PRIMEBALL
Parallel Processing Framework Benchmark, for BigData Applications in the Cloud

Jaume Ferrarons, Mulu Adhana, Carlos Colmenares, Sandra Pietrowska, Fadila Bentayeb and Jérôme Darmont

Université Lumière Lyon 2

August 26th, 2013

Outline

Outline

Introduction

Introduction

Context

- Massive increasing volume of data produced and stored worldwide.
- Traditional approaches and techniques for data analysis proved limited because:
  - lack of parallelism.
  - lack of fault tolerance.

Introduction

Context

- Massive increasing volume of data produced and stored worldwide.
- Traditional approaches and techniques for data analysis proved limited because:
  - lack of parallelism.
  - lack of fault tolerance.
- Recently many platforms for parallel processing have been created to satisfy this need.
- It has become important to fix standards to allow accurate comparisons of these frameworks.

Introduction

Context

- Massive increasing volume of data produced and stored worldwide.
- Traditional approaches and techniques for data analysis proved limited because:
  - lack of parallelism.
  - lack of fault tolerance.
- Recently many platforms for parallel processing have been created to satisfy this need.
- It has become important to fix standards to allow accurate comparisons of these frameworks.

PRIMEBALL is a complete and unified cloud benchmark, two main axes are involved:

- Parallel processing frameworks.
- Cloud computing service providers.

Related work

Related work

PRIMEBALL is the first benchmark for BigData in the cloud, but similar ones are:

Related work

PRIMEBALL is the first benchmark for BigData in the cloud, but similar ones are:

- **TPC-DS:** a decision support benchmark.
- **MalStone:** benchmark for data intensive computing and analysis.
- **Cloud Harmony:** measures performance of cloud providers as black boxes.
- **YCSB:** framework to facilitate performance comparisons among cloud database systems.
- **SWIM:** is a Statistical Workload Injector for MapReduce.

Related work

PRIMEBALL is the first benchmark for BigData in the cloud, but similar ones are:

- **TPC-DS:** a decision support benchmark.
- **MalStone:** benchmark for data intensive computing and analysis.
- **Cloud Harmony:** measures performance of cloud providers as black boxes.
- **YCSB:** framework to facilitate performance comparisons among cloud database systems.
- **SWIM:** is a Statistical Workload Injector for MapReduce.

	TPC-DS	MalStone	Cloud Harmony	YCSB	SWIM	PRIMEBALL
Real data	$\sim$	No	$\sim$	No	No	Yes
Real workload	Yes	No	$\sim$	No	Yes	Yes
Parallel processing	Yes	Yes	Yes	Yes	Yes	Yes
Hardware-oriented	No	No	Yes	No	No	No
MapReduce-oriented	No	No	No	No	Yes	$\sim$
Cloud properties	$\sim$	Yes	No	Yes	Yes	Yes
Complex data	No	No	Yes	No	Yes	Yes
Big data	$\sim$	Yes	Yes	Yes	Yes	Yes
Technology-indep	Yes	Yes	Yes	No	No	Yes

Table: Comparison of Benchmark Features

Application Model

Application Model

PRIMEBALL aims to provide a solution for an accurate comparison among parallel processing frameworks in the context of BigData applications in the cloud.

Application Model

PRIMEBALL aims to provide a solution for an accurate comparison among parallel processing frameworks in the context of BigData applications in the cloud. PRIMEBALL is set in the context of a fictitious BigData application, named "**New Pork Times' News Hub**".

Application Model

PRIMEBALL aims to provide a solution for an accurate comparison among parallel processing frameworks in the context of BigData applications in the cloud. PRIMEBALL is set in the context of a fictitious BigData application, named "**New Pork Times' News Hub**".

- New Pork Times would be an international news site.

Application Model

PRIMEBALL aims to provide a solution for an accurate comparison among parallel processing frameworks in the context of BigData applications in the cloud. PRIMEBALL is set in the context of a fictitious BigData application, named "**New Pork Times' News Hub**".

- New Pork Times would be an international news site.

PRIMEBALL:

- is a cloud and BigData oriented benchmark.
- does provide meaningful properties for comparing frameworks.
- does provide objective and repeatable experiments.

