



Object-oriented Enterprise Application
Programming Model for In-Memory Databases

Minimal Projection Enforcement

Minimal Projections

2

- AKA “How to avoid SELECT *”
- As a developer, it’s convenient to write SELECT * and don’t think about which attributes are going to be used
 - See analysis of SAP ERP code
- ORMs use SELECT * / full projections by default
- Problem:
 - It’s slower, since unused attributes are transmitted to the client
- Solutions:
 - Educate developers about the issue (?)
 - Make ORMs smarter


How serious is this?

3

- Tested on MySQL 5
 - TODO: test on column-based DB
- 50 columns (mixed varchar(255) & integer)
 - Average row length = 850 Bytes
- Vary number of SELECTed columns
 - Fastest case: SELECT column1 FROM table
 - Worst case: SELECT * FROM table
- Use MySQL profiler to analyze where time is spent

SELECT 50 rows

4

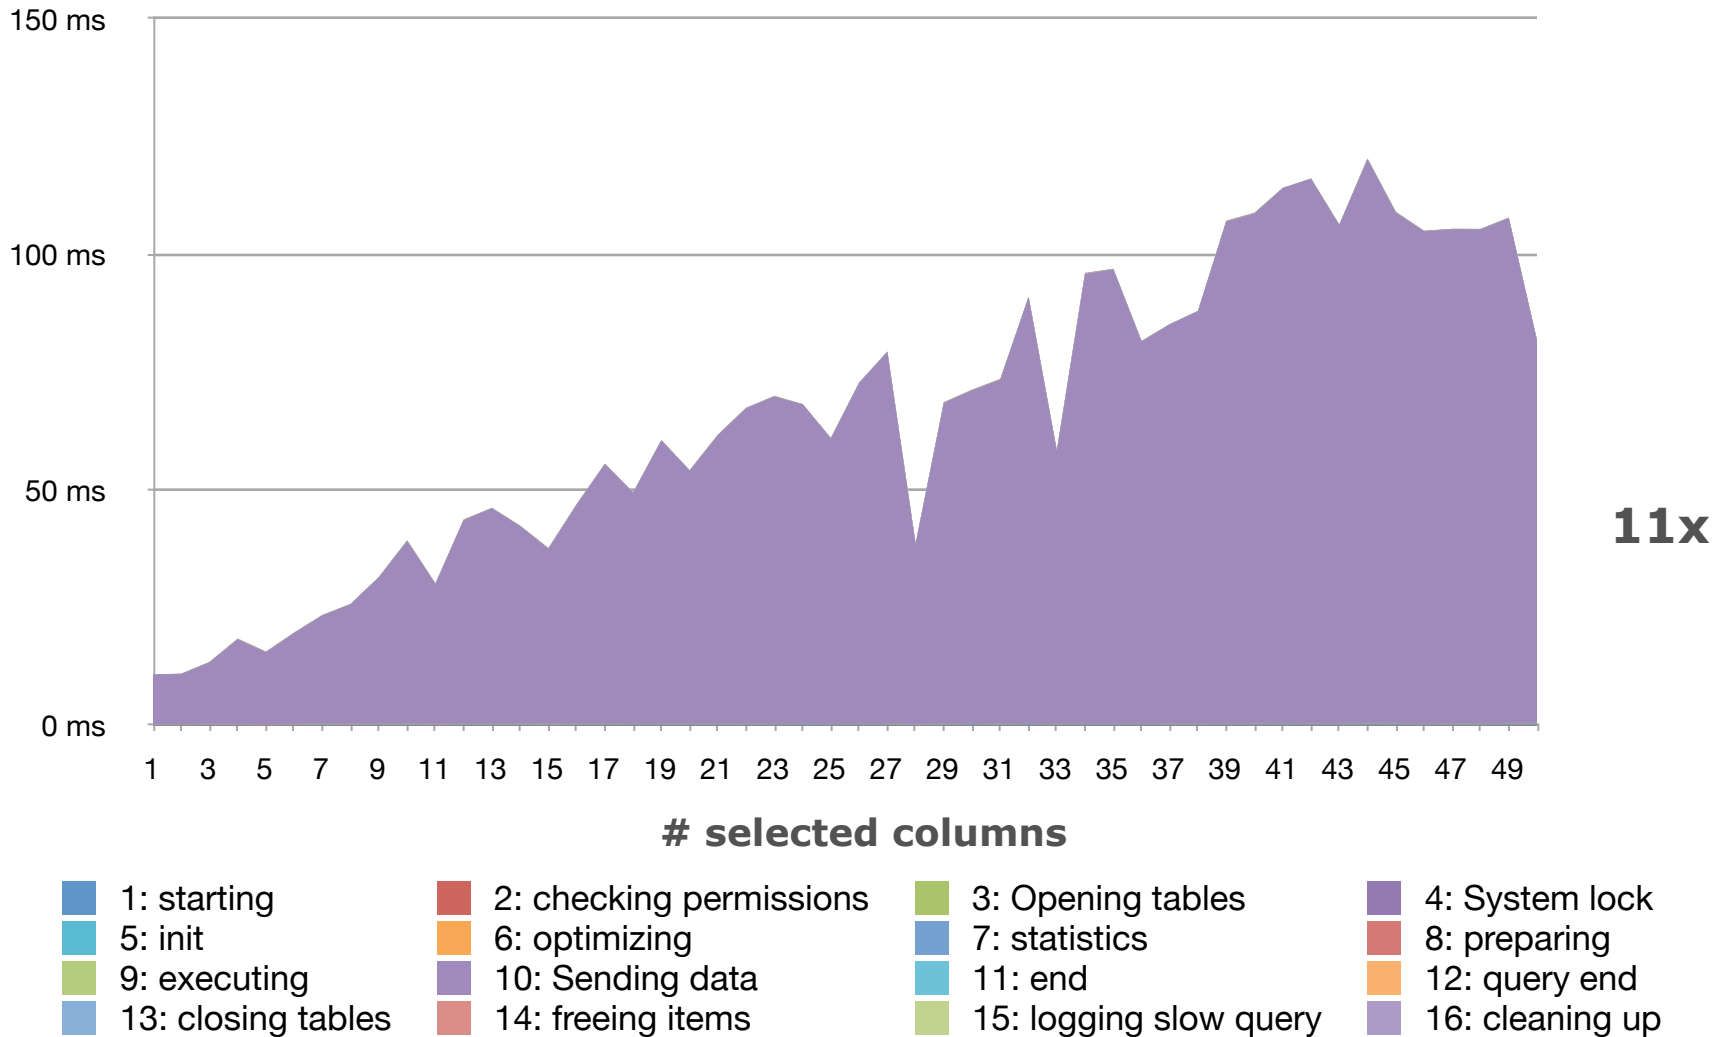


5x

- 1: starting
- 2: checking permissions
- 3: Opening tables
- 4: System lock
- 5: init
- 6: optimizing
- 7: statistics
- 8: preparing
- 9: executing
- 10: Sending data
- 11: end
- 12: query end
- 13: closing tables
- 14: freeing items
- 15: logging slow query
- 16: cleaning up

