistan	capital:Kabul,continent:Asia,pop:31108077#2011
Albania	capital:Tirana,continent:Europe,pop:3011405#2013
...	...

Tabular = Multi-dimensional Maps

Countries				
Primary Key	capital	continent	pop-value	pop-year
Afghanistan	Kabul	Asia	31108077	2011
Albania	Tirana	Europe	3011405	2013
...

Bigtable: The Original Whitepaper

Why did they write another paper?
MapReduce solves everything, right?

MapReduce
authors

Bigtable: A Distributed Storage System for Structured Data

Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach
Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber

{fay,jeff,sanjay,wilsonh,kerr,m3b,tushar,fikes,gruber}@google.com

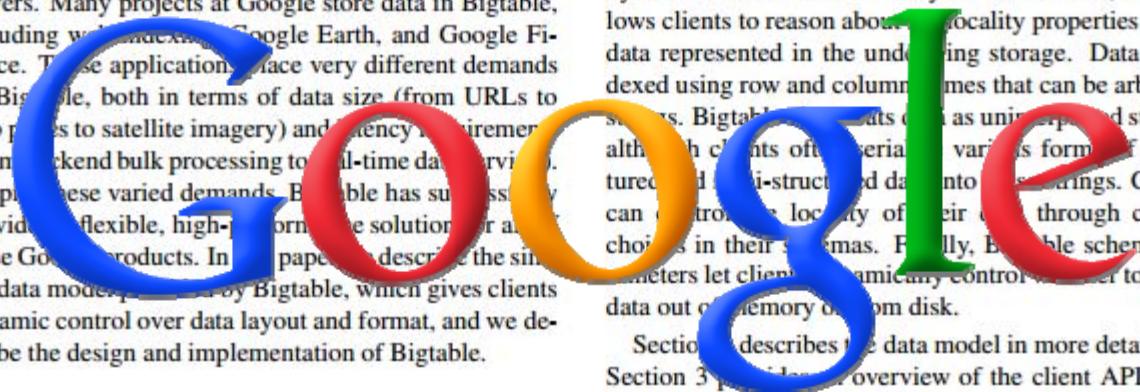
Google, Inc.

Abstract

Bigtable is a distributed storage system for managing structured data that is designed to scale to a very large size: petabytes of data across thousands of commodity servers. Many projects at Google store data in Bigtable, including web indexing, Google Earth, and Google Finance. These applications place very different demands on Bigtable, both in terms of data size (from URLs to web pages to satellite imagery) and latency requirements (from weekend bulk processing to full-time data services). Despite these varied demands, Bigtable has successfully provided a flexible, high-performance solution for all of these Google products. In this paper we describe the simple data model provided by Bigtable, which gives clients dynamic control over data layout and format, and we describe the design and implementation of Bigtable.

achieved scalability and high performance, but Bigtable provides a different interface than such systems. Bigtable does not support a full relational data model; instead, it provides clients with a simple data model that supports dynamic control over data layout and format, and allows clients to reason about the locality properties of the data represented in the underlying storage. Data is indexed using row and column names that can be arbitrary strings. Bigtable stores data as uninterpreted strings, although clients often serialize various forms of structured and semi-structured data into these strings. Clients can control the locality of their data through careful choices in their schemas. Finally, Bigtable schema parameters let clients dynamically control whether to serve data out of memory or from disk.

Section 2 describes the data model in more detail, and Section 3 provides an overview of the client API. Sec-



Bigtable used for ...

A screenshot of a Google search results page for the query "chikoo". The results include a mix of text and image links. A red arrow points to the "Search" link in the sidebar, while a yellow arrow points to the "Images" link. A green arrow points to the first image result, which shows four small images of dogs.

Google earth

chikoo

Search

Everything

Images

Maps

Videos

News

Shopping

Recipes

More

Mountain View, CA

Change location

All results

Sites with images

Related searches

More search tools

Manilkara zapota - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/Manilkara_zapota

Sapodilla is known as chikoo ("चिकू," or chiku, "ચિકૂ,") in India and Pakistan and sapota in some parts of India (Tamil Nadu, Kerala, Karnataka, Andhra ...)

Description - Other names - See also - References

You've visited this page 3 times. Last visit 12/4/11

Images for chikoo - Report images

Chikoo - a simple file organizer for the Mac
codingturtle.com/chikoo/

Chikoo. a simple file organizer for the Mac. Download (30-day trial). Version 0.9.1 ...



orkut by Google™



Bigtable: Data Model

“a sparse, distributed, persistent, multi-dimensional, sorted map.”

- *sparse*: not all values form a dense square
- *distributed*: lots of machines
- *persistent*: disk storage (GFS)
- *multi-dimensional*: values with columns
- *sorted*: sorting lexicographically by row key
- *map*: look up a key, get a value

Bigtable: in a nutshell

$(\text{row}, \text{column}, \text{time}) \rightarrow \text{value}$

- **row**: a row id string
 - e.g., “Afghanistan”
- **column**: a column name string
 - e.g., “pop-value”
- **time**: an integer (64-bit) version time-stamp
 - e.g., 18545664
- **value**: the element of the cell
 - e.g., “31120978”

Bigtable: in a nutshell

(row, column, time) → value

(Afghanistan, pop-value, t₄) → 31108077

	Primary Key	capital	continent	pop-value	pop-year				
Afghanistan	t ₁	Kabul	t ₁	Asia	<td>t₁</td> <td>31143292</td> <td>t₁</td> <td>2009</td>	t ₁	31143292	t ₁	2009
Albania	t ₁	Tirana	t ₁	Europe	<td>t₂</td> <td>31120978</td> <td>t₄</td> <td>2011</td>	t ₂	31120978	t ₄	2011
...	<td>t₄</td> <td>31108077</td> <td>t₄</td> <td>2010</td>	t ₄	31108077	t ₄	2010
				<td>t₁</td> <td>2912380</td> <td>t₁</td> <td>2013</td>	t ₁	2912380	t ₁	2013	
				<td>t₃</td> <td>3011405</td> <td>t₃</td> <td>...</td>	t ₃	3011405	t ₃	...	
				

Bigtable: Sorted Keys

	Primary Key	capital		pop-value		pop-year	
S O R T E D 	Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
				t ₂	31120978		
				t ₄	31108077	t ₄	2011
	Asia:Azerbaijan

	Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
				t ₃	3011405	t ₃	2013
	Europe:Andorra

Benefits of sorted keys vs.
hashed keys?

Bigtable: Tablets

Primary Key	capital		pop-value		pop-year	
Asia:Afghanistan	t_1	Kabul	t_1	31143292	t_1	2009
			t_2	31120978		
			t_4	31108077	t_4	2011
Asia:Azerbaijan
...
Europe:Albania	t_1	Tirana	t_1	2912380	t_1	2010
			t_3	3011405		
		
Europe:Andorra
...

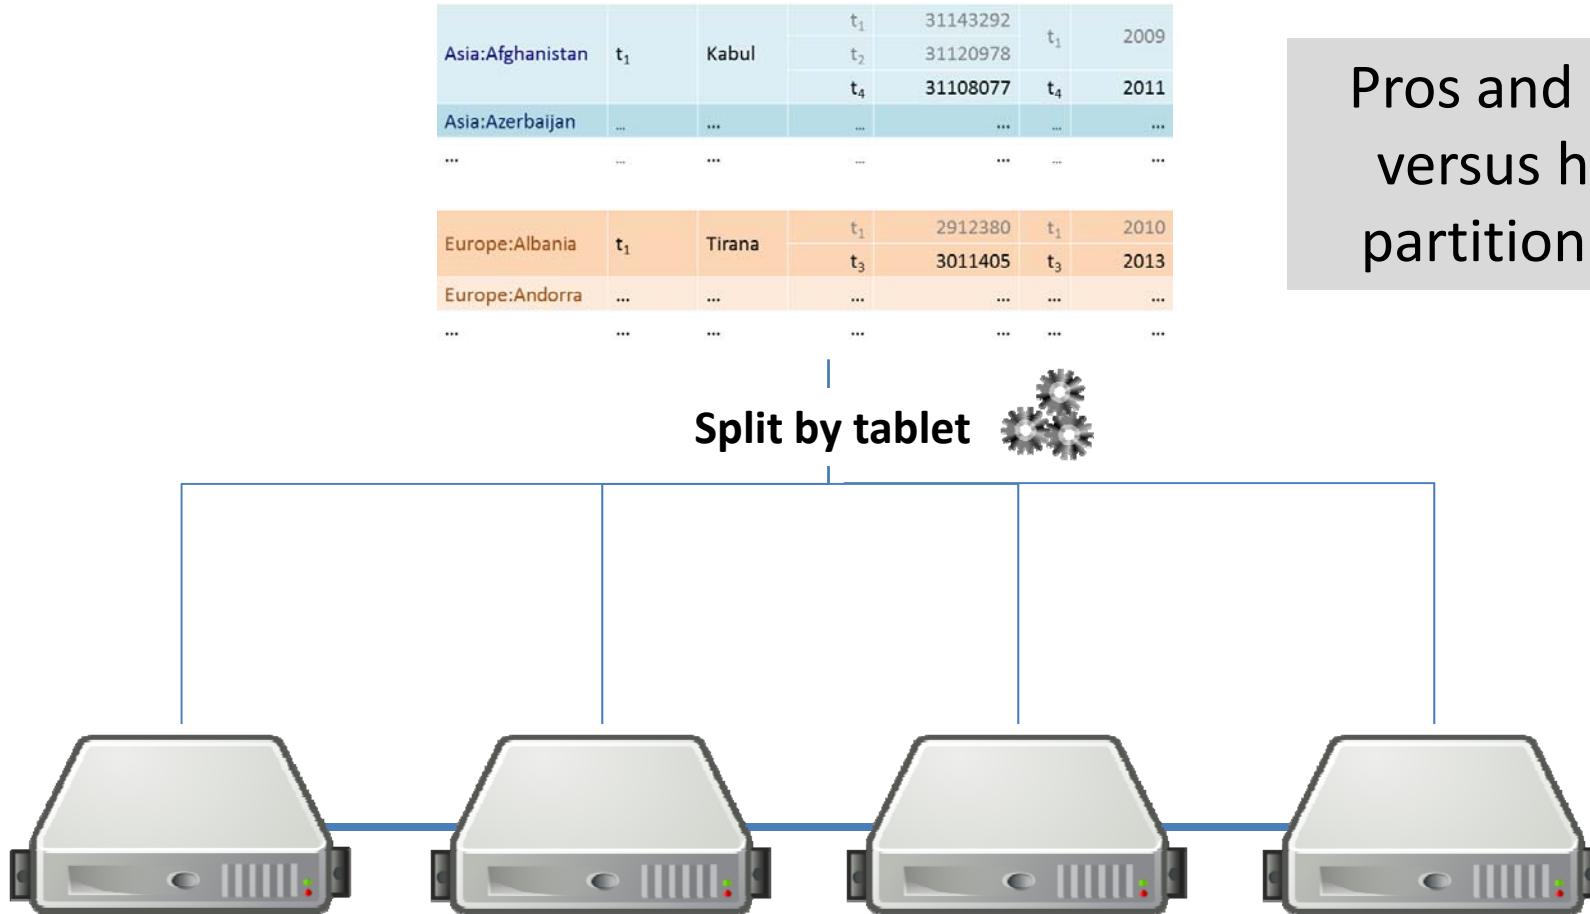


Take advantage of locality of processing!

A real-world example of locality/sorting

Primary Key	language	title		links
com.imdb	t ₁	en	t ₁	IMDb Home
			t ₂	IMDB - Movies
			t ₄	IMDb
com.imdb/title/tt2724064/	t ₁	en	t ₂	Sharknado
com.imdb/title/tt3062074/	t ₁	en	t ₂	Sharknado II
...
org.wikipedia	t ₁	multi	t ₁	Wikipedia
			t ₃	Wikipedia Home
org.wikipedia.ace	t ₁	ace	t ₁	Wikipèdia bahsa Acèh
...

Bigtable: Distribution



Pros and cons
versus hash
partitioning?

Split by tablet

Horizontal range partitioning

Bigtable: Column Families

Primary Key	pol:capital		demo:pop-value		demo:pop-year	
Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
			t ₂	31120978		
			t ₄	31108077	t ₄	2011
Asia:Azerbaijan
...
Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
			t ₃	3011405	t ₃	2013
Europe:Andorra
...

- Group logically similar columns together
 - Accessed efficiently together
 - Access-control and storage: column family level
 - If of same type, can be compressed

Bigtable: Versioning

- Similar to Apache Dynamo
 - Cell-level
 - 64-bit integer time stamps
 - Inserts push down current version
 - Lazy deletions / periodic garbage collection
 - Two options:
 - keep last n versions
 - keep versions newer than t time

Bigtable: SSTable Map Implementation

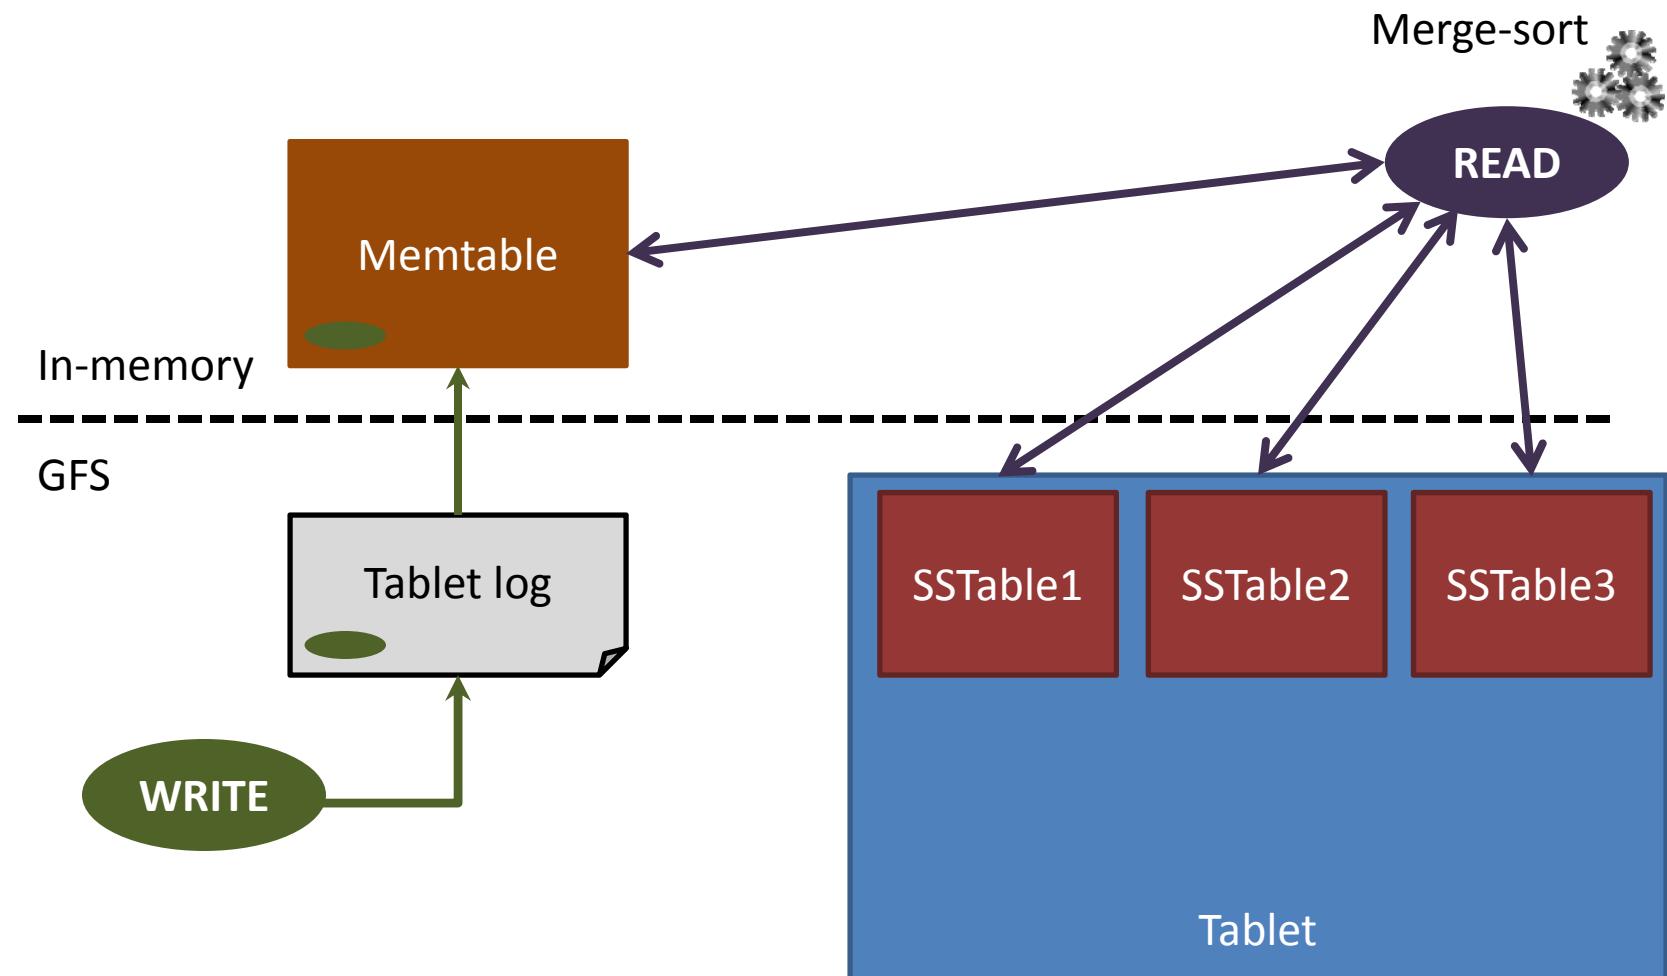
- 64k blocks (default) with index in footer (GFS)
- Index loaded into memory, allows for seeks
- Can be split or merged, as needed

How to handle writes?

