

Fast and Scalable Inequality Joins

-- for Data Cleansing on Scale --

Zuhair Khayyat

PhD Candidate @ InfoCloud group
King Abdullah University of Science and Technology (KAUST)

Data Cleansing

- Two customers having the same zip cannot be in different cities

| Name | Zip | City |
|---------|-------|---------------|
| Winnie | 91340 | San Francisco |
| Robbert | 91340 | New York |
| Emma | 91340 | San Francisco |

Data Cleansing

- Two customers having the same zip cannot be in different cities

| Name | Zip | City |
|---------|-------|---------------|
| Winnie | 91340 | San Francisco |
| Robbert | 91340 | New York |
| Emma | 91340 | San Francisco |

- “inaccurate data has a direct impact ... the average company **losing 12% of its revenue**” -- Ben Davis (Econsultancy)

Big

Data Cleansing

- Two customers having the same zip cannot be in different cities

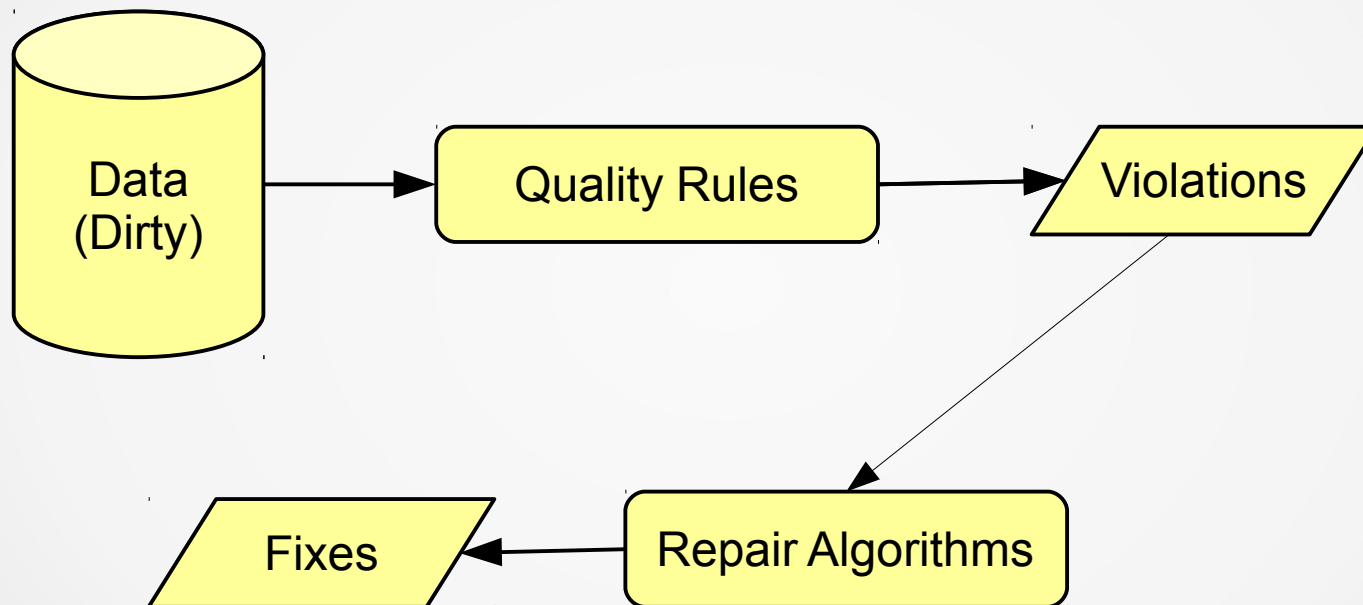
| Name | Zip | City |
|---------|-------|---------------|
| Winnie | 91340 | San Francisco |
| Robbert | 91340 | New York |
| Emma | 91340 | San Francisco |

- “inaccurate data has a direct impact ... the average company **losing 12% of its revenue**” -- Ben Davis (Econsultancy)
- “This is the digital universe. It is **growing 40%** a year into the next decade” -- EMC²

