Preview of Award 1320357 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Grant or Other Identifying Number Assigned by Agency:	1320357
Project Title:	III: Small: Cumulon: Easy and Efficient Statistical Big-Data Analysis in the Cloud
PD/PI Name:	Jun Yang, Principal Investigator Shivnath Babu, Co-Principal Investigator Sayan Mukherjee, Co-Principal Investigator Michael D Ward, Co-Principal Investigator
Recipient Organization:	Duke University
Project/Grant Period:	09/15/2013 - 08/31/2016
Reporting Period:	09/01/2014 - 08/31/2015
Submitting Official (if other than PD\PI):	Jun Yang Principal Investigator
Submission Date:	07/16/2015
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	Jun Yang

Accomplishments

*** What are the major goals of the project?**

"Big data" have been growing in volume and diversity at an explosive rate, bringing enormous potential for transforming science and society. Driven by the desire to convert messy data into insights, analysis has become increasingly statistical, and there are more people than ever interested in analyzing big data. The rise of cloud

computing in recent years, exemplified by the popularity of services such as Amazon EC2, offers a promising possibility for supporting big-data analytics. Its "pay-as-you-go" business model is especially attractive: users gain on-demand access to computing resources while avoiding hardware acquisition and maintenance costs.

However, it remains frustratingly difficult for many scientists and statisticians to use the cloud for any nontrivial statistical analysis of big data. First, developing efficient statistical computing programs requires a great deal of expertise and effort. Popular cloud programming platforms, such as Hadoop, require users to code and think in low-level, platform-specific ways, and, in many cases, resort to extensive manual tuning to achieve acceptable performance. Second, deploying such programs in the cloud is hard. Users are faced with a maddening array of choices, ranging from resource provisioning (e.g., type and number of machines to request on Amazon EC2), software configuration (e.g., number of parallel execution slots per machine for Hadoop), to execution parameters and implementation alternatives. Some of these choices can be critical to meeting deadlines and staying within budget, but current systems offer little help to users in making such choices.

This project aims to build Cumulon, an end-to-end solution for making statistical computing over big data easier and more efficient in the cloud. When developing data analysis programs, users will be able to think and code in a declarative fashion, without being concerned with how to map data and computation onto specific hardware and software platforms. When deploying such programs, Cumulon will present users with best "plans" meeting their requirements, along with information that is actually helpful in making decisions---in terms of completion time and monetary cost. For example, given a target completion time, Cumulon can suggest the best plan on Amazon EC2 that minimizes the expected total cost. A plan encodes choices of not only implementation alternatives and execution parameters, but also cluster provisioning and configuration choices. This project will develop effective cost modeling and efficient optimization techniques for the vast search space of possible plans. Once a plan is chosen, Cumulon automatically takes care of all details, including reserving hardware, configuring software, and executing the program.

*** What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?**

Major Activities:

1. As of the end of Year 2 of this project, we have completed two iterations of system development for Cumulon and expect to finish the third iteration in 6 months.

1.1. The first version demonstrates efficient implementation and intelligent cost-based optimization of matrix-based data analysis workloads on Amazon EC2, using Hadoop/HDFS for underlying execution/storage. The optimization problem is formulated using two criteria clearly understandable to end users: expected execution time and expected monetary cost.

1.2. The second version of Cumulon further demonstrates the ability to leverage auction-based markets (in our case, Amazon spot instances) of cloud resources to lower execution costs. Cumulon makes intelligent choices of bidding strategies, and recovers gracefully from sudden, massive departure of transient machines acquired through bidding. The optimization problem is extended to additionally include a user-specified risk tolerance requirement, because the expected cost alone is unable to capture the risk due to the volatility of auction-based markets.

To support cost-based optimization, we worked on statistical methods for performance prediction and uncertainty quantification. Besides performance, we also modeled the market price of transient resources using the price history of Amazon spot instances.

1.3. The third iteration of Cumulon, currently under active development, aims at making the system adaptive at runtime in a more principled manner. Previously, Cumulon makes a one-time decision taking future uncertainty into account, but assumes that this decision is the only one to make. While the system can always re-optimize at runtime, we want to study how knowing the prospects of future

decisions influences the current decision. The new optimization framework is based on the Markov decision process model and requires re-architecting multiple aspects of the system.

2. As a potential additional application of Cumulon, we have been investigating the use of bilinear mixed-effects modeling for dyadic data. Specially, using machine-coded news events from Eurozone countries between 2001 and 2011, we analyze them with latent network models, which summarize how cooperative or conflictual the public relationships among societal groups are. On the computational side, we have been working on identifying scalability bottlenecks that prevent the analysis to be carried out on larger social networks.

