Smarter Computing & ICT for Sustainable Development of Human Society and Industry in Big Data Environment

Engelbert Mephu Nguifo
LIMOS, CNRS & UBP

NICST 2014 workshop, Weihai (China)
Panel session

Smarter Computing model

IBM Co.

- Based on three key principles:
  - designing systems for data,
  - optimizing systems for specific workloads and
  - managing systems in using a cloud computing architecture

Smarter Computing model

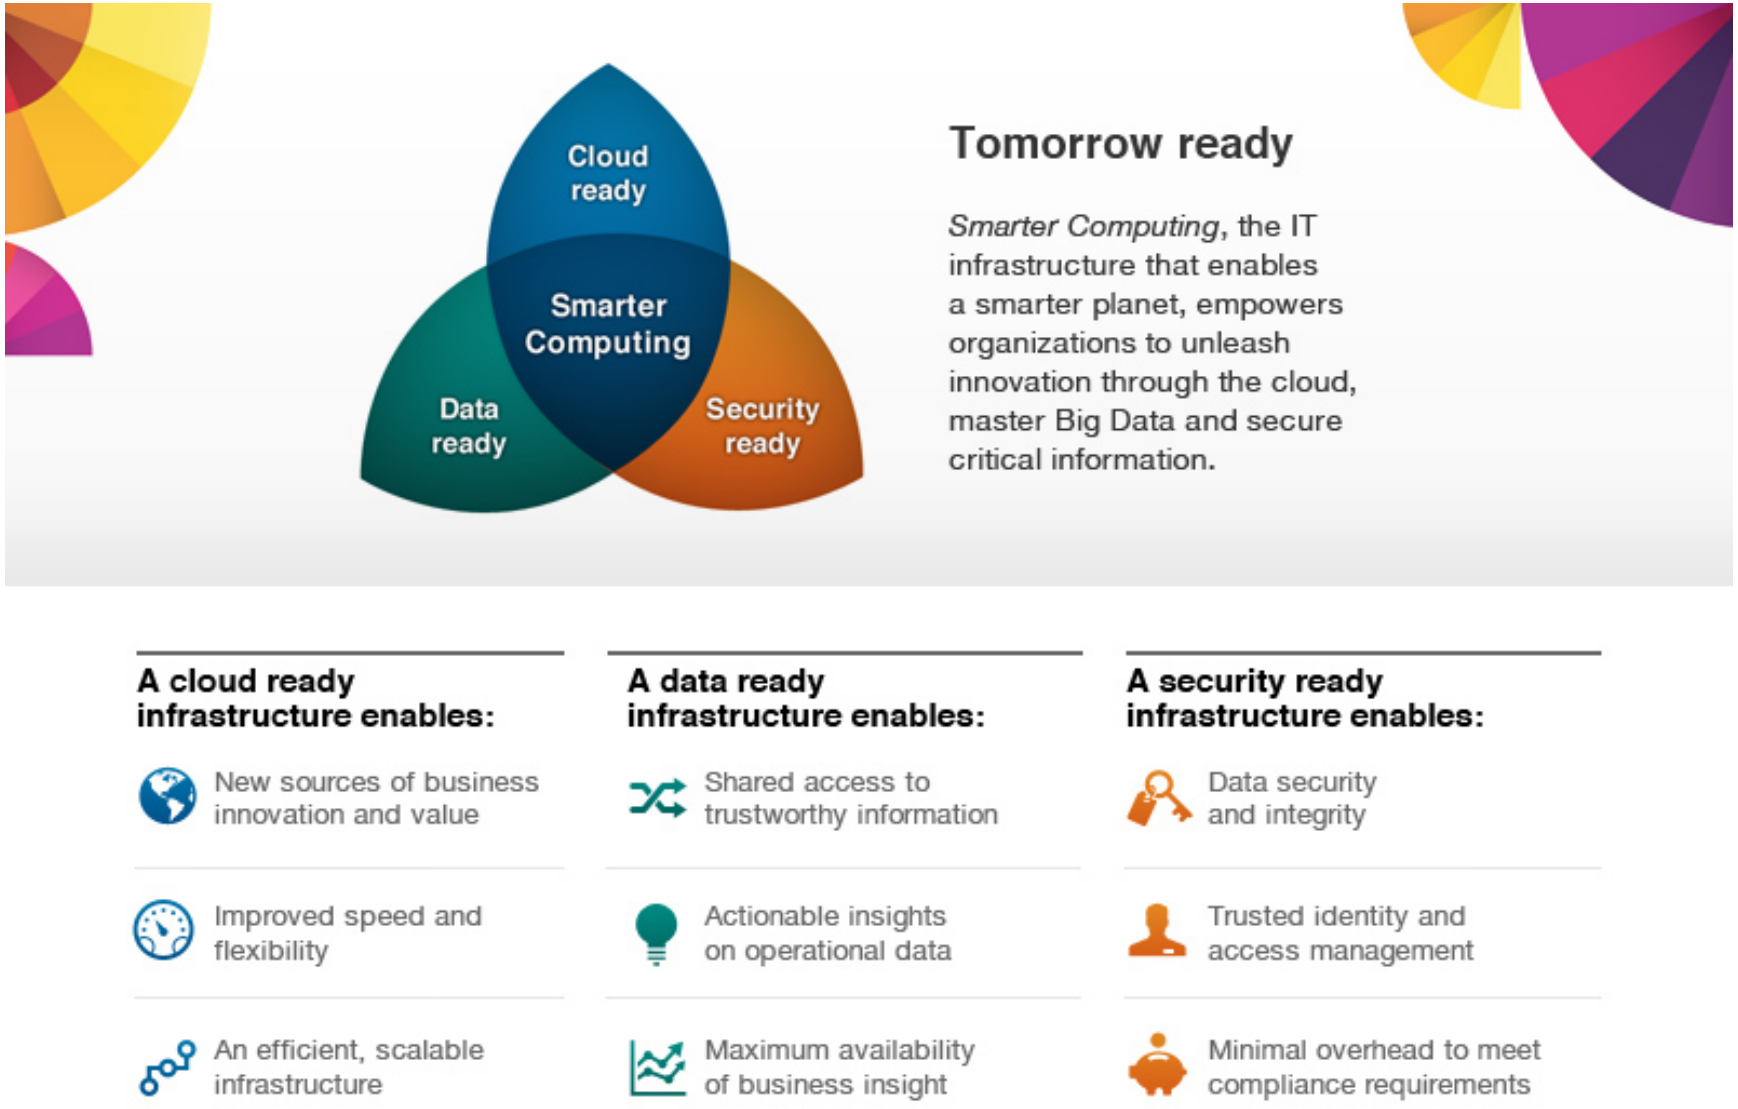

Issues

- **Big Data problem**
- **Skew problem**

Big Data Environment

- **3 V's** Doug Laney , 2001
  - **Volume.**
  - **Velocity.**
  - **Variety.**

- Also :
  - Variability
  - Complexity

Big Data difficulty

Takeaki Uno'slides

- **noisy**
  - not only low accuracy, but big lack, incorrect (lie)…
- **non-uniform**
  - very different granularity, assorted by different objective…
- **elementarity**
  - directly taken from sensors, difficult to understand directly
- **sparsity**
  - Small degree, items, slightly structured
- **diversity**
  - composed of many small groups (minorities), so include locally dense structures

Big Data

Haseeb Budhani, 2008

"Big Data" caught on quickly as a blanket term for any collection of **data set so large** and **complex** that it becomes ***difficult to process*** using on-hand database management tools or traditional data processing applications.

Big Data

http://www.ibm.com/big-data/us/en/,
2014/05/26

Big data is being generated by everything around us at all times. Every digital process and social media exchange produces it. Systems, sensors and mobile devices transmit it.

Big data is arriving from multiple sources at an alarming **velocity, volume and variety**.

To extract meaningful value from big data, you need **optimal processing power**, **analytics capabilities** and **skills**.

Cloud and Big Data
on Smarter computing

- Search **for powerful and cost effective** approach for **massively parallel** analytics

a Google solution

  - **Map-Reduce** paradigm
    - Hadoop implementation

MapReduce framework

- A framework for processing huge datasets.
- Large number of computers and task/node failures.