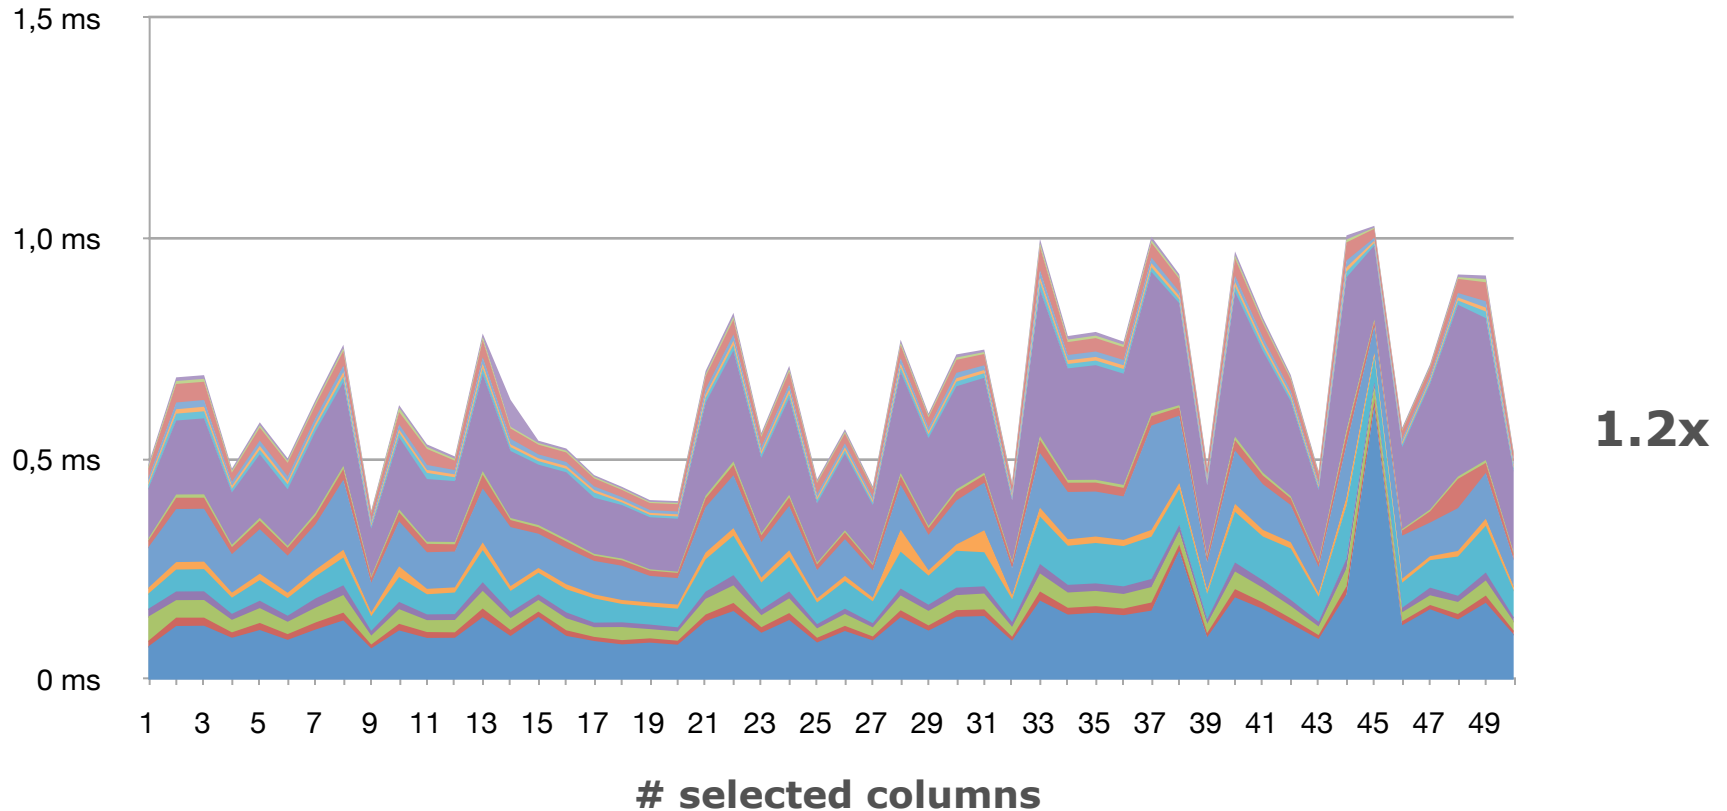
SELECT 5000 rows

5



SELECT 5 rows

6



- 1: starting
- 2: checking permissions
- 3: Opening tables
- 4: System lock
- 5: init
- 6: optimizing
- 7: statistics
- 8: preparing
- 9: executing
- 10: Sending data
- 11: end
- 12: query end
- 13: closing tables
- 14: freeing items
- 15: logging slow query
- 16: cleaning up

Optimized ORM SELECT

7

- 1) Stats Collection Phase
 - Use query backtrace as unique identifier
 - SELECT *
 - record which attributes have been accessed
- 2) Optimized Execution Phase
 - Lookup stats using query identifier
