Twister4Azure : Iterative MapReduce for Azure Cloud

A Decentralized, Dynamically Scalable, Fault Tolerant Iterative MapReduce Framework Built Using Cloud Services for Microsoft Azure Cloud.

Introduction

There exists many algorithms that rely on iterative computations, where each iterative step can be easily specified as a MapReduce computation. MapReduceRoles for Azure (MR4Azure) is a decentralized, dynamically scalable MapReduce runtime we developed for Windows Azure Cloud platform using Microsoft Azure cloud infrastructure services as the building blocks. Twister4Azure extends MR4Azure to support optimized iterative MapReduce executions, enabling a wide array of large scale iterative data analysis and scientific applications to utilize Azure platform easily and efficiently.

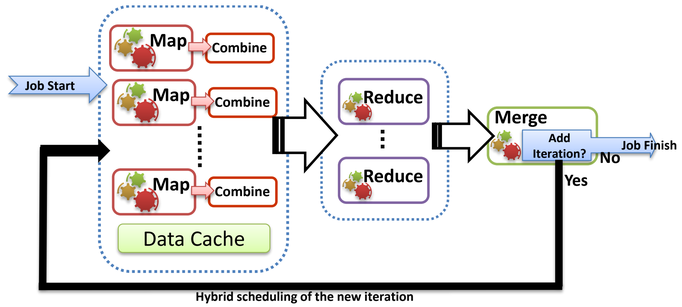

Twister4Azure Computation Flow

MapReduceRoles for Azure

- Use distributed, highly scalable and highly available cloud services as the building blocks.
  - Azure Queues for task scheduling.
  - Azure Blob storage for input, output and intermediate data storage.
  - Azure Tables for meta-data storage and monitoring
- Utilize eventually-consistent , high-latency cloud services effectively to deliver performance comparable to traditional MapReduce runtimes.
- Minimal management and maintenance overhead

Thilina Gunarathne, Judy Qiu, Geoffrey Fox
{tgunarat, xqiu, gcf}@indiana.edu

Twister4Azure

- Familiar MapReduce programming model
- Fault Tolerance features similar to traditional MapReduce.
- No single point of failure.
- Combiner step
- Supports dynamically scaling up and down of the compute resources.
- Web based monitoring console
- Easy testing and deployment using Azure local development fabric.
- Iterative extensions
  - Merge Step
  - In-Memory Caching of static data
  - Cache aware hybrid scheduling using Queues as well as a bulletin board (special table)