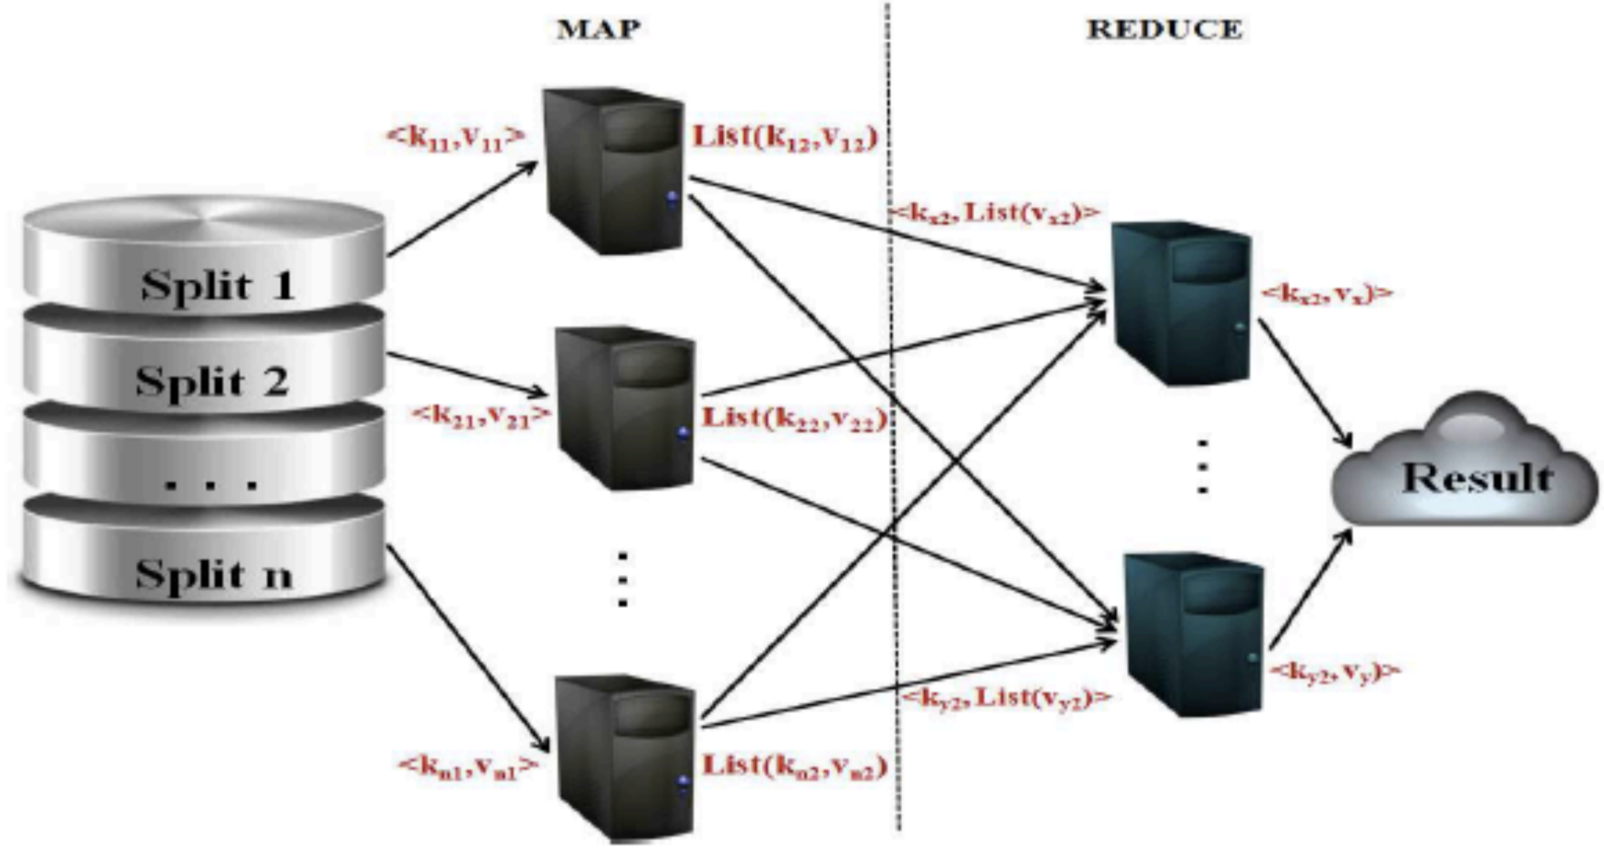

MR skew problem

- Two main phases: Map & Reduce
- In each phase, a subset of the input data is processed by **distributed tasks**
  - When a map task completes, the reduce tasks are notified to pull newly available data.
  - This transfer process is referred to as a **shuffle**. All map tasks must complete before the shuffle part of the reduce phase can complete, allowing the reduce phase to begin.

MR skew problem

- Load imbalancing
  - case where computational load is imbalanced among map or reduce tasks
  - Refer as map-skew and reduce-skew respectively.

- **Skew** can lead to **significantly longer job execution times** and **significantly lower cluster** throughput.

http://nuage.cs.washington.edu/pubs/opencirrus2011.pdf

A Study of Skew in MapReduce Applications

YongChul Kwon, Magdalena Balazinska, Bill Howe
University of Washington, USA
Email:{yongchul,magda,billhowe}@cs.washington.edu

Jerome Rolia
HP Labs
Email: jerry.rolia@hp.com

***Abstract*—This paper presents a study of skew — highly variable task runtimes — in MapReduce applications. We describe various causes and manifestations of skew as observed in real world