		Primary Key	pol:capital	demo:pop-value	demo:pop-year		
0				t ₁	31143292	t ₁	2009
65536	Asia:Afghanistan	t ₁	Kabul	t ₂	31120978	t ₄	2011
				t ₄	31108077		
	Asia:Azerbaijan
...	
65536		Asia:Japan
...		Asia:Jordan
Index:		Block 0 / Offset 0 / Asia:Afghanistan					
		Block 1 / Offset 65536 / Asia: Japan					

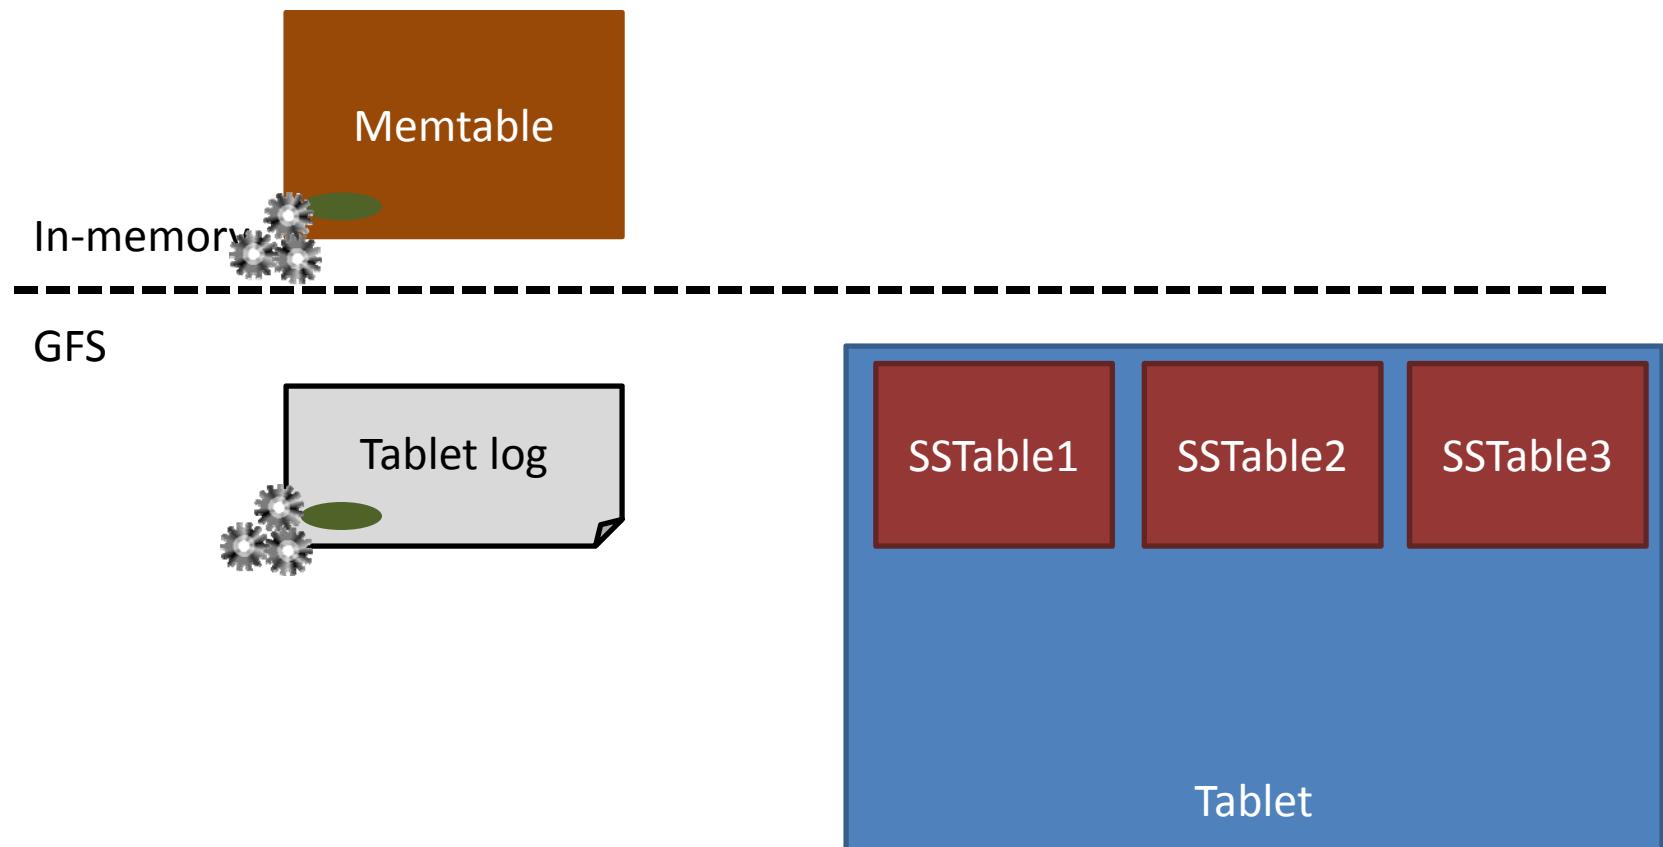
Bigtable: Buffered/Batched Writes

What's the danger here?



Bigtable: Redo Log

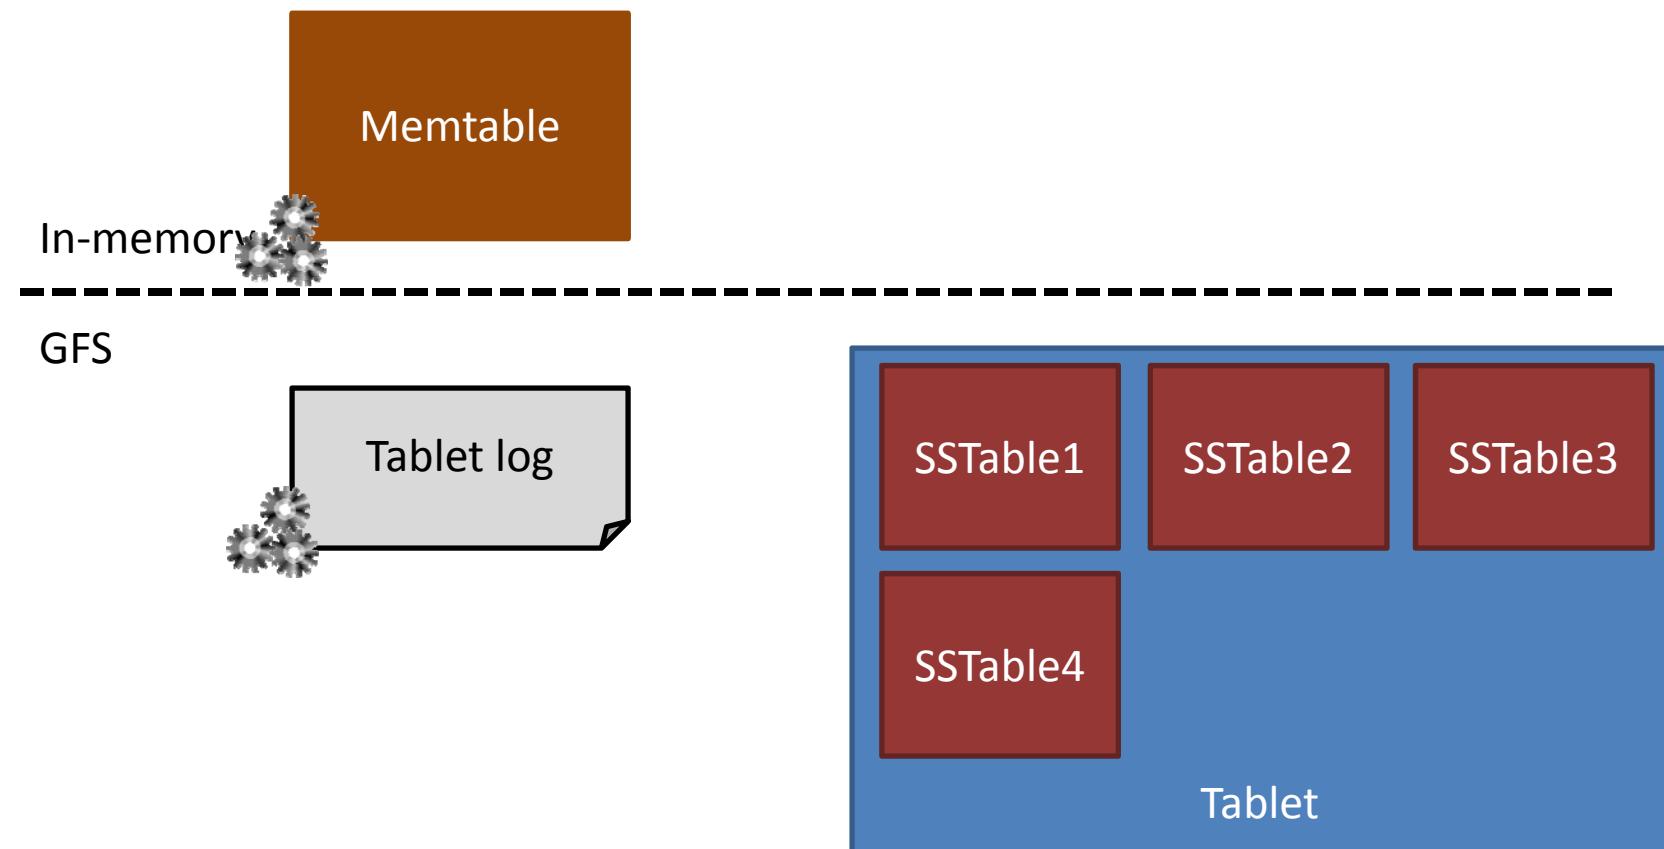
- If machine fails, Memtable redone from log



Bigtable: **Minor** Compaction

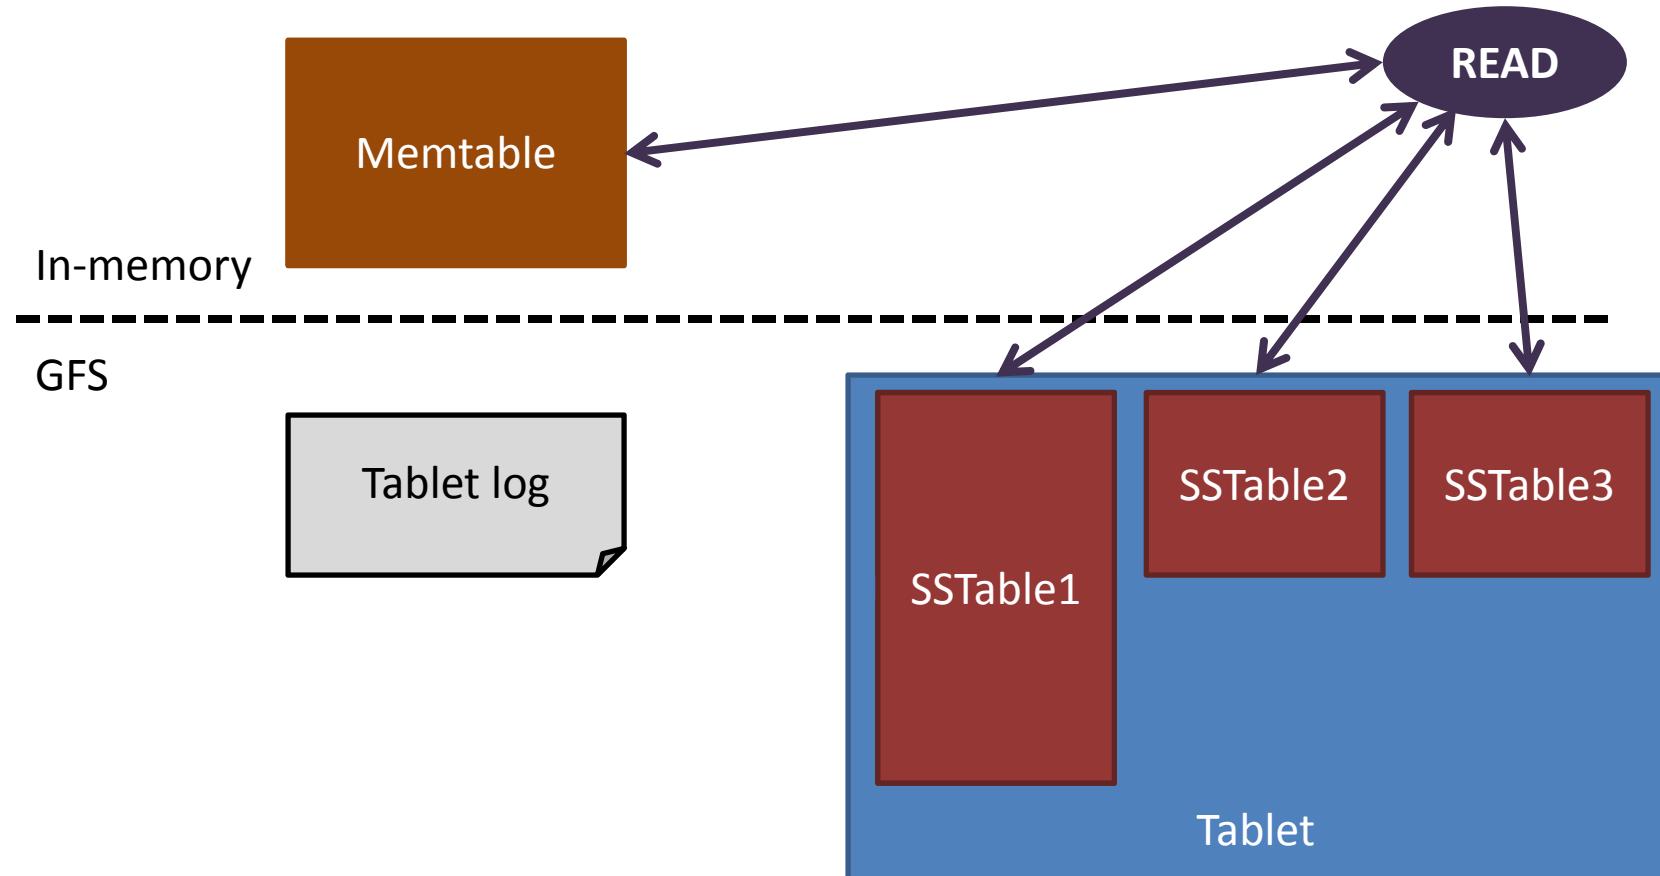
- When full, write Memtable as SSTable

What's the problem here?



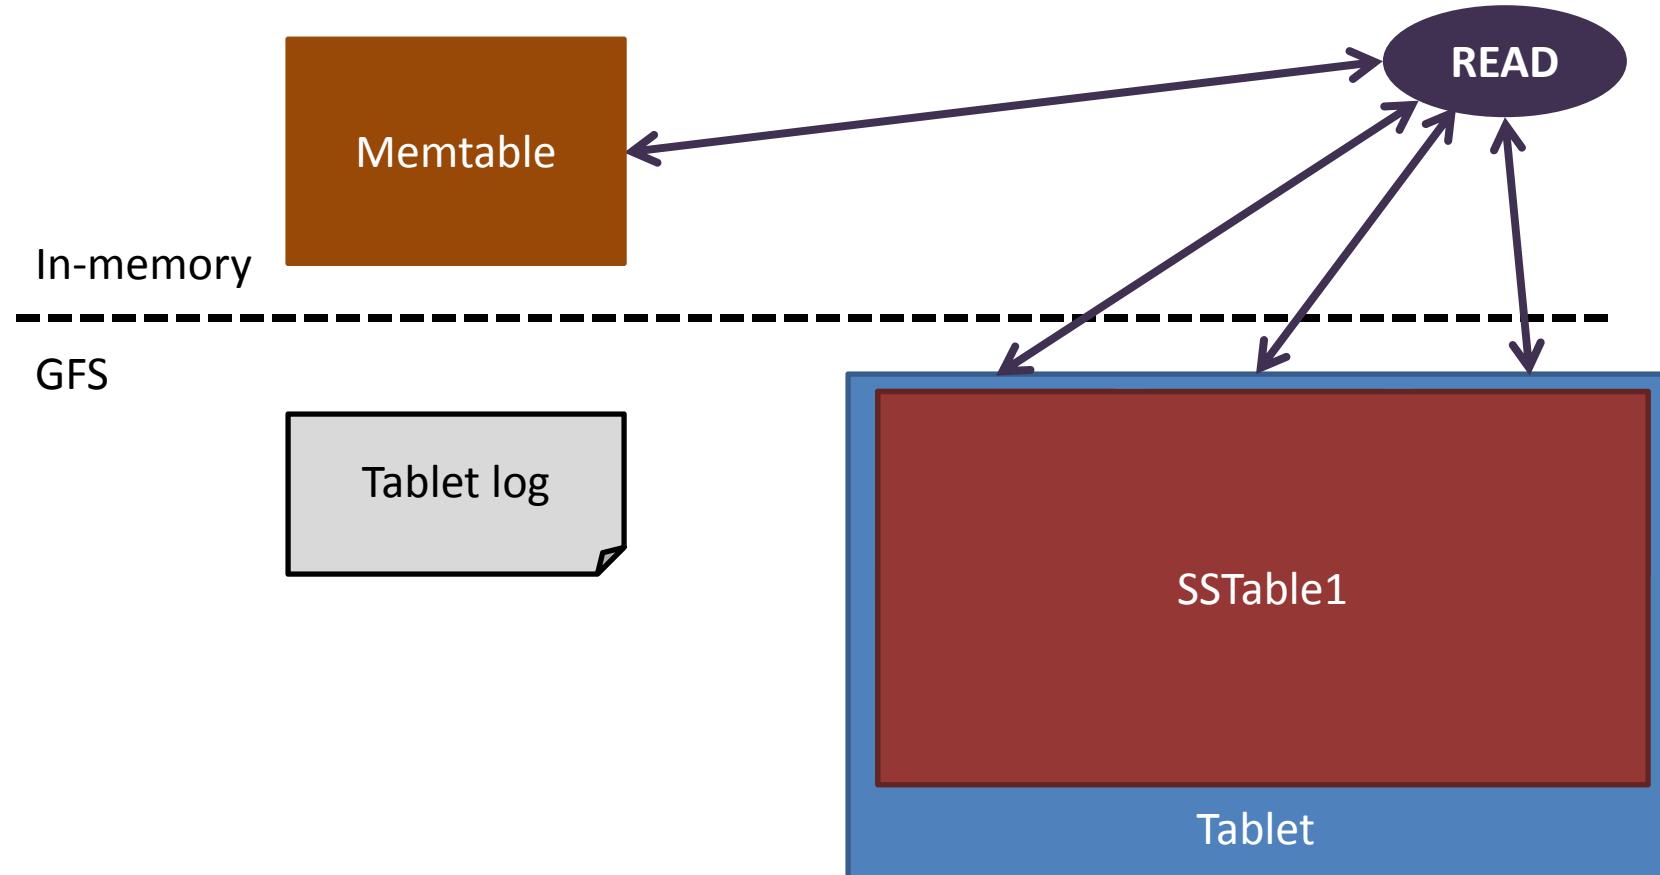
Bigtable: **Merge** Compaction

- Merge **some of the SSTables** (and the Memtable)



