On Development and Optimization of Computing Workflows for Big Data Analytics in Brain Injury Research

Chase Wu[†], Yi Chen[‡], and Namas Chandra[§]
[†‡]Department of Computer Science
[‡]School of Management
[§]Department of Biomedical Engineering
New Jersey Institute of Technology
Newark, NJ 07102
Email: {chase.wu, yi.chen, namas.chandra}njit.edu

***Abstract*—Many scientific applications are producing a colossal amount of data that must be processed and analyzed for scientific exploration and knowledge discovery. The brain injury research at the Center for Injury Biomechanics, Materials and Medicine (CIBM3) at New Jersey Institute of Technology (NJIT) is one such application. We propose to develop a workflow-based solution that integrates various big data analytics approaches to support the brain injury research. Within the workflow framework, the data-intensive analytics modules are implemented and executed on big data platforms such as Hadoop using cloud computing resources in virtualized environments. The proposed workflow system is generic and can be adapted to other scientific domains with similar computing and networking needs.**

***Index Terms*—Big Data Analytics; Brain Injury Research; Scientific Workflow**

I. Introduction

The rapid advance in computing technology is expediting the transition in various basic and applied sciences from traditional laboratory-controlled experimental methodologies to modern computational paradigms involving complex numerical model analyses and extreme-scale simulations of physical phenomena, chemical reactions, climatic changes, and biological processes. These computation-based analyses and simulations have become an essential research and discovery tool in next-generation scientific applications and are producing colossal amounts of data, now frequently termed as “Big Data”, on the order of terabyte at present and petabyte or even exabyte in the predictable future [6]. Other data of similar scales generated in broad science, engineering, and business domains include environmental observation data (satellite data [1], multimodal sensor data, etc.), experimental measurement data (Spallation Neutron Source [2], Large Hadron Collider [3], etc.), astronomical imaging data (One Degree Imager [11], Dark Energy Camera [10], Large Synoptic Sky Survey [13], etc.), bioinformatics and medical data, financial trading and transaction data, and Internet-based social network data. No matter which type of data is considered, a high-performance computing solution is required for geographically distributed users to store, transfer, process, visualize, analyze, and synthesize the data for collaborative research and discovery.

At the core of these applications is the mission to achieve “Big Impact Through Big Data”. However, their success requires the use of a wide range of expensive and powerful resources including supercomputers, PC clusters, high-end workstations, experimental facilities, large-screen display devices, high-performance network infrastructures, and massive storage systems [9]. Typically, these resources are deployed at various research institutes and national laboratories, and are provided to application users through wide-area network connections that may span across the nation or even over several countries, hence inevitably exhibiting an inherent dynamic nature in their accessibility, availability, capacity, and reliability. As new computing and networking technologies rapidly emerge, enabling functionalities and services are progressing at an ever-increasing pace, but unfortunately, so are the dynamics, scale, heterogeneity, and complexity of the data analytics problems and computing environments. Application users, who are primarily domain experts, need to manually search for the data of their interest and suitable toolkits for data analytics, and then configure and run their computing tasks over networks using software tools or systems they are familiar with based on their own empirical studies, oftentimes resulting in unsatisfactory performance in such diverse and dynamic environments.

At the Center for Injury Biomechanics, Materials and Medicine (CIBM3) at New Jersey Institute of Technology (NJIT), a multidisciplinary faculty team from biology, engineering, mathematics, neuroscience and clinical medicine work together to address the holistic questions in brain injury, from protection, diagnostics, therapeutics, to rehabilitation. The researchers have developed unique experimental and computational facilities that can simulate injury causing conditions accurately and are used to test live tissues, small/large animal models, human cadavers, and manikins under realistic conditions. These experimental and computer models could help the researchers understand how the injuries are caused and then develop predictive models useful to engineers and medical professionals. The research attempts to find both biomechanical and biochemical pathways of how injury is caused and how it is manifested as acute and chronic neurological dysfunctions with poor medical outcomes. It is important to not only understand these mechanisms in the animal model but also translate these results to humans. Sometimes, it is not always clear what are the biological, mechanical and clinical parameters that need to be controlled or measured, and significant challenges still remain in processing the large amounts of data for knowledge discovery and model-based prediction.

