



A study on Parallelization of Scientific Workflows and Integration of Swift and Galaxy Systems

Dr. M. Suresh Babu¹, A. Vijay Kumar², M.N.L. Anuradha³

Professor, Department of CSE, Nalla Malla Reddy Engineering College, Hyderabad¹

Associate Professor, Department of IT, Nalla Malla Reddy Engineering College, Hyderabad²

Associate Professor, Department of CSE, Vidya Jyothi Institute of Technology, Hyderabad³

Abstract: Galaxy is a scientific workflow, data integration, and data and analysis persistence and publishing platform that aims to make computational biology accessible to research scientists that do not have computer programming experience. Although it was initially developed for genomics research, it is largely domain agnostic and is now used as a general bioinformatics workflow management system. While user-friendly and intuitive for doing small to medium-scale computations, it currently has limited support for large-scale parallel and distributed computing. The Swift parallel scripting framework is capable of composing ordinary applications into parallel scripts that can be run on multiscale distributed and performance computing platforms. Swift is an implicitly parallel programming language that allows the writing of scripts that distribute program execution across distributed computing resources, including clusters, clouds, grids, and super computers. Swift implementations are open source under the Apache License, version 2.0. In complex distributed environments, and often the user end of the application lifecycle slows because of the technical complexities brought in by the scale, access methods, and resource management nuances. Galaxy offers a simple way of designing, composing, executing, reusing, and reproducing application runs. Integration between the Swift and Galaxy systems can accelerate science as well as bring the respective user communities together in an interactive, user friendly, parallel and distributed data analysis environment enabled on a broad range of computational infrastructures.

Keywords: Swift, Galaxy, Big Data, scientific applications.

1. INTRODUCTION

The appearance and effect of huge information are obvious from the constantly expanding development of information era. Tremendous measures of experimental information lie unexplored even as more calculations and codes are created to dissect enormous information in novel ways. Computational science today progressively is focalizing to huge information investigation. Moreover, exploratory applications are extending their spaces, bringing about an increment in the era and handling of information and in the blossoming requests on computational assets. Electronic coordinated effort, internet handling, and sharing of uses through entrance based situations have ended up ordinary. Moreover, re-Permission to make computerized or printed versions of all or a portion of this work for individual or classroom use is allowed without charge gave that duplicates are not made or conveyed for profit or business point of preference and that duplicates bear this notification and the full reference on the first page. To duplicate generally, to republish, to present on servers or on redistribute to records, requires former specific consent and/or an expense. Prerequisites, for example, investigation of results with visual guides, post processing of execution follows, and tweaked perception have added to the complexity.

The Galaxy [3] system addresses many of these demands, offering a web-based, interactive platform for data analysis via reusable and reproducible workflows. While excelling in the usability aspects, the Galaxy environment has however a limited parallelism capabilities and limited capabilities for interfacing with diverse, scalable computing infrastructures. A Swift script describes strongly typed data, application components, invocations of applications components, and the inter-relations in the dataflow between those invocations. The program statements will automatically run in parallel unless there is a data dependency between them, given sufficient computing resources. The design of the language guarantees that results of a computation are deterministic, even though the order in which statements executes may vary. A special file data type is built into Swift and Swift allows command-line programs to be integrated into a program as typed functions. This allows programmers to write programs that treat command-line programs and files in the same way as regular functions and variables. A concept of "mapping" is used to store and



exchange complex data structures using a file system structure with files and directories. Rapid dispatch of parallel tasks to a wide range of resources is implemented through a mechanism called "Coasters task dispatch". A Message Passing Interface based implementation of the language supports very high task execution rates on large clusters and supercomputers.

