



Renato Javier Marroquín Mogrovejo

Experimental Statistical Analysis Of
MapReduce Joins

DISSERTAÇÃO DE MESTRADO

Dissertation presented to the Programa de Pós-graduação em
Informática of the Departamento Informática PUC-Rio as
partial fulfillment of the requirements for the degree of Mestre
em Informática

Advisor: Prof. Edward Hermann Haeusler

Rio de Janeiro
June 2011



Renato Javier Marroquín Mogrovejo

**Experimental Statistical Analysis Of
MapReduce Joins**

Dissertation presented to the Programa de Pós-graduação em Informática of the Departamento Informática PUC-Rio as partial fulfillment of the requirements for the degree of Mestre em Informática. Approved by the following commission:

Prof. Edward Hermann Haeusler

Advisor

Departamento de Informática — PUC-Rio

Prof. Sérgio Lifschitz

Departamento de Informática—PUC-Rio

Prof. Daniel Schwabe

Departamento de Informática—PUC-Rio

Prof. Javam de Castro Machado

UFC

Prof. José Eugenio Leal

Coordinator of the Centro Técnico Científico — PUC-Rio

Rio de Janeiro — June 28, 2011

All rights reserved.

Renato Javier Marroquín Mogrovejo

Renato Marroquín Mogrovejo graduated from the Universidad Católica de Santa María (Arequipa, Peru) in System Engineering. He then worked at the Banco de Crédito del Perú, and then continued his graduate studies at the Pontifícia Universidade Católica do Rio de Janeiro. While persuing his Master's degree in Informatics at the PUC-Rio, he participated actively in research projects at the BioInformatics Laboratory from PUC-Rio.

Bibliographic data

Marroquín, Renato Javier Mogrovejo

Experimental Statistical Analysis Of MapReduce Joins / Renato Javier Marroquín Mogrovejo; advisor: Edward Hermann Haeusler. — Rio de Janeiro : PUC-Rio, Department of Informática, 2011.

v., 8 f: il. ; 29,7 cm

1. MSc Dissertation - Pontifícia Universidade Católica do Rio de Janeiro, Department of Informática.

Bibliography included.

1. Informática – Dissertation. Statistical Analysis; Join operations; MapReduce; Pig-Latin; Cloud Data Management. I. Haeusler, Edward Hermann. II. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Informática. III. Title.

CDD: 004

Acknowledgments

To God, for everything I have.

To my family, for being pillars of support throughout my life.

To my friends, for helping me with their support in difficult times,

To the PUC-Rio and CAPES, for the chance of pursuing graduate studies
in Brasil and also,

to all the people that help me even without noticing it.

Resumo

Marroquín, Renato Javier Mogrovejo; Haeusler, Edward Hermann(Orientador). **Experimental Statistical Analysis Of MapReduce Joins**. Rio de Janeiro, 2011. 8p. Dissertação de Mestrado — Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro.

Existem muitas aplicações que poderiam conseguir aproveitar melhor o modelo de negócio da computação em nuvem. No entanto, muitas dessas aplicações enfrentam dificuldades na hora de gerenciar grandes quantidades de dados pois a maioria dos bancos de dados tradicionais, como é o caso dos SGBDs relacionais, não foram projetados para conseguir aproveitar todas as características dos ambientes de nuvens computacionais. Neste trabalho, algumas técnicas estatísticas são exploradas para a predição do desempenho de consultas de junções, no contexto de execuções com o paradigma MapReduce. Além disso, o cálculo da seletividade de junções é estudado neste contexto, para que possa ser parametrizado no modelo estudado e proposto nesta dissertação. São realizados testes computacionais simulando um benchmark TPC-DS para avaliar métodos de junção implementados no framework Pig, que trabalha usando Hadoop como motor de execução. O modelo proposto e os resultados obtidos permitem orientar o usuário sobre o método mais adequado e é um passo fundamental na construção de otimizadores nestes ambientes.

Palavras-chave

Análisis Estadístico; Operações de Junção; MapReduce; Pig-Latin; Gerenciamento de dados na nuvem.

Abstract

Marroquín, Renato Javier Mogrovejo; Haeusler, Edward Hermann(Advisor). **Experimental Statistical Analysis Of MapReduce Joins**. Rio de Janeiro, 2011. 8p. MSc Dissertation — Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro.

There are different scalable data management solutions which can take advantage of cloud features making them more attractive for a deployment in such environments. One of the most critical operations in data processing is joining datasets, but this operation is most of the time the one that takes more time, and one of the hardest to optimize. In this work, we explore statistical methods in order to predict join queries execution times. In addition to that, join selectivity estimation is explored in a MapReduce environment in order to use it as another parameter in our model.

Keywords

Statistical Analysis; Join operations; MapReduce; Pig-Latin; Cloud Data Management.

Summary

I	Introduction	9
II	Background Information	11
II.1	Cloud Computing	11
II.2	Cloud Data Management	12
II.3	The Map Reduce programming model	14
II.4	High level computing languages for data intensive applications	17
II.5	Join types in the PIG framework	19
II.6	Statistical Modeling	25
III	Problem Statement	28
III.1	Background	28
III.2	Problem Statement	30
III.3	Motivation	32
IV	Statistical Modelling of MapReduce Joins	36
IV.1	Data Set	36
IV.2	Model Construction	39
IV.3	Join Selectivity	40
IV.4	Experimental setup	51
IV.5	Results	54
V	Conclusions	65
V.1	Related work	65
V.2	Conclusions	69
V.3	Summary of Contributions	71
V.4	Future Work	72
A	Pig-Latin example queries	87

List of Figures

II.1	The MapReduce Model [106]	15
II.2	Hadoop Computing Model	17
II.3	MapReduce Hash Join	20
II.4	MapReduce Fragment-Replicated Join	22
II.5	MapReduce Merge Join	25
III.1	System Architecture	31
III.2	System behavior for determining the correct type of join.	32
IV.1	Relationships from some TPC-DS' tables.	38
IV.2	Sampling file blocks.	43
IV.3	Cardinality sampling using $L = 1000$.	44
IV.4	Cardinality sampling using $L = 1500$.	45
IV.5	Average error for sampled ranges (L being 1000 and 1500).	45
IV.6	Selectivity Percentage Error varying according to table IV.3.	52
IV.7	Graphical Three-Fold validation.	53
IV.8	Workload structure.	53
IV.9	Query naming convention.	54
IV.10	Query execution times against time estimates for fragment replicated joins.	55
IV.11	Query execution times against time estimates for the 80th percentil of the executions.	56
IV.12	Error percentage for the 80th percentil of the executions.	56
IV.13	Query execution times against time estimates for merge joins.	57
IV.14	Query execution times against time estimates for the 80th percentil of merge join executions.	58
IV.15	Error percentage for the 80th percentil of the merge join executions.	59
IV.16	Query execution times against time estimates for hash joins.	59
IV.17	Query execution times against time estimates for the 80th percentil of hash join executions.	60
IV.18	Error percentage for the 80th percentil of hash join executions.	60
IV.19	Hash join query execution variability.	61
IV.20	Query execution times against time estimates for the simulation workload of fragment-replicated joins.	62
IV.21	Query execution times against time estimates for the simulation workload of merge joins.	63
IV.22	Query execution times against time estimates for the simulation workload of hash joins.	64