

S Q L
OR NoSQL



@majidfn

Majid Fatemian

Ubisoft Montreal

BIG DATA

NoSQL

NoSQL



1

No SQL

2

Cluster

3

Schemaless

YOUR SHAPE

Fitness Evolved



YOUR SHAPE

Fitness Evolved



YOUR SHAPE

Fitness Evolved




YOUR SHAPE

Fitness Evolved

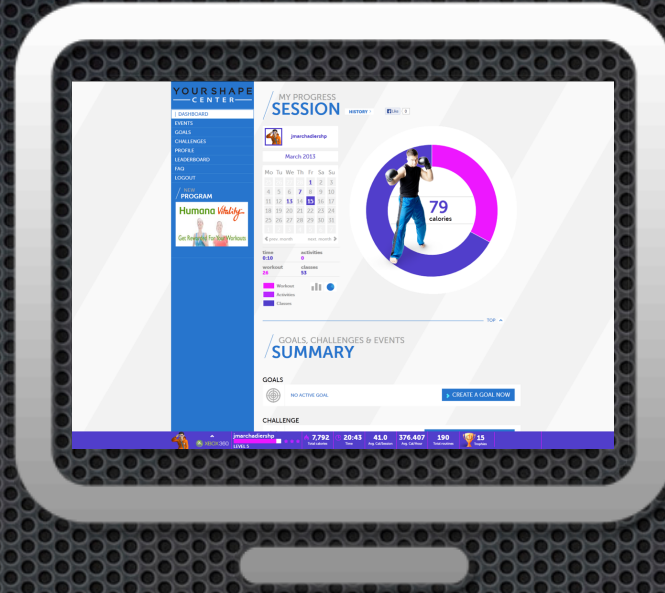
YOUR SHAPE

Fitness Evolved



YOUR SHAPE

Fitness Evolved



STATS

Users

2.2M

Concurrent


10K

Daily Activities


**160K(Max)
40K (Avg)**



2012



2011



2013



2012



2011



2013

Architecture

2011

.....



Architecture

2011











SQL / JOIN
ACID
Transaction
Partitioning
Replication
Tools
OPs
Data Modeling






2012



2011



2013

DATA

Global

Profile Info.

Aggregated stats

Community Events

Specific

Details

Activities

Leaderboards



SQL / JOIN
ACID
Transaction
Partitioning
Replication
Tools
OPs
Data Modeling





SQL / JOIN
ACID
Transaction
Partitioning
Replication
Tools
OPs
Data Modeling





SQL / JOIN
ACID
Transaction
Partitioning
Replication
Tools
OPs
Data Modeling



Volume
No Details
Variant Schema
Partitioning
Scale out
Sharding



SQL / JOIN

xbox v.1

```
{  
  player_id: 1234,  
  game_id: 1002,  
  platform: wiiu,  
  * calories: 75,  
  * duration: 120  
}
```

xbox v.2

```
{  
  player_id: 1234,  
  game_id: 4001,  
  platform: xbox360,  
  * calories: 40,  
  * duration: 75,  
  difficulty_level: medium,  
  jumping_ropes: 178,  
  score: 450,  
  endurance:3,  
  category: 4  
}
```

Volume

No Details

Variant Schema

Partitioning

Scale out

Sharding





Volume
No Details
Variant Schema
Partitioning
Scale out
Sharding



Data Modeling

Data Access Layer



Volume
No Details
Variant Schema
Partitioning
Scale out
Sharding



Data Modeling



SQL / JOIN
ACID
Transaction
Partitioning
Replication
Tools
OPs
Data Modeling



Volume
No Details
Variant Schema
Partitioning
Scale out
Sharding



No
SQL

No SQL



...

mongoDB



①

OPs

②


Support

③

Community







Schema-Less
Aggregated Data
Object Mapping
Large Scale Data
Sharding
Map/Reduce
Eventual Consistency
Memory Storage
Journaling





Schema-Less

Aggregated Data

Object Mapping

Large Scale Data

Sharding

Map/Reduce

Eventual Consistency

Memory Storage

Journaling

```
{  
  xbox v.1  
  player_id: 1234,  
  game_id: 1002,  
  platform: wiiu,  
  * calories: 75,  
  * duration: 120  
}
```

```
{  
  WiiU  
  player_id: 1234,  
  game_id: 4001,  
  platform: xbox360,  
  * calories: 40,  
  * duration: 75,  
  difficulty_level: medium,  
  jumping_ropes: 178,  
  score: 450,  
  endurance:3,  
  category: 4  
}
```




Schema-Less Aggregated Data Object Mapping

Large Scale Data

Sharding

Map/Reduce

Eventual Consistency

Memory Storage

Journaling

```
{  
  player:  
  {  
    name: example, gender: M, country:CA  
  }  
  game_id: 4001,  
  date: "2013-11-08",  
  calories: 40,  
  duration: 75,  
}
```