We propose a workflow-based solution that consists of various computing modules of big data analytics techniques to support brain injury research. Within the workflow framework, the data-intensive analytics modules are implemented and executed on big data platforms such as Hadoop system using cloud computing resources in virtualized environments. We would like to point out that the proposed workflow system is generic and can be adapted to other scientific domains with similar computing and networking needs.

The rest of the paper is organized as follows. In Section II, we provide a brief introduction to brain injury research. In Section III, we sketch a number of data analytics approaches to address the big data challenges in this application domain. In Section IV, we propose a workflow-based solution to big data analytics in high-performance computing and networking environments. In Section V, we describe some technical challenges in implementing and deploying the workflow system in cloud environments.

II. Big Data Challenges in Brain Injury Research

In brain injury research, both experimental and computational data have been generated. In each experiment, we measure pressure, acceleration, strain, and images at multiple locations and collect data at the rate of 1 MHz for 100 msec. A dataset may consist of at least 15 to 20 of these measurements. Currently, all the numerical data are stored in raw data form in native and excel formats. Two or three camera shots of the event are recorded at a resolution of 20000 frames per second and are stored in video format. A single experiment may produce data of about 3 GB.

Biological tissues of the animal post-injury are sliced (200 slices) and the histological information is stored at 100 MB per sample. A similar size data is expected when we do protein, antibody or other biomarker analyses. The computer models for the shock-structure interaction involve about 5 million finite elements with a simulation time of 5 msec (64 processors run for 12 hour clock time) at the rate of 1 picosecond. Typical raw data may be of 3 or 5 GB. We conduct about 300 experiments per year.

All the experimental data are not used in most of the analysis. Currently, we only use a very small subset of data consistent with the testing of our hypothesis. Since our test specimen is a rat of very specific kind (Sprague-dawley male, specific age), we are interested in knowing how to use that data if we do some additional tests on other variations. Also, we want to find out what other information is contained in such big data that can be made useful by following proper data analytics.

III. Big Data Analytics in Brain Injury Research

To address the challenges and make sense of the data in brain injury research, we propose to develop various data analytics techniques, as outlined below.

Currently, all the numerical data are stored in plain text or excel files. To link data of different types, such as experiment setting, position of shock, type of shock, mechanical and medical conditions, special programs have to be developed for efficient data synthesis and analysis. Also, it is necessary to link and correlate the information at different time points to better understand the condition change of an individual animal. To facilitate analysis and reduce development efforts, we will employ database systems to abstract the common modules of analysis needs, enable the expression of user analysis as declarative queries built upon the common modules, and create indices offline to speed up the data processing. In particular, we will investigate traditional relational databases, column-based databases, and NoSQL systems for storage backend options based on the study of analysis needs. We will provide a query processor with a user-friendly interface to allow users to easily query the data. The user input will be converted to a declarative query, which will be processed and visualized. Some example queries may include: which experiment setting leads to identified brain injury, what position is more likely to result in injury, what is the difference in the injury caused by the same shock wave to different subjects, and so on.

The grand challenge is to develop predictive models. Since it is infeasible to conduct all types of experiments on all available species (especially human subject), one solution is to predict a possible impact on a certain subject based on the historical experiments and observations of other subjects (e.g. animals) of possibly different species. Human brain injury data can also be accessed through our collaboration with the emergency clinical medical database. These data provide actual injury scenarios (e.g. falling from height, motor vehicle accidents, blast injury, etc.), treatment regimens, initial prognostics and progression of diseases. One main research goal is to integrate animal and human data in a seamless manner using data analytics. Towards this goal, we plan to study latent models for prediction, e.g. Singular Value Decomposition (SVD) and Factorization Machines (FM). The rationale behind these models is that different combinations of individual animals and experiment settings may lead to different impacts on the brains. However, there are a large amount of features in experiment settings, and even many more features in individual observations, from biological tissue,to protein, antibody or other biomarkers. Since it is challenging to select features, latent models can learn these impact patterns hidden in the data. For SVD, we plan to build an individual-experiment-setting matrix, in which the value of each element is the level of brain impact. The missing values, i.e. the brain impact of experiments that are not conducted, will be predicted by SVD. By adding more features that are known to have big impact based on the suggestion of medical experts, we can build an FM model, which not only computes the internal model using factorized interactions between individual animals and experiment settings, but also allows more real-value features, like those in linear regressions. Given any combination of individual animals and experiment settings, the trained FM model may be able to predict the corresponding impact level.

