UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Performance Comparison of Technological Solutions for Spark Applications in AWS

Riccardo Lancellotti, Stefano Rossi
University of Modena and Reggio Emilia

Giuseppe Calogero Miano, Fabio Miselli
Doxee

Motivation

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Big data analysis
- Cloud-based approach
  - Available VM images
  - Large scale storage
- Case study
- Doxee
  - Data analysis applications
  - Stream of data and events
  - AWS as cloud platform
- Performance analysis of technologies

doxee

aws

Motivation

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Problem of performance evaluation
- Several approaches
  - Benchmarks
  - Theoretical models
  - Simulation
- Huge differences in:
  - Applications to consider
  - Workload characteristics
- Contribution
  - → Specific case study of big data application

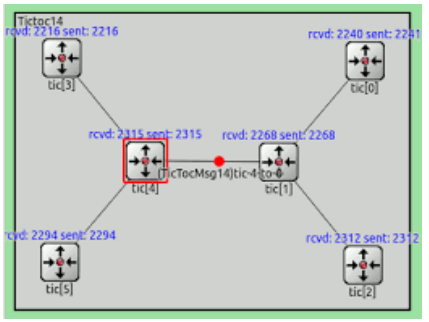

Research questions

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Goal of analysis: compare design alternatives
- RQ1:
  Model for processing time vs. data size in ingestion?
- RQ2:
  Benefits and limit of compression?
- RQ3:
  Best back-end for data for performance?
- RQ4:
  Impact of size on processing time in queries?
- RQ5:
  Cost-effective back-end?

Reference scenario

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Application structure
  - Ingestion
  - Queries in unrefined area
  - Queries in refined area
- Technologies
  - Spark (+Spark.SQL)
  - Hive metastore
  - Hudi storage
  - Tez compression

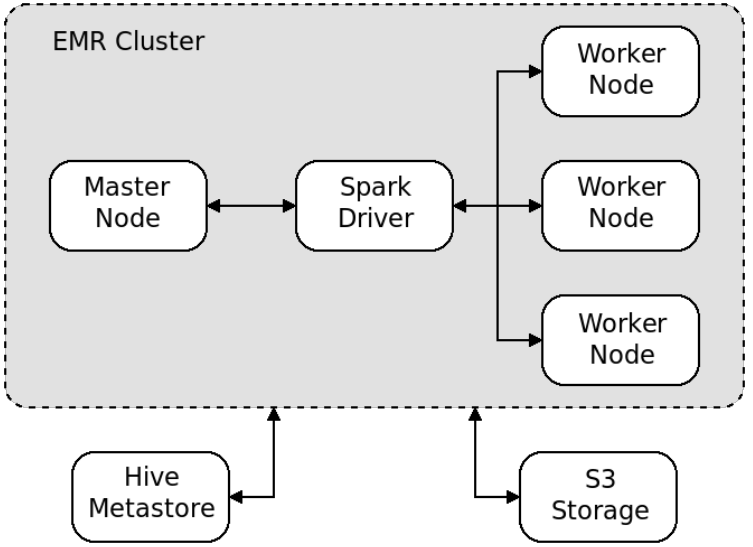

- Use of `m5.xlarge` VMs
- (Anonymized) workload
  - ~40G unrefined area
  - ~10G refined area

Ingestion performance

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Model for processing time vs. data size in ingestion? **RQ1**
- Response time as a function of
  - Workload size
  - # of files
- File of similar size
- Linear model
  - T=20s + Ks × Size
  - T=20s + Kn × NFiles