Application Model

PRIMEBALL aims to provide a solution for an accurate comparison among parallel processing frameworks in the context of BigData applications in the cloud. PRIMEBALL is set in the context of a fictitious BigData application, named "**New Pork Times' News Hub**".

- New Pork Times would be an international news site.

PRIMEBALL:

- is a cloud and BigData oriented benchmark.
- does provide meaningful properties for comparing frameworks.
- does provide objective and repeatable experiments.

PRIMEBALL does not:

- define technical execution details or a storage schema.
- provide expected performance results.
- compare data retrieval or processing algorithms.

Outline

PRIMEBALL Dataset & Schema

PRIMEBALL Dataset & Schema

Mainly three types of files used:
General information (XML), Media files (binary), Metadata (XML).

PRIMEBALL Dataset & Schema

Mainly three types of files used:
General information (XML), Media files (binary), Metadata (XML).
Conceptual Schema

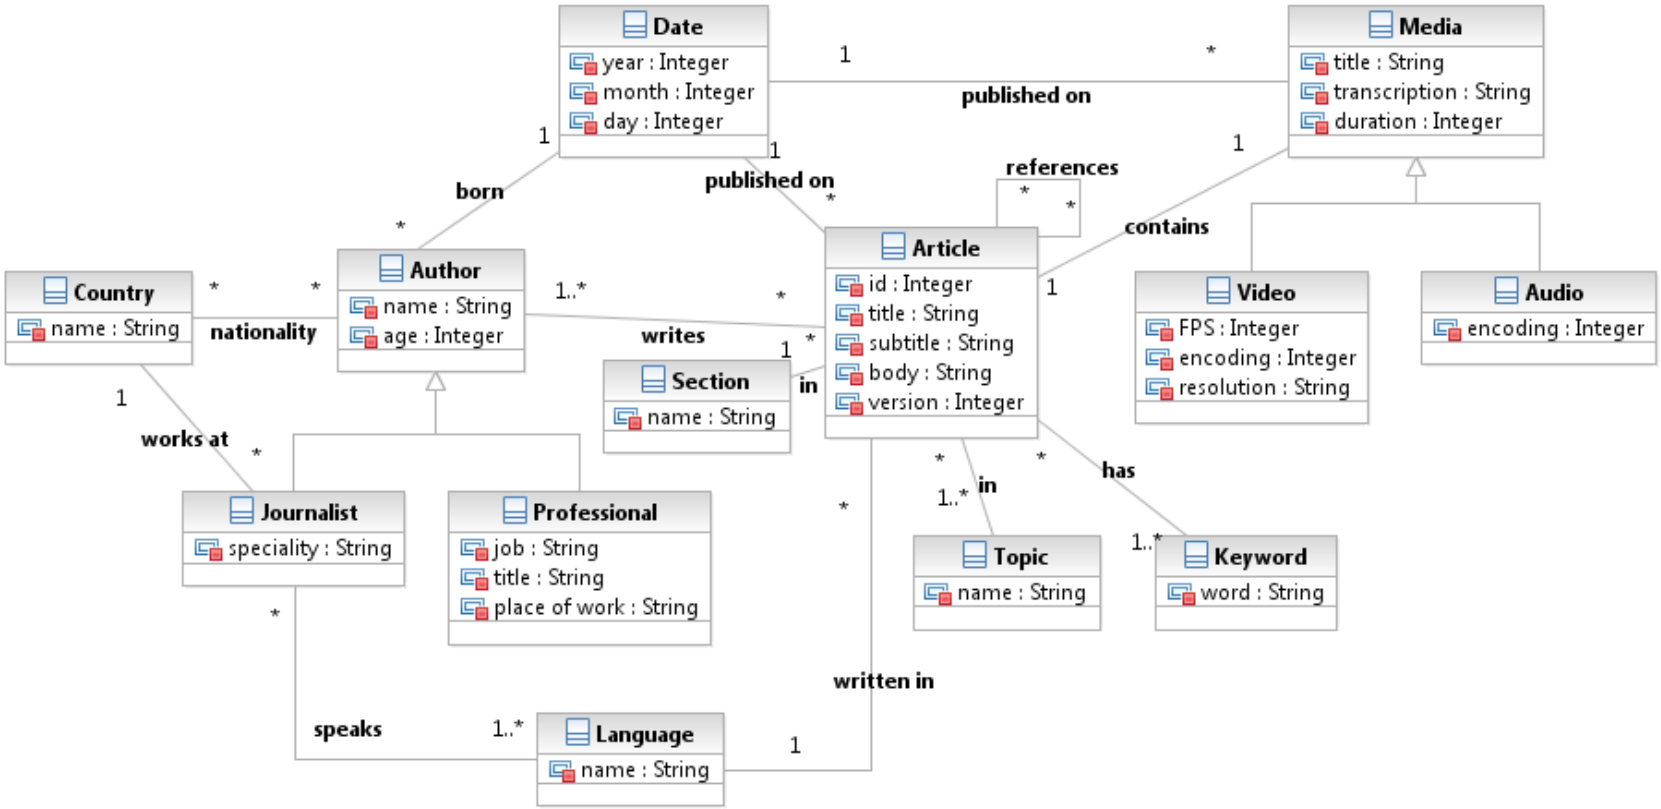

Initial Data and Metadata Processes

Initial Data

Initial Data and Metadata Processes

Initial Data

- Use web crawler to obtain real-world news and related information such: author, media, publisher...
- Long period of time (30-40 years).
- At least 100 TB recommended for running the tests.
- Once the corpus is fetched it can be sliced to select the desired scale.

Initial Data and Metadata Processes

Initial Data

- Use web crawler to obtain real-world news and related information such: author, media, publisher...
- Long period of time (30-40 years).
- At least 100 TB recommended for running the tests.
- Once the corpus is fetched it can be sliced to select the desired scale.

Metadata Processes

Initial Data and Metadata Processes

Initial Data

- Use web crawler to obtain real-world news and related information such: author, media, publisher...
- Long period of time (30-40 years).
- At least 100 TB recommended for running the tests.
- Once the corpus is fetched it can be sliced to select the desired scale.

Metadata Processes

- Extract the metadata.
- Build a structure for Information Retrieval algorithms such as:
  - TF-IDF: to measure the relevance of a word within an article.
  - PageRank: to weight the relevance of the articles for a given search.

Data Scaling and Maintenance

Scaling

Data Scaling and Maintenance

Scaling

1. Increase volume: using the remaining fetched data as progressive slices of PRIMEBALL's corpus.
