Workflow Technologies and CyberInfrastructure: Laying the Foundations for Science

Ewa Deelman
USC Information Science Institute
deelman@isi.edu

Today, many advances in scientific disciplines are achieved not by spending hours in a laboratory or by examining the skies through a microscope atop of a mountain, but rather through sophisticated analyses conducted on large volumes of data. The analyses are frequently composed of several application components, each often designed and tuned by a different researcher. The question that arises is: **How to leverage the knowledge and application components developed by the members of a collaboration to solve ever more complex problems?** Recently, **scientific workflows** have emerged as a means of combining individual application components into large-scale analysis by defining the interactions between the components and the data that they rely on.

Scientific workflows provide a systematic way to capture scientific methodology by providing a detailed trace (provenance) of how the results were obtained. Additionally, workflows are collaboratively designed, assembled, validated, and analyzed. Workflows should be shared in the same manner that data collections and compute resources are shared today among communities. Scientific workflows are a very promising paradigm for large-scale and large-scope scientific inquiry. The scale refers to the complexity and the number of computational steps, whereas the scope refers to the integration of diverse models, phenomena, and disciplines. The scale of the workflows often necessitates that substantial computational and data resources be used to generate the required results.

There are many high-performance resources available to scientists today, for example those of the NSF **TeraGrid** (www.teragrid.org) or the NSF and **DOE Open Science Grid** (www.opensciencegrid.org). Middleware available through the NSF **National Middleware Initiative (NMI)** provides basic functionality of resource discovery, remote job submission, data management, all in a secure fashion. However, there is gap between the specification of the analysis that a scientist is comfortable with, which is often very high-level and the task specification required by the Grid resources in order to conduct the computations. Therefore the following question needs to be answered: **How to bridge the gap between the analysis specifications that are meaningful to the scientists and the resources presented by the CyberInfrastructure?** One way to approach this problem is to provide **mapping technologies** that can interpret the high-level descriptions and map them onto the available computational software, relying on the available software and the available data.

Pegasus, which stands for Planning for Execution for Grids is a software developed and used as part of several NSF ITR projects (GriPhyN, NVO, and SCEC). Pegasus bridges the scientific domain and the execution environment by automatically mapping the high-level workflow descriptions onto distributed resources such as the TeraGrid, the Open Science Grid, and other resources, for example those managed by Condor. Pegasus enables scientists to construct workflows in abstract terms without worrying about the details of the underlying CyberInfrastructure. The software is used day-to-day to map complex, large-scale scientific workflows with thousands of tasks processing TeraBytes of data onto the Grid. Pegasus is used in a variety of applications ranging from astronomy, biology, earthquake science, gravitational-wave physics and others. Pegasus uses Condor's DAGMan for workflow execution, and Globus and Condor to provide the middleware for distributed environments.

Sometimes the workflows as structured by scientists are not tuned for performance. Additionally, given that at the time of the workflow generation, the eventual execution resources are not known, it is impossible to optimize the runtime of the overall workflow. Since Pegasus discovers the available resources and their characteristics and queries for the location of the data (potentially replicated in the environment), it can improve the performance of the application through:

- Data reuse to avoid duplicate computations and provide reliability
- Workflow restructuring to improve resource allocation
- Automated task and data transfer scheduling to improve overall workflow runtime

Pegasus also provides reliability through dynamic workflow remapping when failures during execution are detected.

Scientific workflows are used in astronomy, and in particular in the **Montage** application which delivers science-grade mosaics of the sky. Workflow technologies were used to transform a single-processor Montage code into a complex workflow and parallelized computations to process larger-scale images. Montage workflows mapped by Pegasus to the NSF CyberInfrastructure are characterized by tens of thousands of executable tasks and the processing of thousands of images. Recently, scientists have used Montage to **verify a bar in the spiral galaxy M31,** shown in the image on the right (Beaton et al. *Ap J Lett* in press). Eleven major projects and surveys worldwide, such as the Spitzer Space Telescope Legacy teams have integrated Montage into their pipelines and processing environments to generate science and browse products for dissemination to the astronomy community.