Hadoop applications. Runtime task distributions from these applications demonstrate the presence and negative impact of skew on performance behavior. We discuss best practices recommended for avoiding such behavior and their limitations.**

I. Introduction

MapReduce [1] has proven itself as a powerful and cost-effective approach for massively parallel analytics [2]. A MapReduce job runs in two main phases: map phase and reduce phase. In each phase, a subset of the input data is processed by distributed tasks in a cluster of computers. When a map task completes, the reduce tasks are notified to pull newly available data. This transfer process is referred to as a shuffle. All map tasks must complete before the shuffle

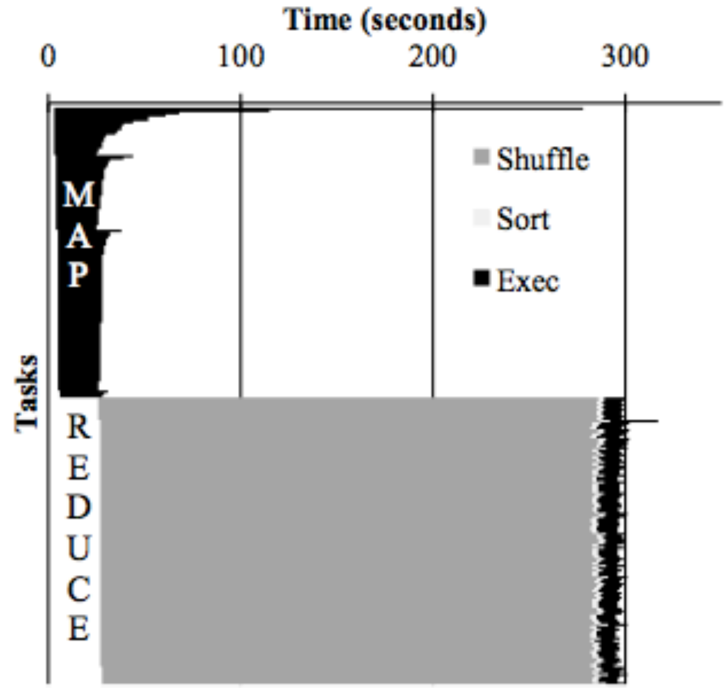

MR skew Best Practices

1. Use **domain knowledge**
2. Try **different partitioning schemes**
3. Implement **a combiner to reduce the amount of data** going into the reduce-phase
4. Use a pre-processing MapReduce job that **extracts properties** of the input data in the case of a longruning, skew-prone map phase.
5. Design algorithms whose **runtime depends only on the amount of input data** and not the data distribution.