3. Motivated by the recent hardware trend, where most public cloud offerings offer nodes equipped with SSD storage instead of hard drives, we have investigated efficient indexing techniques for SSDs. In particular, we studied the problem of further minimizing writes for the popular log-structured merge approach to indexing. We show how selective partial merge, combined with block-preserving merge, can significantly lower write traffic while preserving other advantages of log-structured merge.

Specific Objectives: At the end of Year 2:

1. We have met our main objectives of building and improving the Cumulon system, evaluating its effectiveness, and learning from this experience.

2. We continue to make progress towards applying Cumulon to practical problems. We have evaluated Cumulon on several well-known, widely applicable techniques, such as Gaussian non-negative matrix factorization and randomize singular value decomposition. Applying Cumulon to the concrete problem in social science we had in mind still requires additional work.

3. We have improved indexing techniques for SSD storage, which has become prevalent in recent years.

Significant Results: 1.1. We have shown that the popular MapReduce programming model is a poor fit for matrix workloads. We have developed a simplified parallel execution model that is more efficient, flexible, and easier for automatic optimization. Instead of reinventing the wheel, we have demonstrated how to implement this new model on top of MapReduced-based Hadoop and HDFS, but in a way that avoids the limitations of MapReduce. This novel strategy offers large savings over more "traditional" uses of Hadoop for matrix workloads.

1.2. Additionally, we have developed storage and execution techniques capable of leveraging transient machines acquired through bidding on a market. Our techniques handle heterogeneous clusters that result from a combination of upfront provisioning and runtime bidding, and gracefully cope with sudden and massive departure of transient machines to reduce their adverse effects on execution costs.

1.3. We have learned how to obtain good cost estimates for matrix workloads in a cloud by benchmarking, simulation, and modeling. Various sources of uncertainty in the cloud make cost estimation particularly challenging. We have devised a better method for predicting job completion time, by accounting for the variance among the speeds of individual tasks. We have also developed a stochastic market price model based on historical prices of Amazon spot instances, which allows us to predict departures and their effects on overall execution costs, and to

quantify the uncertainty in the predictions.

1.4. We have developed optimization techniques that use time, monetary cost, and user-specified risk tolerance as objectives and constraints, and consider a bigger "plan" space whose dimensions include the choices of not only implementation alternatives and execution parameters, but also hardware provisioning and bidding strategies as well as software configuration settings. We have developed a principled approach towards online, adaptive optimization of bidding and execution strategies, based on the Markov decision process.

2. Through experiments on Amazon EC2, we have demonstrated the benefit of our approach with implementations of several well-known, widely applicable statistical analysis techniques, such as Gaussian non-negative matrix factorization and randomize singular value decomposition.

3. For the popular SSD indexing approach, log-structure merge, we developed a deeper understanding of how different merge policies affect performance, and we proposed a new, intelligent partial merge policy with block preservation, which significantly improve the performance on SSDs.

Key outcomes or Other achievements:

1.1. A paper on the first iteration of Cumulon has been published in SIGMOD 2014. Our paper on the second iteration, focusing on leveraging auction-based markets, is currently under revision for PVLDB 2015. An overview paper summarizing our vision and experience up to the second version of Cumulon was published in a special issue of IEEE Engineering Bulletin in 2014. Our paper on the third iteration, focusing on adaptive optimization, is under preparation and should be ready for submission within 6 months.

1.2. An in-depth effort on modeling execution time distributions in Cumulon is described in the 2015 dissertation of Nick Jarrett. Papers on this execution time model and on the market price model of spot instances, with the statistics community as intended audience, are currently under preparation.

3. A paper describing our results on merging techniques for log-structured merge on SSDs is close to be finished and submitted.

4. Jun Yang co-organized a panel about big data at VLDB 2014 titled "Big and Useful: What's in the Data for Me?" Jun Yang also co-chaired the First International Workshop on Bringing the Value of "Big Data" to Users, held in conjunction with VLDB 2014.

*** What opportunities for training and professional development has the project provided?**

Botong Huang, the lead computer science PhD student on this project, and Nick Jarrett, the lead statistics PhD student on this project, have been working closely with each other; they have learned much from each other's discipline as well as about practical cloud computing skills.

Botong interned at IBM Almaden Research Lab in the summer of 2014, working on SystemML, a system being developed at IBM with similar goals as Cumulon (albeit from a provider's perspective instead of users'). He has a SIGMOD 2015 paper based on this internship work. He is schedule to defend in Fall 2015.

Nick Jarrett obtained his PhD in statistics in Spring 2015.

Risi Thonangi, another PhD student, worked on efficient indexing for SSDs. He is scheduled to defend in Summer 2015.