2. Metadata update: the system has to extract the necessary metadata to ensure that subsequent queries can be performed.

Data Scaling and Maintenance

Scaling

1. Increase volume: using the remaining fetched data as progressive slices of PRIMEBALL's corpus.
2. Metadata update: the system has to extract the necessary metadata to ensure that subsequent queries can be performed.

Maintenance

Data Scaling and Maintenance

Scaling

1. Increase volume: using the remaining fetched data as progressive slices of PRIMEBALL's corpus.
2. Metadata update: the system has to extract the necessary metadata to ensure that subsequent queries can be performed.

Maintenance

It is required to recalculate some structures used for information retrieval. The process involves recomputation of:

- TF-IDF from all the documents.
- Topic modelling for all the articles.

Outline

Query Set

Query Set

A set of queries that are typically performed by NPT-NH users is used for testing the performance of a given parallel data processing framework.

Query Set

A set of queries that are typically performed by NPT-NH users is used for testing the performance of a given parallel data processing framework.

Some examples

Query Set

A set of queries that are typically performed by NPT-NH users is used for testing the performance of a given parallel data processing framework.

Some examples

The hot topics and the most common

- Articles published during an interval of time sorted by topics. The output contains pairs of an article title and a topic.

Query Set

A set of queries that are typically performed by NPT-NH users is used for testing the performance of a given parallel data processing framework.

Some examples

The hot topics and the most common

- Articles published during an interval of time sorted by topics. The output contains pairs of an article title and a topic.

Topic evolution w.r.t. time

- Find the most frequently used keywords in X day on two different years sorted by count decreasingly.

Query Set

A set of queries that are typically performed by NPT-NH users is used for testing the performance of a given parallel data processing framework.

Some examples

The hot topics and the most common

- Articles published during an interval of time sorted by topics. The output contains pairs of an article title and a topic.

Topic evolution w.r.t. time

- Find the most frequently used keywords in X day on two different years sorted by count decreasingly.

Diversity of articles w.r.t source

- Articles written by an author A from a given country C that best match a search term S.