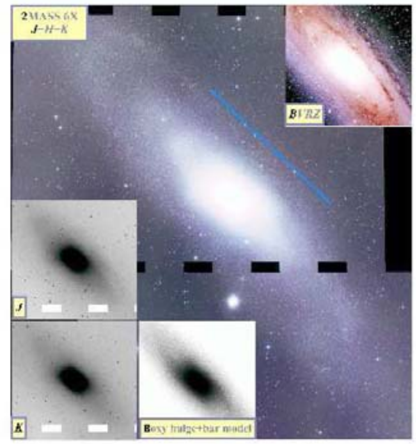

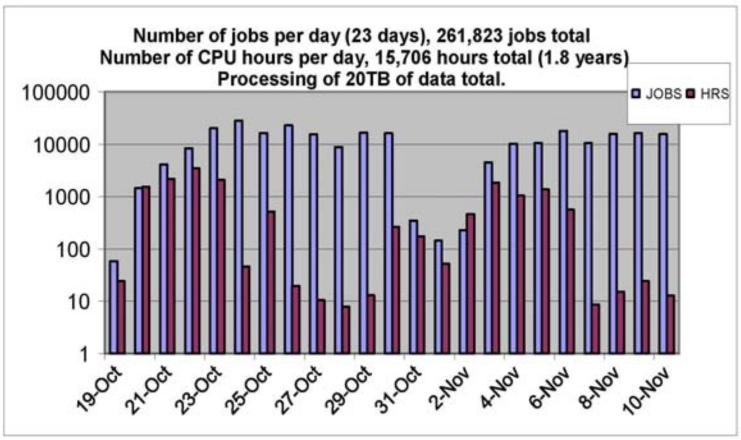

The **Southern California Earthquake Center** (SCEC) uses workflow technologies to produce more accurate seismic hazard maps. These maps, generated as part of the SCEC CyberShake project, indicate the maximum amount of shaking expected at a particular geographic location over a certain period of time, often 50 years. The hazard maps are used by civil engineers to determine building design tolerances. **Pegasus** maps the CyberShake workflows onto SCEC and NSF CyberInfrastructure resources. The figure on the left shows the results of running CyberShake on the TeraGrid in the fall of 2005. The **workflows ran over a period of 23 days and processed 20TB of data using 1.8 CPU Years**. The total number of tasks in all workflows was 261,823. **CyberShake delivers new insights** into how rupture directivity and sedimentary basin effects contribute to the shaking experienced at different geographic locations. As a result **more accurate hazard maps can be created.**

Pegasus is also used in the **Laser Gravitational Wave Observatory** (LIGO) project to map binary inspiral analysis workflows onto the NSF and DOE Open Science Grid. A month of LIGO data requires many **thousands of jobs, running for days on hundreds of CPUs**.

Workflows can have great benefits in the education of students and young researchers. Experts in a particular domain can set up the end-to-end workflows, set the values of the parameters and choose representative data sets. Students can run the sophisticated workflows using training data sets and learn to interpret the results. Young researchers can benefit from the expertise of senior scientists and run the workflows on the data of interest to them and discover new phenomena. Workflows can also be **shared between collaborations, enabling cross-organizational result replication and validation.**

Relevant Links: Pegasus: pegasus.isi.edu, Montage: montage.ipac.caltech.edu, SCEC: www.scec.org, LIGO: www.ligo.caltech.edu/
TeraGrid:www.teragrid.org, Open Science Grid: www.opensciencegrid.org, Globus: www.globus.org, Condor: www.cs.wisc.edu/condor/,
NSF Workshop on Challenges of Scientific Workflows: vtcpc.isi.edu/wiki/, E. Deelman and Y. Gil (chairs)
Workflows for e-Science, Taylor, I.J.; Deelman, E.; Gannon, D.B.; Shields, M. (Eds.), Dec. 2006, *to appear*

Acknowledgments: The Pegasus team consists of Ewa Deelman, Gaurang Mehta, Mei-Hui Su, and Karan Vahi (ISI). Thanks to all collaborators: scientific workflow issues: Yolanda Gil; Montage: Bruce Berriman, John Good, Dan Katz, and Joe Jacobs; SCEC: Tom Jordan, Robert Graves, Phil Maechling, David Okaya, Li Zhao; LIGO: Kent Blackburn, Duncan Brown, and David Meyers. Thanks to the National Science Foundation for the support of this work.