We have redesigned the undergraduate database curriculum ("Introduction to Database Systems" and "Everything Data") at Duke with heavy use of cloud-based virtual machines. The students have been introduced to cloud computing concepts and skills.

*** How have the results been disseminated to communities of interest?**

A paper on the first iteration of Cumulon was published in SIGMOD 2014; Botong Huang delivered a well-received talk at the conference, and was invited to visit the SystemML group at IBM Almaden. Botong interned with the group in the summer of 2014, and has a SIGMOD 2015 paper on his internship work.

Our paper on the second iteration, focusing on leveraging auction-based markets, is currently under revision for PVLDB 2015. An overview paper summarizing our vision and experience up to the second iteration of Cumulon was published in a special issue of IEEE Engineering Bulletin in 2014.

Jun Yang co-organized a panel about big data at VLDB 2014 titled "Big and Useful: What's in the Data for Me?" In the panel, Jun argued that "democratizing data analysis" is as important as "democratizing data," and that we need more systems and research like Cumulon that are more user-facing and user-friendly. To further promote such research, Jun is co-chaired the First International Workshop on Bringing the Value of "Big Data" to Users, held in conjunction with VLDB 2014.

Cumulon's poster has been used in Duke Computer Science Department's graduate recruiting events.

*** What do you plan to do during the next reporting period to accomplish the goals?**

The next reporting period will see the graduation of some of the key students on this project. A new PhD student is scheduled to join the project in Fall 2015 and we are exploring various possible directions that may be best suited to the student. The possibilities include a repository of execution traces and market prices, "deep extensibility" (extending not only Cumulon's functionality with new operators but also their optimizability), and support for tensors, a generalization of the matrix concept that has seen a resurgence of interest in machine learning in recent years.

Products

Books

Book Chapters

Conference Papers and Presentations

Inventions

Journals

Botong Huang, Matthias Boehm, Yuanyuan Tian, Berthold Reinwald, Shirish Tatikonda, and Frederick R. Reiss (2015). Resource Elasticity for Large-Scale Machine Learning. *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. . Status = PUBLISHED; Acknowledgment of Federal Support = No ; Peer Reviewed = Yes

Botong Huang, Nicholas W. D. Jarrett, Shivnath Babu, Sayan Mukherjee, and Jun Yang (2014). Cumulon: cloud-based statistical analysis from users perspective. *IEEE Data Engineering Bulletin*. 37 (3), 77. Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Botong Huang, Shivnath Babu, and Jun Yang (2013). Cumulon: optimizing statistical data analysis in the cloud. *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. . Status = PUBLISHED; Acknowledgment of Federal Support = Yes ; Peer Reviewed = Yes

Licenses

Other Products

Other Publications
Botong Huang, Nicholas W.D. Jarrett, Shivnath Babu, Sayan Mukherjee, and Jun Yang (2014). *Cumulon: matrix-based data analytics in the cloud with spot instances*. Technical report, Duke University, July 2014.. Status = UNDER_REVIEW; Acknowledgement of Federal Support = Yes

Risi Thonangi and Jun Yang (2015). *Optimizing log-structured merge for solid state drives*. Technical report, Duke University. Status = SUBMITTED; Acknowledgement of Federal Support = Yes

Rada Chirkova and Jun Yang (2014). *Proceedings of the 2014 International Workshop on Bringing the Value of Big Data to Users*. Held in conjunction with VLDB 2014.. Status = PUBLISHED; Acknowledgement of Federal Support = No

Patents

Technologies or Techniques

Thesis/Dissertations
Nicholas W.D. Jarrett. *Nonlinear Prediction in Credit Forecasting and Cloud Computing Deployment Optimization*. (2015). Duke University. Acknowledgement of Federal Support = No

Websites
Cumulon
http://db.cs.duke.edu/projects/cumulon

This is the project website for Cumulon.

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Yang, Jun	PD/PI	4
Babu, Shivnath	Co PD/PI	1
Mukherjee, Sayan	Co PD/PI	1
Ward, Michael	Co PD/PI	1
Huang, Botong	Graduate Student (research assistant)	12
Jarrett, Nicholas	Graduate Student (research assistant)	4
Thonangi, Risi	Graduate Student (research assistant)	12

Full details of individuals who have worked on the project:

Jun Yang
Email: junyang@cs.duke.edu
Most Senior Project Role: PD/PI

Nearest Person Month Worked: 4

Contribution to the Project: As the PI, Jun Yang leads the team in building Cumulon, and is responsible for data management for the project.