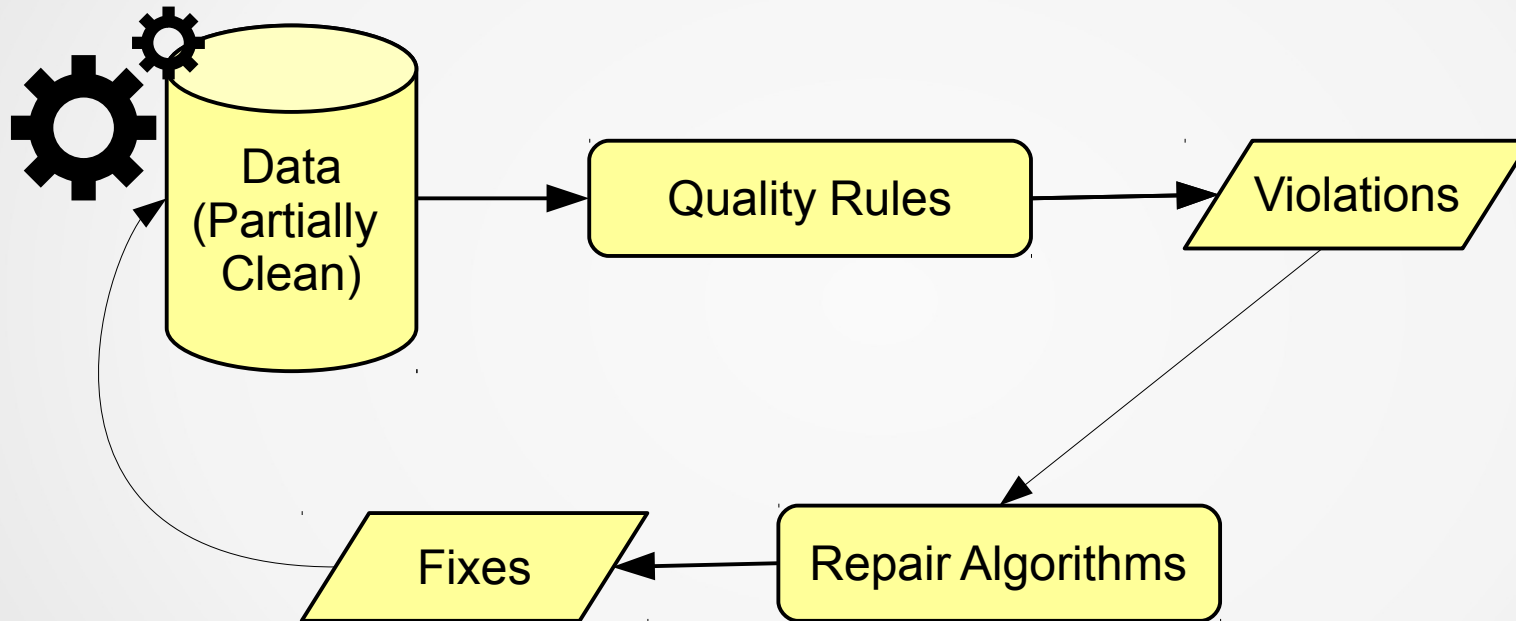
Data Cleansing System