Test Protocol

Test Protocol

Scenarios

Test Protocol

Scenarios

Different scenarios help to benchmark the system.

Each scenario has: context, initial state, operations and measures.

Test Protocol

Scenarios

Different scenarios help to benchmark the system.
Each scenario has: context, initial state, operations and measures.
Seven scenarios have been defined with the following objectives:

- Scenario 1: scale-up and horizontal scalability.
- Scenario 2: consistency and version handling.
- Scenario 3: data availability using failure simulation.
- Scenario 4: system concurrency and consistency.
- Scenario 5: analysis procedures performance.
- Scenario 6: system loading and initialization.
- Scenario 7: system upgrade.

Example scenario

Example scenario

Scenario 2

Example scenario

Scenario 2

- **Context**: a very popular article has been published with many mistakes and publishers are correcting it constantly. Their main concern is to deliver a consistent view of the article to people.
- **Initial state**: the system contains a dataset of the specified scale factor.
- **Operations**:
  1. Initiate a thread performing 100 queries per second to retrieve the given article.
  2. Start another thread updating the same article every 5 seconds.
- **Measure**: how many times the article is read in the older versions after being read once in the new version.

Outline

Performance metrics

Performance metrics

Metrics to evaluate the performance of the SUT.

Performance metrics

Metrics to evaluate the performance of the SUT.

Throughput

- Total time required to execute a scenario.

Performance metrics

Metrics to evaluate the performance of the SUT.

Throughput

- Total time required to execute a scenario.

Price performance

$$Price\ performance = \frac{throughput}{price}$$

Performance metrics

Metrics to evaluate the performance of the SUT.

Throughput

- Total time required to execute a scenario.

Price performance

$$Price\ performance = \frac{throughput}{price}$$

Pricing depends mainly on: cloud provider, number of instances, storage space, platform inherent costs and execution time.

Properties and Metrics

Generic cloud properties

Properties and Metrics

Generic cloud properties

- Scale up: more data can be processed when adding more computers.
- Elastic speedup: higher throughput when adding more computers.
- Horizontal scalability: even distribution of workload (data & computation).
- Latency: time to execute a set of operations.
- Durability: retain information for a long time.
- Consistency and version handling: homogeneous vision of data.
- Availability: data is accessible even when there are failures.
- Concurrency: offer service to multiple clients.

Properties and Metrics

Complex data properties

Properties and Metrics

Complex data properties

- Path traversals: linking data from different parts of schema
- Construction of complex results: Ability to generate structured output.
- Polymorphism: ability to deal with inheritance.

Properties and Metrics

Complex data properties

- Path traversals: linking data from different parts of schema
- Construction of complex results: Ability to generate structured output.
- Polymorphism: ability to deal with inheritance.

Big data properties

Properties and Metrics

Complex data properties

- Path traversals: linking data from different parts of schema
- Construction of complex results: Ability to generate structured output.
- Polymorphism: ability to deal with inheritance.

Big data properties

- Analysis: summarize data and reporting.

Properties and Metrics

Complex data properties

- Path traversals: linking data from different parts of schema
- Construction of complex results: Ability to generate structured output.
- Polymorphism: ability to deal with inheritance.

Big data properties

- Analysis: summarize data and reporting.

Information retrieval properties

Properties and Metrics

Complex data properties

- Path traversals: linking data from different parts of schema
- Construction of complex results: Ability to generate structured output.
- Polymorphism: ability to deal with inheritance.

Big data properties

- Analysis: summarize data and reporting.

Information retrieval properties

- Full text: searching for single words.

Outline

Conclusions

- Propose specifications for PRIMEBALL.
- PRIMEBALL is a guideline to build an integral solution for benchmarking could platforms.
- Uses a real-life model called New Pork Times.
- Workload compound by: queries and data-intensive batch processes.
- Tests split in scenarions: each assessing the frameworks performance from different points of view.
- Integrated benchmark:
  - parallelization capabilities.
  - cloud features.
  - big data analysis ability.
  - uses real data.

Future work

- Implement the crawler for fetching real data from the Web.
- Create and distribute the built dataset online.
- Implement the benchmark in several cloud environments.
- Obtain experimental results and publish comparison results.

Thank you very much!
Do you have any questions?