Brain injury researchers are also interested in identifying similar cases. For instance, a user may want to find similar experiment settings, or similar statuses of biomarkers. However, this is very challenging due to the high dimensionality of the data under consideration. Furthermore, since the biomarker data are represented as sequences/time series, we will study similarity measure for sequences and develop effective techniques for similarity measurement.

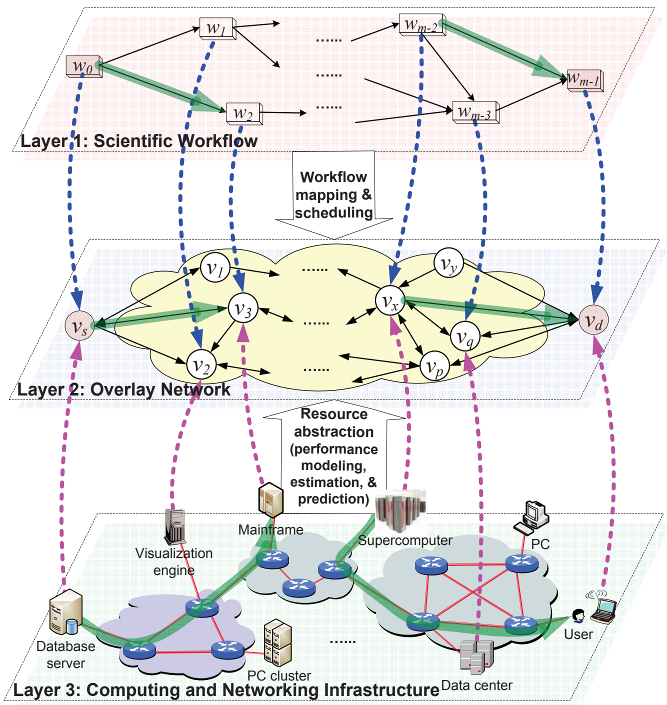

Fig. 1: A generic three-layer workflow architecture for executing and optimizing big data scientific applications in distributed environments.

IV. A Workflow-based Solution to Big Data Analytics

Many of the aforementioned techniques to be developed for data processing and model-based prediction (including model construction, training, and testing) in brain injury research are computationally intensive and hence require the use of high-performance computing resources. Considering the sheer volume of data to be processed, we implement and execute data-intensive analytics modules on big data platforms such as Hadoop in virtualized environments. Also, since the execution of one computing module might depend on the result of another, we propose to develop a generic workflow-based solution to support big data analytics in brain injury research.

As shown in Fig. 1, we consider a layered workflow architecture that consists of three interrelated layers: (i) a top layer – abstract scientific workflow, (ii) a middle layer – virtual overlay network, and (iii) a bottom layer – physical computing and networking infrastructure. The interactions and transformations between these three layers produce an integrated and intelligent workflow solution to optimize the performance of scientific applications including brain injury research in resource sharing environments.

The top layer defines abstract scientific workflows comprised of computational and computing modules with high-level functional and I/O descriptions. This layer provides a unified web interface for users to compose, dispatch, and monitor domain-specific workflows while the rest of the system is made completely transparent to them. The simplest workflow may include only two modules for a point-to-point data transfer while a complex one may involve as many as thousands of modules with intricate execution dependencies.

The bottom layer defines underlying physical system resources including large data repositories storing (simulated, observational, or experimental) high-resolution multimodal scientific datasets, high-speed network infrastructures provisioning high bandwidth for fast data transfer, and HPC facilities generating countless CPU cycles for expeditious data processing.

The top and bottom layers meet at the middle layer that defines a virtual overlay network through the following two operations, which are interleaved when dynamic workflow reconfigurations are required:

i) Resource abstraction: Build a virtual overlay network from a given scientific workflow and the underlying networked computing resources via performance modeling and prediction. Each overlay node with estimated processing power corresponds to a physical machine/site or a virtual machine as in clouds, and each overlay link with estimated bandwidth corresponds to a network path consisting of a set of physical links. In most scientific applications, inter-task data exchanges are performed through overlay links at the IP or logical level.

ii) Workflow mapping and scheduling: Determine a workflow mapping scheme that assigns each module in the scientific workflow to an overlay node in the overlay network, and decide a task scheduling policy on each mapping node to optimize end-to-end workflow performance such as latency, throughput, or reliability. The resultant mapping scheme and scheduling policy can be utilized by any existing workflow engine to dispatch individual modules for distributed execution. We will leverage the functionalities of existing workflow systems such as Condor/DAGMan [4], [5], SWAMP [15], Pegasus [8], and Oozie/Hadoop to reduce our implementation efforts so we can focus on the research aspects of the proposed workflow solution.

