

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

A Multi-Dimensional Classification Model for Scientific Workflow Characteristics

Permalink

<https://escholarship.org/uc/item/8fj0s88m>

Author

Ramakrishnan, Lavanya

Publication Date

2010-06-03

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

A Multi-Dimensional Classification Model for Scientific Workflow Characteristics

Lavanya Ramakrishnan
Lawrence Berkeley National Lab
Berkeley, CA
lramakrishnan@lbl.gov

Beth Plale
Indiana University
Bloomington, IN
plale@cs.indiana.edu

ABSTRACT

Workflows have been used to model repeatable tasks or operations in manufacturing, business process, and software. In recent years, workflows are increasingly used for orchestration of science discovery tasks that use distributed resources and web services environments through resource models such as grid and cloud computing. Workflows have disparate requirements and constraints that affects how they might be managed in distributed environments. In this paper, we present a multi-dimensional classification model illustrated by workflow examples obtained through a survey of scientists from different domains including bioinformatics and biomedical, weather and ocean modeling, astronomy detailing their data and computational requirements. The survey results and classification model contribute to the high level understanding of scientific workflows.

Categories and Subject Descriptors

A.1 [Introductory]: Survey

General Terms

Design

Keywords

Cloud, Grid, Scientific Workflows

1. INTRODUCTION

Workflows and workflow concepts have been used to model a repeatable sequence of tasks or operations in different domains including the scheduling of manufacturing operations, business process management [Taylor et al. 2006], inventory management, etc. The advent of internet and web services has seen the adoption of workflows as an integral component of cyberinfrastructure for scientific experiments [Deelman and Gil 2006, Atkins 2002]. The availability of distributed resources through grid and cloud computing models has enabled users to share data and resources using workflow tools and other user interfaces such as portals.

Workflow tools are in use today in various cyberinfrastructure projects to satisfy the needs of a specific science problem [Deelman et al. 2003, Altintas et al. 2004, Condor DAGMan, Taylor et al. 2004, Oinn et al. 2006]. This has resulted in innovative solutions for workflow planning, resource management, provenance generation, etc.

Each scientific workflow has different resource requirements and constraints associated with them. Workflows

can vary in their size, resource requirements, constraints, amount of user intervention, etc. For example, application workflows with stringent timeliness constraints such as for weather prediction or economic forecasting are now increasingly run in distributed resource environments. These workflows require a clear understanding of their workflow requirements to manage user's deadline constraint with the variability of underlying resources. Additionally the parallelism of the workflow and its storage needs often affect design choices. However, there is a limited description and understanding of usage, performance and characteristics of scientific workflows. This is a major road block to reuse of existing technologies and techniques and innovation of new workflow approaches.

In this paper, we present a qualitative classification model of workflow characteristics. These characteristics help classify or "bin" workflow types enabling broader engagement and applicability of solutions. We discuss workflow examples from different domains: bioinformatics and biomedicine, weather and ocean modeling, astronomy, etc and cast them in the context of our model. The workflow examples have been obtained through surveying domain scientists and computer scientists who composed and/or run these workflows. Each of these workflows has been modeled using one of several workflow tools and/or through scripts. For each workflow we specify the running time of applications and input and output data sizes associated with each task node. Running time of applications and data sizes for a workflow depend on a number of factors including user inputs, specific resource characteristics and run-time resource availability variations [Kramer and Ryan 2003]. Thus our numbers are approximate estimates for typical input data sets that are representative of the general characteristics of the workflow.

The remainder of the paper is organized as follows. Section 2 discusses related work. Section 3 presents the classification model that is drawn from the survey of workflows presented in Section 4. In section 5 we tie the model and survey together with a taxonomy. The paper concludes in Section 6.

2. RELATED WORK

New workflow tools have been developed to represent and run scientific processes in a distributed grid environment. Workflow tools such as Kepler [Altintas et al. 2004, Ludscher et al. 2005], Taverna [Oinn et al. 2006], Pegasus [Deelman et al. 2003, Deelman et al. 2003], Triana [Churches et al. 2006] allow users to compose tasks (i.e., analysis, modeling, synthesis, mapreduce, and data-driven) and services into a

logical sequence. These tools often are developed in the context of one or more specific application domains and have various features to allow users to compose and interact with workflows through a graphical interface, provides seamless access to distributed data, resources and web services. Yu and Buyya provide a taxonomy for scientific workflow systems that classify systems based on four elements of a grid workflow systems - a) workflow design, b) workflow scheduling, c) fault tolerance and d) data movement [Yu and Buyya 2005]. Thain et al. characterize a collection of scientific batch-pipelined workloads on processing, memory, and I/O demands, and the sharing characteristics of the workloads. However there has been no detailed study of characteristics of complex scientific workflows and representing a qualitative classification model to capture the features of a workflow.