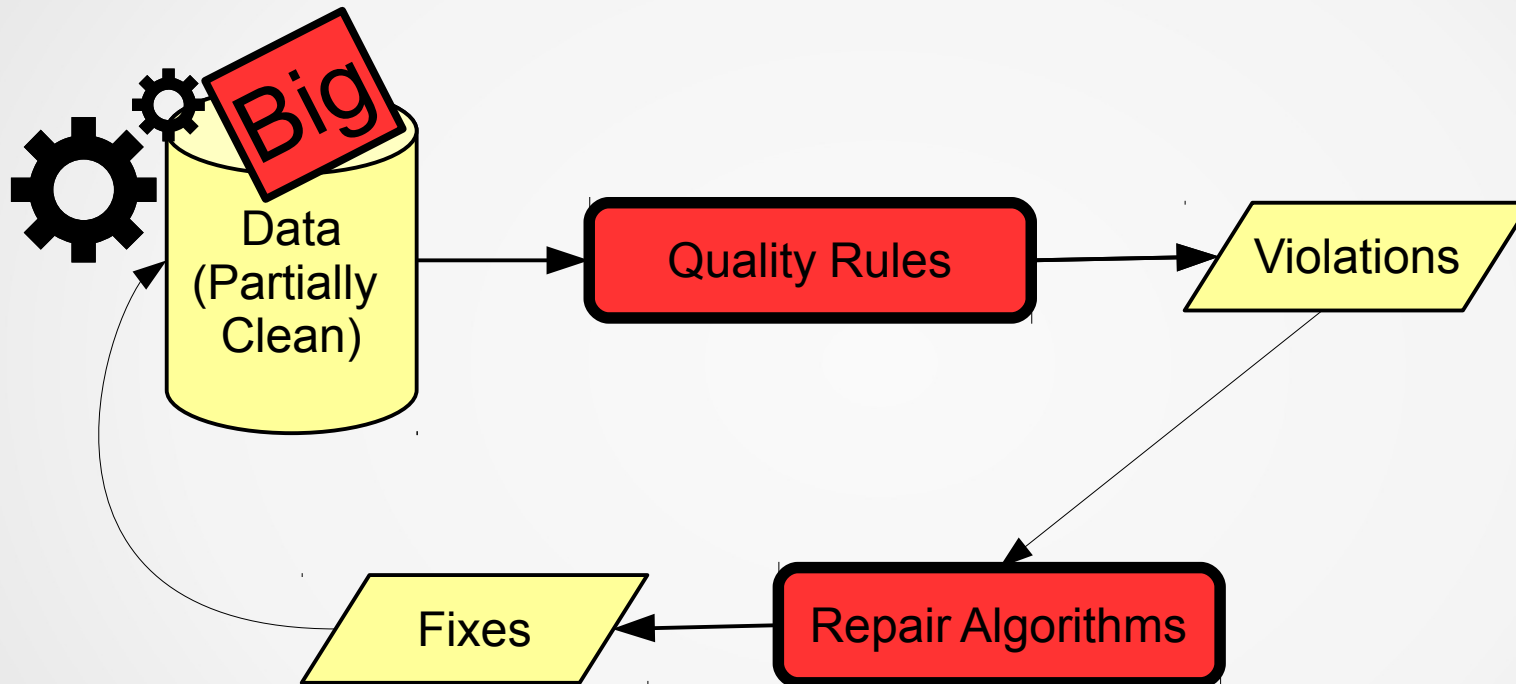
Data Cleansing System



Data Cleansing System