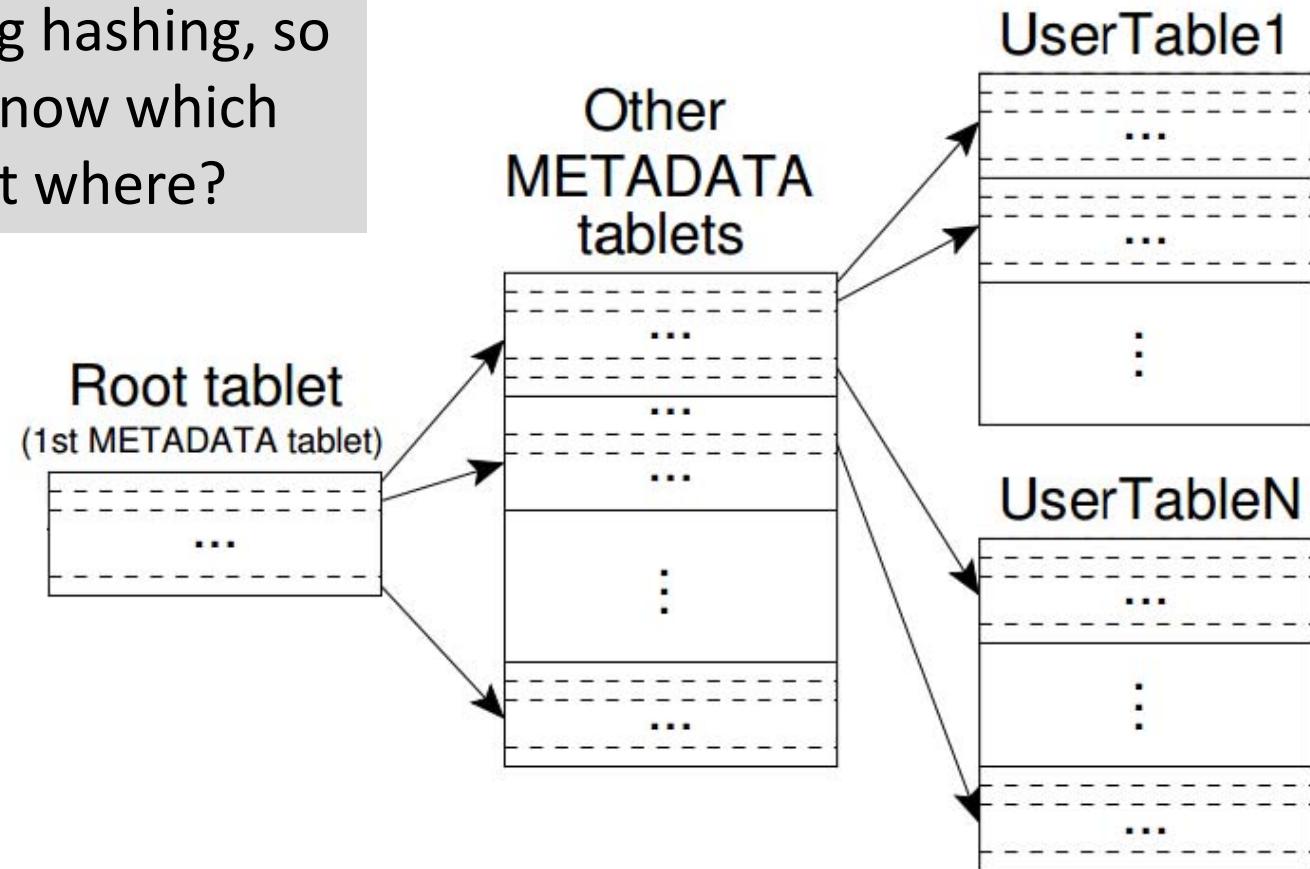
Bigtable: **Major** Compaction

- Merge **all** SSTables (and the Memtable)
- Makes reads more efficient!



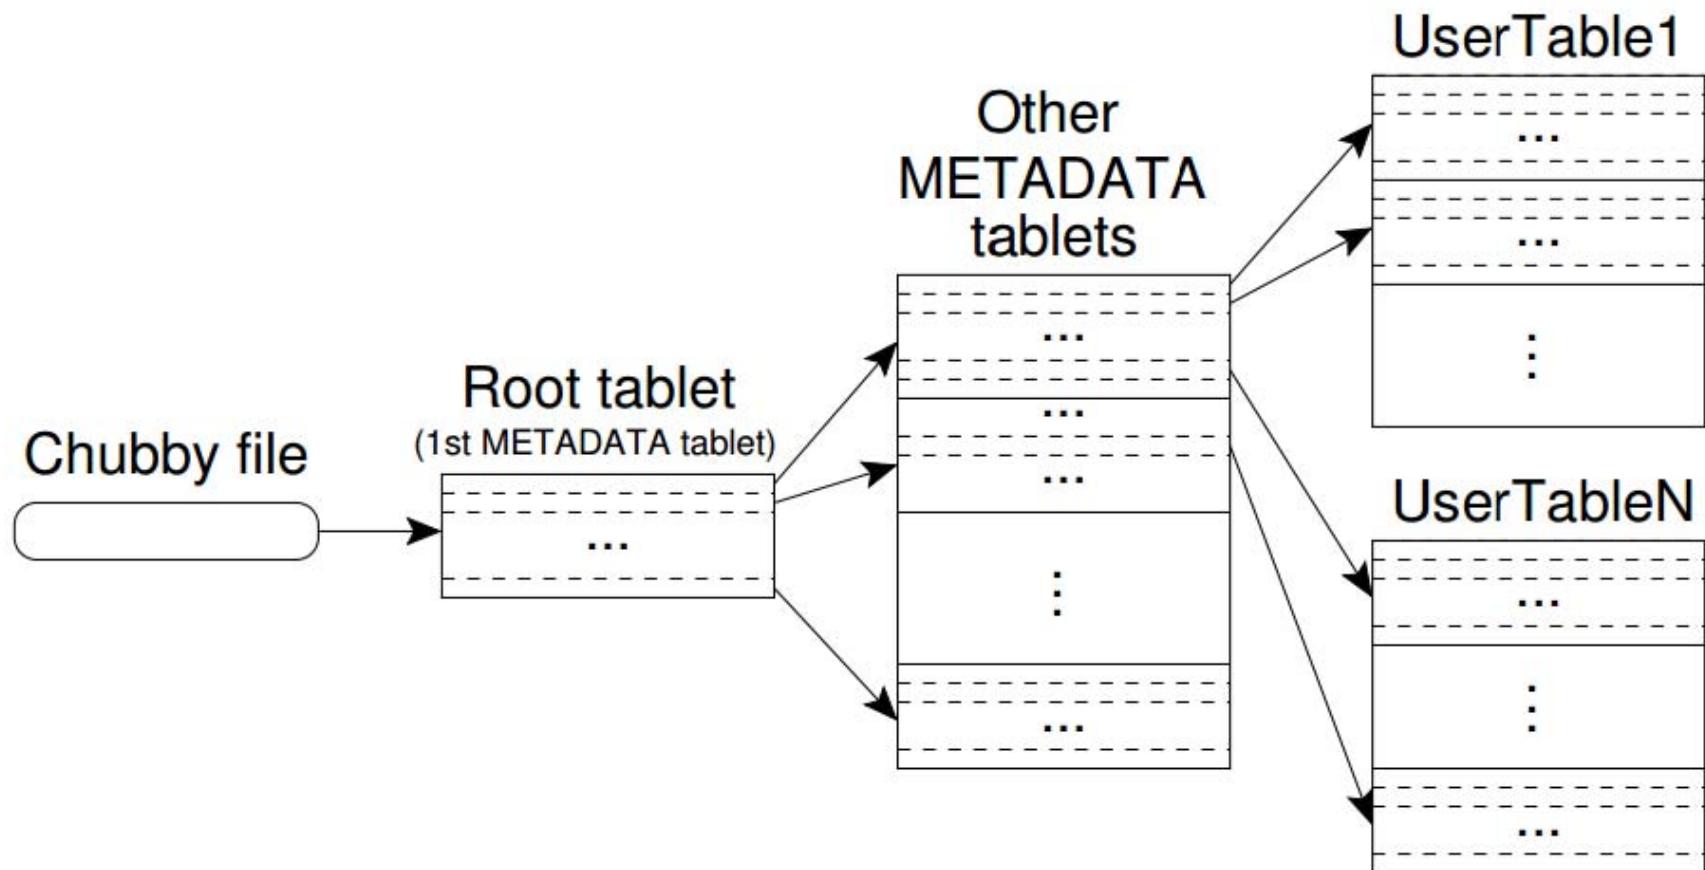
Bigtable: Hierarchical Structure

We're not using hashing, so
how do we know which
tablet went where?



How do we choose the location of the Root tablet?

Bigtable: Hierarchical Structure

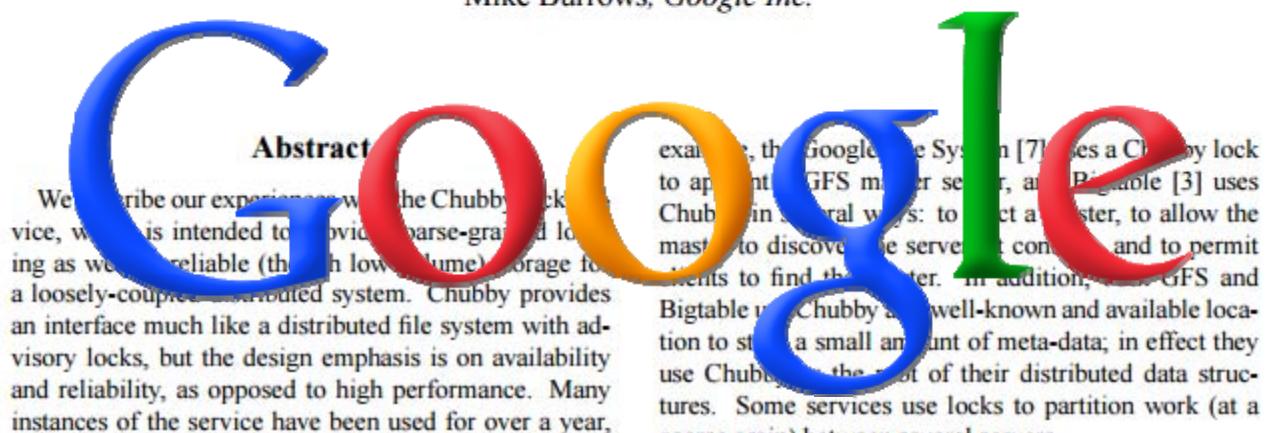


Bigtable: Consistency

- CHUBBY: Distributed consensus tool based on PAXOS
 - Quorum of five:
 - one master (chosen automatically) and four slaves
 - Co-ordinates distributed locks
 - Stores location of main “root tablet”
 - Holds other global state information (e.g., active servers)

The Chubby lock service for loosely-coupled distributed systems

Mike Burrows, Google Inc.



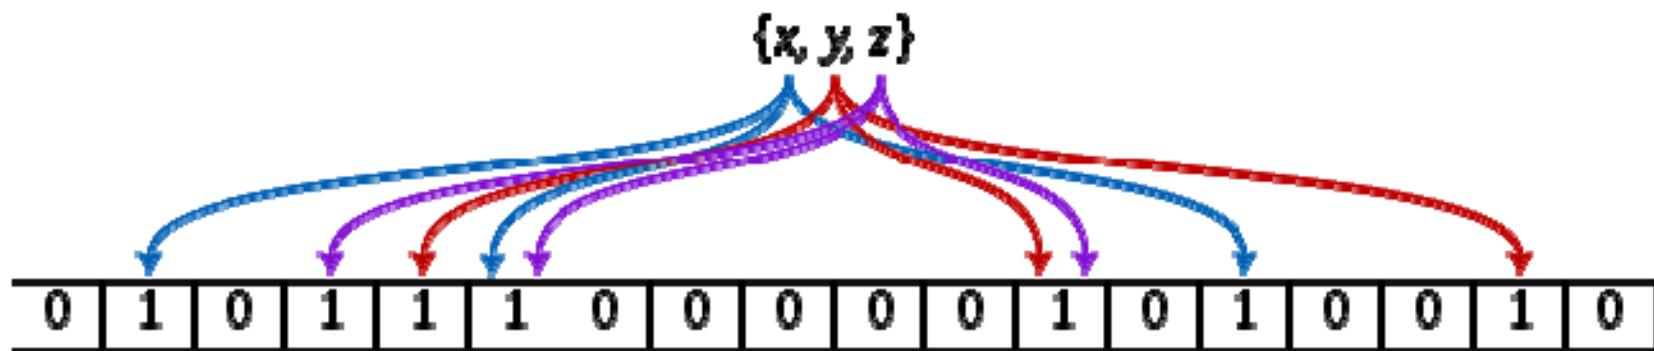
Bigtable: A Bunch of Other Things

- **Locality groups:** Group multiple column families together; assigned a separate SSTable
- **Select storage:** SSTables can be persistent or in-memory
- **Compression:** Applied on SSTable blocks; custom compression can be chosen
- **Caches:** SSTable-level and block-level
- **Bloom filters:** Find negatives cheaply ...

Aside: Bloom Filter

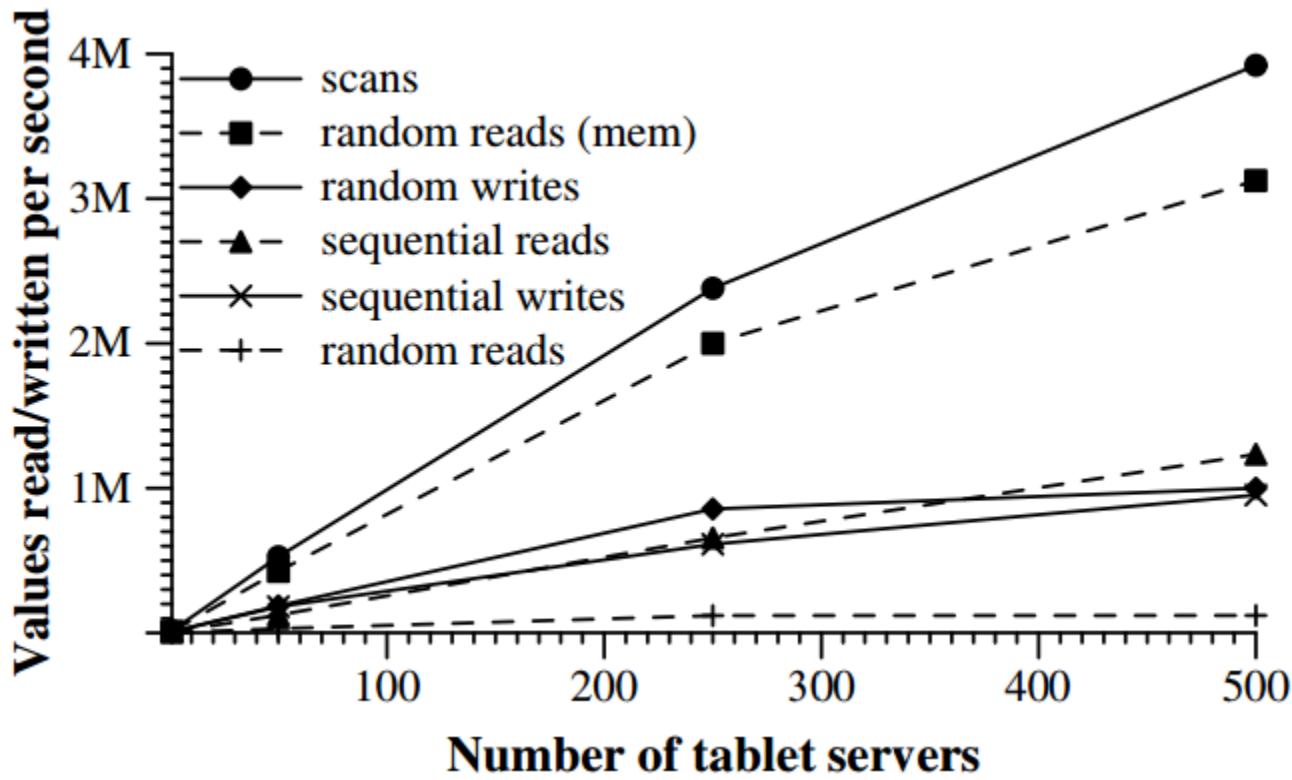
Reject “empty” queries using very little memory!

- Create a bit array of length m (init to 0’s)
- Create k hash functions that map an object to an index of m (even distribution)
- Index o : set $m[\text{hash}_1(o)], \dots, m[\text{hash}_k(o)]$ to 1
- Query o :
 - **any** $m[\text{hash}_1(o)], \dots, m[\text{hash}_k(o)]$ set to **0** = not indexed
 - **all** $m[\text{hash}_1(o)], \dots, m[\text{hash}_k(o)]$ set to **1** = **might** be indexed



Bigtable: an idea of performance [2006]

- Values are 1 Kb in size (blocks 64 Kb)
- Values are aggregate over all machines



Why are
random (disk)
reads so slow?