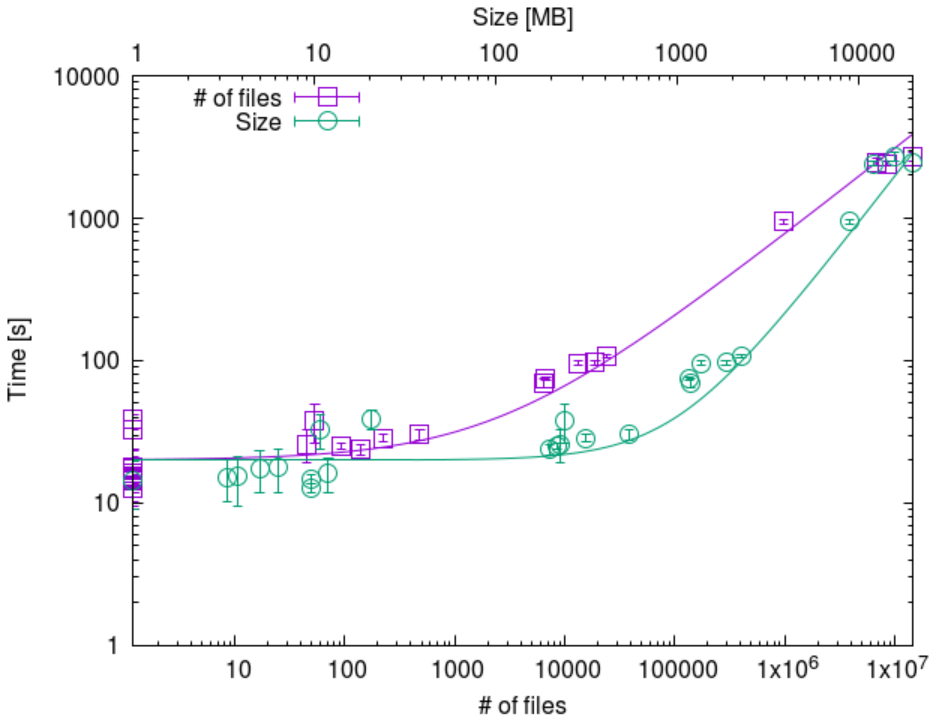

Ingestion performance

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Breakdown of time
  - CPU-intensive tasks
  - I/O-intensive tasks
- Both tasks depend on amount of data processed
  - Explanation of performance model
- How to best manage storage?

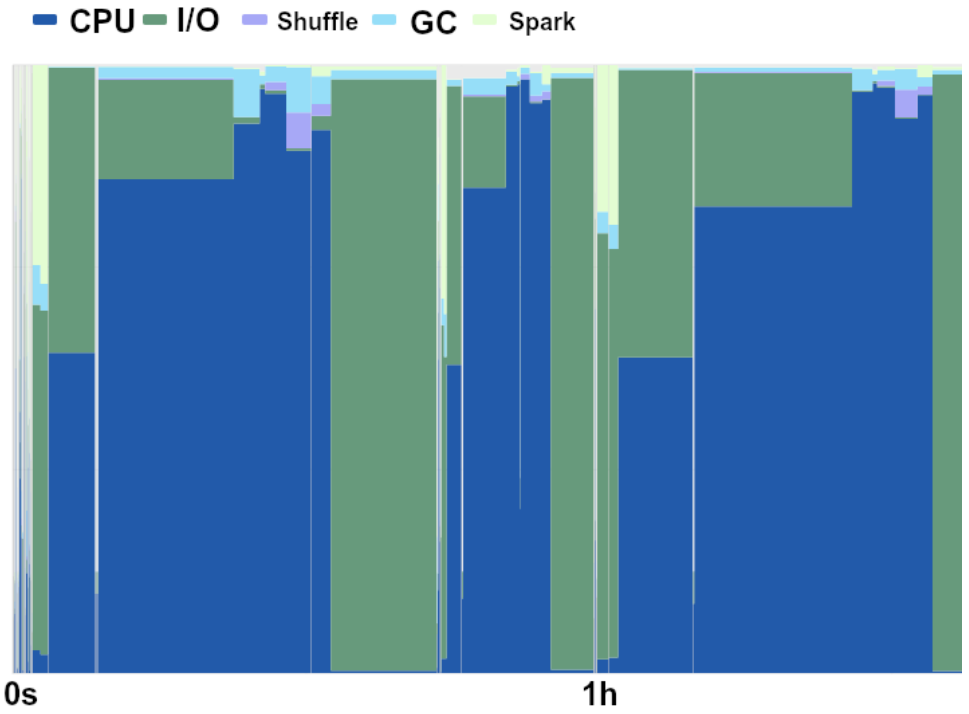

Impact of compression

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Benefits and limit of compression? **RQ2**
- Use of data compression
  - Benefit for I/O tasks + reduced storage costs
  - Higher CPU demand
- Several algorithms considered
  - Compression with gzip, lzw, tez, ...
- Similar results for several algorithms
  - Focus just on tez compression
- Comparison of compressed/uncompressed results
  - Working set size
  - Ingestion time