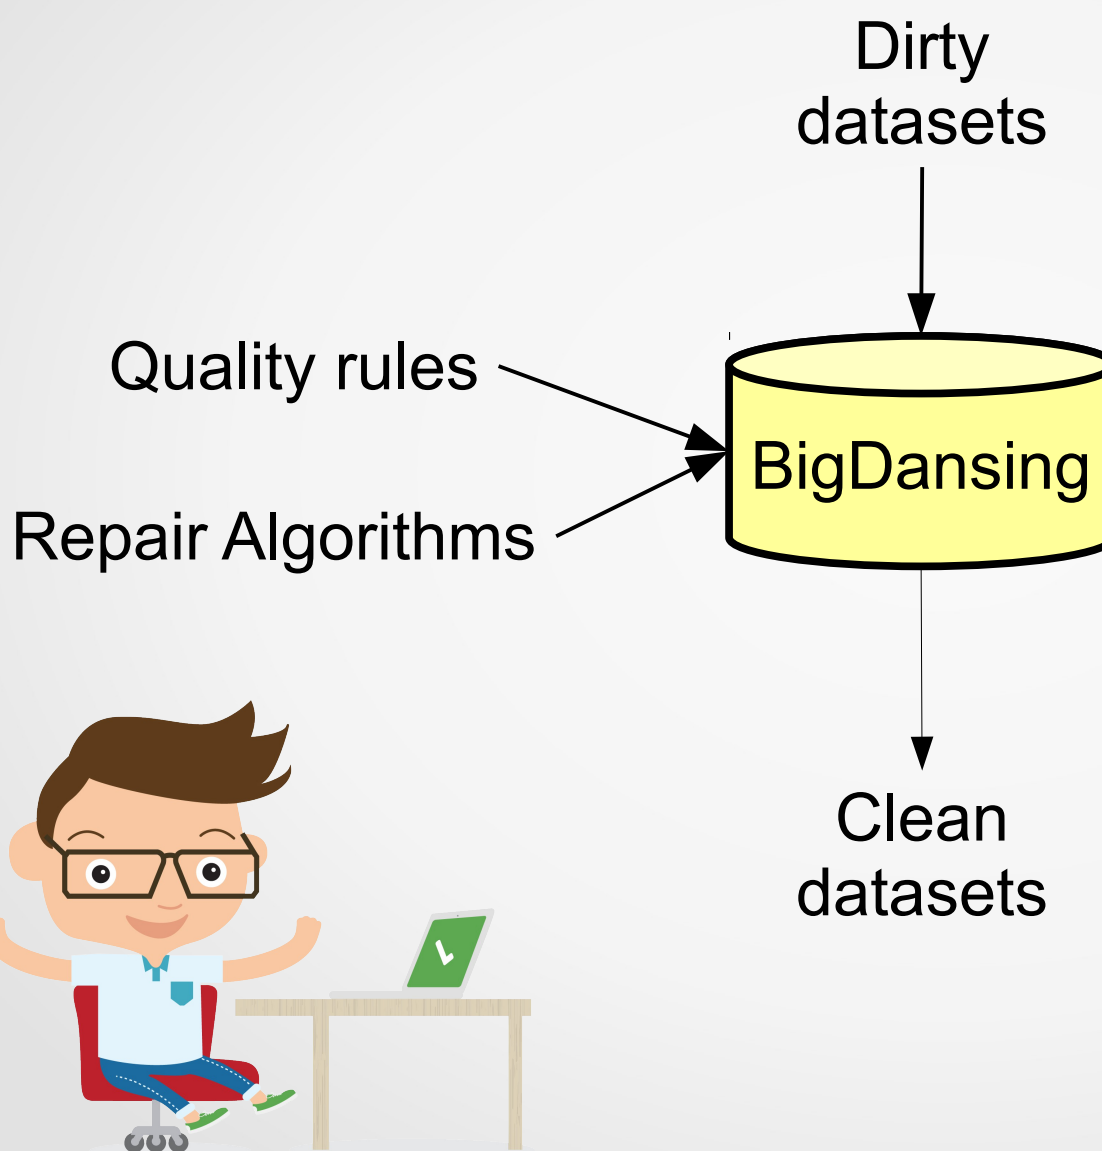


Big Data Cleansing System