The read sizes are 1 kb,
but a different 64 kb
block must be sent over
the network (almost)
every time

Bigtable: an idea of performance [2006]

- Values are 1 Kb in size (blocks 64 Kb)
- Average values/second per server:

Experiment	# of Tablet Servers			
	1	50	250	500
random reads	1212	593	479	241
random reads (mem)	10811	8511	8000	6250
random writes	8850	3745	3425	2000
sequential reads	4425	2463	2625	2469
sequential writes	8547	3623	2451	1905
scans	15385	10526	9524	7843

- Adding more machines does add a cost!
- But overall performance does increase

Bigtable: examples in Google [2006]

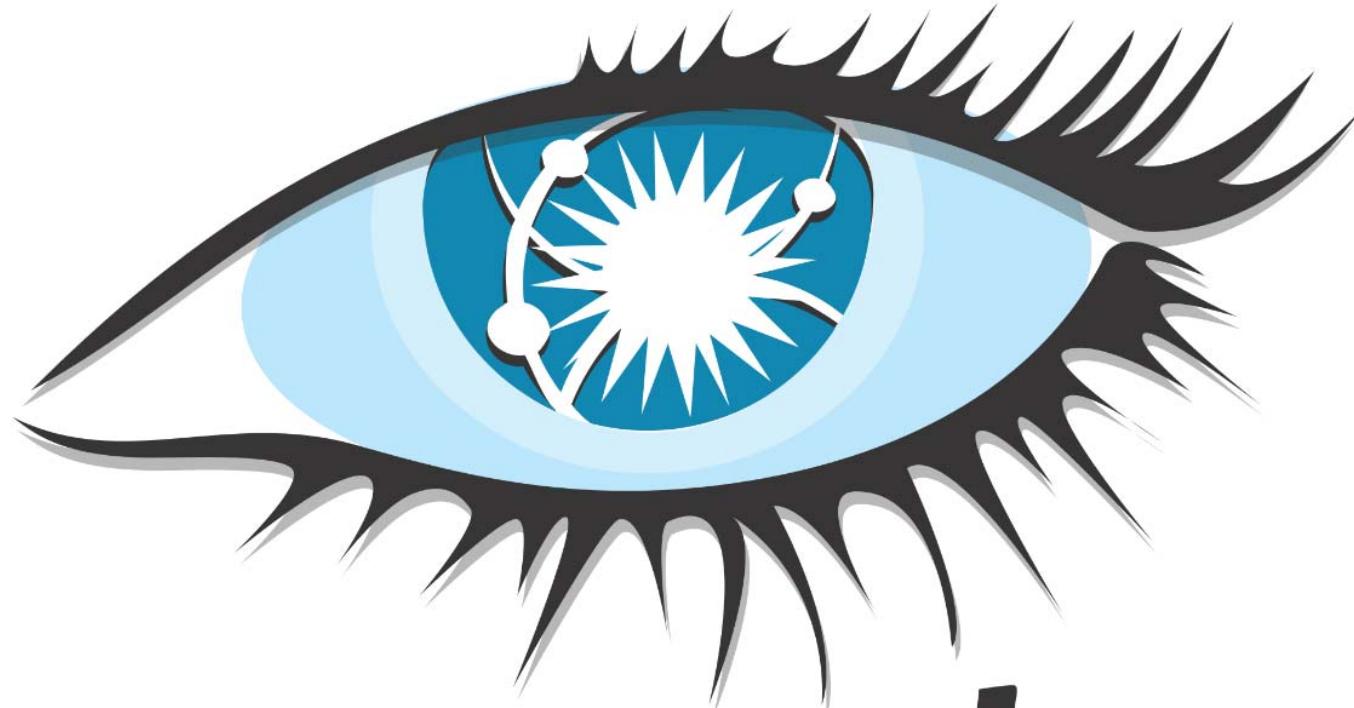
Project name	Table size (TB)	Compression ratio	# Cells (billions)	# Column Families	# Locality Groups	% in memory	Latency-sensitive?
<i>Crawl</i>	800	11%	1000	16	8	0%	No
<i>Crawl</i>	50	33%	200	2	2	0%	No
<i>Google Analytics</i>	20	29%	10	1	1	0%	Yes
<i>Google Analytics</i>	200	14%	80	1	1	0%	Yes
<i>Google Base</i>	2	31%	10	29	3	15%	Yes
<i>Google Earth</i>	0.5	64%	8	7	2	33%	Yes
<i>Google Earth</i>	70	–	9	8	3	0%	No
<i>Orkut</i>	9	–	0.9	8	5	1%	Yes
<i>Personalized Search</i>	4	47%	6	93	11	5%	Yes

Bigtable: Apache HBase



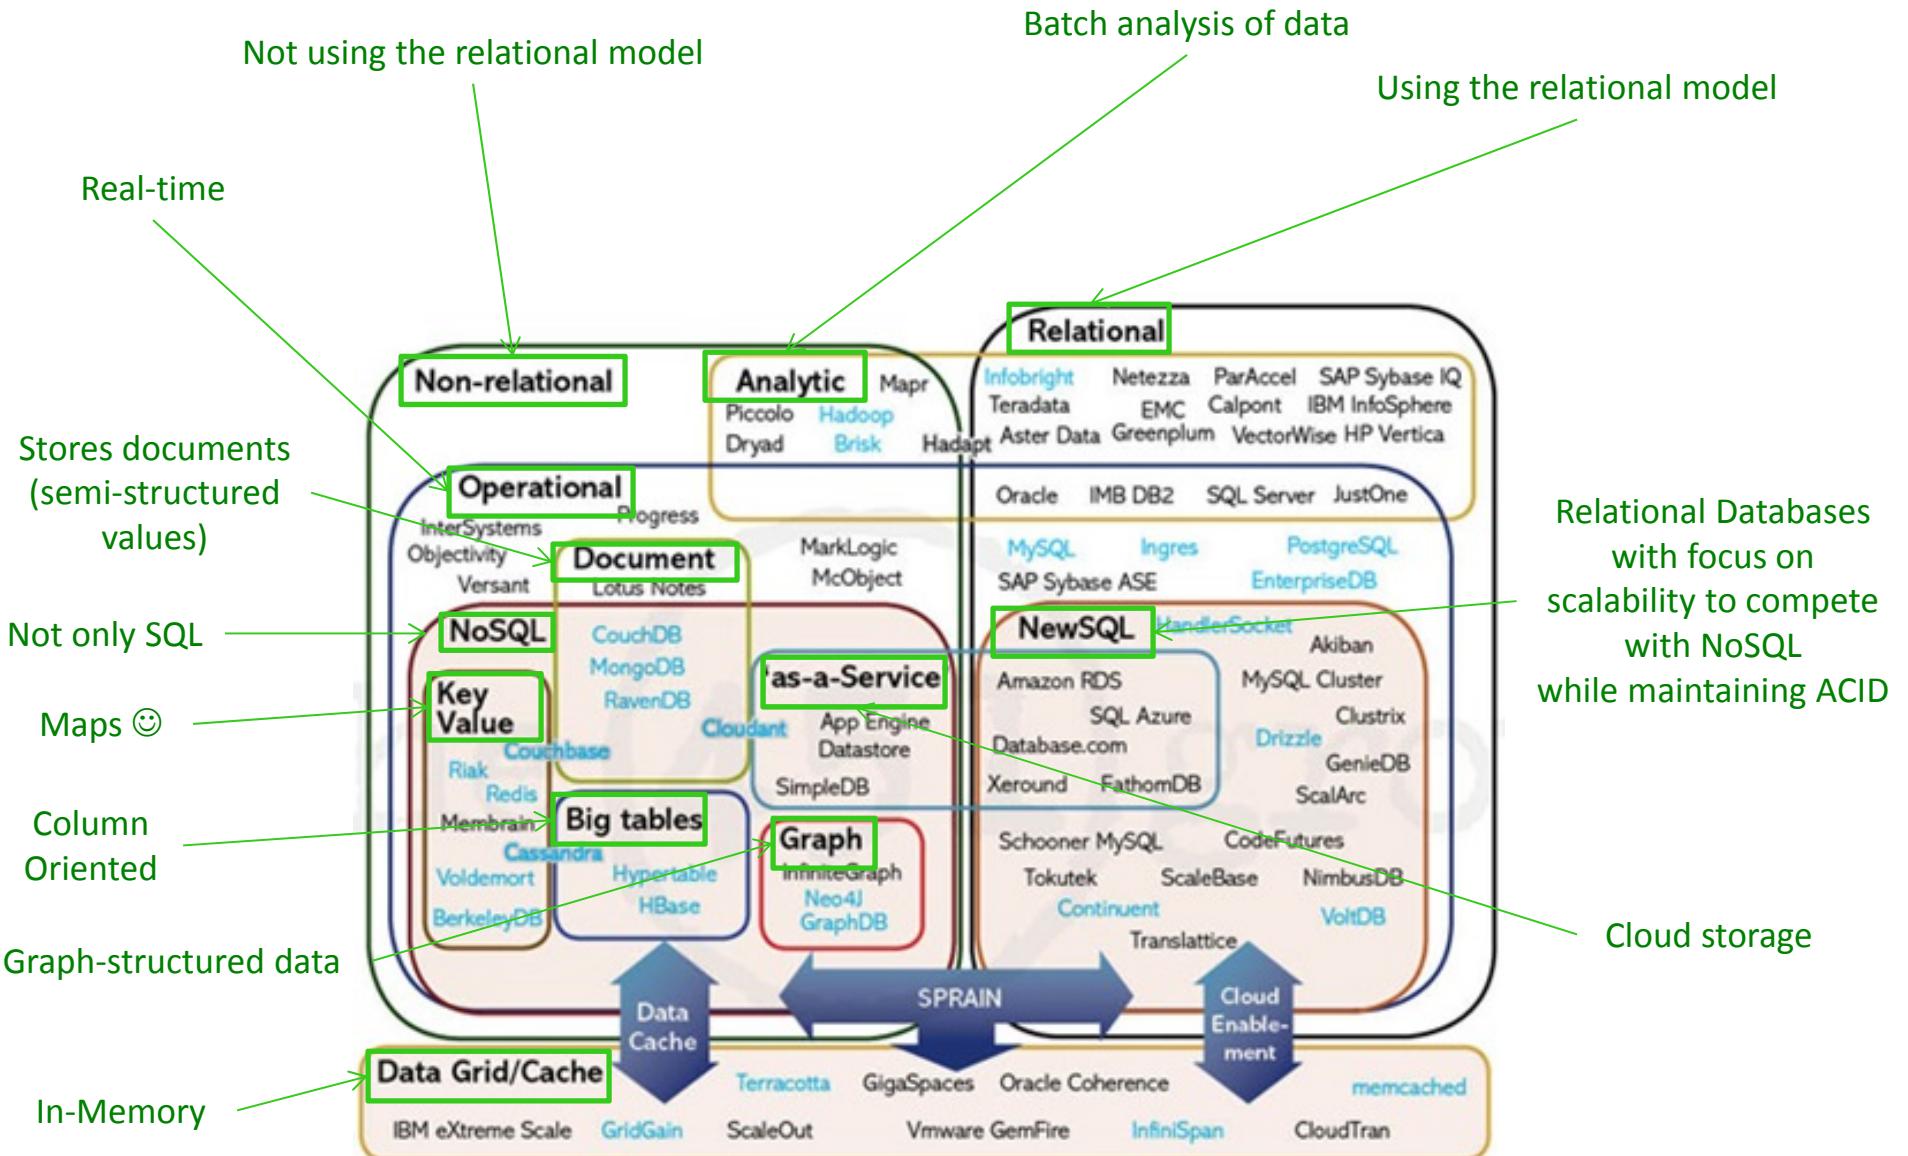
Open-source implementation of Bigtable ideas