Thanks

MR skew study

- Best Practice 1.

Use **domain knowledge** when choosing the map output partitioning scheme if the reduce operation is expensive: Range partition or some other form of explicit partition may be better than the default hash-partition.

MR skew study

- Best Practice 2.

Try **different partitioning schemes** on sample workloads or collect the data distribution at the reduce input if a MapReduce job is expected to run several times.

MR skew study

- Best Practice 3.

Implement **a combiner to reduce the amount of data** going into the reduce-phase and, as such, significantly dampen the effects of any type of reduce-skew.

MR skew study

- Best Practice 4.

Use a pre-processing MapReduce job that **extracts properties** of the input data in the case of a longruning, skew-prone map phase. Appropriately partitioning the data before the real application runs can significantly reduce skew problems in the map phase.

MR skew study

- Best Practice 5.

Design algorithms whose **runtime depends only on the amount of input data** and not the data distribution.

Map Reduce

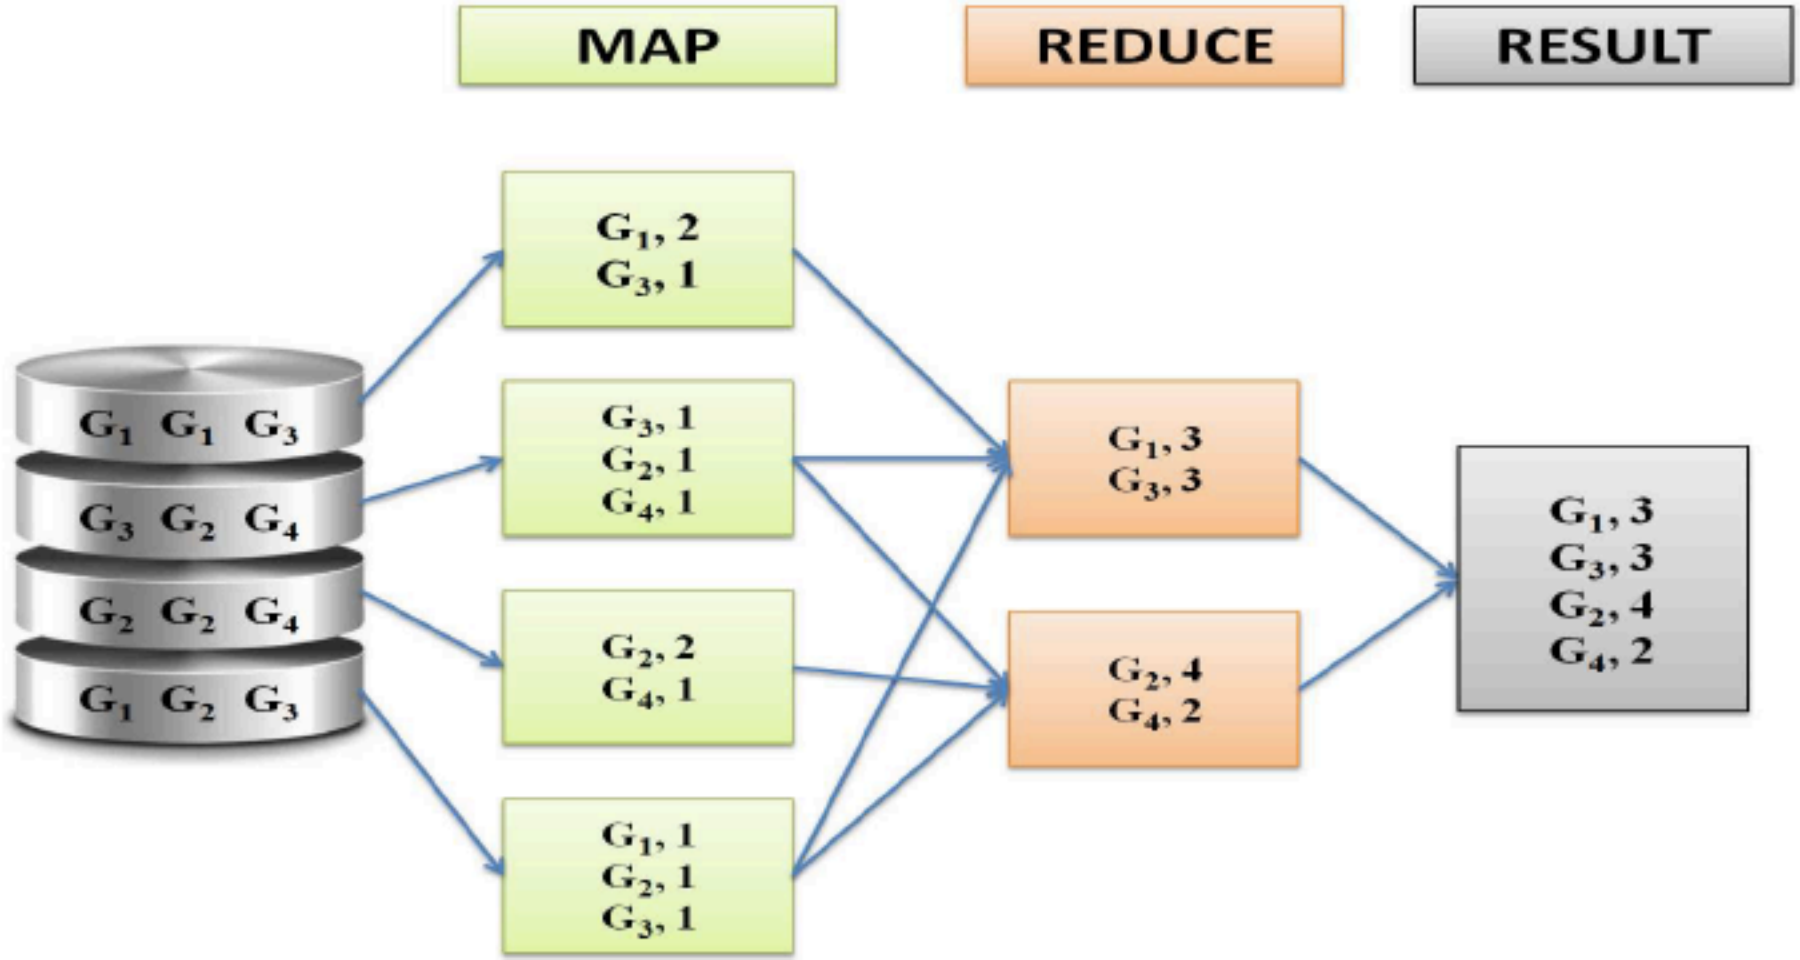

MAP
REDUCE
RESULT
G_1 G_1 G_3
G_3 G_2 G_4
G_2 G_2 G_4
G_1 G_2 G_3
G_1, 2
G_3, 1
G_3, 1
G_2, 1
G_4, 1
G_2, 2
G_4, 1
G_1, 1
G_2, 1
G_3, 1
G_1, 3
G_3, 3
G_2, 4
G_4, 2
G_1, 3
G_3, 3
G_2, 4
G_4, 2

Big Data

Doug Laney (currently with Gartner) 2001

- **Volume.**
- **Velocity.**
  - Data is streaming in at unprecedented speed and must be dealt with in a timely manner. RFID tags, sensors and smart metering are driving the need to deal with torrents of data in near-real time. Reacting quickly enough to deal with data velocity is a challenge for most organizations.
- **Variety.**
  - Data today comes in all types of formats. Structured, numeric data in traditional databases. Information created from line-of-business applications. Unstructured text documents, email, video, audio, stock ticker data and financial transactions. Managing, merging and governing different varieties of data is something many organizations still grapple with.

Big Data

- **Variability.**
  - In addition to the increasing velocities and varieties of data, data flows can be highly inconsistent with periodic peaks. Is something trending in social media? Daily, seasonal and event-triggered peak data loads can be challenging to manage. Even more so with unstructured data involved.
- **Complexity.**
  - Today's data comes from multiple sources. And it is still an undertaking to link, match, cleanse and transform data across systems. However, it is necessary to connect and correlate relationships, hierarchies and multiple data linkages or your data can quickly spiral out of control.