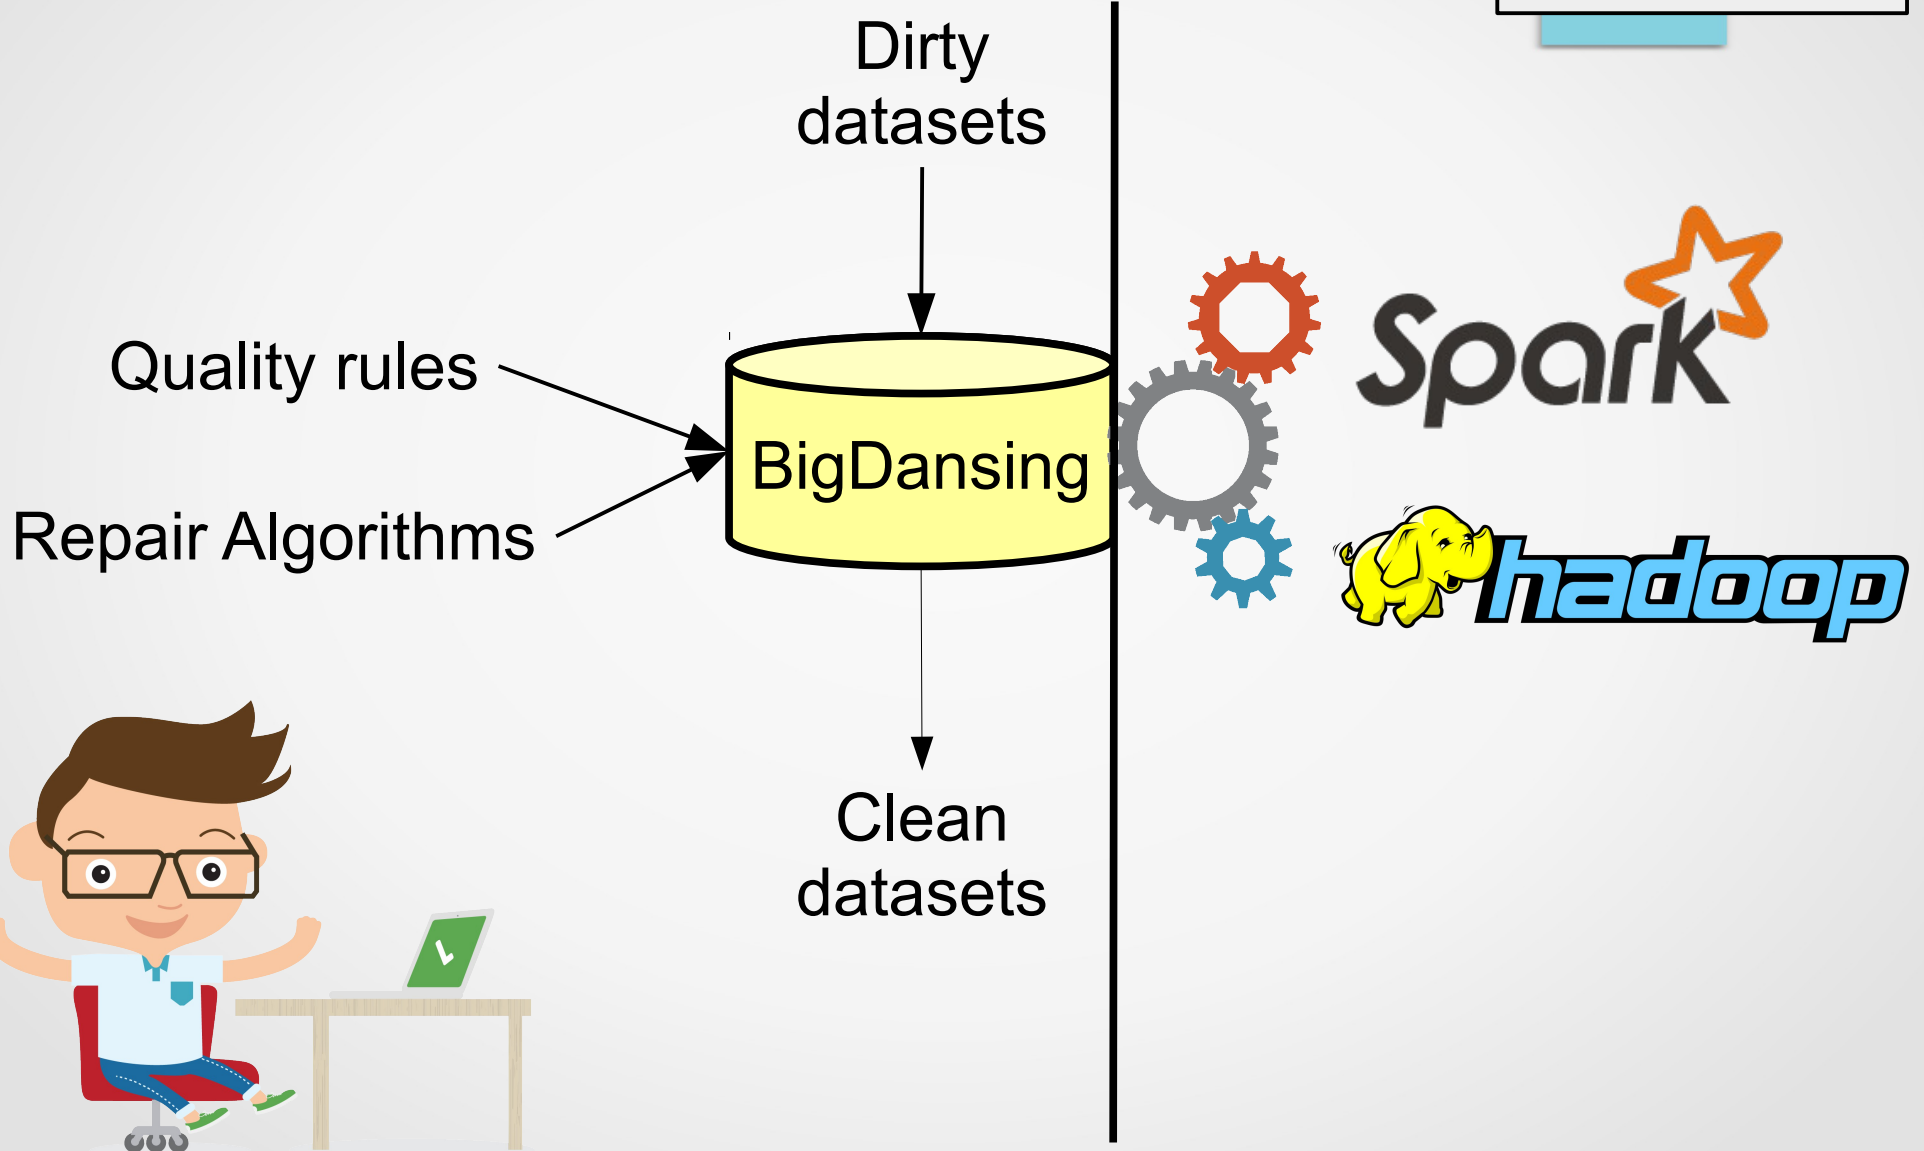
BigDancing: A System for Big Data Cleansing