Impact of compression

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

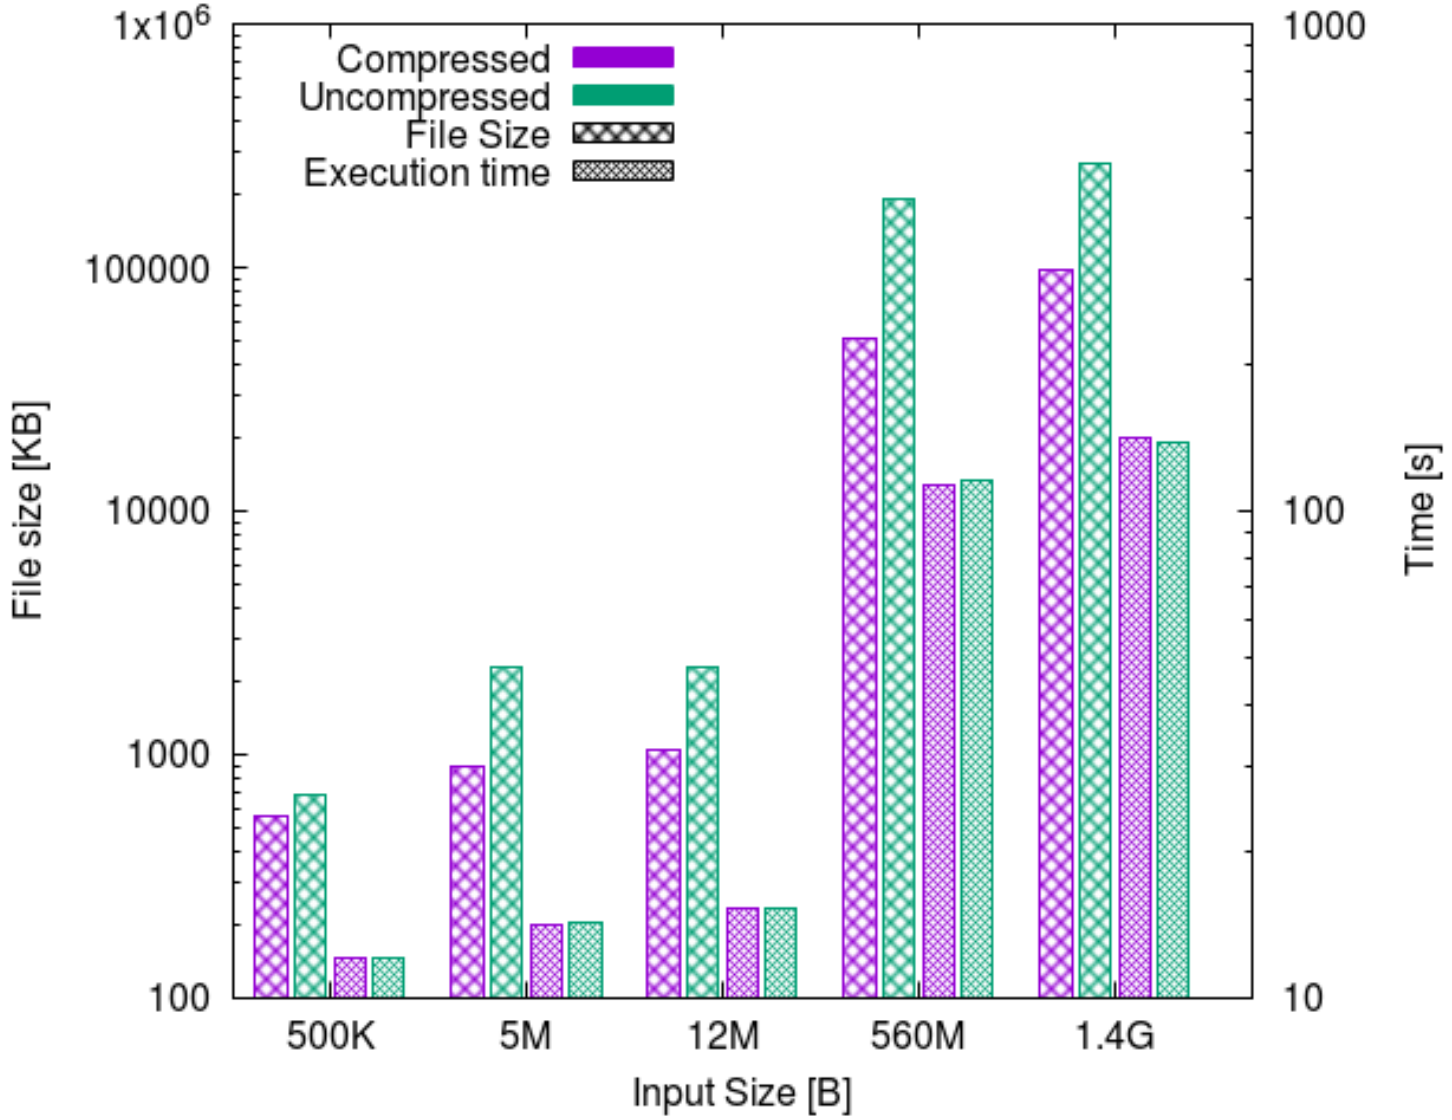

Queries in unrefined area

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Best back-end for data for performance?
- Queries on ingested data
- Several queries in sequence
  - Single query time
  - Total execution time (sequence of query)