Similar ideas in Cassandra

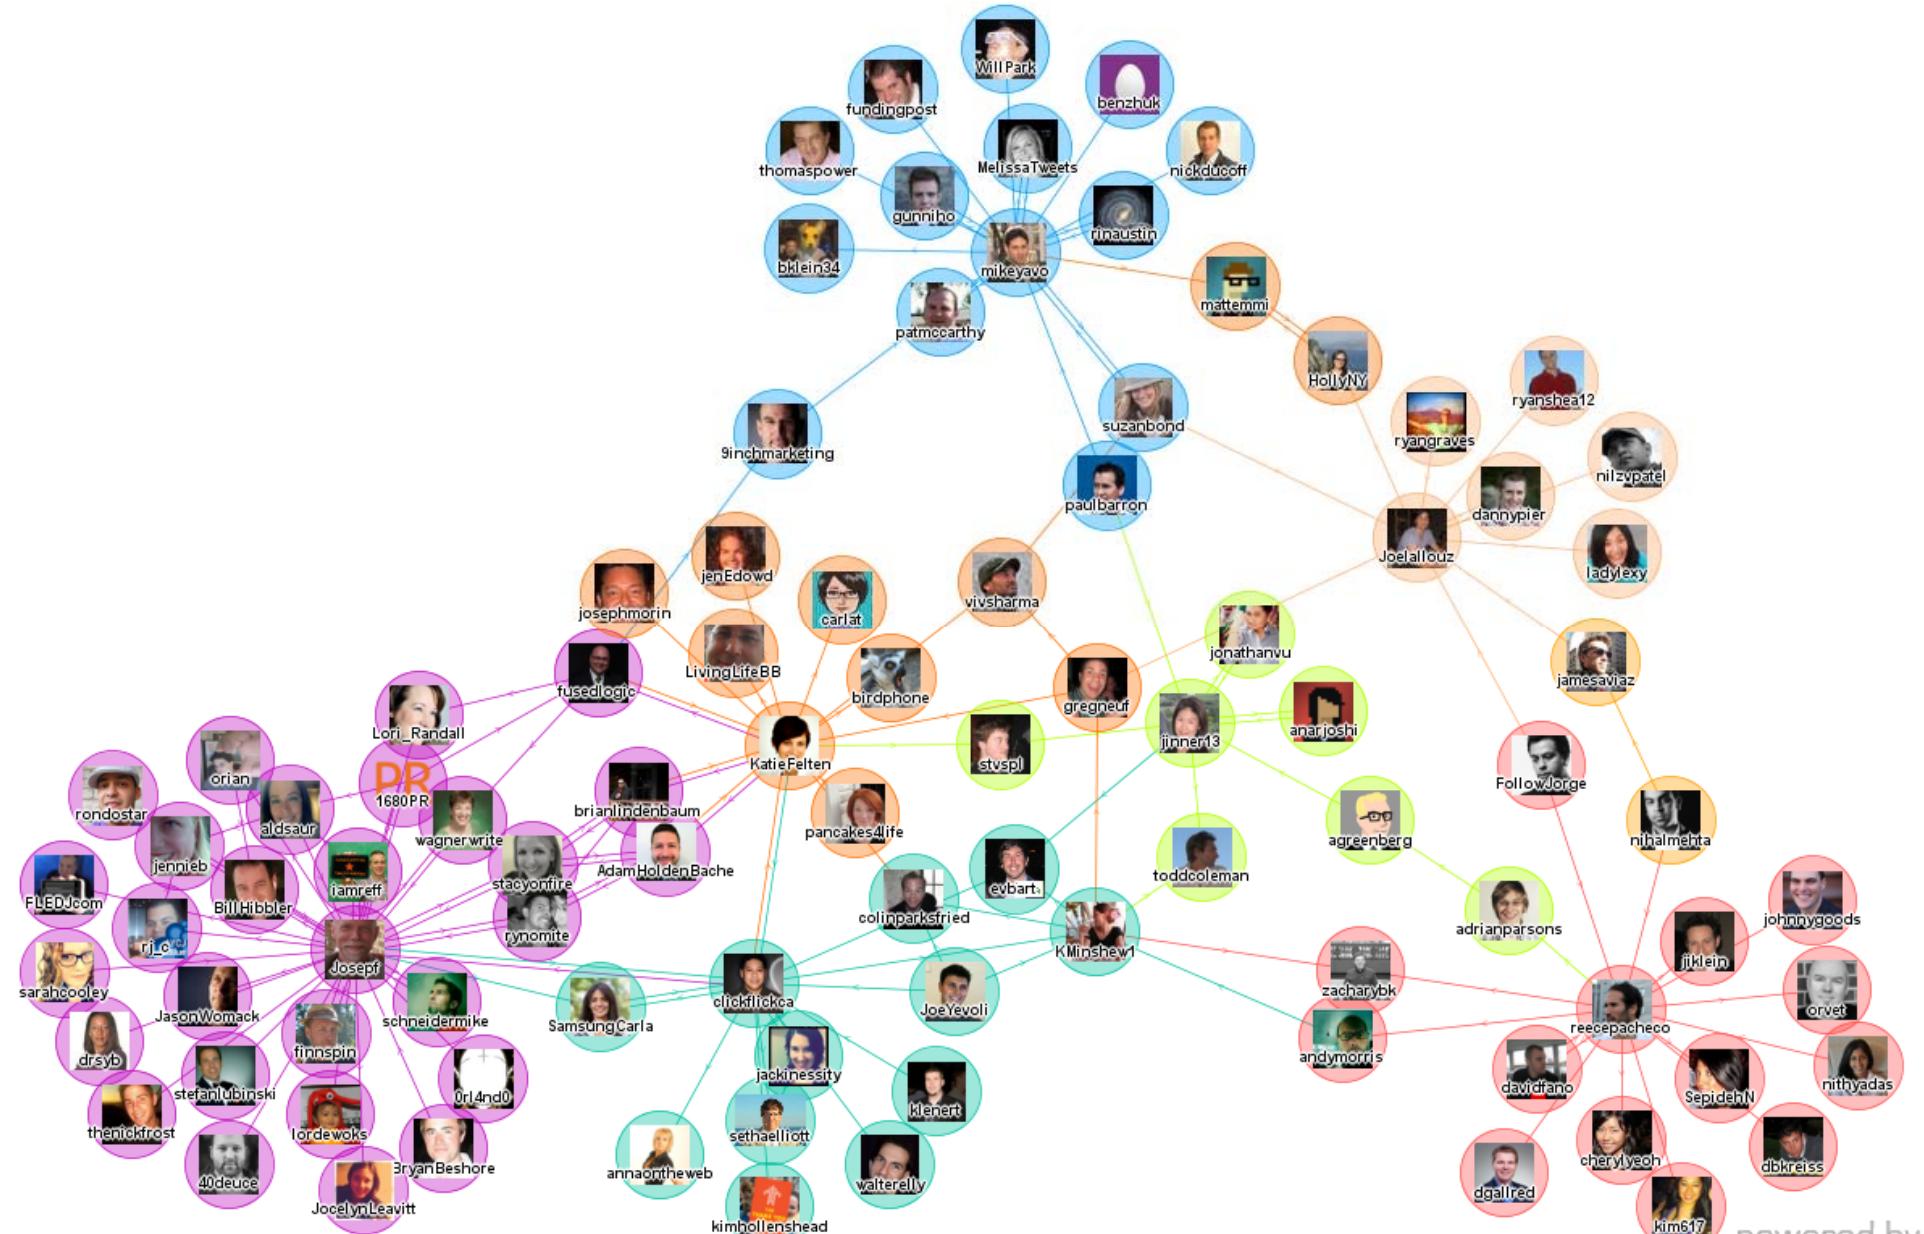


cassandra

The Database Landscape

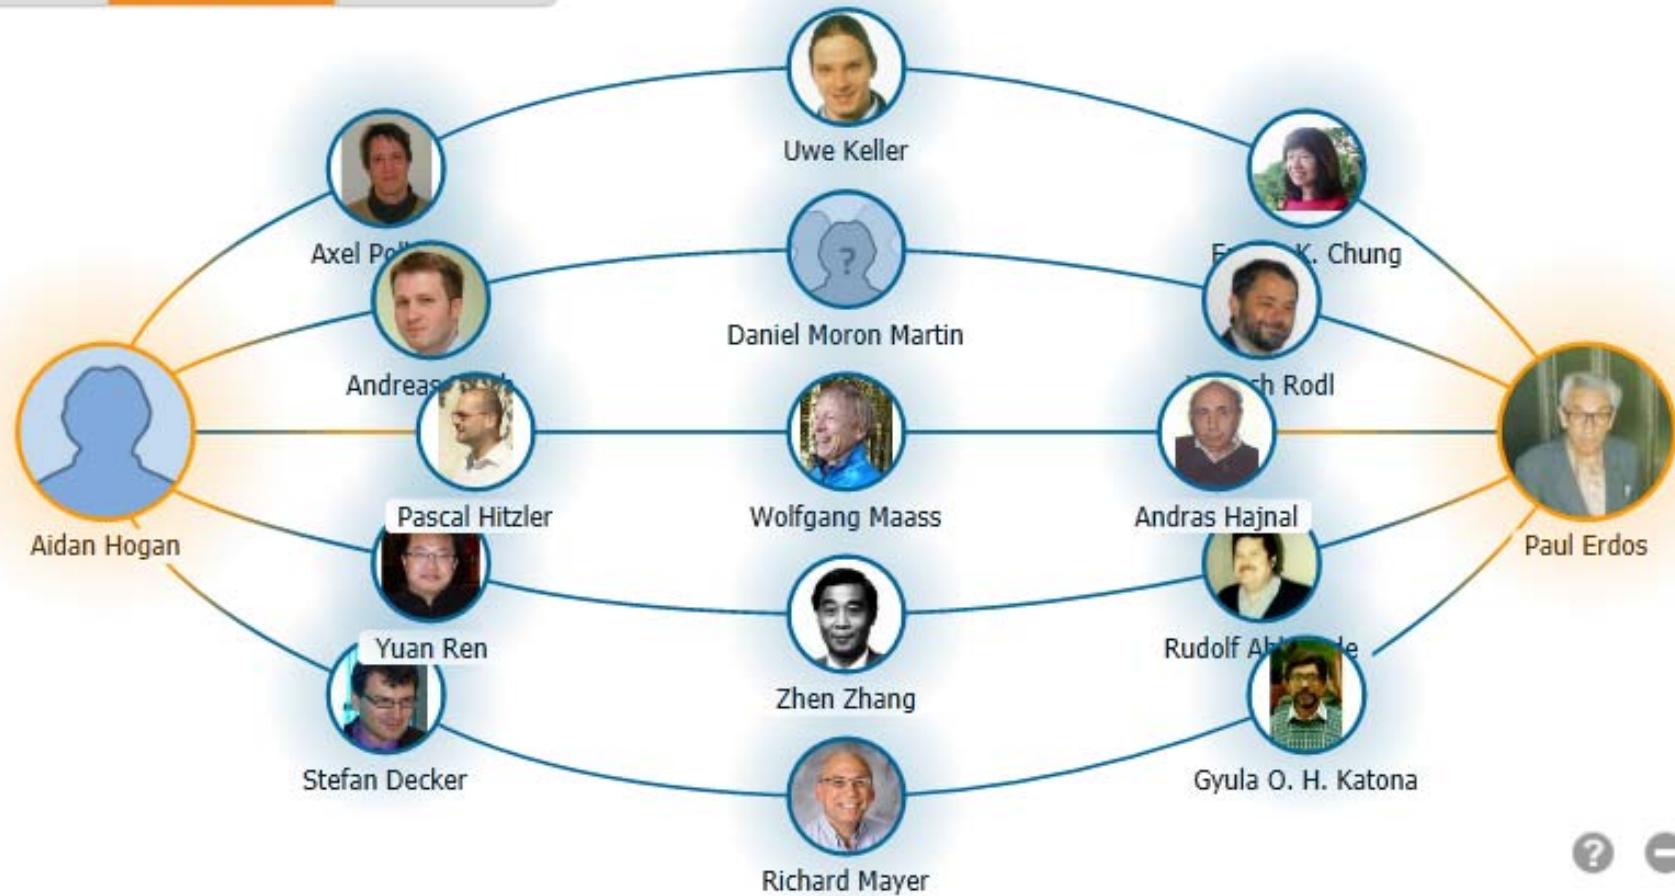


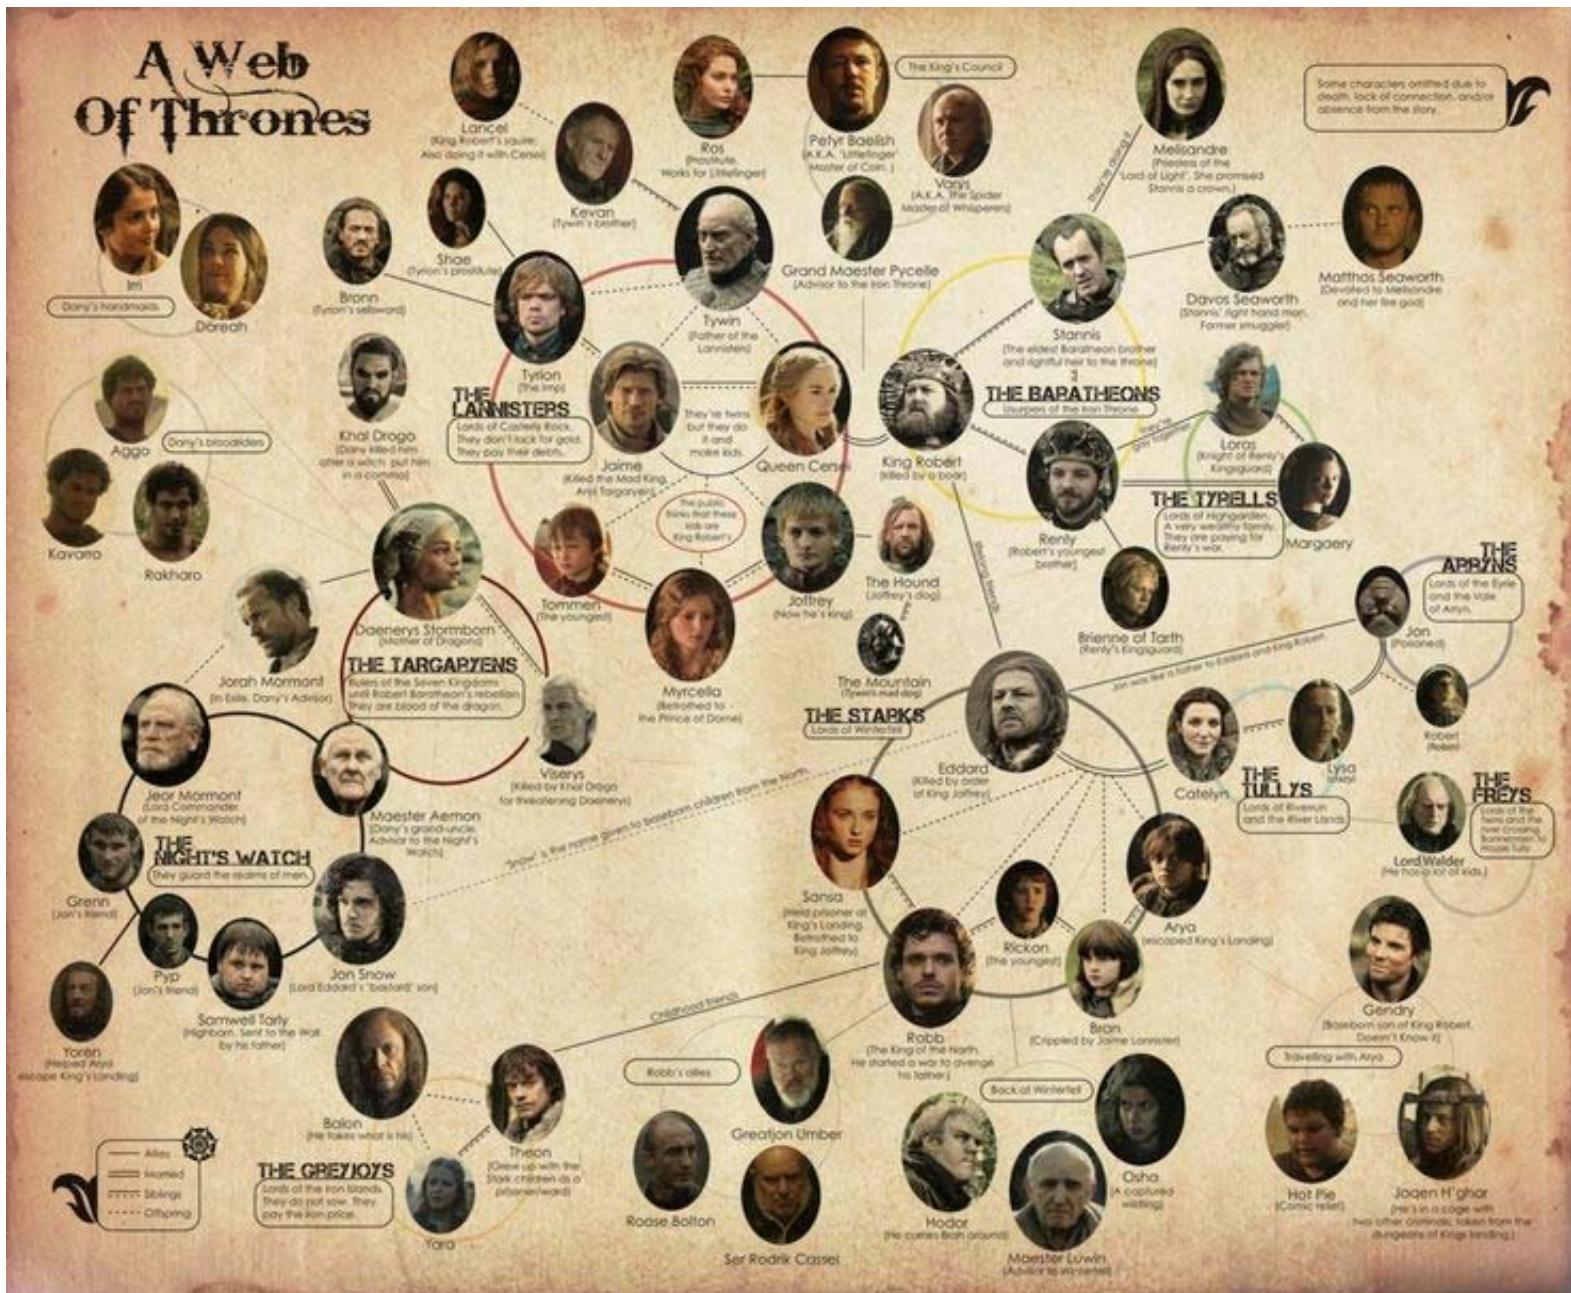
GRAPH DATABASES

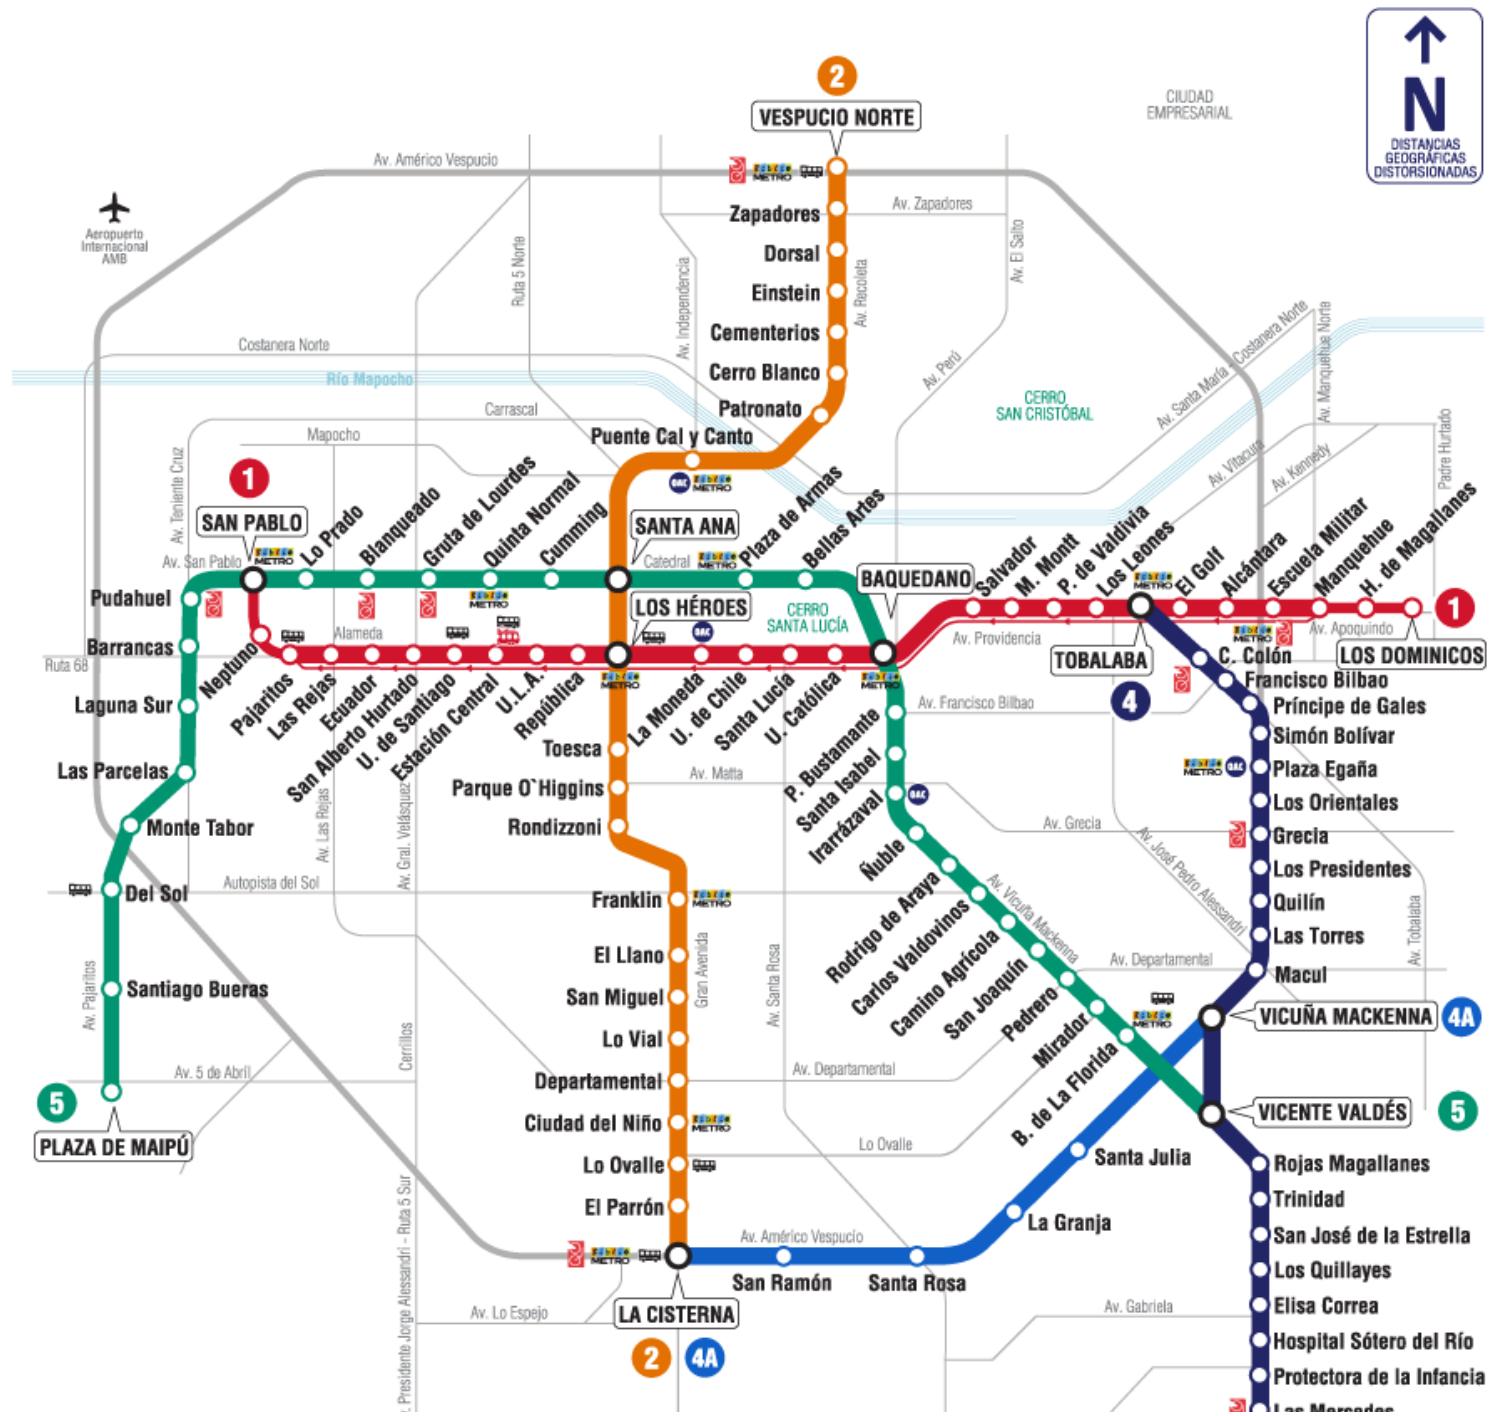


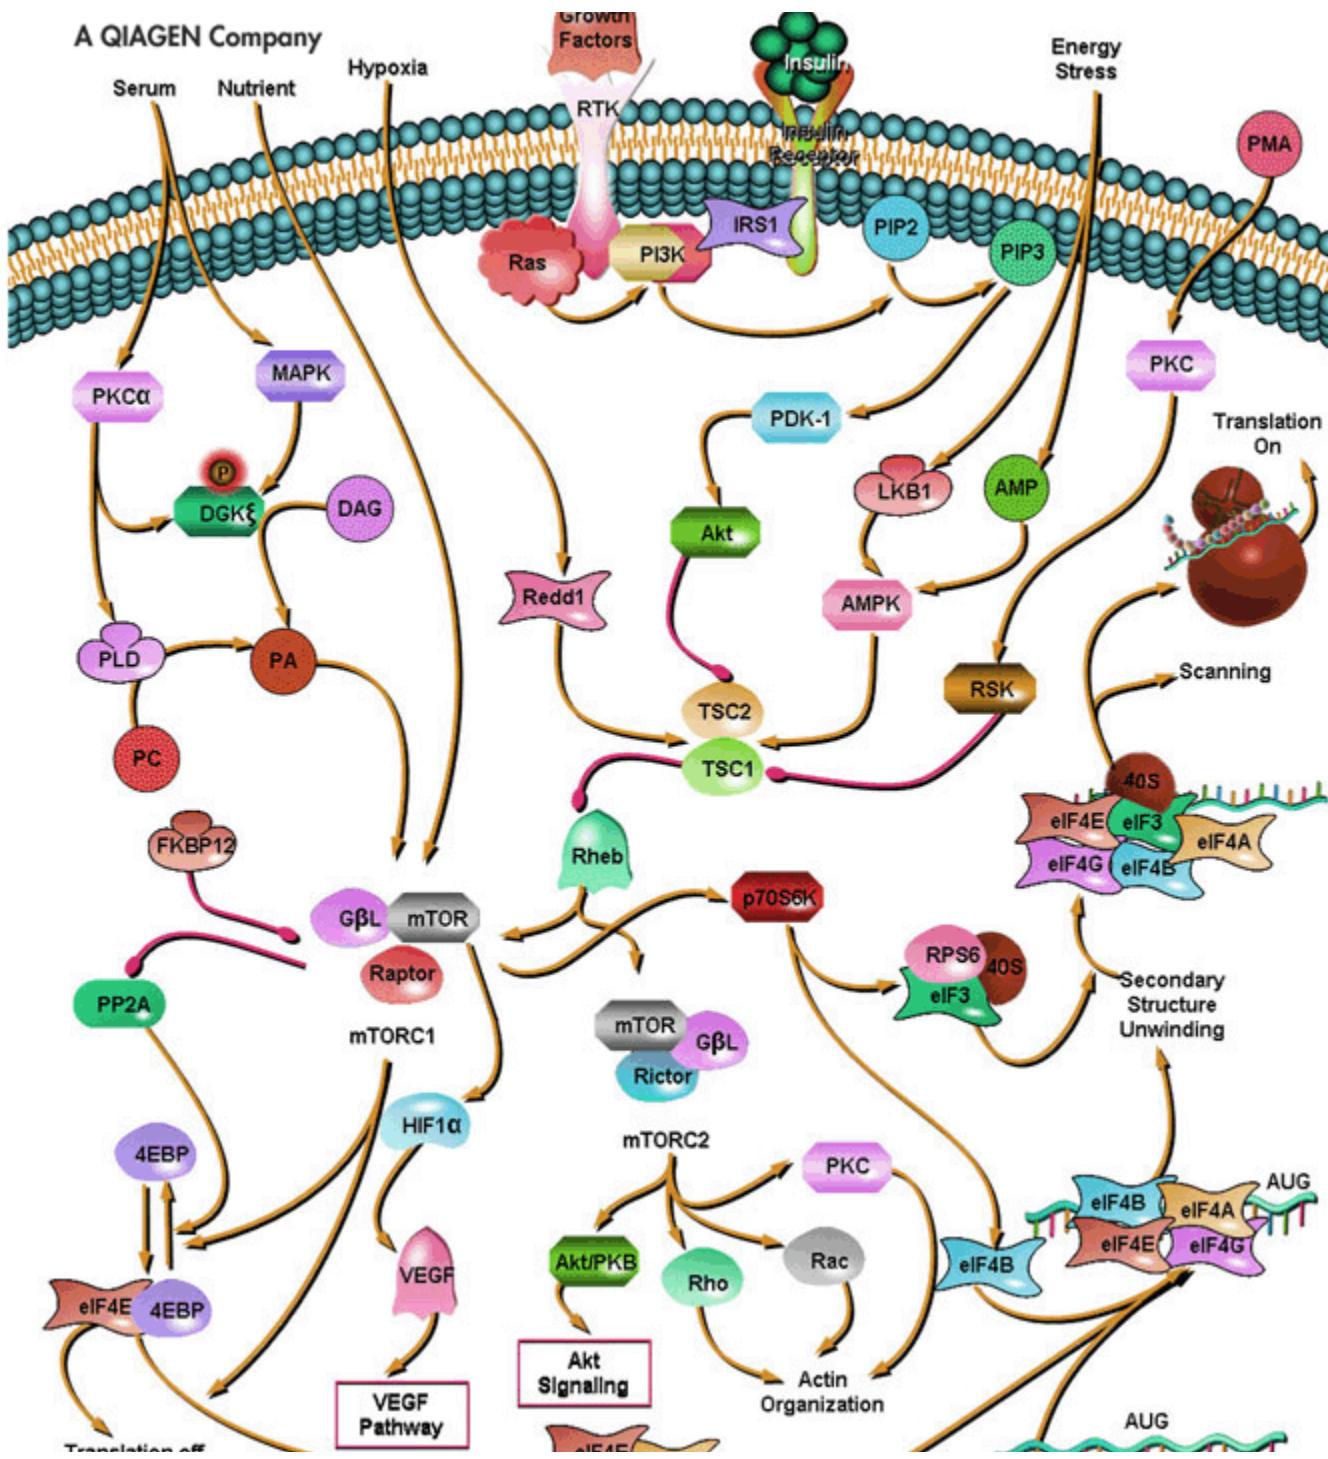
powered by
TouchGraph

Co-author Graph Co-author Path Citation Graph





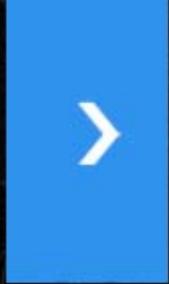






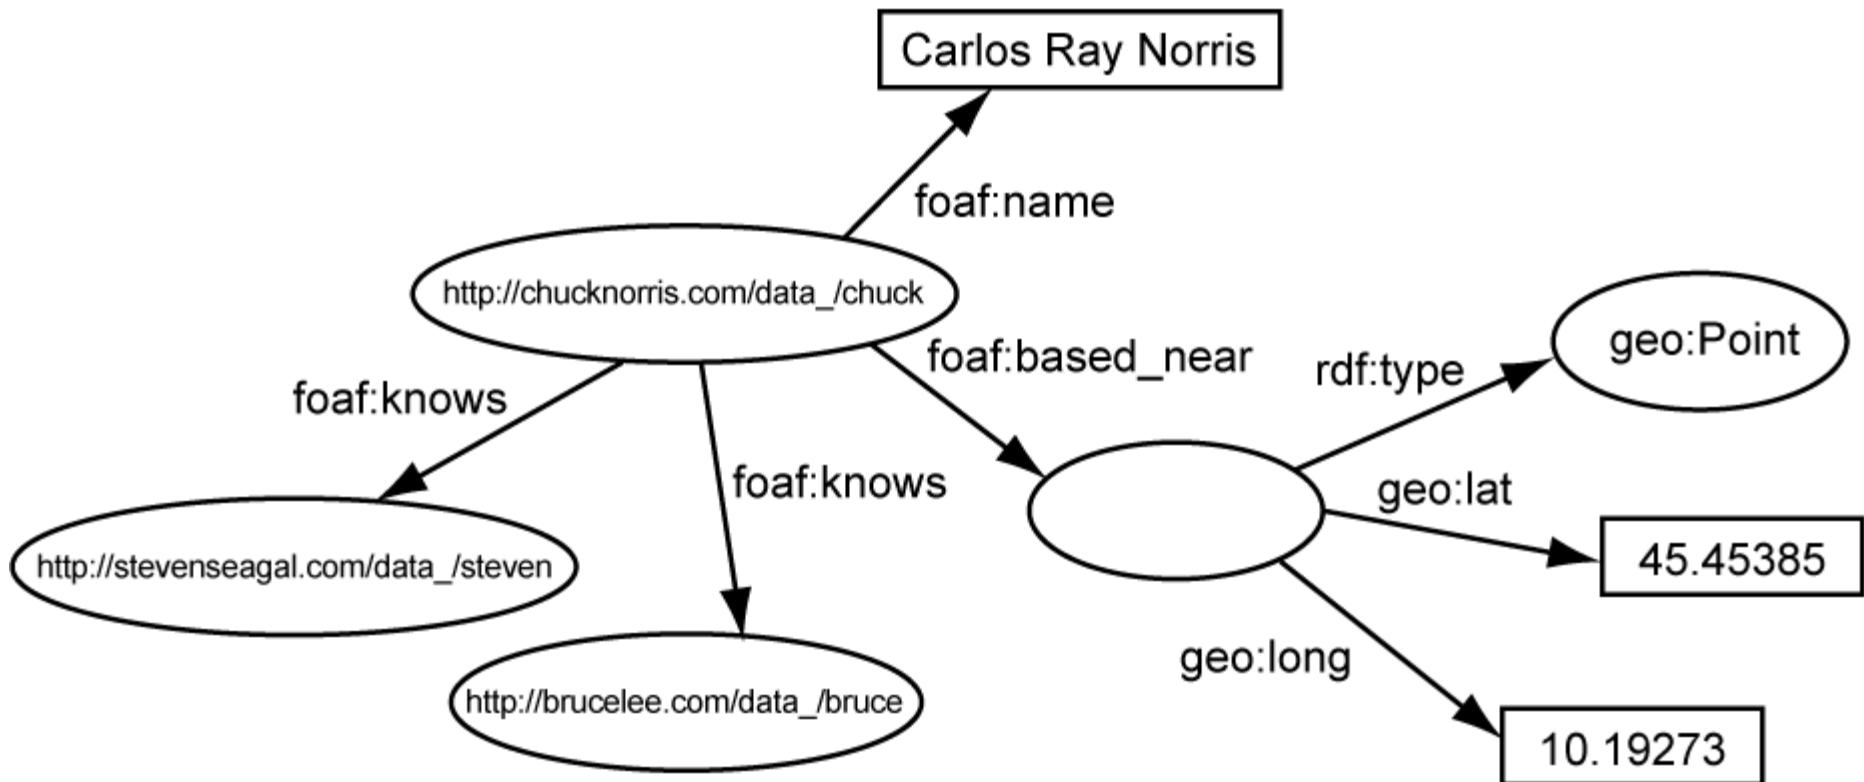
The image shows a network graph visualization centered around the Google logo. The graph consists of nodes (represented by small circles) connected by lines, forming a complex web. Several prominent nodes are highlighted with larger circles containing images: a portrait of Leonardo da Vinci, a Canadian flag, a landscape, a portrait of another historical