The user can draw workflows in which multiple copies of a task are created to operate on multiple inputs at once, but this is not a scalable approach. Thus, integration with an existing highly distributed, portable system is desirable. Swift [8] is a scripting framework for concurrent execution of ordinary programs. The Swift runtime contains a powerful platform for running user programs on a broad range of computational infrastructures, such as clouds, grids, clusters, and supercomputers, out of the box. Users rapidly prototype complex application flows on small, local resources and then deploy them on a diverse range of remote computation systems. Compact, C-like Swift scripts provide powerful semantics to exploit non-obvious concurrencies in complex application flows. The concise and powerful foreach loop can be used to generate thousands of concurrent tasks with just two lines of code. Thus, highly concurrent applications may be constructed in a small amount of script logic without explicitly managing process ids, batch queues, graph construction, or other tedious details.

Moreover, the Swift framework is specialized to run scientific computations on multiple, remote resources while being well interfaced with diverse schedulers, transport protocols, and security measures. Swift handles the data movement among distributed file systems and can leverage powerful data movement mechanisms. Swift works well with ssh, PBS, Condor, SLURM, LSF, SGE, Cobalt, and Globus to run applications and with scp, http, ftp, and GridFTP to move data.

In this paper, we report on our efforts to extend the Galaxy environment with parallel and distributed execution systems via the Swift framework. The result is an integrated environment with the same user-friendly Galaxy frontend and Swift as a backend. We motivate the integration by arguing that a combination of the best of both worlds will result in a powerful solution useful to the wider user community (x2).

We support our argument with a description of various integration schemes and resulting benefits (x3). We describe our experience with some real-world application implementations (x4). We provide perspectives on related systems by comparison with similar work (x5). Finally, we conclude with a summary and a description of current and future work (x6).

2. MOTIVATION

More and more computational resources are being acquired by organizations such as universities, national laboratories and research consortia to meet an increasing demand for computational and storage capabilities for big data analytics. Similarly, private owners such as Amazon, Google, and Microsoft have started offering resources in the form of clouds. Science users are likely to have allocations from many of these resources with heterogeneous architectures, accessibility modes, and resource managers. Thus, to effectively use each system in isolation or in combination, users need to expend additional efforts that can be a significant distraction and may require special expertise. Consequently, just the provision of resources is not sufficient. While the demand for computational resources is being met, a wide gap still exists between effective utilization of these systems by end-users. Portal-based systems such as Galaxy specialize in providing user-friendly computational analysis environments. Galaxy workflows largely run on local resources.

More computational assets are being gained by associations, for example, colleges, national labs and examination consortia to take care of an expanding demand for computational and stockpiling abilities for enormous information investigation. Also, private proprietors, for example, Amazon, Google, and Microsoft have begun offering assets as clouds. Science clients are liable to have allotments from large portions of these assets with heterogeneous architectures, openness modes, and asset chiefs. Subsequently, to viably utilize every framework in detachment or in blend, clients need to consume extra endeavors that can be a huge diversion and may require unique aptitude. Thus, simply the procurement of assets is not adequate. While the interest for computational assets is being met, a wide crevice still exists between successful uses of these frameworks by end-clients. Entrance based frameworks, for example, Galaxy represent considerable authority in giving easy to use computational examination situations. Galaxy work processes generally keep running on local resources.

Galaxy providers offer their own servers as execution platforms. As the number of users increases however, servers face spikes in demands. As a result, Galaxy offers a cloud based solution such that users run their own Galaxy servers in cloud and manage cloud resources (e.g. via ad hoc Condor pools). Although community developments are aimed at running Galaxy tools on clusters, using them requires special tool configurations and constraints, for example, a shared file system between Galaxy server and the target cluster.

The following points summarize the motivation behind our conceptualization of the Swift-Galaxy integration.