- Breakdown
  - Setup/Execution
- Comparison of back-ends
  - Spark/Spark+Hive

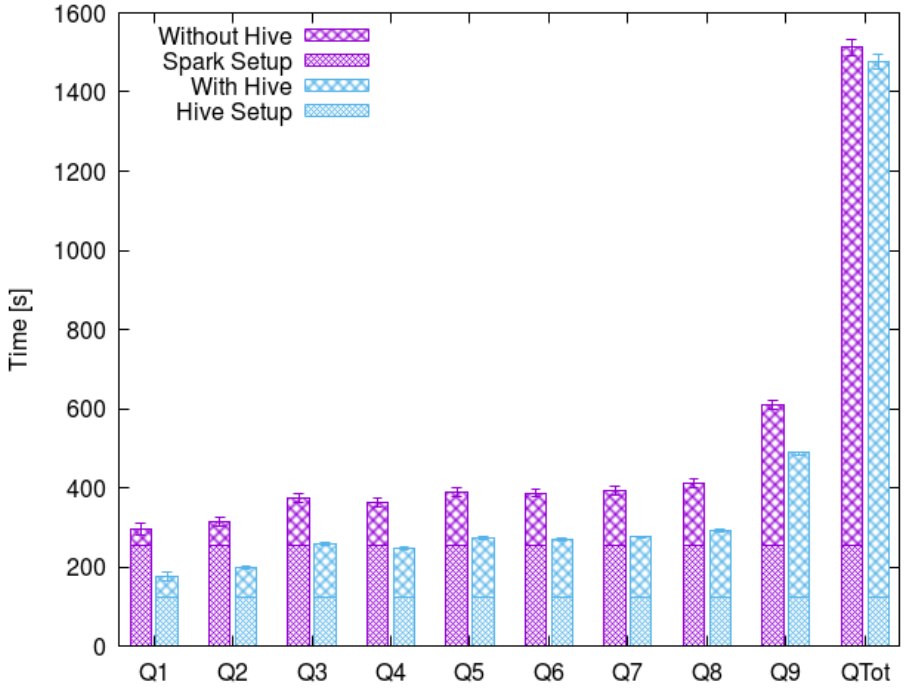

- Hive
  - Faster
  - Long setup time
- Long sequence of queries balances setup overhead

Queries in unrefined area

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Impact of size on processing time in queries?