 - SELECT only attributes that are likely to be used & PKs
 - When unfetched attributes are accessed, load them using the record's PK & update stats

ORM Test Case

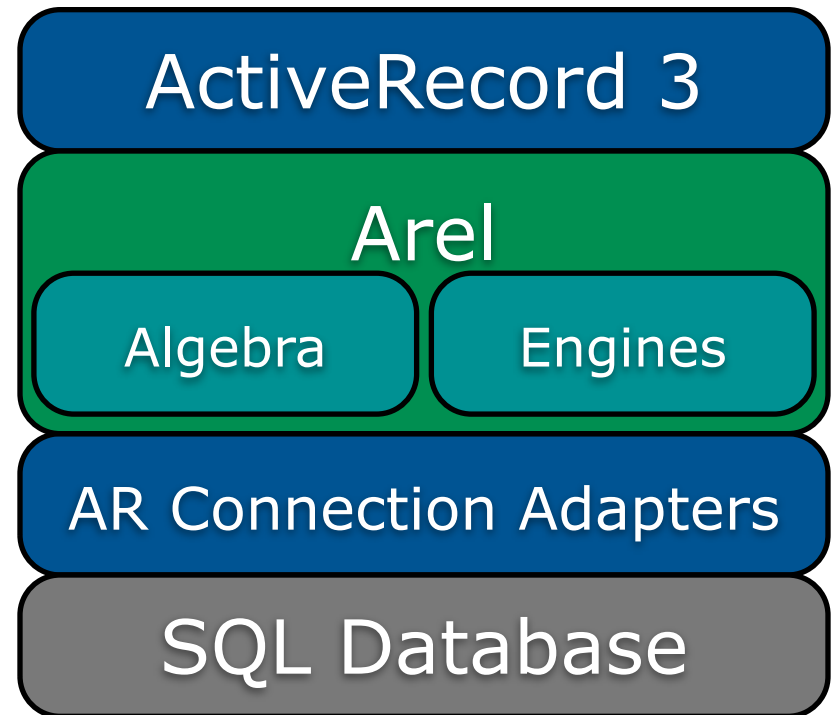
8

- ActiveRecord: default “ORM” used in Ruby on Rails
 - Active record pattern + inheritance + associations
- Arel: An object-oriented interpretation of relational algebra in Ruby

```
unis = University.where(:city  
=> "Potsdam").order("size  
DESC").limit(10)
```

```
unis = unis.select(:name)
```

Ruby on Rails 3:



<http://www.slideshare.net/brynary/arel-ruby-relational-algebra>