3. CLASSIFICATION MODEL

Workflows vary significantly in their characteristics and their computational and data requirements. Workflows vary in their size, structure and usage. A number of the bioinformatics workflows often have tasks that are based on querying large databases in order of minutes for the task execution. In other cases we see each of the tasks of the workflow require computation time on the order of hours or days on multiple processors. In some cases sub-parts of the workflow might also present different characteristics. It is critical to understand these characteristics to effectively manage these workflows. We present a multi-dimensional workflow characterization model that considers the following a) Size b) Resource Usage c) Structural Pattern d) Data Pattern and e) Usage Scenarios.

3.1 Size

The size of the workflow is an important characteristic since a workflow might have from a small number of tasks to thousands of tasks. The size of the workflows that are deployed today in most production environments are relatively small. The largest workflows in our survey contain about a couple of hundred independent tasks. The Avian Flu (Figure 11) and PanSTARRS (Figures 13 and 14) workflows has over a thousand nodes but the computation at each node is expected to take only a few minutes to an hour. Scientists express a need to run larger sized workflows but are often limited by available resources or workflow tool features that might be needed to support such large-scale workflows.

In addition to the total number of tasks in a workflow it is also important to consider the width and length of the workflows. The width of the workflow (i.e. maximum number of parallel branches) determines the concurrency possible and the length of the workflow characterizes the makespan (or turnaround time) of the workflow. We observe that in our workflow examples, the larger sized workflows such as the Motif workflow (Figure 8) and the astronomy workflows (Figures 13 and 14) the width of the workflow is significantly larger than the length of the workflow.

Thus we define three properties in the size dimension -

- **Total Number of Tasks.** defines the total number of tasks in the workflow.
- **Number of Parallel Tasks.** defines the maximum number of parallel tasks in any part of the workflow.

- **Longest Chain.** defines the number of tasks in the longest chain of the workflow.

3.2 Resource Usage

In addition to the structure and pattern of a workflow it is important to understand the computational requirements. In the presented workflow examples we observe that computational time required by the workflows can vary from a few seconds to several days. A number of the bioinformatics workflows depend on querying large databases and have small compute times. Some examples include the Glimmer workflow (Figure 6), Gene2Life (Figure 7), caDSR (Figure 12). Similarly the initial parts of the LEAD forecast workflow (Figures 1 and 2) and the LEAD data mining workflows (Figure 3) have small computational load. A number of the workflows including the forecasting parts of the LEAD workflow, Pan-STARRS workflows (Figures 13 and 14), SCOOP (Figure 4), SNS (Figure 15), Motif (Figure 8), NCFS (Figure 5) have medium to large sized compute requirements.

- **Max task processor width.** is the maximum concurrent number of processors required by the workflow.
- **Computation time.** is the total computational time required by the workflow.
- **Data Sizes.** is the data size of the workflow inputs, outputs and intermediate data products.

3.3 Structural Pattern

Each workflow might include one or more patterns. Our goal is to capture the dominant pattern seen in the workflow. The workflows that we surveyed depict the basic control flow patterns such as sequence, parallel split, synchronization [van der Aalst et al. 2003]. The parallel split-synchronization pattern has similarities to the map-reduce programming paradigm. A number of workflows divide the work units into distinct work units and the results are then combined - Motif workflow (Figure 8), Pan-STARRS workflows (Figures 13 and 14). Thus we classify our workflows into the following patterns:

- **Sequential.** consists of tasks that follow one after the other.
- **Parallel.** consists of multiple tasks that can be run at the same time.
- **Parallel-split.** one task's output feeds to multiple tasks.
- **Parallel-merge.** multiple tasks merge into one task.
- **Parallel-merge-split.** both parallel-merge and parallel-split.
- **Mesh.** task dependencies are interleaved.