RQ4

- Analysis as a function of dataset size
- Hive setup
  - Constant time
- Spark setup
  - Depends on data size
- Spark setup
  - Faster for small dataset

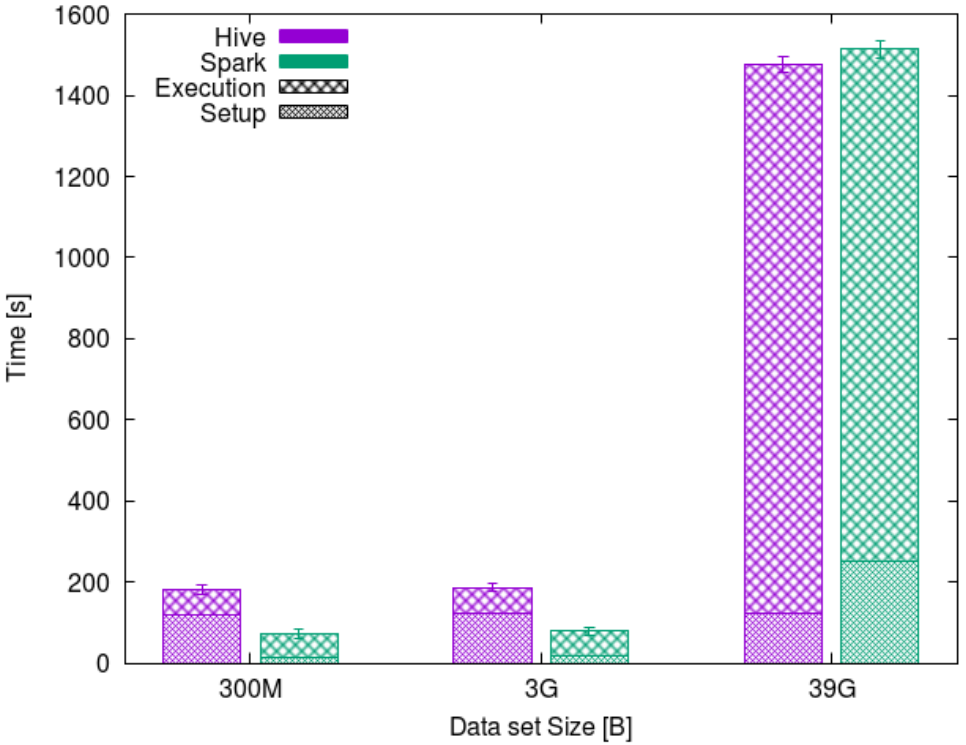

- Hive faster query execution overweights setup only for large working set

Queries in refined area

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Impact of size on processing time in queries?
- Refined area
  - Smaller working set
  - Complex queries
- Impact of cluster size
- Setup time → constant
- Execution time → inversely proportional
  - Amdahl's law