figure, a cityscape, and a portrait of a man in a red cap. The word "GOOGLE" is written in its signature multi-colored, 3D-style font, with each letter connected to the network graph.

GOOGLE



The Knowledge Graph

Learn more about one of the key breakthroughs
behind the future of search.



Data = Graph

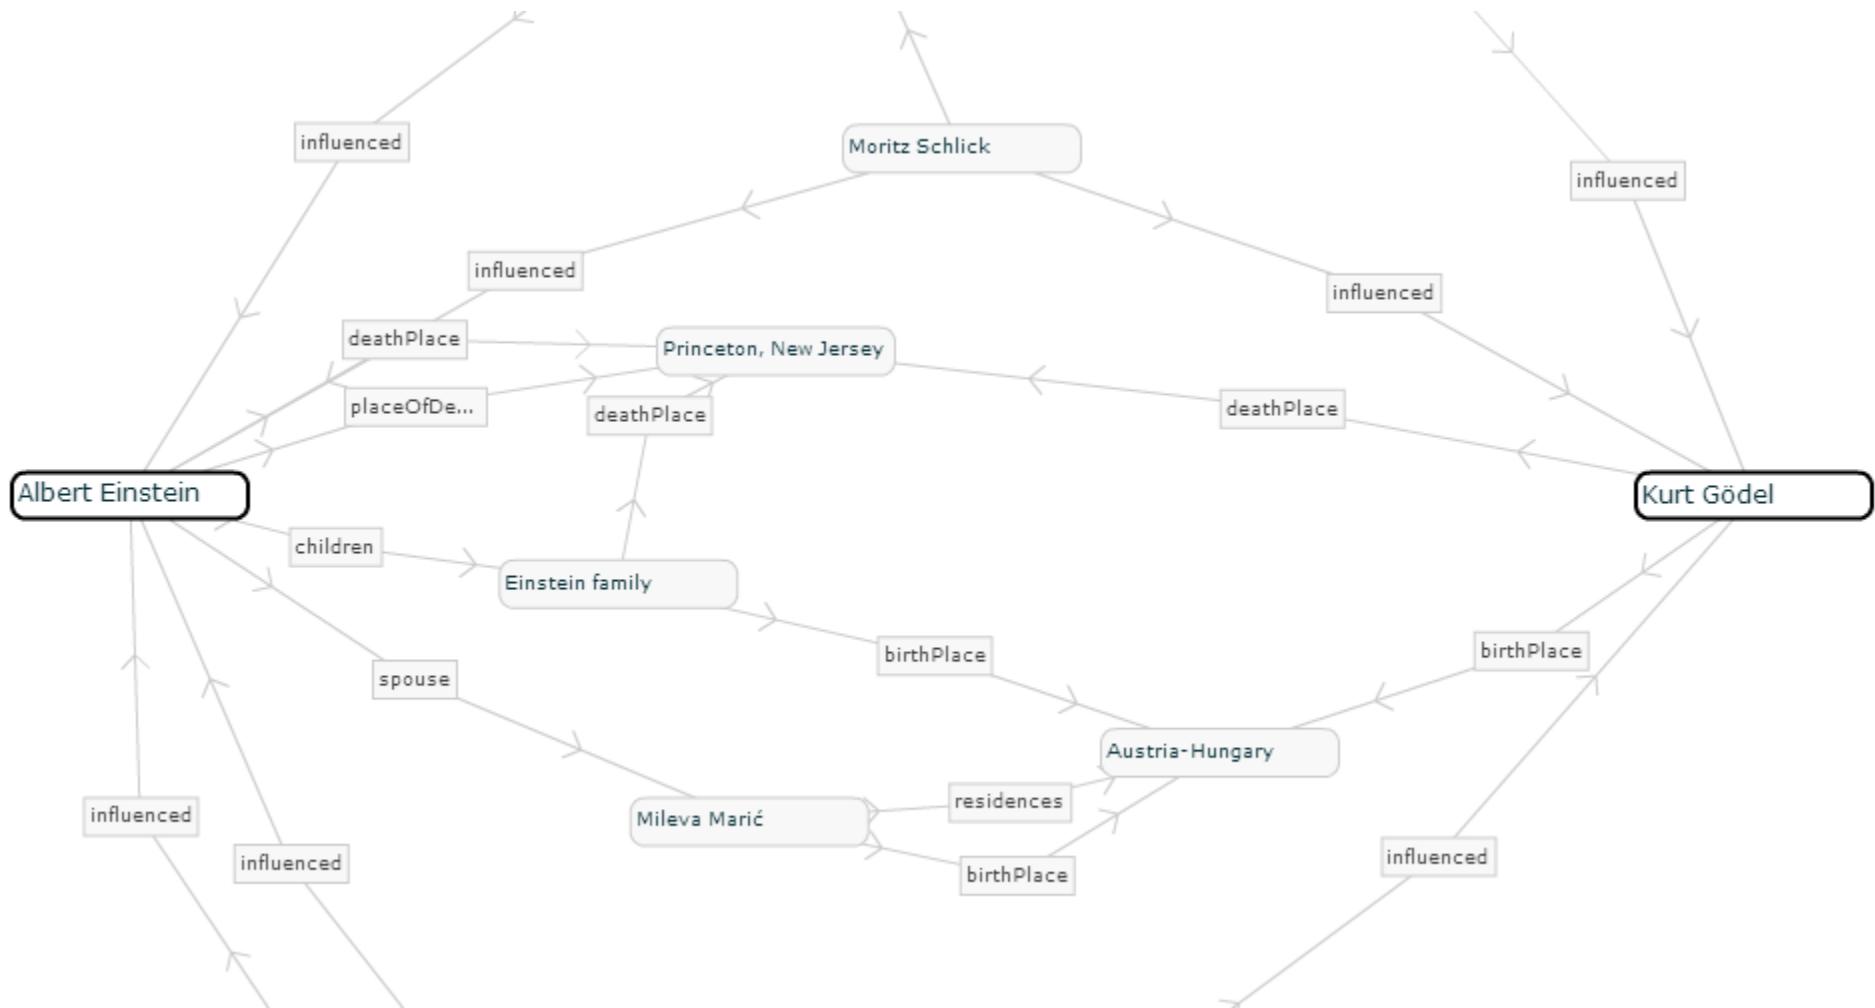
- Any data can be represented as a directed labelled graph (not always neatly)]

When is it a good idea to consider data as a graph?

When you want to answer questions like:

- How many social hops is this user away?
- What is my Erdős number?
- What connections are needed to fly to Perth?
- How are Einstein and Godel related?

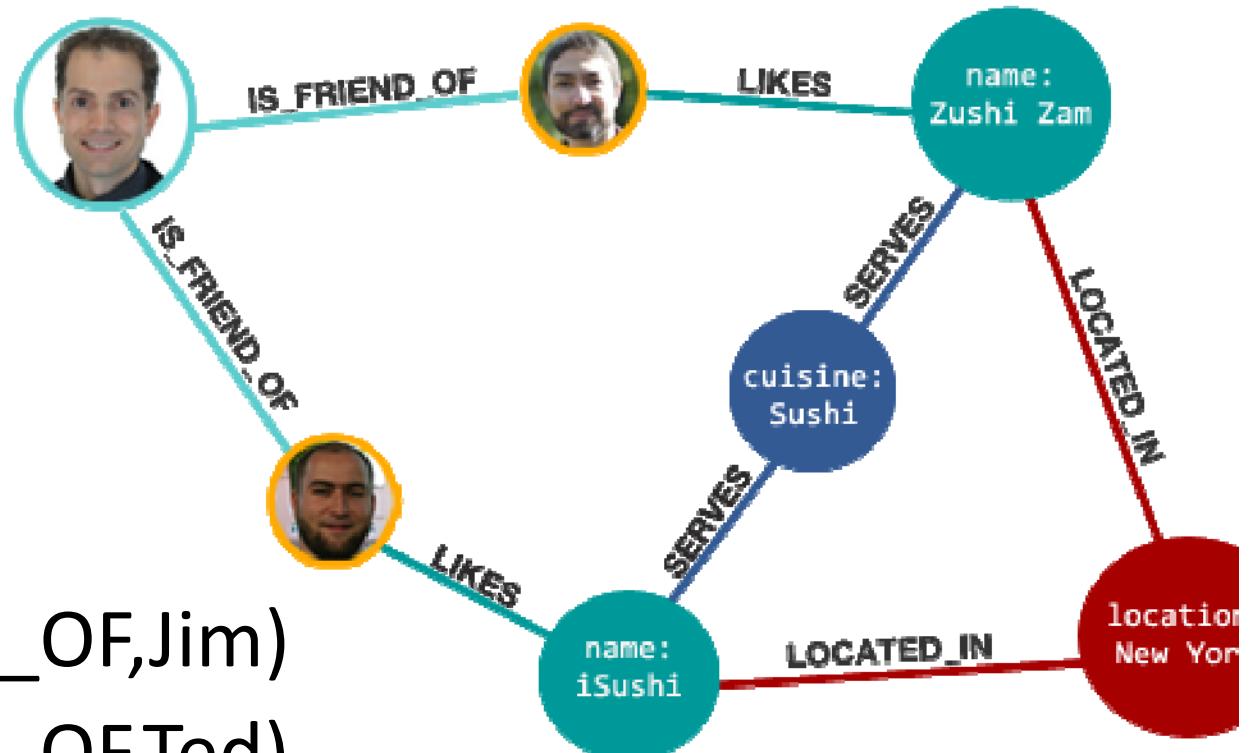
RelFinder



Leading Graph Database



Graph Databases



(Fred,IS_FRIEND_OF,Jim)