3.4 Data Pattern

The workflows are associated with different types of data including input data, backend databases, intermediate data products, output data products. A large number of the bioinformatics applications often have small input and small data products but often rely on huge backend databases that are queried as part of task execution. These workflows require that the databases be pre-installed on various sites and

resource selection is often based on selecting the resources where the data might be available. Workflows such as LEAD (Figures 1 and 2), SCOOP (Figure 4), NCFS (Figure 5) and Pan-STARRS workflows (Figures 13 and 14) have fairly large sized input, intermediate and output data products. The Glimmer workflow (Figure 6) has similar sized input and output data products but its intermediate data products are smaller. In today’s production environments workflows often compress data products to reduce transfer times through intermediate scripts etc. When scheduling workflows on resources, a number of data issues need to be considered including the availability of the required data as well as the data transfer time of both input and output products. We classify the workflows as

- **Data reduction.** where the output data is smaller than the input data of the workflows.
- **Data production.** where the output data is larger than the input data of the workflow.
- **Data processing.** where input data is processed but data sizes do not change dramatically.

3.5 Usage scenarios

It is often important to understand the use case scenarios for the workflows. Workflows are used in a number of different scenarios - a new workflow might be initiated in response to dynamic data or a number of workflows might be launched as part of an educational workshop. In addition, the user might want to *specify constraints* to adjust the number of workflows to run based on resource availability [Ramakrishnan et al. 2007].

Interactive Workflows Scientific explorations often require a “human-in-the-loop” as part of the workflow. The typical mode of usage of science cyberinfrastructure is where a user logs into the portal and launches a workflow for some analysis. The user selects a pre-composed workflow and supplies the necessary data for the run. The user might also want the ability to pause the workflow at the occurrence of a predefined event, inspect intermediate data and make changes during workflow execution.

Event-driven Workflows A number of scientific workflows get triggered by newly arriving data. Multiple dynamic events and their scale might need priorities between users for appropriate allocation of limited available resources. Resources must be allocated to meet deadlines. Additionally, to ensure successful completion of tasks, we might need to replicate some of the workflow tasks for increased fault tolerance. It is possible that with advance notice of upcoming weather events, we might want to anticipate the need for resources and try to procure them in advance. The weather forecasting, storm surge modeling (Figure 4), flood-plain mapping (Figure 5) and the astronomy workflows (Figures 13 and 14) are launched with the arrival of data.

User Constrained Workflows An advanced user might want to provide a set of constraints (e.g. time deadline or budget) on a workflow. Scientific processes such as weather prediction, financial forecasting have a number of parameters and computing an exact result is often impossible. To improve confidence in the result, it is often necessary to run a minimal number of the workflows. There is a need to run multiple workflows (i.e. *workflow sets*) that need to be scheduled together. Thus for workflow sets, users specify

that they minimally require M out of N workflows to complete by the deadline.

4. WORKFLOW SURVEY EXAMPLES

Table 1 presents an overview of the workflow survey. The survey includes workflows from diverse scientific domains and cyberinfrastructure projects. In the following sections, we identify the project from which the workflow is drawn, the workflow and usage model as available at the time of the survey. For each of the workflows, we also provide a DAG representation of the workflow annotated with computation and data sizes. Sections 4.1 describes the weather and ocean modeling workflows and Sections 4.2 describes the bioinformatics and biomedicine workflows. Sections 4.3 describe the astronomy and neutron science.

4.1 Weather and Ocean Modeling

In the last few years the world has witnessed a number of severe natural disasters such as hurricanes, tornadoes, floods, etc. The models used to study weather and ocean phenomenon use considerable and diverse real-time observational data, static data, and parameters that are varied to study the possible scenarios for prediction. In addition the models must be run in a timely manner and information disseminated to groups such as disaster response agencies. This creates the need for *large scale modeling* in the areas of meteorology and ocean sciences, coupled with an *integrated environment* for analysis, prediction and information dissemination.