- Benefits to users from the Galaxy community in interfacing their computations to large-scale parallel systems will include remote clusters, clouds, and supercomputers.
- An interactive distributed and parallel computing environment resulting in powerful capabilities of defining computations, orchestrating tasks to remote machines and monitoring the executions and results from a portal-based platform will give unprecedented power to Galaxy users.
- Swift users will be able to easily manage their applications and workflows. Swift users can share, reuse, and publish their distributed computations with Galaxy style conventions and practices that are well adapted by a large user base. This can result in a wider community uptake of new computational codes and expansion of existing capabilities.
- The integration will result in a clear separation of concerns. Task coordination and orchestration will be handled by Galaxy while parallel execution on remote nodes and related data management will be handled by the Swift engine. Application developments benefit by this clean separation of responsibilities; distribution will be readily available without additional configuration on either side.
- Currently, the Galaxy server is operated on a single cluster at the University of Pennsylvania. With an expanding user community and increasing demand for computational resources, a broader capacity of resources and capability of execution will help meet this demand.

This will result in a productive cycle where the growing demand in computation is met by deploying Galaxy servers on diverse systems, thus reaching wider science and engineering domains and resulting in more such deployments.

3. INTEGRATION SCHEMES

Both the Swift and Galaxy systems are powerful in different and complementary aspects, since they were designed with different goals in mind. Galaxy is a web-based portal for visual workflow composition, execution, and data analysis.

Swift is a distributed computing framework for text-based workflow composition. The Galaxy environment is easy to adapt, simple, with limited control over resources. Swift is flexible, implicitly parallel, and robustly interfaced to a wider range of execution platforms. Clearly they form two different paradigms of computing for addressing the challenges faced in data science.

Swift and Galaxy workflows can interoperate because of their similar model of basic interactions with external systems. A unit of execution in Galaxy, called a "tool", is defined by an XML document describing the command line, executable, interpreter to execute the command line, and any required inputs and outputs. The user controls tool inputs and outputs through the Galaxy workspace interface. Tools are composed together into workflows described internally as JSON [1] documents. The Galaxy engine processes the connections between the tools, identifies dependencies among tools, and schedules them for execution accordingly. Workflow execution results are saved as reproducible "histories", which may be shared among multiple users. Both Swift and Galaxy operate on executable by invoking them via operating system utilities. The overarching goal of our integration is to enable a powerful environment by combining the strengths of each of the systems. We aim to achieve this with minimal modification of either systems. In other words, we take both systems "as is" and explore ways to integrate them. The basic mechanism we employ is to use Swift as the internal, low-level execution engine and to use Galaxy as the higher-level, external, user-visible framework.

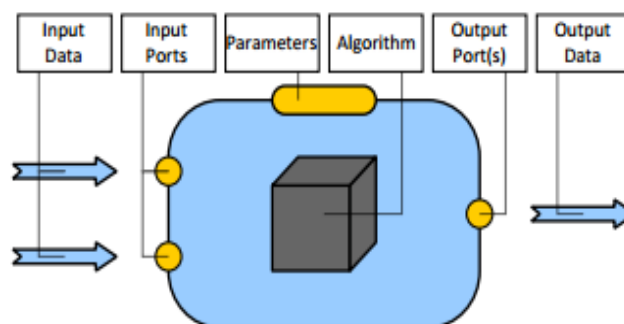


Figure 1 : The Structure of Single workflow task.

Figure 1: Swift acting as an interface between locally installed Galaxy tools and remotely located resources. Figure 1 shows a general scheme of Swift's role in extending Galaxy to a wide range of scientific tools and computational resources. Under this scheme, we are developing modalities through which a variety of Swift-enabled applications will be readily used as Galaxy tools. The execution starts from a local host as a local Galaxy run. The Galaxy tool invokes



Swift application, using preconfigured application libraries written in Swift. The Swift engine in turn orchestrates and sends tasks to remote resources. In the rest of this section we describe our experience in integrating the Swift and Galaxy gateways. We develop different schemes of integration and discuss the benefits of each.




Figure 2: Swift-Galaxy integration view: custom Galaxy tools call Swift.