Funding Support: NSFIIS0916027 NSFIIS1408846 Research grants froms Amazon Web Services and Google Cloud

International Collaboration: No
International Travel: No

Shivnath Babu
Email: shivnath@cs.duke.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: As co-PI, Shivnath Babu leads the design and implementation of the Cumulon system together with Jun Yang, and coadvises the lead computer science PhD student on the project.

Funding Support: None other.

International Collaboration: No
International Travel: No

Sayan Mukherjee
Email: sayan@stat.duke.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: As co-PI, Sayan Mukherjee oversees the statistical aspects of the project. He supervised work on statistical models of runtime as well as those of spot auction pricing.

Funding Support: None other.

International Collaboration: No
International Travel: No

Michael D Ward
Email: michael.d.ward@duke.edu
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: As co-PI, Mike Ward works on statistical methods in political and social sciences, and will apply the proposed research and evaluate the system in problems from these domains.

Funding Support: None other.

International Collaboration: No
International Travel: No

Botong Huang

Email: bhuang@cs.duke.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 12

Contribution to the Project: Botong is the lead computer science student on the project, and the main developer of the Cumulon system.

Funding Support: NSF-IIS-09-16027

International Collaboration: No
International Travel: Yes, Australia - 0 years, 1 months, 0 days

---

Nicholas W.D. Jarrett
Email: nwj2@stat.duke.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 4

Contribution to the Project: Nick Jarrett is the lead student on the statistical aspects of the projec. He worked on developing the various statistical models used in Cumulon.

Funding Support: NSF DMS-12-09155: Collaborative: Numerical Algebra and Statistical Inference

International Collaboration: No
International Travel: No

Risi Thonangi
Email: rvt@cs.duke.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 12

Contribution to the Project: Risi is leading the effort of investigating the implications of the prevalence of SSDs in modern cloud hardware.

Funding Support: None other.

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?
Nothing to report.

What other collaborators or contacts have been involved?
Nothing to report

Impacts

What is the impact on the development of the principal discipline(s) of the project?

Cumulon has helped to demonstrate the power of declarative languages and automatic, cost-based optimization---stables of the database systems---in the new setting of matrix-based data analysis workloads in the cloud. Although

there has been work on automatic optimization of data-parallel workloads in the cloud, what distinguishes Cumulon from others is its focus on the users' perspective, and its ability to provide an end-to-end solution. Cumulon's search space includes not only the traditional dimensions of an execution plan (e.g., alternative implementations and execution parameters) but also novel ones (e.g., software configuration parameters, cluster provisioning and bidding strategies). Cumulon's optimization criteria are also user-centric (time, monetary cost, and risk tolerance) as opposed to provider-centric (e.g., overall throughput). We are hopeful that this new perspective will open up new research directions.

What is the impact on other disciplines?

Cumulon makes principled use of statistics towards modeling and handling of uncertainty. As a result, it has motivated new applications and research questions of interest to the statistics community. A significant portion of Nick Jarrett's dissertation is devoted to problems that arise in Cumulon.

We have been working on testing Cumulon in application domains of other disciplines, and hope to report on its impact soon.

What is the impact on the development of human resources?

So far, the project has supported training and development of three PhD students---two from computer science and one from statistics. One has graduated with a PhD and two others are expected to do so soon. One new PhD student in computer science is expected to join in Year 3. We have also revamped the undergraduate curriculum at Duke to incorporate elements of cloud computing.

What is the impact on physical resources that form infrastructure?

In conjunction with this project, we have obtained research grants from Amazon in the form of Amazon Web Services credits, and from Google in the form of Google Cloud credits, both of which can be used to rent hardware resources as needed from these public cloud providers.

What is the impact on institutional resources that form infrastructure?
Nothing to report.

What is the impact on information resources that form infrastructure?

None so far. However, one plan is to develop a public repository of execution traces, historical market prices, as well as performance and price models that are derived from them. This information resource will be valuable to a cloud-based computing infrastructure.

What is the impact on technology transfer?

Botong Huang, the lead computer science PhD student on the project, interned with the SystemML group at IBM Almaden Research Center in the summer of 2014. SystemML is an IBM project with similar aims as Cumulon. Through this connection, we hope that some of Cumulon's technology can be incorporated into commercial products.

What is the impact on society beyond science and technology?

None so far. We plan to pursue applications of Cumulon to problems in quantitative political science, with impact to the society beyond science and technology.

Changes/Problems

Changes in approach and reason for change
Nothing to report.

Actual or Anticipated problems or delays and actions or plans to resolve them
Nothing to report.

Changes that have a significant impact on expenditures
Nothing to report.

Significant changes in use or care of human subjects
Nothing to report.

Significant changes in use or care of vertebrate animals
Nothing to report.

Significant changes in use or care of biohazards
Nothing to report.