



HEMERA - The Map-Red Challenge:

Scalable Distributed Processing Using the Map-Reduce Paradigm

G. Antoniu (KerData), G. Fedak (GRAAL)

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE



centre de recherche
RENNES - BRETAGNE ATLANTIQUE

Why Map-Reduce?

Map-Reduce appears as THE emerging programming model for scalable data-intensive, large-scale computing

- Advertised by Google (initially for internal needs)
- Adopted by the major cloud PaaS service providers
- Focus of many research efforts in the scientific cloud computing community

Why Map-Reduce@HEMERA?

- A challenging experimental context for scalable data-intensive computing on Aladdin-Grid'5000
- Typical orders of magnitude at Google
 - Map tasks: 200,000
 - Reduce tasks: 4,000
 - Workers: 2,000



What is Map-Reduce?

Map-Reduce: a **simple programming model** for **data-intensive** computing

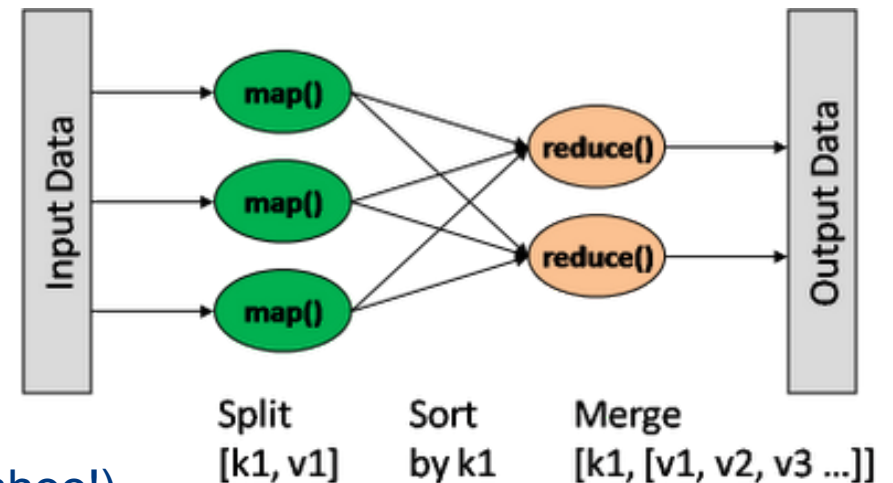
Typical problem solved by Map-Reduce

- Read a lot of data
- **Map**: extract something you care about from each record
- Shuffle and Sort
- **Reduce**: aggregate, summarize, filter, or transform
- Write the results

Approach: **hide messy details** in a runtime library

- Automatic parallelization
- Load balancing
- Network and disk transfer optimization
- Transparent handling of machine failures

Implementations: Google MapReduce, Hadoop (Yahoo!)



Map-Red Challenge: Goal, Open Topics

State of the art

- Current platform still in their initial phase
- E.g. despite his wide adoption, Hadoop's file system has limited support for concurrent access to shared data

Goal: overcome the limitations of current Map-Reduce implementations to enable scalable MapReduce-based processing on ALADDIN-G5K

Issues

- Infrastructure: cope with various environments
 - Clouds, dektop grids, hybrid infrastructures
- Storage architecture: support for highly-concurrent I/O
- Scheduling: heterogeneous platforms, fair execution of multiple M/R applications
- Fault tolerance: specifically optimized mechanisms (e.g. for checkpoint/restart)
- Security
- Others?

The « MapReduce » ANR Project (2010-2013)

Goal: an optimized Map-Reduce platform for cloud infrastructures

- INRIA - KerData team (Rennes) – coordinator (Gabriel Antoniu)
- INRIA - GRAAL team (Lyon), France
- Joint UIUC/INRIA Laboratory for Petascale Computing
- University of Illinois at Urbana Champaign, USA
 - Cloud Computing Testbed
- Nimbus team, Argonne National Lab/University of Chicago, USA
 - FutureGrid
- IBM Products and Solutions Center, Montpellier, France
- Institute of Biology and Chemistry of Proteins, Lyon, France
- MEDIT (SME), Palaiseau, France



Map-Reduce: Related Actions

MAPREDUCE workshop (HPDC) chaired by Gilles Fedak

- A success in 2010, renewed for 2011 (PC ready)

Project proposal to the joint INRIA-Microsoft call

- Experimental platform: Grid'5000 (first step), Microsoft Azure (second step)
- Involves the KerData (Rennes) and PARIETAL (Saclay) INRIA teams

EQUIPEX proposal: CGB (Rennes)

- CGB: Symbiose + KerData INRIA Teams (Rennes)

Other projects that you know?

Mailing list: mapreduce@inria.fr