3.1 Scheme 1: Swift as Galaxy tool

As a first integration scheme, we developed a generic tool with a capability of executing user-provided arbitrary Swift scripts. The tool is set up such that users can select various configuration parameters from available choices. This capability allows users to provide arbitrary parameters specific to the application and the execution-specific parameters such as target compute site.

Figure 2 presents the user view of the Galaxy interface with preset, static Swift-wrapped tools that are available for use individually or as part of workflows. In addition to the tool's own configurable properties, Swift-specific properties allow users to choose the remote resource on which to run the tools with custom configuration such as desired degree of parallelism, and job distribution among sites. With this scheme users can run individual tools as Swift scripts, as well as stitch tools together running a workflow made up of other Galaxy tools and Swift tools.

3.2 Scheme 2: Translating Galaxy to Swift




Figure 3: Interoperability between Galaxy workflowsto Swift scripts.

Figure 3 shows a scheme where predefined Galaxy workflowsare transliterated into Swift scripts via a post processingscript. This development will enable interoperability between the two platforms. Galaxy workflows are currentlyexpressed as JSON [1] documents linking the Galaxy tools,which are expressed as XML documents. Swift scripts areexpressed as C-like programs. A two-way translator willenable the workflows to be expressed in either format thusbeing able to run under two environments. Additionally, thefact that Swift scripts are translated into intermediate XML representation makes this process easier in that no effort isrequired to write new modules in either of the workflow engines.

3.3 Scheme 3: Data splitting within tool




Figure 4: Swift wrapper to Galaxy tool.

Input datafrom the Galaxy tool input is split, the base tool isrun under a Swift foreach on each fragment, and theresults are merged and presented as the Galaxy tooloutput.



Figure 4 shows a scheme of running Galaxy tools wrapped into a Swift foreach parallel loop. The scheme enables running ordinary sequential tools that perform homogeneous processing on a large dataset to run in parallel on a split dataset. The results of individual operations then get merged and emerge as output for consumption downstream. This scheme resembles the popular MapReduce computational paradigm. The implementation of this scheme is straightforward thanks to Swift's indexed arrays that map to files.

The splits and merge on data can be organized into arrays that map to file fragments and can be processed in the Swift script as variables. For this tool to be successfully applied, the operations on split data must be associative.

4. EXPERIENCE

We describe our experience in running applications encoded as Swift scripts running on large-scale remote resources from within the Galaxy environment. In our work to date, we have run Swift-Galaxy applications on two institutional clusters and one XSEDE system: the University of Chicago Midway and UC3 clusters and the XSEDE Stampede cluster. RCC Midway (rcc.uchicago.edu) is the University of Chicago Research Computing Center cluster supporting university-wide high-end computational needs. The cluster has multiple resource partitioning dedicated to specialized computing such as HPC, HTC, and GPU computing and runs a SLURM batch queue scheduler. The UC3 cluster (uc3.uchicago.edu) at the University of Chicago is an open computing framework for connecting users to shared distributed high-throughput computing resources, both on and off campus. The cluster runs Condor DRM capable of flexibly extending compute resources to the OSG environment. XSEDE (www.xsede.org) is an NSF-funded, national cyberinfrastructure comprising multiple large-scale computation systems on sites across the United States. Stampede is one of the supercomputing systems offered by XSEDE. Stampede runs the SLURM scheduler for submitting user jobs.

Figure 5 shows a screen capture of an application workspace interface. The drop-down dialog allows users to select resources for a run. The rest of the text dialog provides an interface to enter application-specific values. Parameter values provided through this interface are passed to Swift via the Galaxy tool runner. Galaxy runs the tool that in turn calls a Swift wrapper that in turn runs the application and relinquishes control back to Galaxy after the execution is completed. Figure 6 shows a screen capture of a post-execution view for a Swift tool. The plot shows the number of parallel active jobs during the run of a 2000-task application on Midway. More than 120 active parallel jobs run at once over a period of 55 seconds to complete the 2,000 tasks. On the right pane is the history of this run captured by Galaxy that shows the standard output and an HTML page that has links to the results (not shown in figure) of the execution. A brief description of applications enabled and ongoing work with the Swift-Galaxy interface follows.