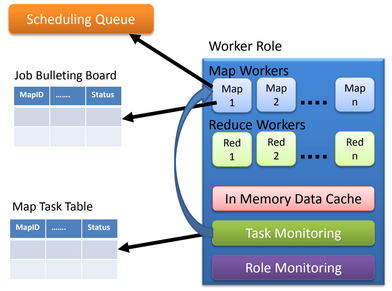

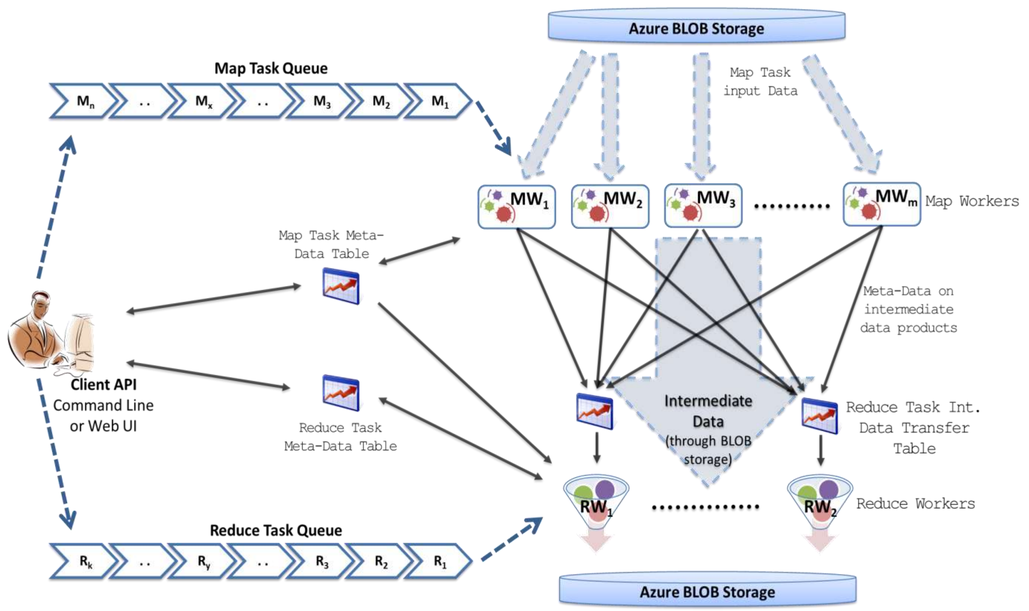

Performance Comparison

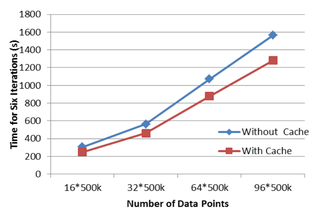

K-Means Clustering

KMeans iterative MapReduce performance. 16 Azure Small instances, 6 iterations, 8 to 48 million 20-D data points.

Left: Performance with and without data caching. **Right**: Speedup obtained from using the data cache

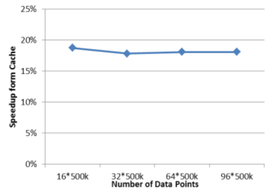

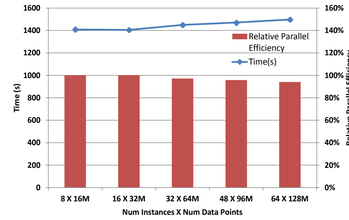

Left: Scaling speedup with increasing number of instances (Azure Small) & data for 10 iterations.
Right: Increasing number of iterations using 16 million data points with caching using 16 Azure Small instances.

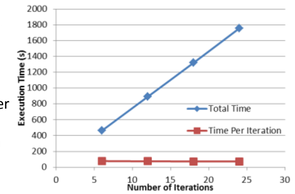

BLAST

NCBI BLAST+ sequence searching using NR protein sequence database (~8.7 GB) using 128 CPU cores. 16 Extra-Large Azure instances were used for Twister4Azure. 8-core 16GB memory per node cluster was used Apache Hadoop. 16-core 16GB memory per node Windows HPC cluster was used for DryadLINQ testing. Each query file == 100 sequences.

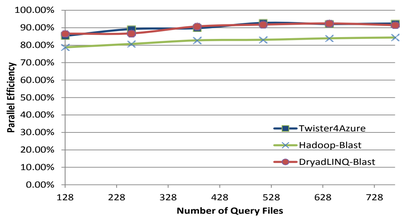

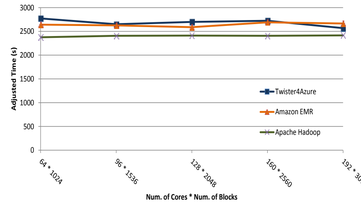

Smith Waterman-GOTOH All-Pairs Sequence Alignment

Azure Small instances for Twister4Azure, Amazon EC2 High-CPU-Extra-Large instances for Amazon Elastic MapReduce and 8 core 16GB per node cluster for Apache Hadoop. Block == 200 sequences.

CAP3 Sequence Assembly

Using Azure Small instances for Twister4Azure, Amazon EC2 High-CPU-Extra-Large instances for Amazon Elastic MapReduce and 8 core 16GB per node cluster for Apache Hadoop. Each file contained 458 reads.

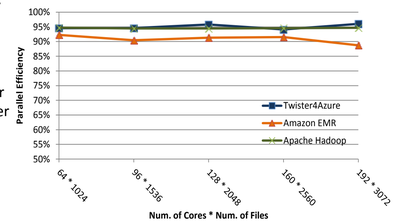

REFERENCES

Gunarathne, T., Wu,T.L., Qiu, J., and Fox, G.C. 2010. MapReduce in the Clouds for Science. In *Proceedings of CloudCom 2010* Conference (Indianapolis,December 2010)

Ekanayake, J., Li, H., Zhang B., *et al.*, 2010. Twister: A Runtime for iterative MapReduce, in *Proceedings of the First International Workshop on MapReduce and its Applications* of ACM HPDC 2010 conference (Chicago, June 2010)

Gunarathne, T., Wu, T. L., Qiu, J., *et al.*, 2010. Cloud computing paradigms for pleasingly parallel biomedical applications. Submitted for publication in Concurrency and Computation: Practice and Experience journal.

ACKNOWLEDGEMENTS

Microsoft Exploratory Research in Clouds and Platforms Grant, FutureGrid and Salsa Group.

SALSA HPC

SCHOOL OF INFORMATICS AND COMPUTING
INDIANA UNIVERSITY

PERVASIVE TECHNOLOGY INSTITUTE
INDIANA UNIVERSITY