In SIGMOD 2015



BigDancing: A System for Big Data Cleansing

In SIGMOD 2015



BigDancing: A System for Big Data Cleansing




In SIGMOD 2015

Functional
dependencies

Denial
constraints

Entity
resolution

Inclusion
dependencies



Domain Specific Language



Optimized execution plan


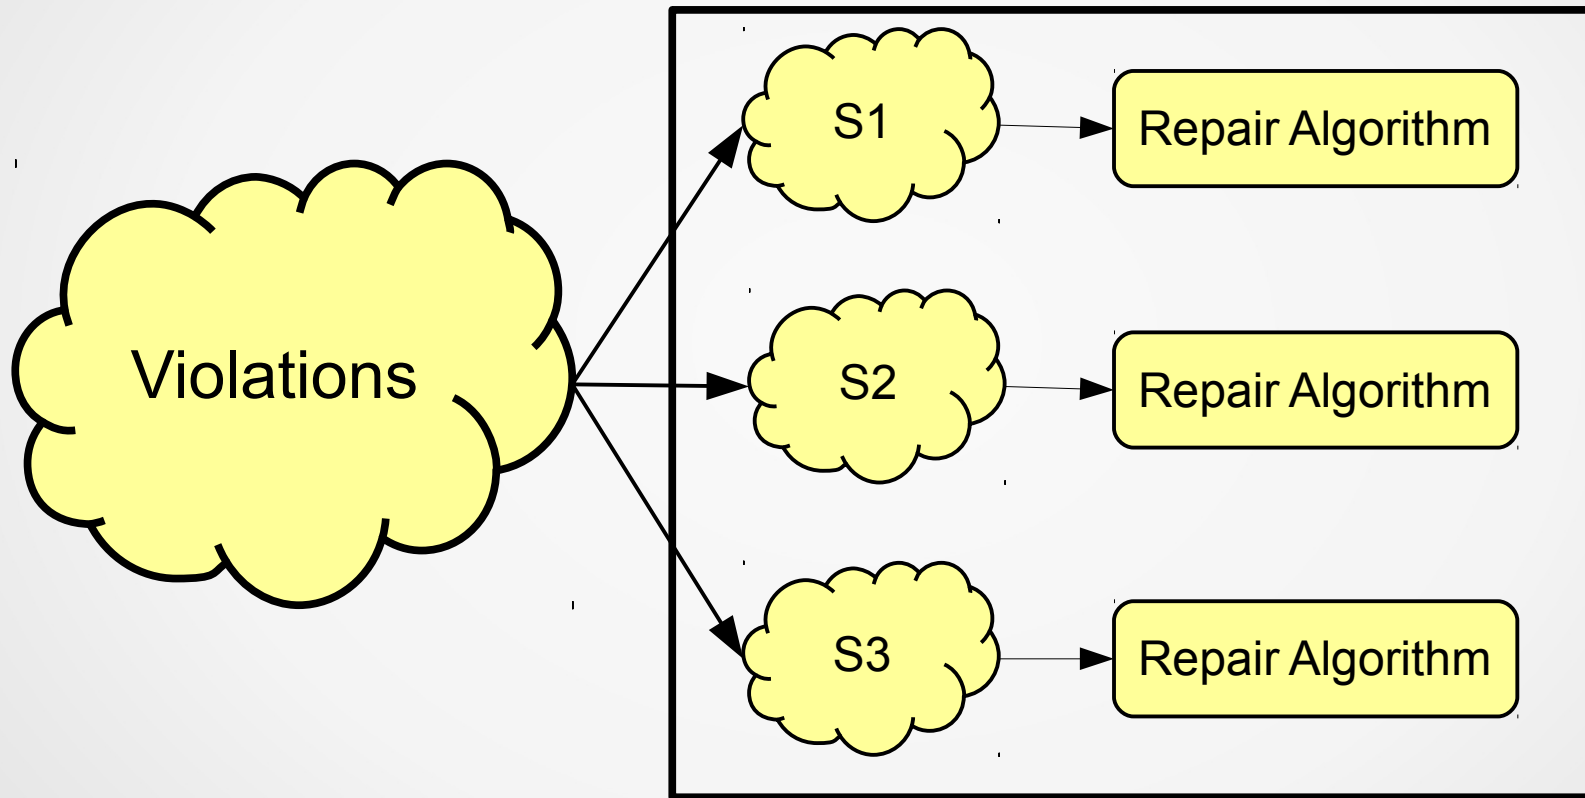


 **Spark**

 **hadoop**

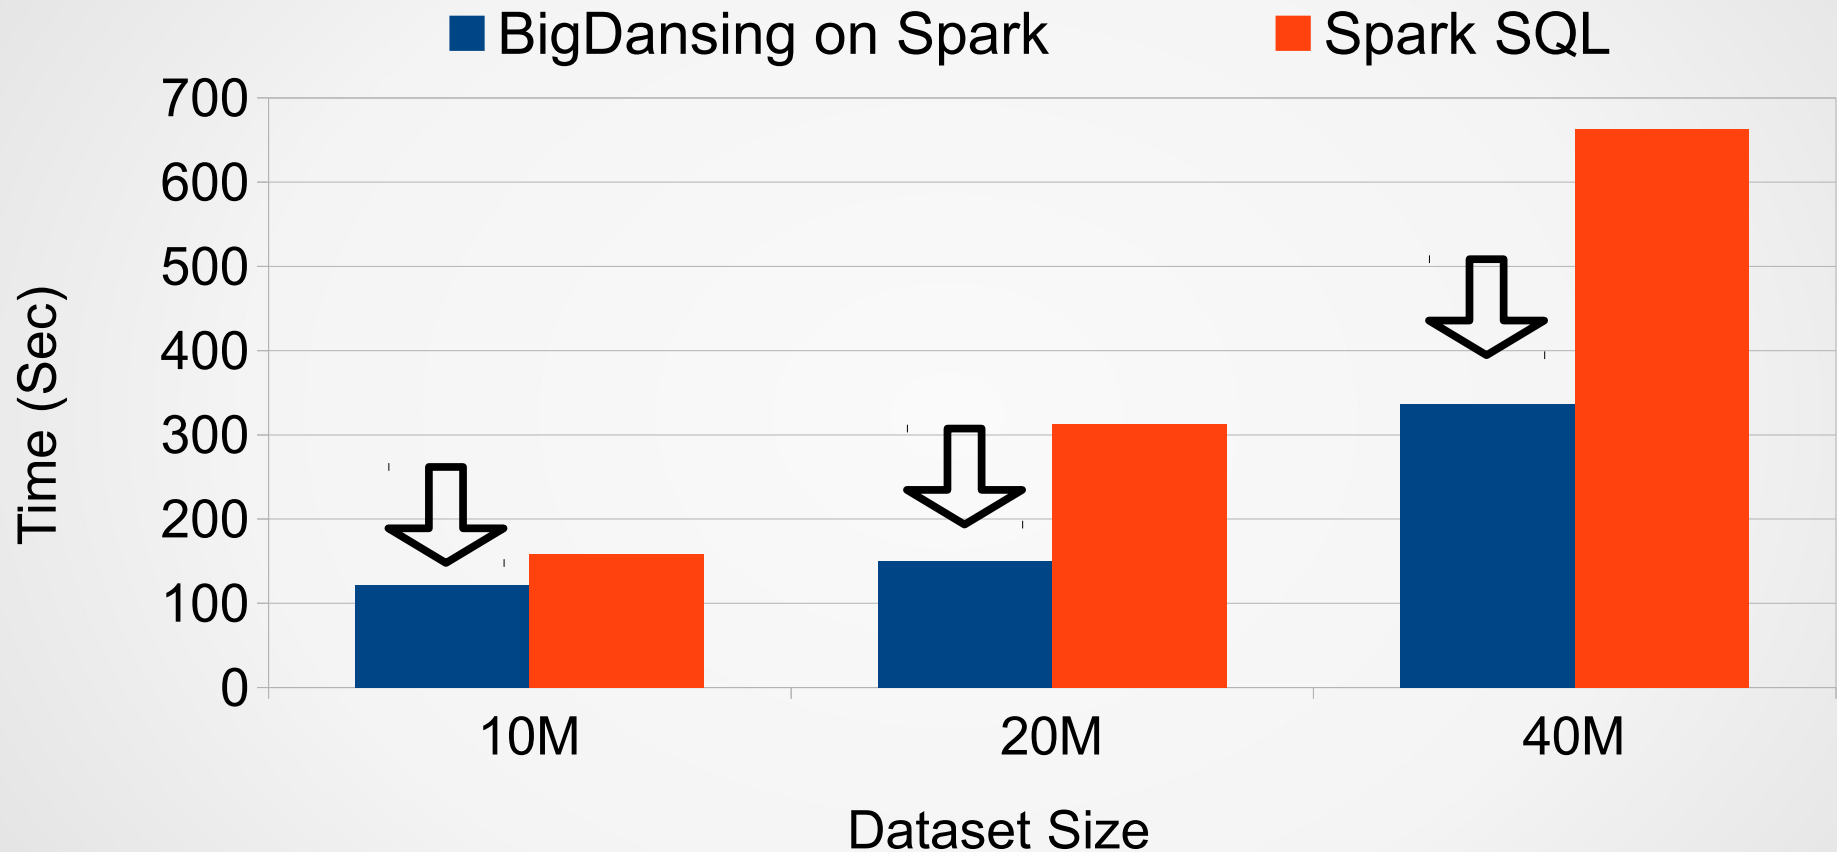
BigDancing: A System for Big Data Cleansing

In SIGMOD 2015



BigDancing: A System for Big Data Cleansing

In SIGMOD 2015



Two customers having the same zip cannot be in different cities
(FD: Zip \rightarrow City)

Complex Quality rule: Inequality joins

- If a person has a higher salary, he must pay more taxes compared to others

Complex Quality rule: Inequality joins

- If a person has a higher salary, he must pay more taxes compared to others
- ```
Select * from D t1 JOIN D t2 on
 t1.Salary > t2.Salary AND
 t1.Tax < t2.Tax;
```
- Processed as a Cartesian product:  $O(n^2)$

# Lightning Fast and Space Efficient Inequality Joins

In VLDB 2015

- Sort on Salary:

|         |         |         |        |
|---------|---------|---------|--------|
| t3(150) | t4(120) | t1(100) | t2(90) |
|---------|---------|---------|--------|



Permutation array:

|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 2 | 3 |
|---|---|---|---|

- Sort on Tax:

|        |        |       |       |
|--------|--------|-------|-------|
| t3(15) | t4(10) | t2(9) | t2(5) |
|--------|--------|-------|-------|



|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 3 | 2 |
|---|---|---|---|



- Bit-array:

|   |   |   |   |
|---|---|---|---|
| 0 | 0 | 0 | 0 |
|---|---|---|---|

| Tuple | Salary | Tax |
|-------|--------|-----|
| t1    | 100    | 5   |
| t2    | 90     | 9   |
| t3    | 150    | 15  |
| t4    | 120    | 10  |



# Lightning Fast and Space Efficient Inequality Joins

In VLDB 2015

- Sort on Salary:

|         |         |         |        |
|---------|---------|---------|--------|
| t3(150) | t4(120) | t1(100) | t2(90) |
|---------|---------|---------|--------|