Schema-Less
Aggregated Data
Object Mapping

**Large Scale Data
Sharding
Map/Reduce**


Eventual Consistency
Memory Storage
Journaling






Schema-Less
Aggregated Data
Object Mapping
Large Scale Data
Sharding
Map/Reduce
Eventual Consistency
Memory Storage
Journaling






Schema-Less
Aggregated Data
Object Mapping
Large Scale Data
Sharding
Map/Reduce
Eventual Consistency
Memory Storage
Journaling





Schema-Less
Aggregated Data
Object Mapping
Large Scale Data
Sharding
Map/Reduce
Eventual Consistency
Memory Storage
Journaling



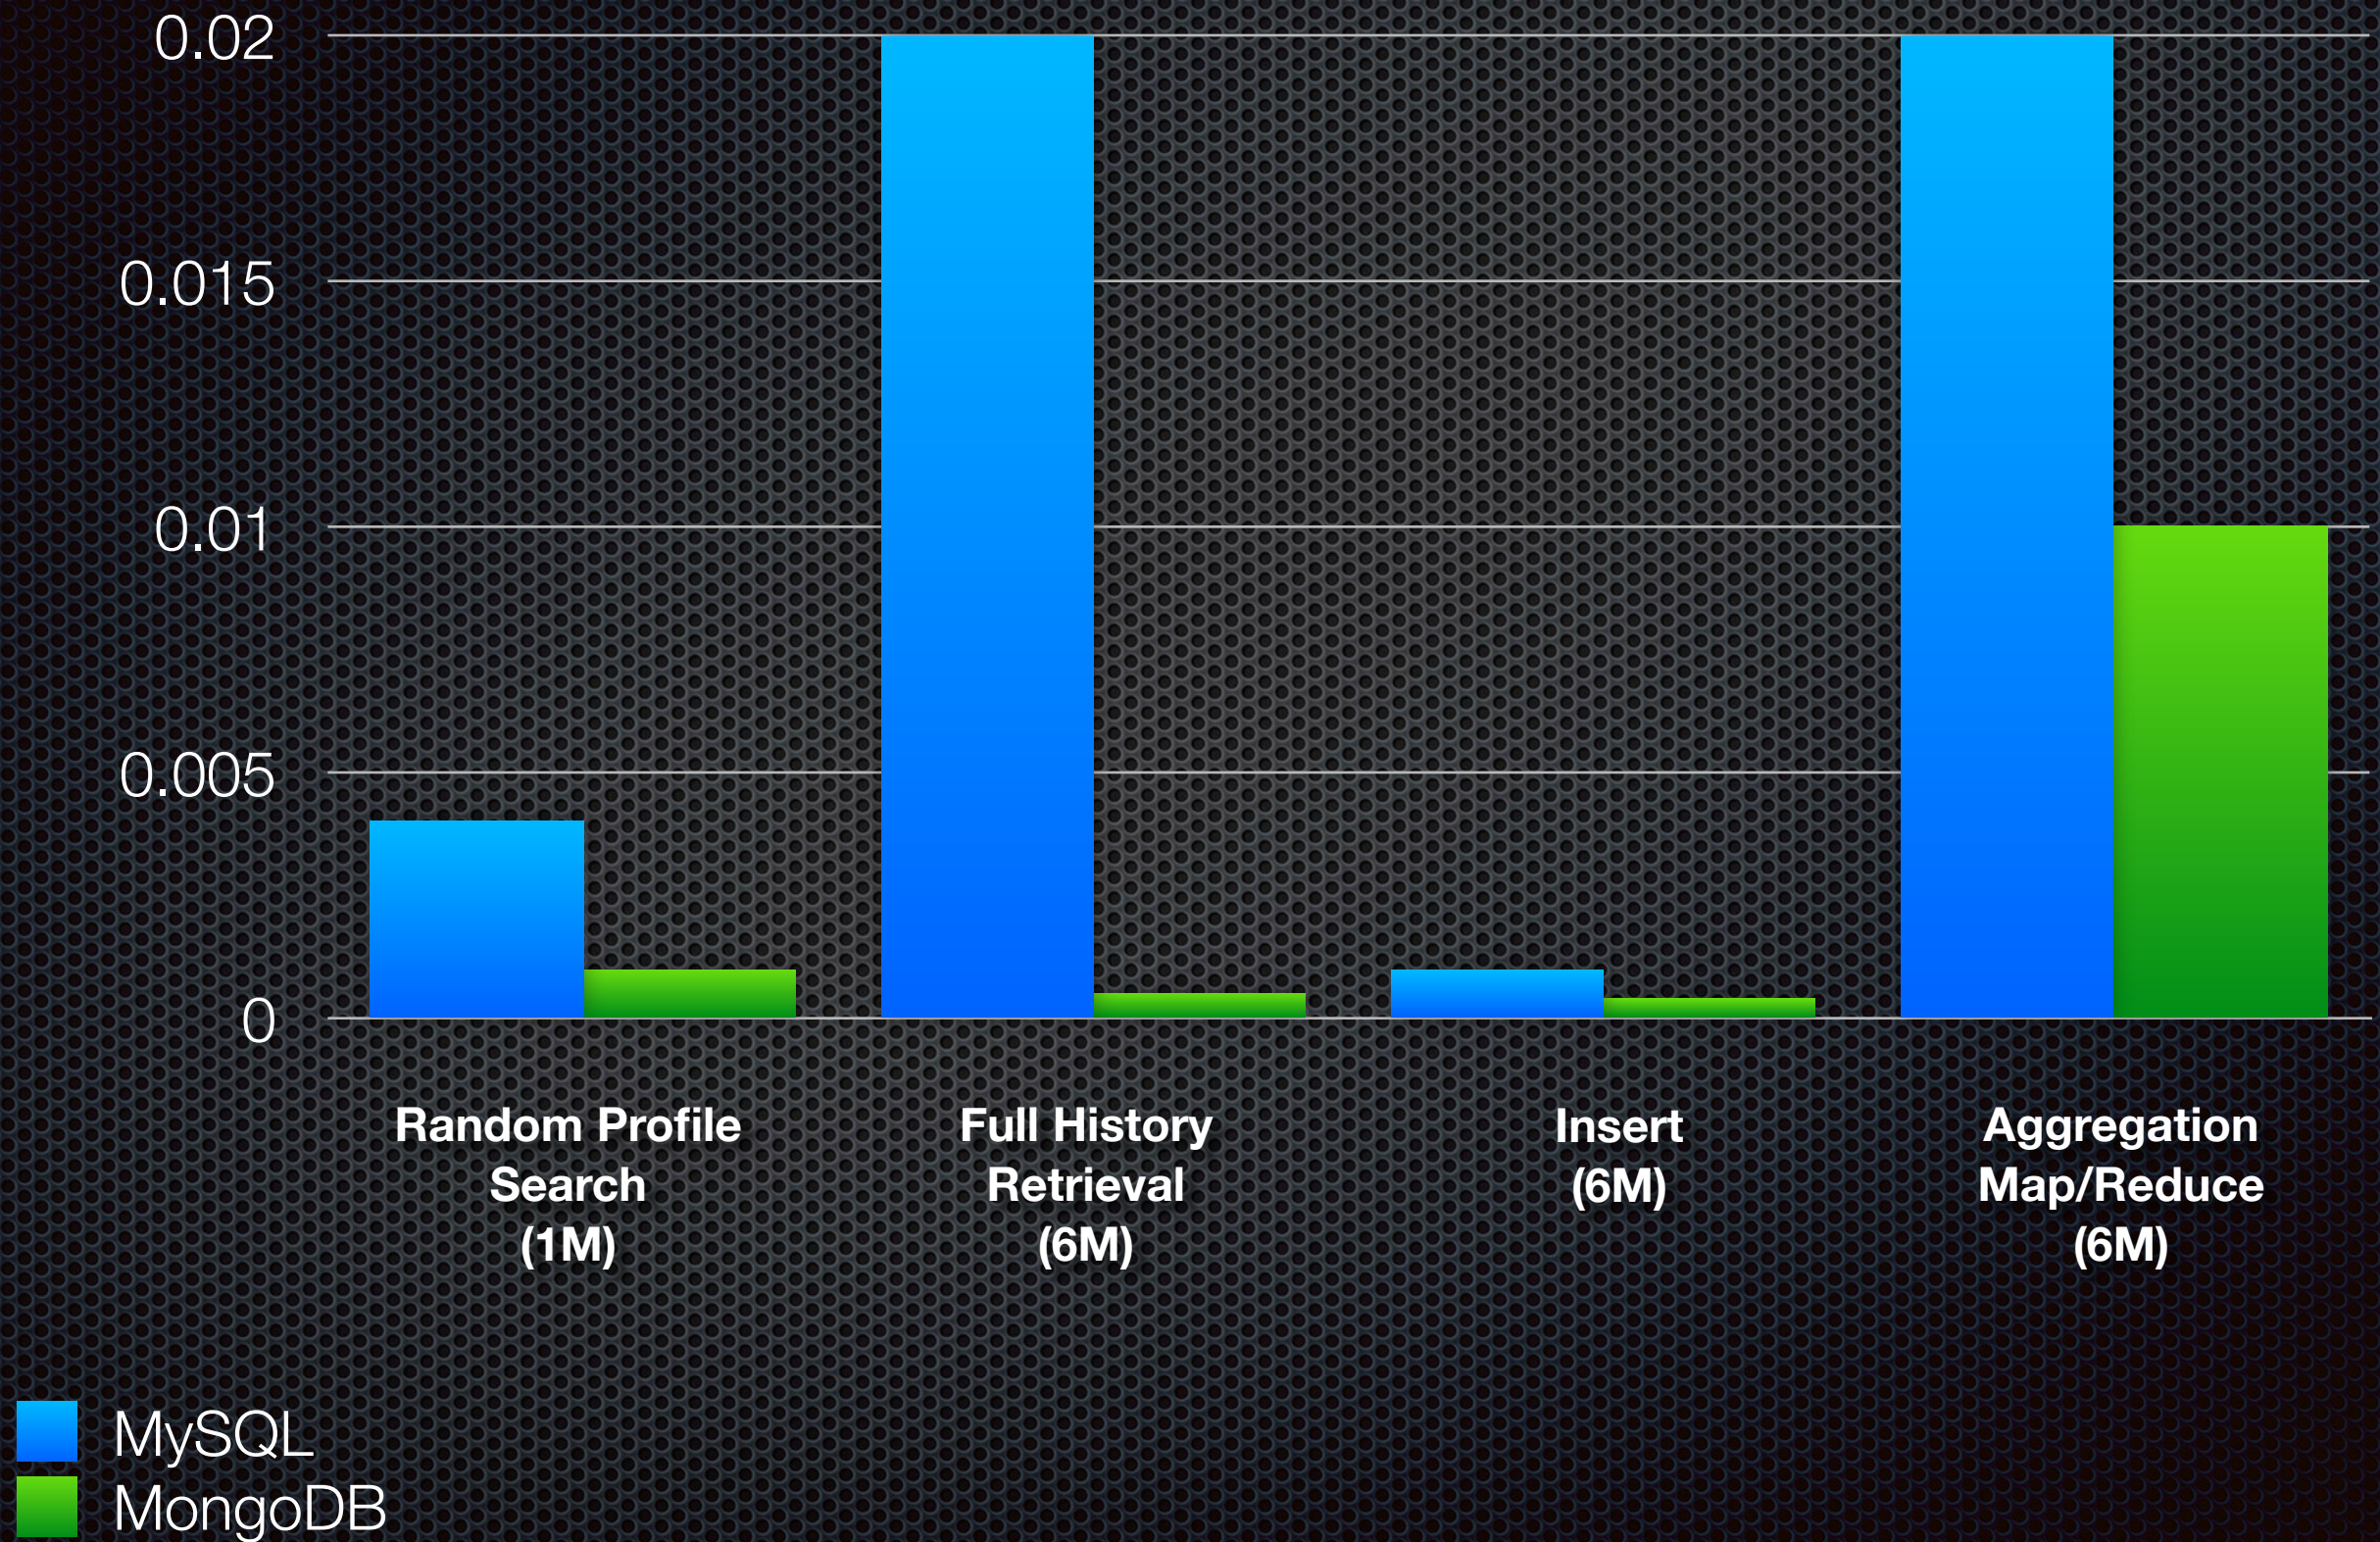


Schema-Less