Power Grid Analysis.

Large-scale analysis of power grid quantities is emerging as an important engineering application of national interest. Applications have many stakeholders, including the government, power companies, ISOs, and domestic and industrial consumers. In this work we enable an application concerning inference analysis of stochastic unit commitment models.

The application involves hierarchical stochastic computation over a range of samples and batches. The code in Listing 1 shows a Swift script representation of the core application logic. A three-level-deep nested foreach loop generates about 100K concurrent tasks for a sample size of 5 and a batch comprising sizes (10, 20, and 30). Currently we implement the application under Scheme 1 described above running on institutional and XSEDE resources.

```

1 foreach S, idxs in nS{
2   o = gensample( wind_data );
3   foreach B, idxb in [10:30:10]{
4     foreach k in [0: B]{
5       o = gensample (S, wind_data );
6       obj_out_l[ idxs ][ idxb ][k] = ampl_app_L ( params );
7       o = gensample (S, wind_data );
8       obj_out_u[ idxs ][ idxb ][k] = ampl_app_U ( params );
9     } } }

```

Listing 1: Swift code for power grid analysis application. Climate Models.

The Decision Support System for Agrotechnology Transfer (DSSAT) application analyzes the effects of climate change on agricultural production [4]. Projections of crop yields at regional scales are carried out by running simulation ensemble studies on available datasets of land cover, soil, weather, management, and climate. The computational framework starts with the DSSAT crop systems model and evaluates this model in parallel. One study involves analysis of climate impact for a single crop (maize) across the conterminous USA (120K cells) with daily weather data and climate model output spanning 120 years (1981-2100) and 16 different configurations of fertilizer, irrigation, and cultivar choice.



Currently, the climate model runs as a single, generic Galaxytool wrapped via a shell script and callable by Galaxy as atool.

Genomics.

A large section of the Galaxy community is geared towardgenomics data analysis, mainly next-generation sequencingdata (NGS). Many tools working on all aspects of genomicsanalysis are currently used, and many new tools are beingdeveloped. The analysis of NGS data can be partitioned into a series of steps common among the multipleNGS applications (i.e., transcriptome, whole-exome, wholegenome). One of the most popular applications for accuratevariant calling of exome data is the Genome AnalysisToolkit (GATK) from the Broad Institute [2], which providesa best practices pipeline with optimal parameters ineach step of the process. The pipeline includes running anumber of GATK applications that can be computationallydemanding and in some cases can be parallelized. As anongoing development, a Swift-enabled GATK best practicespipeline is being integrated into Galaxy. We are currentlyporting the application using a combination of Scheme 1and Scheme 3 described in x3.The Unified Genotyper (UG) tool analyzes multiple chromosomes.

By splitting the analysis by chromosome, we intendto run foreach loop over chromosomes. The workflowconsists of 123 input human genome files containing 23 chromosomeseach. Thus UG can be run 23 times for individualchromosomes.

5. RELATED WORK

In this section, we briefly present our perspectives on relatedwork. Many scientific gateways and portals are in use globally.Historically, problem-solving environments (PSEs) canbe considered as ancestors of portal-based science gateways.With the advent of fast networks and sophisticated web programmingtechnologies, PSEs have evolved into portals.The driving force behind many separate developments canbe attributed to user communities bonded by a scientificdomain or adapted to a particular set of tools.

The GPSI [7] portal is a prototype portal using Swift as itsworkflow engine. A Python-driven web environment allows users to submit their Swift scripts to distributed resourcesand to manage their data from within the web-based environment.Integrationof the Vbrowser portal with the MOTEUR workflowengine [6] is another example of similar work gearedtoward combining network-based portals with a task orchestrationengine. The target execution infrastructure in thiscase is the European Grid environment, and the target communityis the medical imaging community. Nanohub [5] isanother recent portal focused on the nanotechnology community.