RQ4

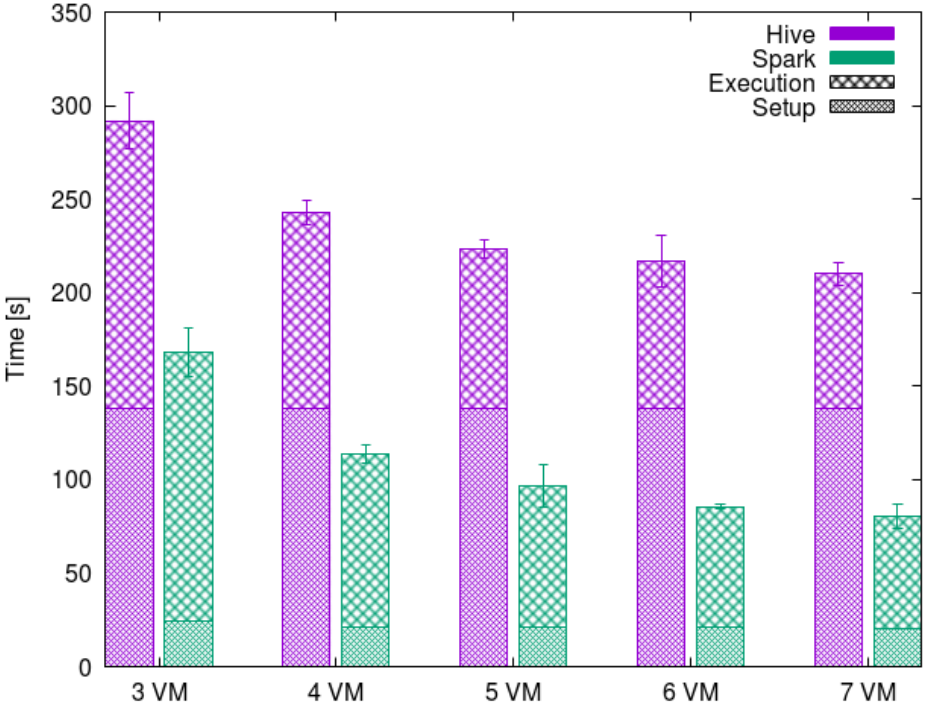

- Setup time remains dominating

Queries in refined area

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Cost-effective back-end?
- Cost of VMs by second
  - Consider multiple executions on same cluster
- Setup cost grows with cluster size
- Execution cost can be reduced
- Hive → no economic gain
- Spark → sweet spot for 4 VMs

RQ5

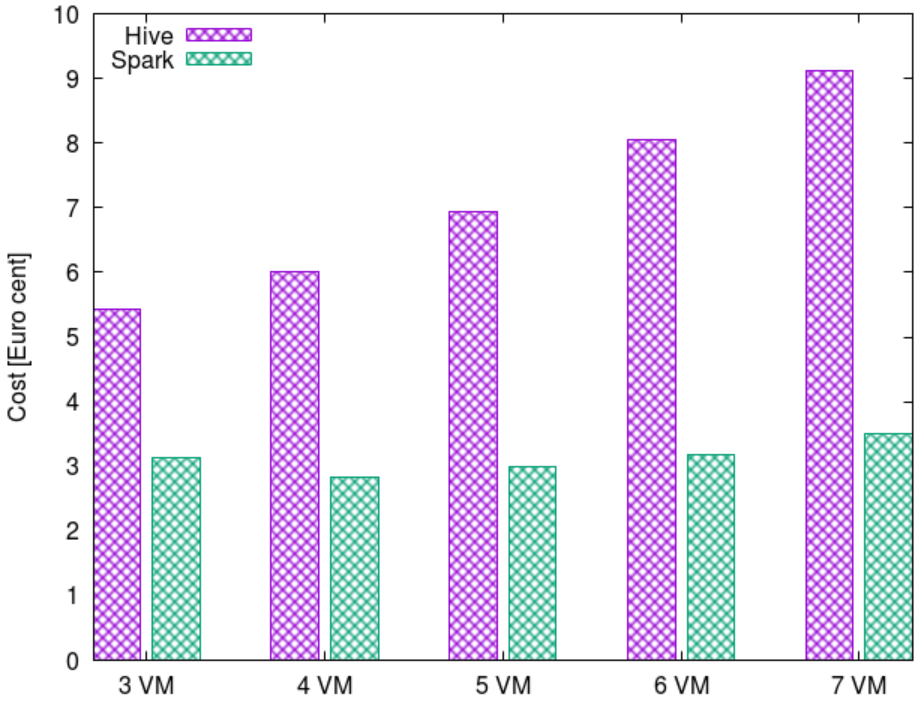

Conclusion

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

- Performance analysis
  - Realistic big data analysis application
  - Comparison of alternatives
- Critical impact of setup phases
  - Significant result of experiments
  - Negligible for “many TB” applications
  - Major effect for smaller scenarios

UNIMORE
UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

Performance Comparison of Technological Solutions for Spark Applications in AWS

Riccardo Lancellotti, Stefano Rossi
University of Modena and Reggio Emilia

Giuseppe Calogero Miano, Fabio Miselli
Doxee