Permutation array:


|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 2 | 3 |
|---|---|---|---|

- Sort on Tax:

|        |        |       |       |
|--------|--------|-------|-------|
| t3(15) | t4(10) | t2(9) | t2(5) |
|--------|--------|-------|-------|



|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 3 | 2 |
|---|---|---|---|



- Bit-array:

|   |   |   |   |
|---|---|---|---|
| 0 | 0 | 0 | 0 |
|---|---|---|---|



| Tuple | Salary | Tax |
|-------|--------|-----|
| t1    | 100    | 5   |
| t2    | 90     | 9   |
| t3    | 150    | 15  |
| t4    | 120    | 10  |

# Lightning Fast and Space Efficient Inequality Joins

In VLDB 2015

- Sort on Salary:

|         |         |         |        |
|---------|---------|---------|--------|
| t3(150) | t4(120) | t1(100) | t2(90) |
|---------|---------|---------|--------|

Permutation array:

|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 2 | 3 |
|---|---|---|---|

- Sort on Tax:


|        |        |       |       |
|--------|--------|-------|-------|
| t3(15) | t4(10) | t2(9) | t1(5) |
|--------|--------|-------|-------|

|   |   |   |   |
|---|---|---|---|
| 0 | 1 | 3 | 2 |
|---|---|---|---|



- Bit-array:


|   |   |   |   |
|---|---|---|---|
| 0 | 0 | 0 | 0 |
|---|---|---|---|




$O(n \log n)$

| Tuple | Salary | Tax |
|-------|--------|-----|
| t1    | 100    | 5   |
| t2    | 90     | 9   |
| t3    | 150    | 15  |
| t4    | 120    | 10  |

# IEJoin vs. DBMS (Single machine)

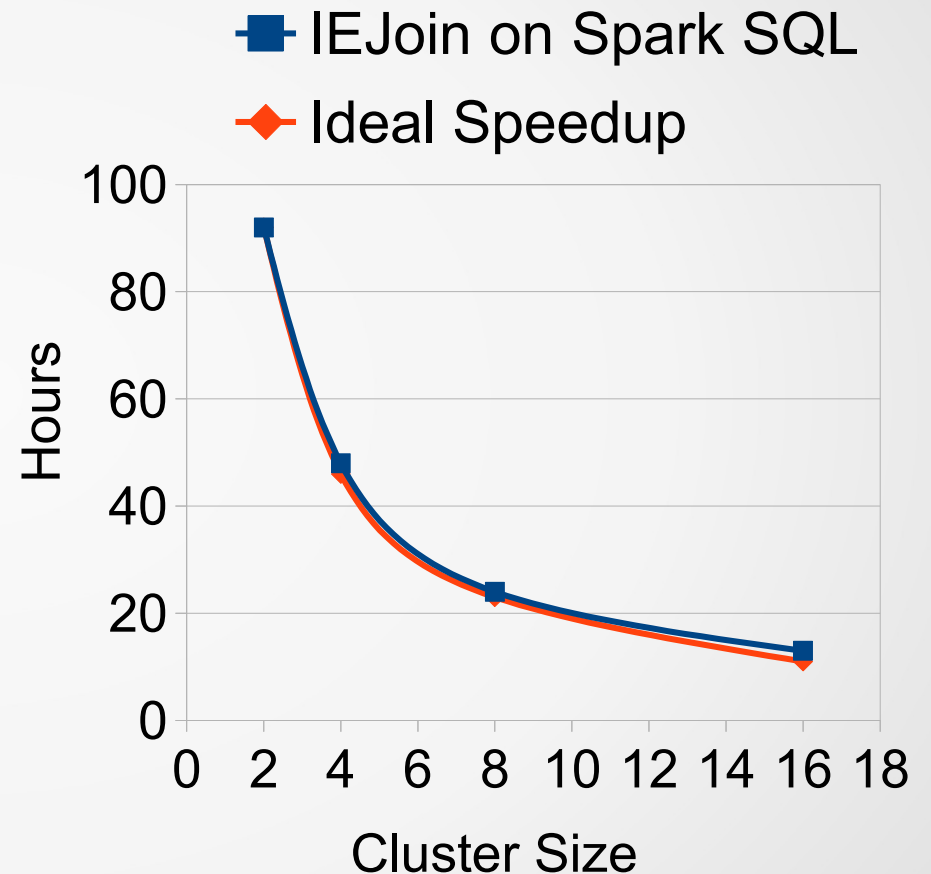


# IEJoin vs. Spark SQL (Distributed)



# IEJoin on 8B Rows

- A cluster of 16 workers
- 8B rows, 287 GB on Disk
- Runtime in 13 hours
- Close to the ideal speedup



# Visit us!

- **Zuhair Khayyat**
  - [cloud.kaust.edu.sa](http://cloud.kaust.edu.sa)
- SIGMOD 15 – BigDancing paper
- VLDB 15 – IEJoin Paper -----> to be presented in VLDB 16
- SIGMOD 16 – Demo Paper