Aggregated Data
Object Mapping
Large Scale Data
Sharding
Map/Reduce
Eventual Consistency
Memory Storage
Journaling



Denormalized data
Disk Space
Expertise
Complex Querying
Resource Usage
DB-Level Locking

BENCH
MARK






Relational Data Model
Ease of Use



Schema Less
Performance
Sharding

GO!

Hybrid



Concerns



Inconsistency



Consistency vs. Latency



Data Duplication



Complexity

Challenges



Availability



Time



Expertese

Data Access Layer

OR



Plans



Backup



Failure


Plans

Backup

Failure



DATA Migration



1

One Shot Migration

2

Profile Versioning

1

One Shot Migration

4 hrs !!

2

Profile Versioning

1


One Shot Migration

4 hrs !!

2

Profile Versioning

Live!



TEST!

1

Unit

2


Functional

3

Integration

Final Results

Scalability



&...

Details



User



B.I.

Response Time

0.02"  0.00002"


Architecture

2013



Architecture

2013



> 1 year uptime

1

Based on Requirements

2

Benchmark!

3

Choose Wisely!


Thank **YOU!**




<https://joind.in/9977>

References


NoSQL Distilled



MongoDB The definitive guide



MySQL High Performance



we need
YOU!

<http://montreal.ubisoft.com/en/video-games-jobs-my-career>

Got
Q?



@majidfn