The workflow engine converts abstract workflows to concrete workflows after resolving the location of data files and executables and establishing network channels required by component modules. The actual workflow execution occurs in the resource layer where computing modules are executed on computer nodes. Such a layered approach coupled with modularized design makes the proposed workflow system generic in handling the disparities in the file format, analysis program, data transfer, and result display in different application domains and network environments.

V. Targeting Cloud Platforms

With the emergence of cloud computing, an increasing number of scientific workflows have been moved or are in active transition to clouds from traditional computing environments. This shift of computing paradigm has reaped the most significant benefit of resource virtualization by untangling application users from complex management and maintenance of underlying resources, but meanwhile has also brought on many new challenges for workflow execution and performance optimization. We focus on the following aspects targeting cloud platforms.

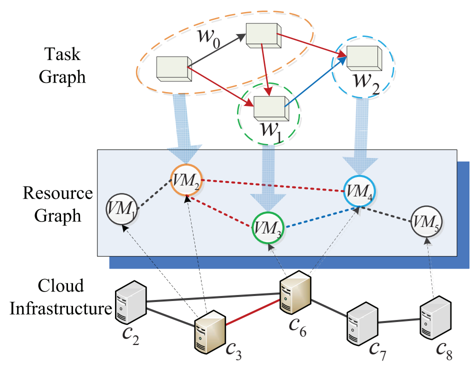

Fig. 2: Modeling workflow execution in cloud environments.

Manageability: As cloud computing makes computing a utility, both the cloud provider and the workflow owner face the challenge to reduce cost in time, finance, and energy. To support cost-effective execution of scientific workflows in clouds, we adapt the general three-layer architecture in Fig. 1 from traditional networks to clouds by focusing on resource virtualization. As shown in Fig. 2, the resource graph layer represents a network of virtual machines (VMs) provisioned from the cloud infrastructure layer. In the task graph layer, we consider scientific workflows preprocessed by an appropriate clustering technique based on the inter-module dependencies and the volumes of inter-module data transfer [12], [14], [7], where a group of modules in the original workflow are bundled together as one aggregate module in the resulted task graph. We will extend the performance models in traditional networks to quantify the performance of big data scientific workflows in Infrastructure as a Service (IaaS) cloud environments, and generalize the workflow optimization problem to minimize the workflow end-to-end delay under a user-specified cost constraint. We will design efficient algorithms for this problem by leveraging the mapping and scheduling algorithms initially targeting traditional computing environments.

Scalability and Elasticity: Cloud computing enables service providers and application users to scale up or down computing resources as needed, which is an important feature especially when experiencing a sudden growth or drop in demand. This scalability/elasticity can lead to significant economic benefit, but may also affect workflow mapping/scheduling as the resources needed by a particular computing module may not be always available on a designated VM. Particularly, for MapReduce-based computing modules, a cluster of VMs need to be provisioned and configured to process big data in a distributed manner. We will refine our workflow mapping/scheduling algorithms to provision a set of appropriate VMs that are able to execute the workflow to its completion with the minimum computing and networking resources. Also, due to the correlation between the resource capacities of those VMs provisioned on the same physical server, the mapping and scheduling algorithms must be revised to explicitly account for resource integration and job migration across VMs.

Reliability: As the scope and scale of the cloud infrastructure are rapidly expanding, node and link failures are inevitable, causing a detrimental impact on the workflow performance. We will formulate and investigate multi-objective problems (MOPs) that take reliability into consideration in addition to those traditional goals such as latency or throughput. Another challenge in clouds arises from the network requirement for big data transfer. Since cloud computing utilizes massively distributed resources on the Internet, there may exist large data movement between VMs during workflow execution, for which, we need to either provision dedicated bandwidth or choose less congested links to guarantee fast and reliable access to remote data in the cloud. We will also refine our workflow mapping algorithms to avoid unnecessary data transfer by mapping workflow modules to the VM(s) provisioned from the same or nearby physical servers.