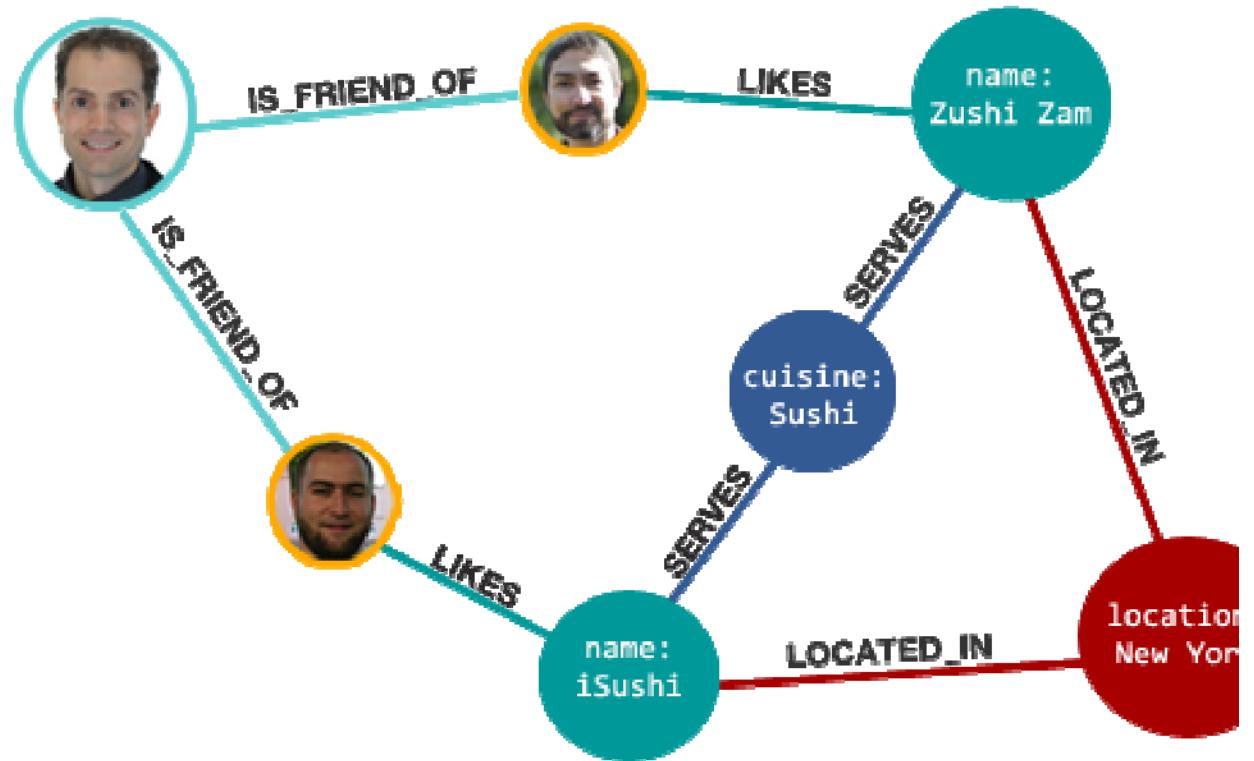
(Fred,IS_FRIEND_OF,Ted)

(Ted,LIKES,Zushi_Zam)

(Zuzhi_Zam,SERVES,Sushi)

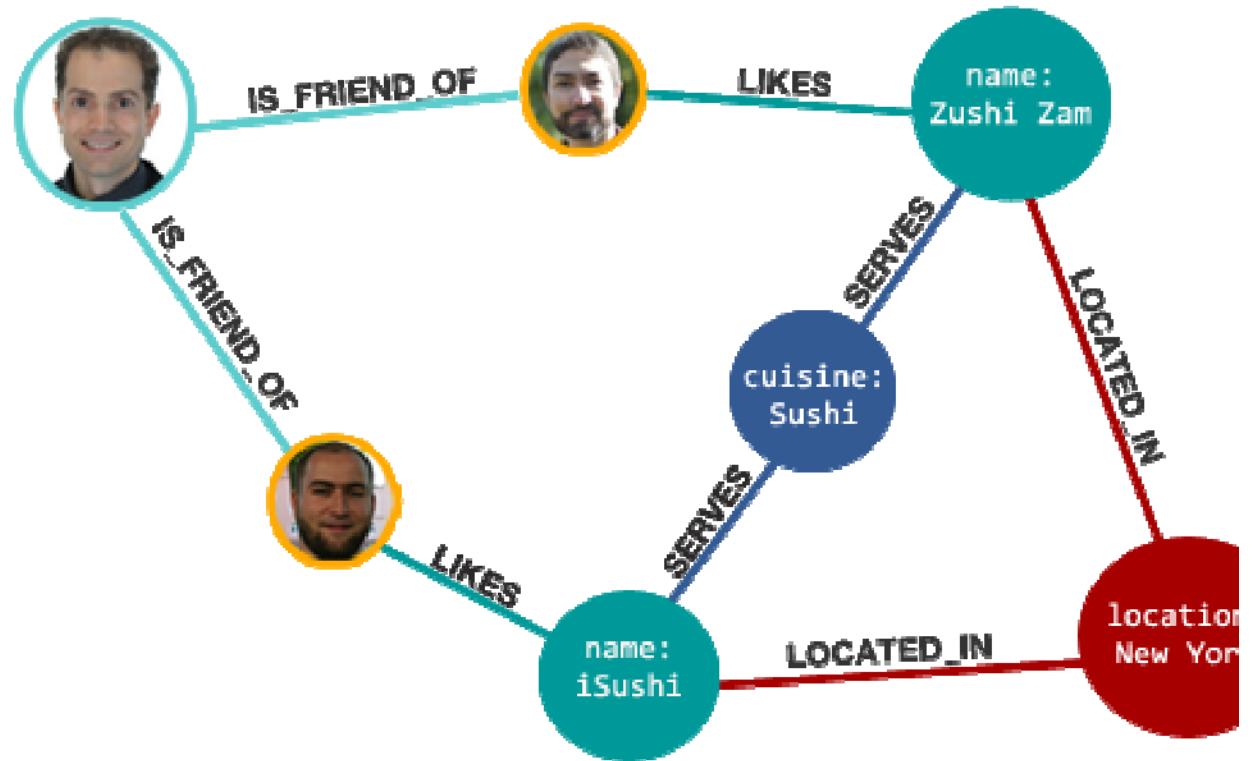
...

Graph Databases: Index Nodes



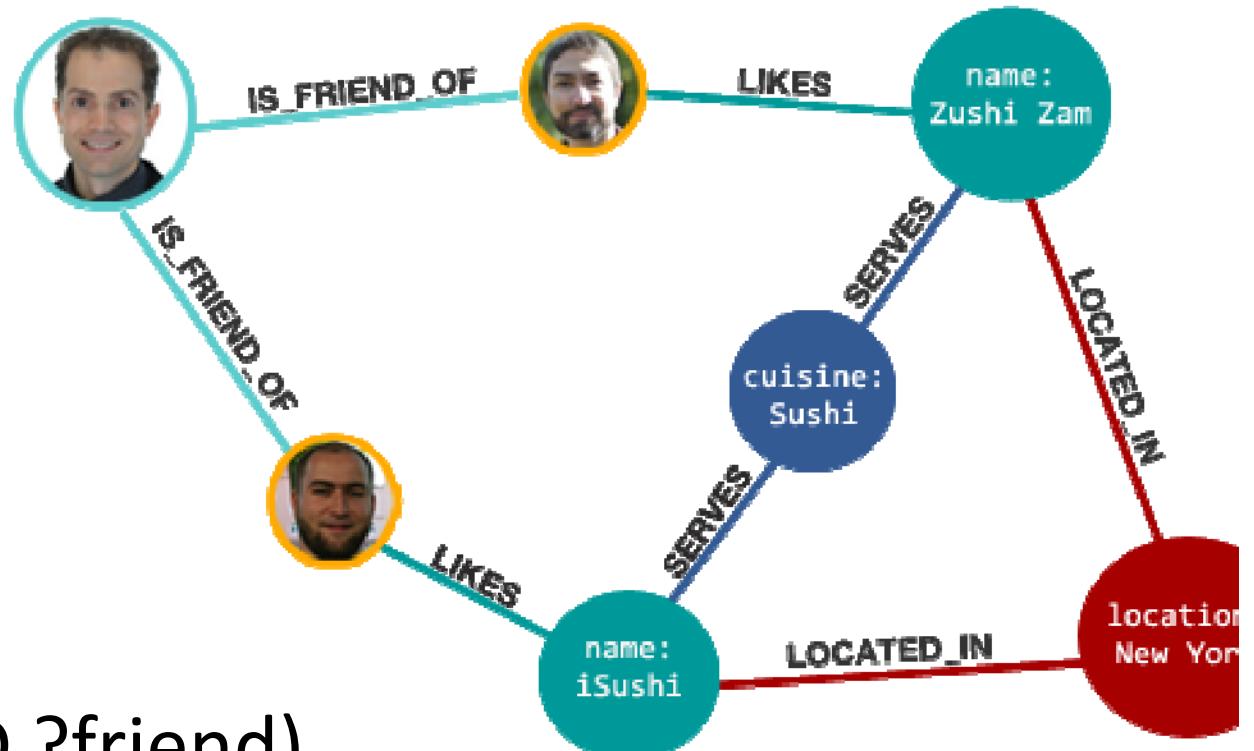
Fred -> (Fred,IS_FRIEND_OF,Jim)
(Fred,IS_FRIEND_OF,Ted)

Graph Databases: Index Relations



LIKES -> (Ted,LIKES,Zushi_Zam)
(Jim,LIKES,iSushi)

Graph Databases: Graph Queries



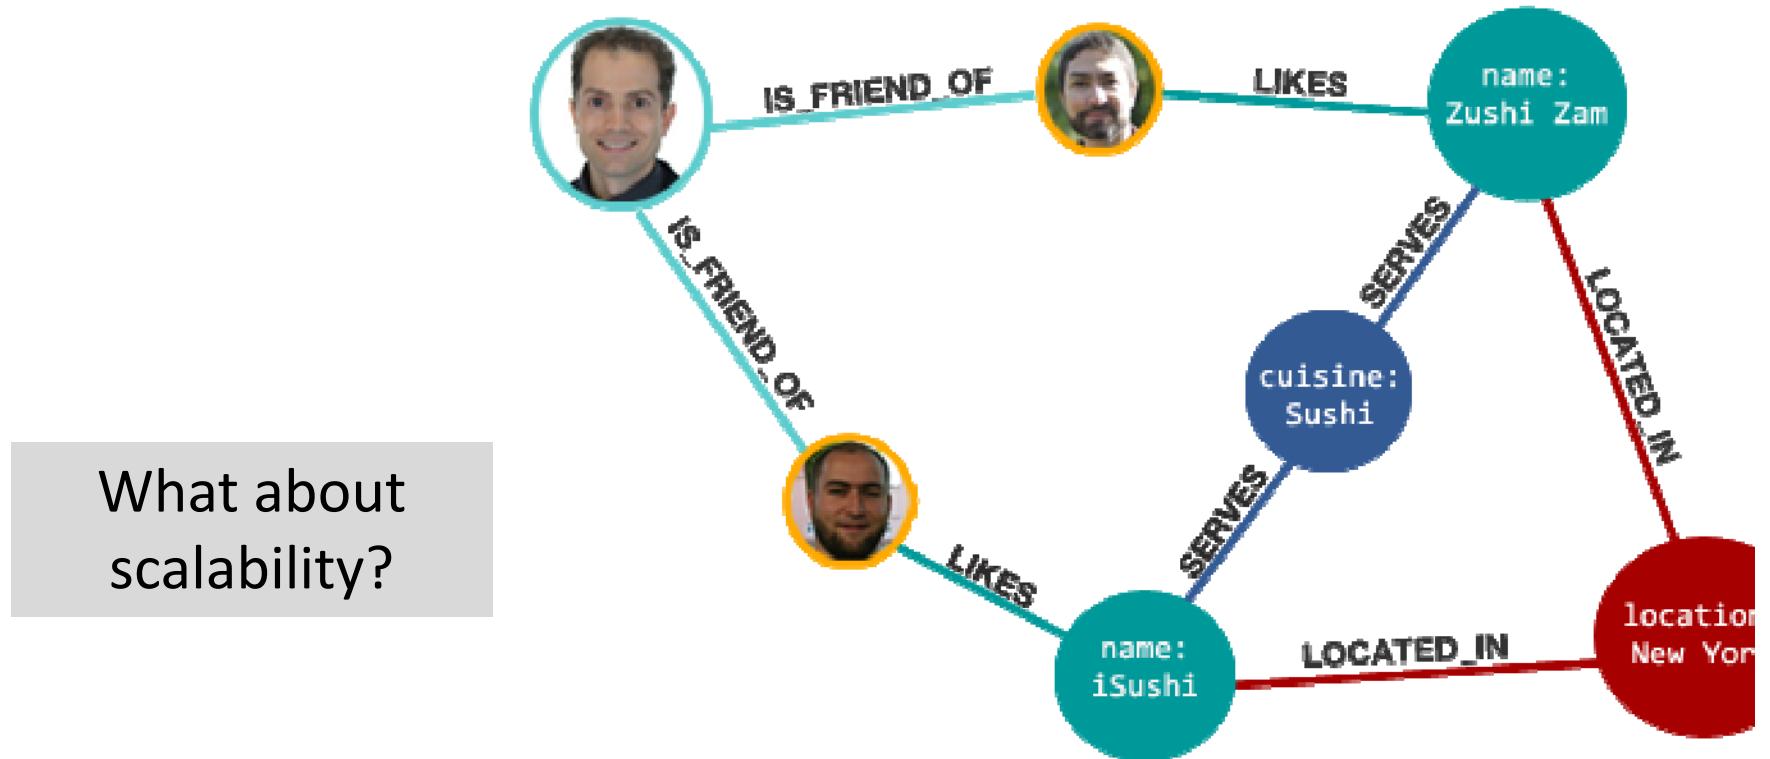
(Fred,IS_FRIEND,?friend)

(?friend,LIKES,?place)

(?place,SERVES,?sushi)

(?place,LOCATED_IN,New_York)

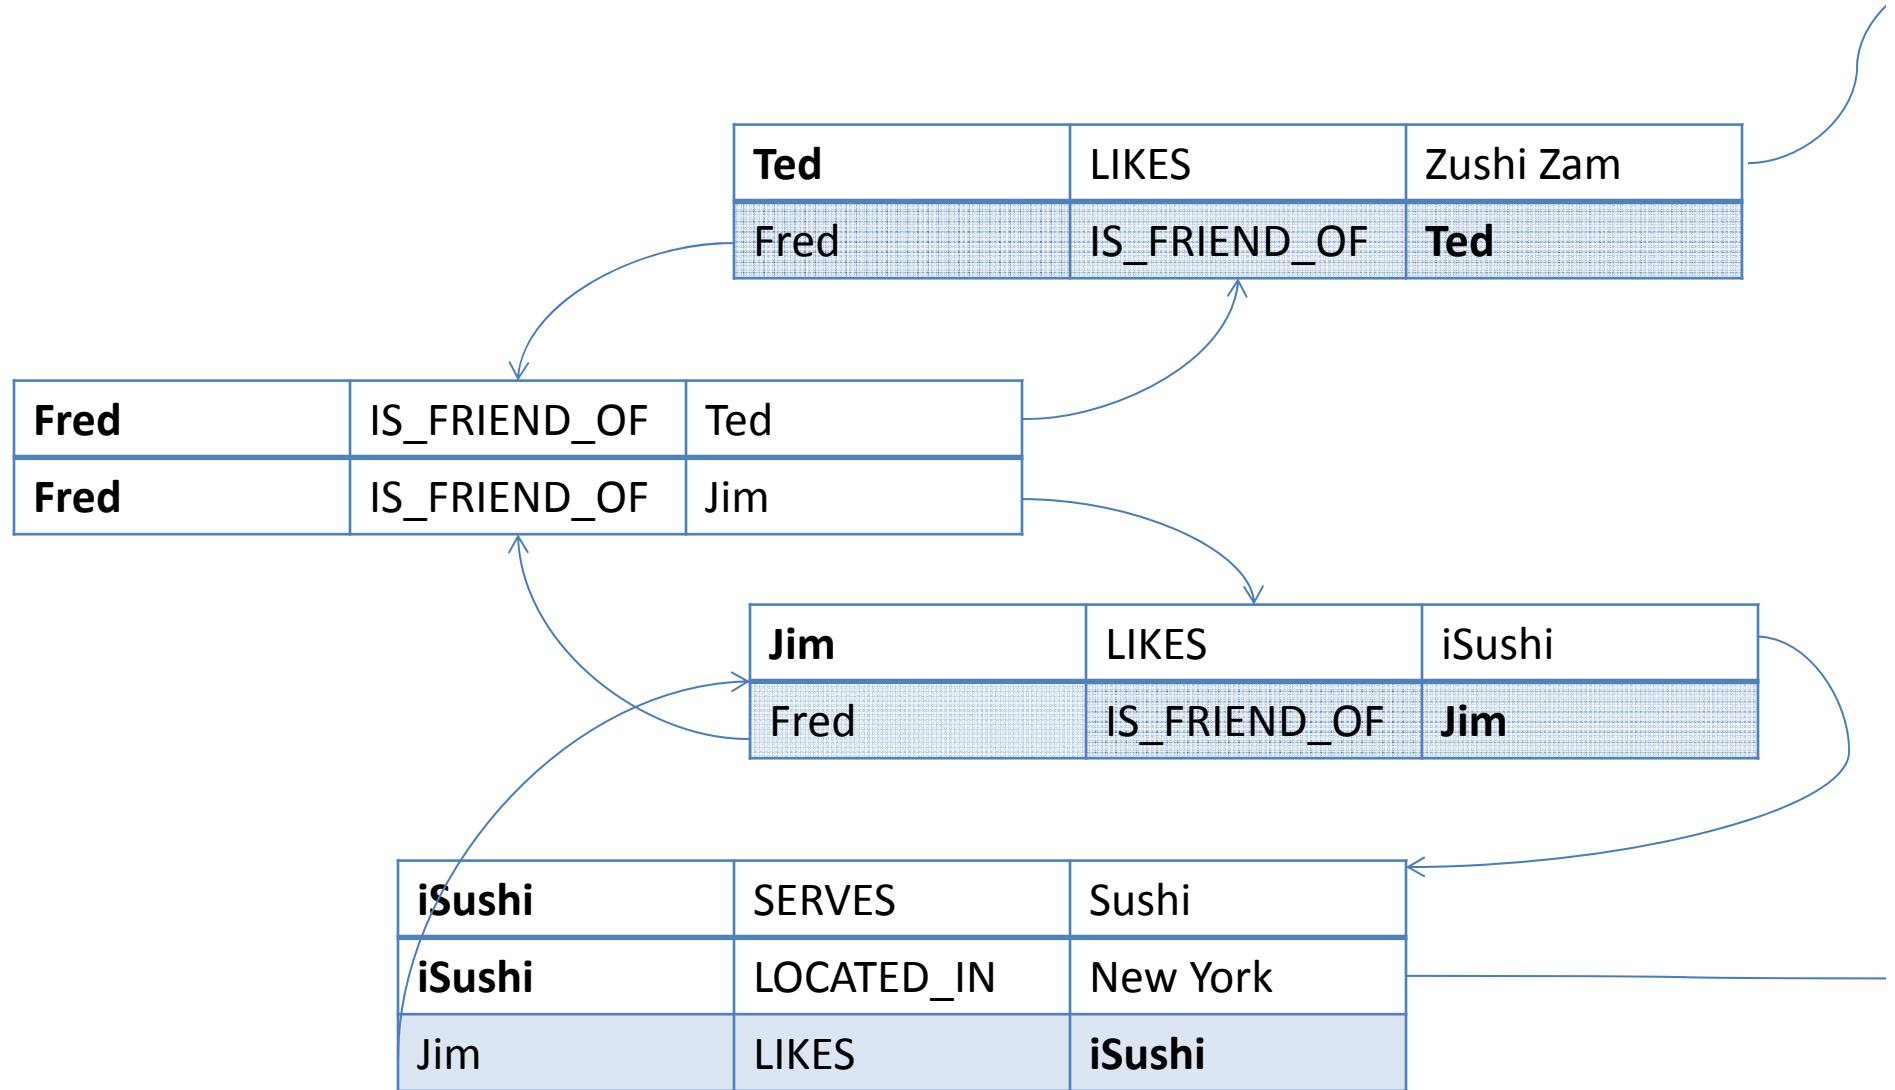
Graph Databases: Path Queries



(Fred,IS_FRIEND*,?friend_of_friend)