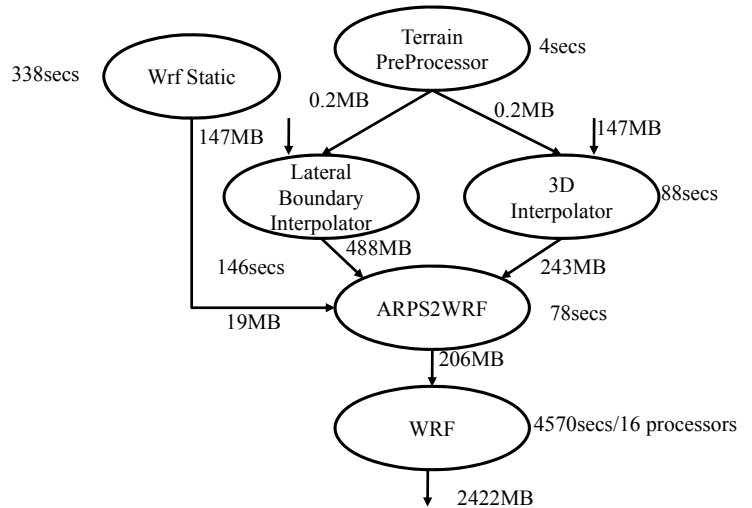


Figure 1: LEAD North American Mesoscale (NAM) initialized forecast workflow. The workflow processes terrain and observation data to produce weather forecasts.

4.1.1 Mesoscale Meteorology

The Linked Environments for Atmospheric Discovery (LEAD) [Droegemeier et al. 2005] is a cyberinfrastructure project in support of dynamic and adaptive response to severe weather. A LEAD workflow is constrained by execution time and accuracy due to weather prediction deadlines. The typical inputs to a workflow of this type are real time observational

Domain	Project	Website	Tool
Weather and Ocean Modeling	Linked Environments for Atmospheric Discovery (LEAD) TeraGrid Science Gateway	http://portal.lead.project.org	xbaya, GPEL, Apache ODE
	Southeastern Coastal Ocean Observing and Prediction Program (SCOOP)	http://www.renci.org/focusareas/disaster/scoop.php	[Scripts]
	North Carolina Floodplain Mapping Program		[Scripts]
Bioinformatics and Biomedical	North Carolina Biportal, TeraGrid Biportal Science Gateway	http://www.renci.org/focusareas/biosciences/motif.php	Taverna
	MotifNetwork	http://www.motifnetwork.org/	Taverna
	National Biomedical Computation Resource (NBCR), Avian Flu Grid, Pacific Rim Application and Grid Middleware Assembly	http://nbcrc.sdsc.edu/mozdev.org http://www.pragma-grid.net/ http://avianflugrid.pragma-grid.net/ http://mgltools.scripps.edu/	Kepler, Gemstone, [Scripts] and Vision
	cancer Biomedical Informatics Grid (caBIG)	http://www.cagrid.org/	Taverna
Astronomy	Pan-STARRS	http://pan-starrs.ifa.hawaii.edu/public/ , http://www.ps1sc.org/	
Neutron Science	Spallation Neutron Source (SNS), Neutron Science TeraGrid Gateway (NSTG)	http://neutrons.ornl.gov/	

Table 1: Workflow Survey Overview. Survey is of workflows used in domains from meteorology and ocean modeling, bioinformatics and biomedical workflows, astronomy and neutron science

data and static data such as terrain information [Droegemeier et al. 2005, Plale et al. 2006] both of which are pre-processed and then used as input to one or more ensemble of weather models. The model outputs are post-processed by a data mining component that determines whether some ensemble set members must be repeated to realize statistical bounds on prediction uncertainty. Figures 1, 2 and 3 show the workflows available through the LEAD portal and include weather forecasting and data mining workflows [Li et al. 2008]. Each workflow task is annotated with computation time and the edges of the directed acyclic graph (DAG) are annotated with file sizes. The weather forecasting workflows are largely similar and vary only in their preprocessing or initialization step. While the data mining workflow can be run separately today, it can trigger forecast workflows and/or steer remote radars for additional localized data in regions of interest [Plale et al. 2006].