6. CONCLUSIONS AND FUTURE WORK

In this paper we motivate the benefits of integrating portalbasedexecution frameworks with a large-scale parallel programmingframework. We demonstrate the integration withGalaxy as a representative of the former and Swift of thelatter.

The following points summarize the features of Swift integrationwith Galaxy:

- A generic Galaxy tool with the ability to run arbitrarySwift scripts from within a Galaxy interface.
- Static application tools for special-purpose executionwith flexibility of user-provided application parameters,and execution site.
- Capability to arbitrarily select remote execution sitesfrom a drop-down list of the tool's visual representation.
- Post execution analysis of runs and plots of job completionrate and parallel tasks.
- Capability to bring results into the Galaxy dataset systemfor reuse and publishing.
- The following activities summarize our future work.
- Extending and combining of the integration schemesinto a toolkit capable of encapsulating arbitrary Galaxytools into Swift.
- Finer integration of external Swift data and managedGalaxy dataset system.
- Deeper integration of the Swift engine into the Galaxyframework. This will enable Galaxy workflows to beorchestrated as a Swift execution, retaining the Galaxyvisual framework.

REFERENCES

1. D. Crockford. RFC 4627: The application/json mediatype for JavaScript Object Notation (JSON). 2006.
2. GATK Best Practices. [2] www.broadinstitute.org/gatk/guide/best-practices.
3. M.Hategan. J.Woznaik and K.Maheswari, Coasters uniform resource provisioning and access for scientific computing on clouds and grids.In Proc Utility and Cloud computing 2011.



4. K.Maheshwari, K.Birman, J.Woznaik and D.V.Zandt. Evaluating cloud computing techniques for smart power grid design using parallel scripting. In cluster, cloud and grid computing (CCGrid), 2013 13th IEEE/ACM International Symposium on, 2013.
5. M. Stonebraker and J. Robertson, "Big data is 'buzzword du jour,' cs academics 'have the best job'," Commun. ACM, vol. 56, no. 9, pp. 10–11, 2013.
6. M. Stonebraker and J. Hong, "Researchers' big data crisis; understanding design and functionality," Commun.ACM, vol. 55, no. 2, pp. 10–11, 2012.
7. R. Ramakrishnan, "Cap and cloud data management," IEEE Computer, vol. 45, no. 2, pp. 43–49, 2012.
8. J. Shen, J. Fang, H. Sips, and A. L. Varbanescu, "Performance Traps in OpenCL for CPUs," in PDP 2013, February 2013, pp. 38–45.
9. J. Shen, A. L. Varbanescu, H. J. Sips, M. Arntzen, and D. G. Simons, "Glinda: a framework for accelerating imbalanced applications on heterogeneous platforms – in conference computing frontiers – 2013 Page 14.

BIOGRAPHIES



Dr. M. Suresh Babu, received Post graduation from Osmania University, M.Phil (Computer Science) from Bharathiar University and PhD (Computer Science & Technology) from Sri Krishnadevaraya University. He is currently working as Professor in Department of Computer Science at NallaMalla Reddy Engineering College. He has published more than 80 publications in the areas of Database, Semantic, Data Mining, Cloud Computing and Network Security. He is associated with various Professional bodies like ISTE, ISCA, IAEngg, IACSIT, UACEE, SOULS, CSTA, SDIWC and EAI. He is editorial board member, reviewer for more than 10 International journals.



A. Vijaya Kumar completed Post Graduation from HCU and M.Tech from Sathyabama University in 2007. His Research interests include Data Mining, MANETs, and cloud computing. Presently he is working as an Associate Professor in the IT Department in NallaMalla Reddy Engineering College, Hyderabad. He has published research papers in various National, International conferences and Journals. He is a life member of various professional societies like ISTE, IAEngg, CSTA, UACEE, SDIWC, IACSIT.