Security and Privacy: Security and privacy are among the most daunting challenges in clouds and must be carefully addressed for workflow execution. For instance, cloud vendors face major issues in data confidentiality, integrity, and availability, and may employ various authentication and authorization mechanisms to protect the client's information by only allowing access to those data and applications pertaining to the submitted job. As a result, some workflow modules can only be assigned to certain VMs, which must be explicitly considered in the mapping process. The mapping problem becomes even more complex if we have to coordinate the workflow execution across multiple domains with different security and privacy policies.

VI. Conclusion

We proposed a workflow-based solution to big data analytics in brain injury research. The proposed workflow system is generic and is applicable to other scientific domains with similar computing and networking needs. We will implement this workflow system in Hadoop environment using cloud computing resources and integrate various MapReduce-based data analytics programs to solve big data problems in different application domains including brain injury research.

Acknowledgments

This research is sponsored by National Science Foundation under Grant No. CNS-1560698, and Oak Ridge National Laboratory, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725/4000141963 with New Jersey Institute of Technology.

References

[1] Atmospheric Sciences Division, Brookhaven National Laboratory. http://www.ecd.bnl.gov.

[2] Spallation Neutron Source. http://neutrons.ornl.gov, http://www.sns.gov.

[3] Relativistic Heavy Ion Collider. http://www.bnl.gov/rhic.

[4] Condor. http://www.cs.wisc.edu/condor.

[5] DAGMan. http://www.cs.wisc.edu/condor/dagman.

[6] Synergistic challenges in data-intensive science and exascale computing. Technical report, March 30 2013. Data Subcommittee Report, Advanced Scientific Computing Advisory Committee (ASCAC).

[7] W. Chen and E. Deelman. Integration of workflow partitioning and resource provisioning. In *Proc. of the 12th IEEE/ACM Int. Symp. on Cluster, Cloud and Grid Computing*, pages 764–768, Washington, DC, USA, 2012.

[8] E. Deelman, G. Singh, M. Su, J. Blythe, A Gil, C. Kesselman, G. Mehta, K. Vahi, G. B. Berriman, J. Good, A. Laity, J. C. Jacob, and D. S. Katz. Pegasus: a framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming Journal*, 13:219–237, 2005.

[9] High-performance networks for high-impact science, Aug. 13-15 2002. Report of the High-Performance Network Planning Workshop, http://DOECollaboratory.pnl.gov/meetings/hpnpw.

[10] B. Flaugher. The dark energy survey camera (decam). *Bulletin of the American Astronomical Society*, 42:406, Jan. 2010.

[11] G.H. Jacoby, J.L. Tonry, B.E. Burke, C.F. Claver, B.M. Starr, A. Saha, G.A. Luppino, and C.F.W. Harmer. Wiyn one degree imager (odi). In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4836, pages 217–27, Dec. 2002.

[12] J.Y. Jung and J. Bae. Workflow clustering method based on process similarity. In *Proc. of the Int. Conf. on Computational Science and Its Applications*, volume 2, pages 379–389, 2006.

[13] A. Rasmussen, K. Gilmore, S.M. Kahn, J. Geary, S. Marshall, M. Nordby, P. O'Connor, S. Olivier, J. Oliver, V. Radeka, T. Schalk, R. Schindler, J. Tyson, R. Van Berg, and LSST Camera Team. The camera for LSST and its focal plane array. *Bulletin of the American Astronomical Society*, 41:221, Jan. 2010.

[14] G. Singh, M.H. Su, K. Vahi, E. Deelman, B. Berriman, J. Good, D. Katz, and G. Mehta. Workflow task clustering for best effort systems with pegasus. In *Proc. of the 15th ACM Mardi Gras Conference*, pages 9:1–9:8, 2008.

[15] Q. Wu, M. Zhu, X. Lu, P. Brown, Y. Lin, Y. Gu, F. Cao, and M.A. Reuter. Automation and management of scientific workflows in distributed network environments. In *Proc. of the 6th Int. Workshop on Sys. Man. Tech., Proc., and Serv.*, Atlanta, GA, April 19 2010.