4.1.2 Storm surge modeling

Southeastern Universities Research Association’s (SURA) Southeastern Coastal Ocean Observing and Prediction (SCOOP) program is creating an open-access grid environment for the southeastern coastal zone to help integrate regional coastal observing and modeling systems [SCOOP Website, Ramakrishnan et al. 2006]. Storm surge modeling requires assembling input meteorological and other data sets, running models, processing the

output and distributing the resulting information. In terms of modes of operation, most meteorological and ocean models can be run in hindcast mode, as an after fact of a major storm or hurricane, for post-analysis or risk assessment, or in forecast mode for prediction to guide evacuation or operational decisions [Ramakrishnan et al. 2006]. The forecast mode is driven by real-time data streams while the hindcast mode is initiated by a user. Often it is necessary to run the model with different forcing conditions to analyze forecast accuracy. This results in a large number of parallel model runs, creating an ensemble of forecasts. Figure 4 shows a five member ensemble run of the tidal and storm-surge ADCIRC [Luettich et al. 1992] model. For increased accuracy of forecast the number of concurrent model runs might be increased. ADCIRC is a finite element model that is parallelized and using the MPI message passing model. The workflow has a predominately parallel structure and the results are merged in the final step.

The SCOOP ADCIRC workflows are launched according to the typical six hour synoptic forecast cycle used by the National Weather Service and the National Centers for Environmental Prediction (NCEP). NCEP computes an atmospheric analysis and forecast four times per day at six hour intervals. Each of the member runs i.e. each branch of the workflow gets triggered when wind files arrive through Local Data Manager (LDM) [Unidata Local Data Manager (LDM)], an event-driven data distribution system that se-

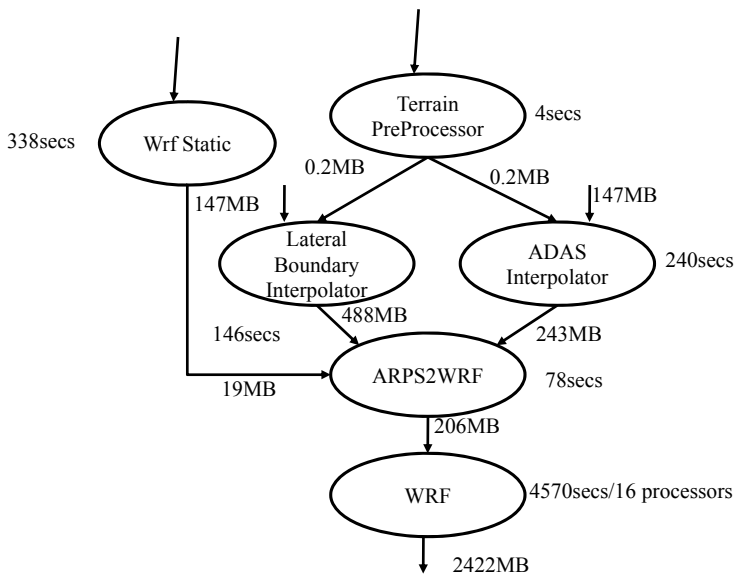


Figure 2: LEAD ARPS Data Analysis System (ADAS) initialized forecast workflow. The workflow processes terrain and observation data to produce complex 4D assimilated volume that initializes a weather forecast model.

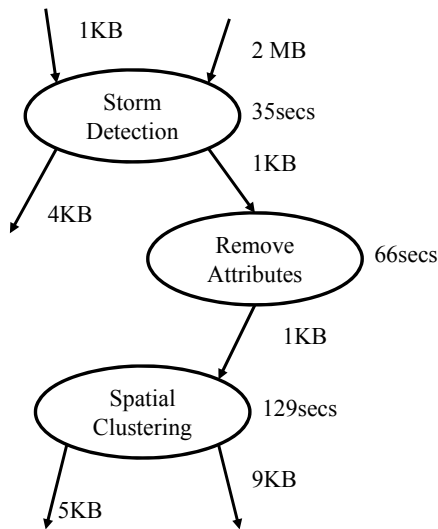


Figure 3: LEAD Data Mining workflow. The workflow processes Doppler radar or model generated forecast data to identify regions where weather phenomenon might exist in the future.

lects, captures, manages and distributes meteorological data products. The outputs from the individual runs are synthesized to generate the workflow output that is then distributed through LDM.

In the system today each arriving ensemble member is handled separately through a set of scripts and Java code [Ramakrishnan et al. 2006]. The resource selection approach [Lander et al. 2008] makes a real-time decision for each model run and uses knowledge of scheduled runs to load-balance across available systems. However this approach does not have any means of guaranteeing desired QoS in terms of completion time.

