

Building scalable data management and analysis infrastructure for metagenomics

Wei Tang, Jared Wilkening, Jared Bischof, Wolfgang Gerlach,
Andreas Wilke, Narayan Desai, Folker Meyer

Argonne National Laboratory
{wtang, jared, jbischof, wgerlarch, wilke, desai, folker}@mcs.anl.gov

ABSTRACT

Next-generation sequencing technology has reduced the cost of DNA sequencing dramatically and shifted the bottleneck of metagenomics from data generation to data analysis. For example, MG-RAST, a free open-public metagenome annotation system, has been experiencing an increasingly large amount of data being submitted for analysis—a situation that threatens to overwhelm efficient production. To address this situation, we developed a pair of open-source software products: a data management system named Shock and a workflow management system named AWE. Shock and AWE can be used to build scalable infrastructure for biological sequence data management and analysis.

Keywords

data analysis infrastructure, data management system, workflow engine, cloud computing, bioinformatics, metagenomics

1. INTRODUCTION

Metagenomics is the study of microbial communities that inhabit the world. Insight into these communities sheds light on the a range of areas, including human health and environmental processes. Next-generation sequencing technology has exponentially reduced the cost of data generation. This has greatly thrived the metagenomics study as increasing amounts of data become available. However the computational costs of analysis have grown unsustainably. Therefore, the bottleneck in metagenomics has moved from sample collection and data generation to data analysis.

MG-RAST [2][4][6] is a web-based set of comparative analysis tools for microbial communities. It provides automated data analysis capabilities and tools for interactive exploration and comparison of user-submitted data sets. MG-RAST was launched in 2007 at Argonne National Laboratory as a free service, and has since grown to become the dominant community resource for metagenomic data analysis. As of middle October 2013, MG-RAST has an-

alyzed 36 TBp (tera-base-pairs) of sequence data in 95,000 metagenome samples for over 10,000 registered users from over 60 different countries. A typical procedure of metagenomics study is shown in Figure 1.




Figure 1: A typical procedure of metagenomics

The MG-RAST project has experienced the data deluge firsthand, as submission volumes have grown exponentially. This data growth has caused two major issues. First, the data deluge is causing longer job turnaround times as the computing capability is limited within one computing facility. Second, with the growing data not managed properly, unnecessary duplicated computation is inevitable, and reproducing of the results becomes harder. To address these problems, we developed Shock data management system and AWE workflow management system in order to build a scalable, reproducible data analysis platform [5].

2. SYSTEMS AND METHODOLOGY

We introduce here the design of our data management and workflow management systems.

2.1 Shock data management system

Shock is a data management system for moderate to large-scale biological data sets. Shock is designed as a layer of system software built on top of backend storage systems, providing object-based storage interface tailored for bioinformatics workloads. An object in Shock encapsulates the actual data file and the metadata of its scientific provenance as well as computational provenance. Shock supports a number of common bioinformatics data formats and can provide

subselection and merge capabilities that minimize I/O requirements for analysis workloads and improves scalability for large numbers of computing clients. These capabilities are implemented by indexing the data in chunks (e.g., one sequence per chunk for biological sequence data).

2.2 AWE workflow management system

AWE is a workflow engine and resource manager. AWE models application related concepts into three hierarchical elements: job, task, and workunit. A job, characterized by its input data and the workflow description, is parsed into tasks representing certain command; a task can be split into multiple workunits running the same command on different parts of data.

AWE is designed as a distributed system that contains a centralized server and multiple distributed clients. The server receives job submissions and parses jobs into tasks, splits tasks into workunits, and manages workunits in a queue. The AWE clients, running on distributed, heterogeneous (in terms of hardware or applications) computing resources, keep checking out workunits from the server queue and dispatch the workunits on the local computing resource. AWE uses Shock to handle input and output data (retrieval, storage, splitting, and merge). REST APIs are used for communication between all components. Also, AWE provides flexible, fault-tolerant queue management and computing resource management.

Collectively, Shock and AWE can build an integrated data analysis ecosystem including task, computing resource, and data management. A working diagram is described in Figure 2.




Figure 2: Working diagram. A centralized AWE server parses submitted jobs and manages the workunit queue. Distributed AWE clients obtain workunits from the server and execute them locally. Shock is used for managing data.

3. EXAMPLE DEPLOYMENT

Figure 3 shows an example deployment of Shock and AWE for MG-RAST. In this diagram, users submit their jobs to MG-RAST as usual. But if the users have computing resources that can play virtual machine (VM) instances installed with MG-RAST applications and the AWE client, workunits in the central task queue can be scaled out to the users' computing resources, including those in local computing facilities or the public clouds such as Amazon EC2.

Input data will be fetched from Shock and output data will be put back to Shock. The final results will be available through MG-RAST website. In this way, MG-RAST users who have computing resources can have their job process accelerated. We can choose to scale out tasks with high computing-to-i/o ratio so that the data movement overhead can be kept at a very low level. This working diagram has been tested with multiple users from several institutes and the throughput of MG-RAST has been improved with the addition of the users' computing cycles.




Figure 3: Example deployment: running MG-RAST tasks on users' computers to accelerate job process

Shock and AWE are both open-source projects [1][3] and can be used to build general data analysis platforms for data analysis services similar to MG-RAST.

4. ACKNOWLEDGMENTS

This work is supported by the U.S. Department of Energy under Contract DE-AC02-06CH11357.

5. REFERENCES

- [1] AWE project. <https://github.com/MG-RAST/AWE>
- [2] MG-RAST website. <http://metagenomics.anl.gov>
- [3] Shock project. <https://github.com/MG-RAST/Shock>
- [4] F. Meyer, D. Paarmann, M. D'Souza, R. Olson, E. M. Glass, M. Kubal, T. Paczian, A. Rodriguez, R. Stevens, A. Wilke, J. Wilkening, and R. Edwards. The metagenomics RAST server - a public resource for the automatic phylogenetic and functional analysis of metagenomes. *BMC Bioinformatics*, 9(386):1-8, 2008.
- [5] W. Tang, J. Wilkening, N. Desai, W. Gerlach, A. Wilke, and F. Meyer. A scalable data analysis platform for metagenomics. In *Proc. IEEE International Conference on Big Data*, 2013.
- [6] A. Wilke, E. Glass, D. Bartels, J. Bischof, D. Braithwaite, M. D'Souza, W. Gerlach, T. Harrison, K. Keegan, H. Matthews, R. Kottmann, T. Paczian, W. Tang, W. Trimble, P. Yilmaz, J. Wilkening, N. Desai, F. Meyer. A metagenomics portal for a democratized sequencing world. *Methods Enzymol.*, 531:487-523, 2013.