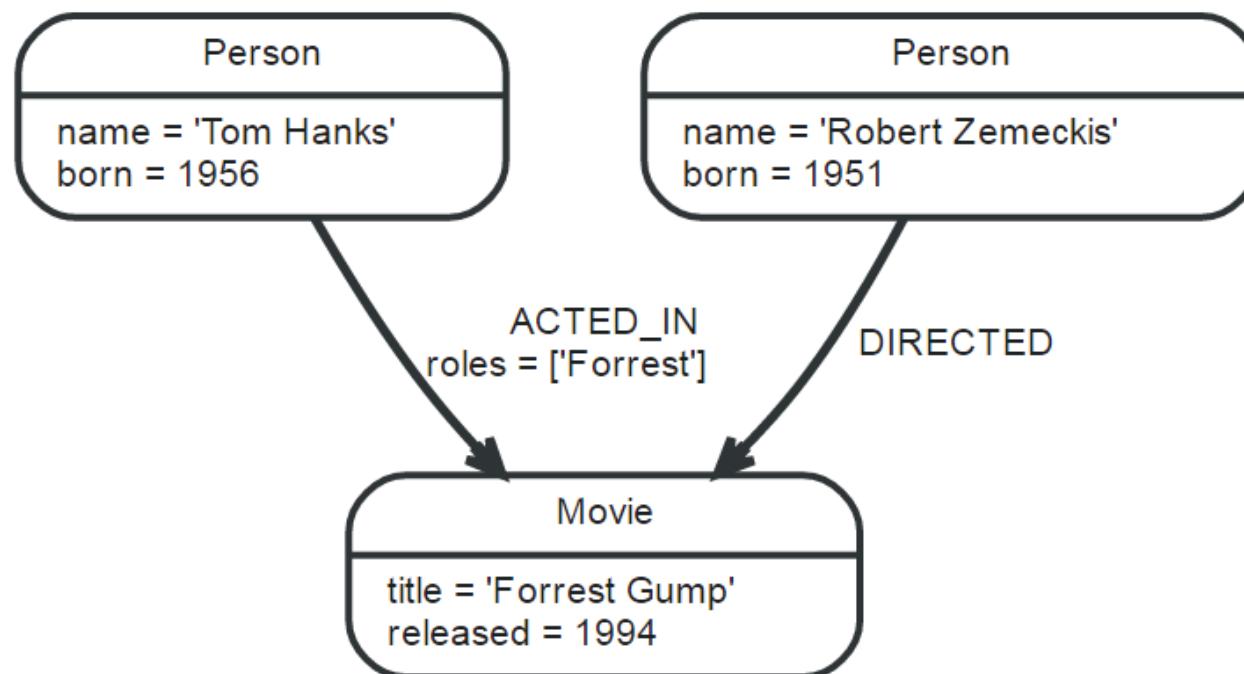
(?friend_of_friend,LIKES,Zushi_Zam)

Graph Database: Index-free Adjacency

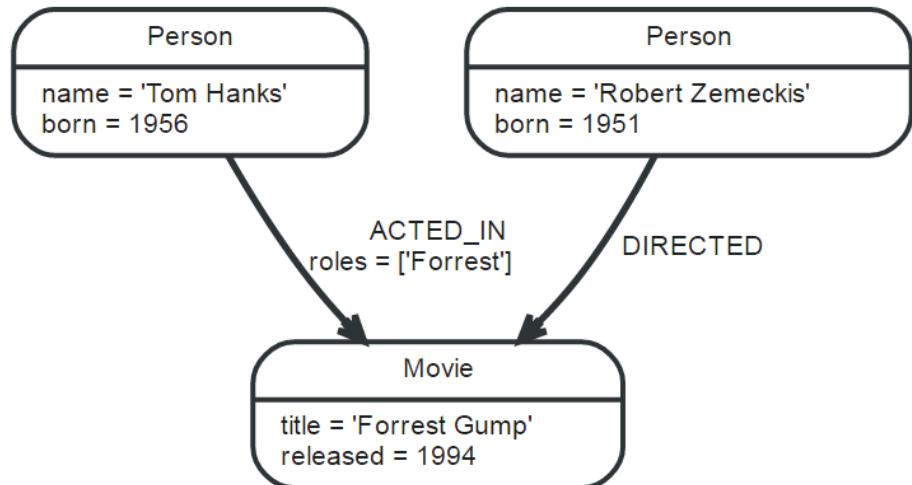


Property Graph Model

- Graph nodes and edges can have
 - Multiple attributes like name, born, roles
 - Multiple labels like Person, Actor, Movie



Query Language: Cypher



```
MATCH (p:Person { name:"Tom Hanks" })-[r:ACTED_IN]->(m:Movie)
RETURN m.title, r.roles
```

What will this return?

m.title	r.roles
"Forrest Gump"	["Forrest"]

1 row

Query Language: Cypher

- Relational query features (UNION, FILTER, ORDER BY, GROUP BY, COUNT, etc.)
- Path query features (*, 1..5, etc.)

<http://neo4j.com/docs/developer-manual/current/#cypher-query-lang>



Table of Contents

+	Introduction
+	Get started
-	Cypher Query Language
-	5. Introduction
-	6. Syntax
-	7. General Clauses
-	8. Reading Clauses
-	9. Writing Clauses

Cypher Query Language

The Cypher part is the authoritative source for details on the Cypher Query Language. For a short introduction, see [What is Cypher?](#). To take your first steps with Cypher, see [Get started with Cypher](#). For the terminology used, see [Terminology](#).

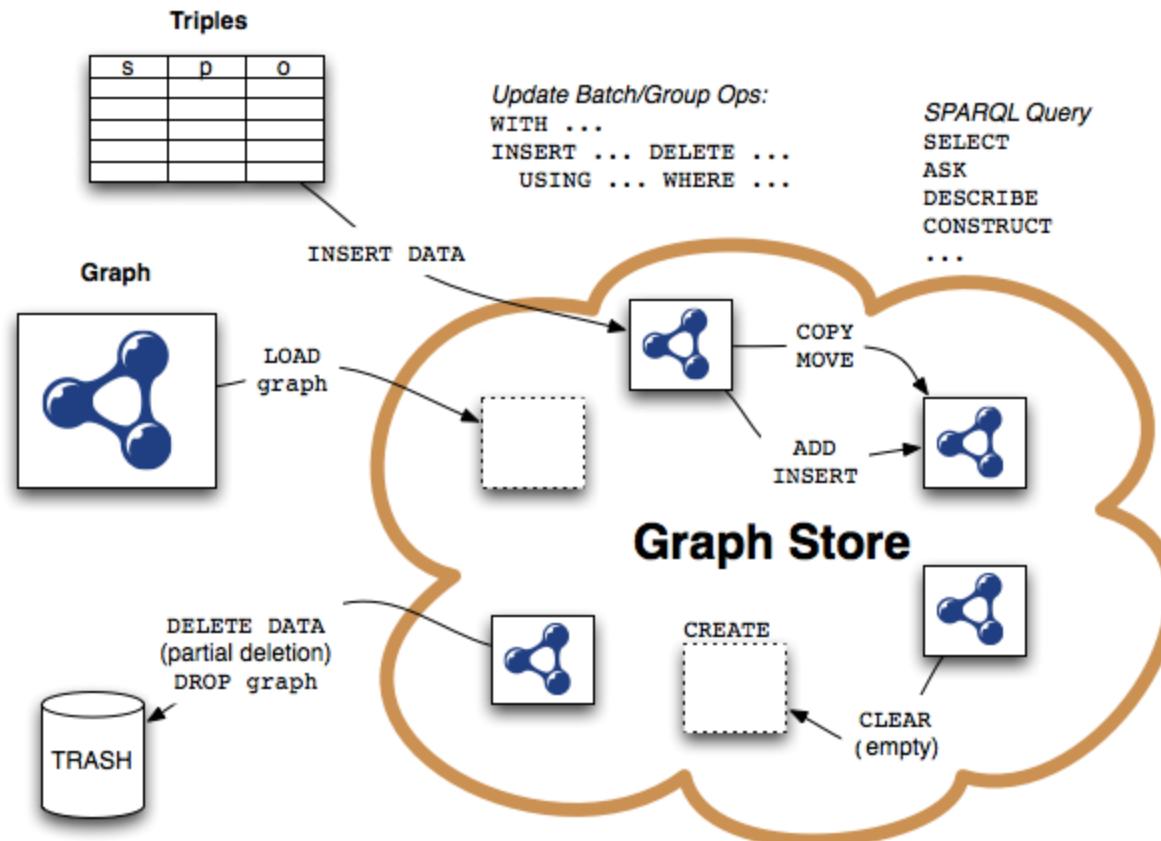
5. Introduction

To get an overview of Cypher, continue reading [What is Cypher?](#). The rest of this chapter deals with the context of Cypher statements, like for example transaction management and how to use parameters. For the Cypher language reference itself see other chapters at [Cypher Query Language](#). To take your first steps with Cypher, see [Get started with Cypher](#). For the terminology used, see [Terminology](#).

5.1. What is Cypher?

Rank			DBMS	Database Model	Score		
May 2016	Apr 2016	May 2015			May 2016	Apr 2016	May 2015
1.	1.	1.	Oracle	Relational DBMS	1462.02	-5.51	+19.93
2.	2.	2.	MySQL 	Relational DBMS	1371.83	+1.72	+77.56
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1142.82	+7.77	+11.79
4.	4.	4.	MongoDB 	Document store	320.22	+7.78	+42.90
5.	5.	5.	PostgreSQL	Relational DBMS	307.61	+3.89	+34.09
6.	6.	6.	DB2	Relational DBMS	185.96	+1.87	-15.09
7.	↑ 8.	↑ 8.	Cassandra 	Wide column store	134.50	+4.83	+27.95
8.	↓ 7.	↓ 7.	Microsoft Access	Relational DBMS	131.58	-0.39	-14.00
9.	9.	↑ 10.	Redis 	Key-value store	108.24	-3.00	+13.51
10.	10.	↓ 9.	SQLite	Relational DBMS	107.26	-0.70	+2.10
11.	11.	↑ 14.	Elasticsearch 	Search engine	86.31	+3.73	+21.48
12.	↑ 13.	↑ 13.	Teradata	Relational DBMS	73.74	+1.48	+3.62
13.	↓ 12.	↓ 11.	SAP Adaptive Server	Relational DBMS	71.48	-1.84	-14.01
14.	14.	↓ 12.	Solr	Search engine	65.62	-0.40	-17.31
15.	15.	15.	HBase	Wide column store	51.84	+0.35	-9.87
16.	16.	↑ 17.	Hive	Relational DBMS	47.51	-1.57	+2.94
17.	17.	↓ 16.	FileMaker	Relational DBMS	46.71	+0.60	-6.19
18.	18.	18.	Splunk	Search engine	44.31	+1.96	+3.59
19.	19.	↑ 21.	SAP HANA 	Relational DBMS	41.37	+1.02	+8.59
20.	↑ 21.	↑ 25.	MariaDB 	Relational DBMS	33.97	+2.38	+10.37
21.	↓ 20.	↑ 22.	Neo4j 	Graph DBMS	32.61	+0.70	+3.97
22.	22.	↓ 19.	Informix	Relational DBMS	30.58	-0.94	-6.16
23.	23.	↓ 20.	Memcached	Key-value store	27.90	-0.11	-5.21
24.	24.	24.	Couchbase 	Document store	24.29	-0.73	-0.95
25.	25.	↑ 30.	Amazon DynamoDB 	Multi-model 	23.60	+0.48	+8.73

SPARQL



Rank			DBMS	Database Model	Score		
May 2016	Apr 2016	May 2015			May 2016	Apr 2016	May 2015
1.	1.	1.	Oracle	Relational DBMS	1462.02	-5.51	+19.93
2.	2.	2.	MySQL +	Relational DBMS	1371.83	+1.72	+77.56
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1142.82	+7.77	+11.79
4.	4.	4.	MongoDB +	Document store	320.22	+7.78	+42.90
5.	5.	5.	PostgreSQL	Relational DBMS	307.61	+3.89	+34.09
6.	6.	6.	DB2	Relational DBMS	185.96	+1.87	-15.09
7.	▲ 8.	▲ 8.	Cassandra +	Wide column store	134.50	+4.83	+27.95
8.	▼ 7.	▼ 7.	Microsoft Access	Relational DBMS	131.58	-0.39	-14.00
9.	9.	▲ 10.	Redis +	Key-value store	108.24	-3.00	+13.51
10.	10.	▼ 9.	SQLite	Relational DBMS	107.26	-0.70	+2.10
11.	11.	▲ 14.	Elasticsearch +	Search engine	86.31	+3.73	+21.48
12.	▲ 13.	▲ 13.	Teradata	Relational DBMS	73.74	+1.48	+3.62
13.	▼ 12.	▼ 11.	SAP Adaptive Server				
14.	14.	▼ 12.	Solr				
15.	15.	15.	HBase				
16.	16.	▲ 17.	Hive				
17.	17.	▼ 16.	FileMaker				
18.	18.	18.	Splunk				
19.	19.	▲ 21.	SAP HANA +				
20.	▲ 21.	▲ 25.	MariaDB +				
21.	▼ 20.	▲ 22.	Neo4j +				
22.	22.	▼ 19.	Informix				
23.	23.	▼ 20.	Memcached				
24.	24.	24.	Couchbase +				
25.	25.	▲ 30.	Amazon DynamoDB +				



RECAP

Recap

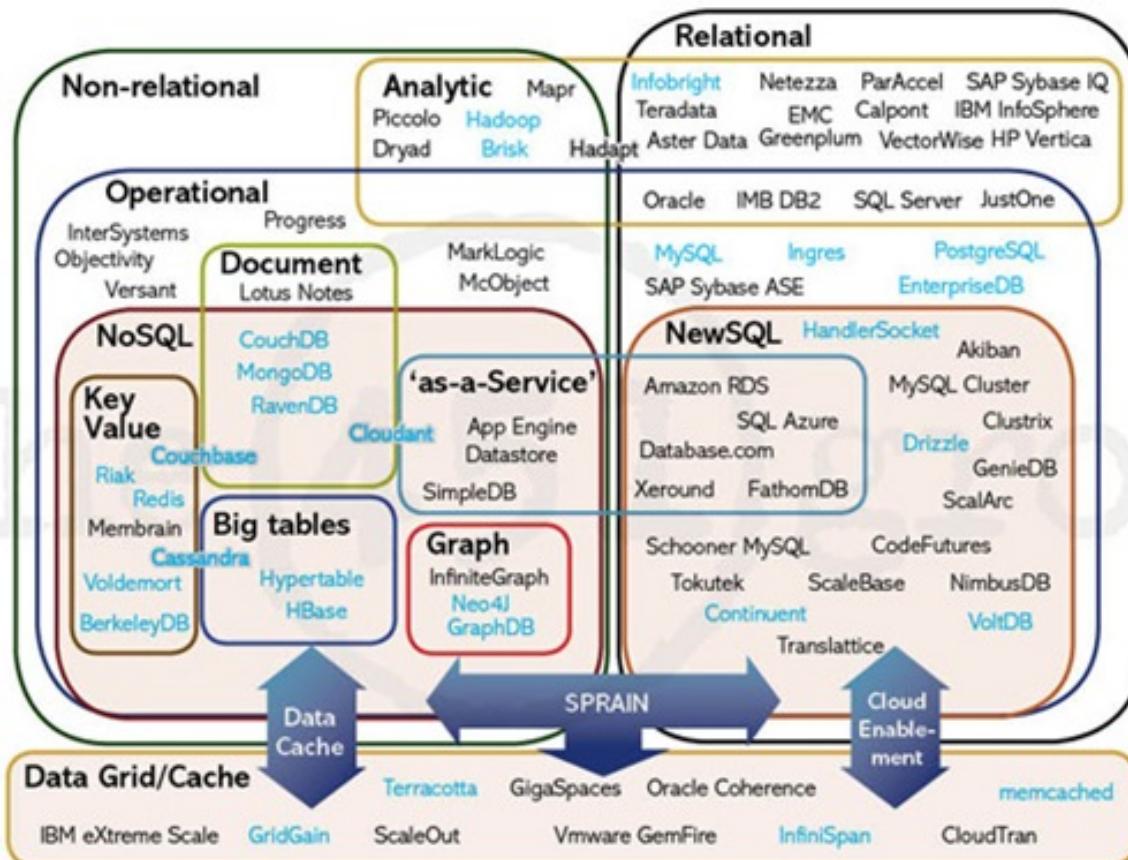
- Relational Databases don't solve everything
 - SQL and ACID add overhead
 - Distribution not so easy
- NoSQL: what if you don't need SQL or ACID?
 - Something simpler
 - Something more scalable
 - Trade efficiency against guarantees

NoSQL: Trade-offs

- **Simplified transactions** (no ACID)
- **Simplified (or no) query language**
 - Procedural or a subset of SQL
- **Simplified query algebra**
 - Often no joins
- **Simplified data model**
 - Often map-based
- **Simplified replication**
 - Consistency vs. Availability

Simplifications enable scale to thousands of machines. But a lot of relational database features are lost!

NoSQL Overview Map



Types of NoSQL Store

- **Key–Value Stores** (e.g., Dynamo):
 - Distributed unsorted maps
 - Some have secondary indexes
- **Document Stores** (e.g., MongoDB):
 - Map values are documents (e.g., JSON, XML)
 - Built-in document functions/indexable fields
- **Table/Column-Based Stores** (e.g., Bigtable):
 - Distributed multi-dimensional sorted maps
 - Distribution by Tablets/Column-families
- **Graph Stores** (e.g., Neo4J)
 - Stores vertices and relations: Index-free adjacency
 - Query languages for paths, reachability, etc.
- **Hybrid/mix/other**

Categories are far from clean. But aside from graph stores, most NoSQL stores are just fancy (distributed) maps. ☺

Bigtable

- Column family store: (row, column, time) → value
- Sorted map, range partitioned
- PAXOS for locks, root table
- Tablets: horizontal table splits
- Column family: logical grouping of columns stored close together
- Locality groups: grouping of column families
- SSTable: sequence of 64k blocks
- Batch writes
- Compactions: merge SSTables

Graph stores

- Indexes on nodes, relations, (attributes)
- Relational-style queries
- Path-style queries
- Index-free adjacency
- Neo4J:
 - Property graph data model
 - Cypher query language

Questions