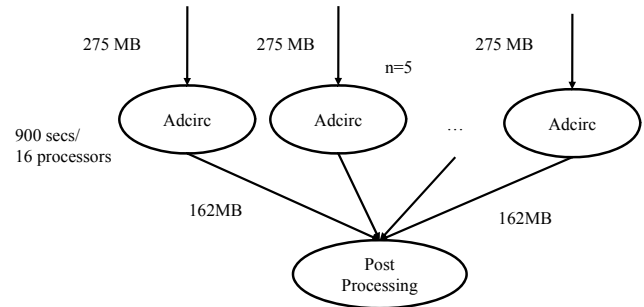


Figure 4: SCOOP workflow. The arriving wind data triggers the ADCIRC model that is used for storm-surge prediction during hurricane season.

4.1.3 Floodplain Mapping

The North Carolina Floodplain Mapping Program [North Carolina Floodplain Mapping Program, Blanton et al. 2008] is focused on developing accurate simulation of storm surges in the coastal areas of North Carolina. The deployed system today consists of a four-model system that consists of the Hurricane Boundary Layer (HBL) model for winds, Wave-Watch III and SWAN for ocean and near-shore wind waves, and ADCIRC for storm surge. The models require good coverage of the parameter space describing tropical storm characteristics in a given region for accurate flood plain mapping and analysis. Figure 5 shows the dynamic portion of the workflow. Forcing winds for the model runs are calculated by the Hurricane Boundary Layer(HBL) model that serve as inputs to the workflow. The HBL model is run on a local commodity Linux cluster. Computational and storage requirements for these workflows are fairly large requiring careful resource planning. An instance of this workflow is expected to run for over a day. The remainder of the workflow runs on a supercomputer.

4.2 Bioinformatics and Biomedical workflows

The last few years have seen large scale investments in cyberinfrastructure for bioinformatics and biomedical research. The infrastructure allows users to access databases and web services through workflow tools and/or portal environments. We surveyed three major projects in the United States - North Carolina Bioportal, cancer Biomedical Informatics Grid (caBIG), and National Biomedical Computational Resource (NBCR) to gain a better understanding of the workflows. Significant number of these workflows involve small computation but involve access to large-scale databases that

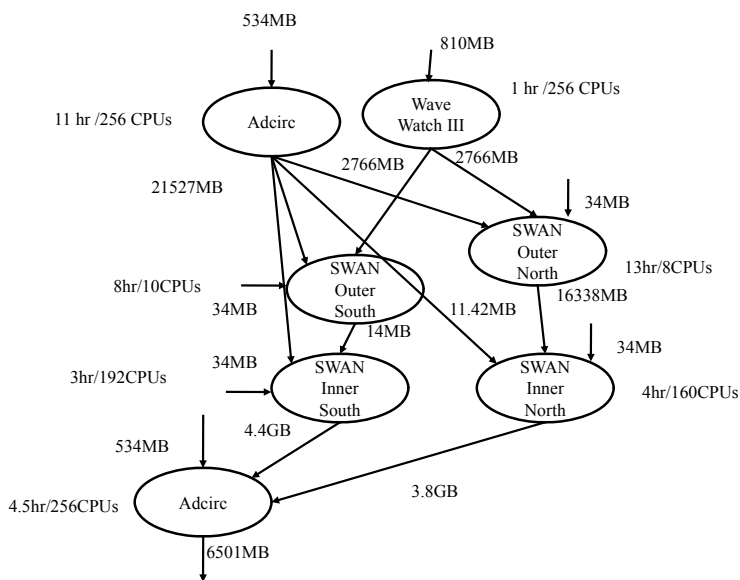


Figure 5: NCFS workflow. A multi-model workflow used to model the storm surges in the coastal areas of North Carolina.

must be preinstalled on available resources. While the typical use cases today have input data sizes in the order of megabytes, it is anticipated that in the future data sizes might be in the gigabytes.

4.2.1 Glimmer

The North Carolina Biportal and The TeraGrid Biportal Science Gateway [Ramakrishnan et al. 2006] provides access to about 140 bioinformatics applications and a number of databases. Researchers and educators use the applications interactively for correlation, exploratory genetic analysis, etc. The Glimmer workflow shown in Figure 6 is one such example workflow that is used to find genes in microbial DNA. The Glimmer workflow is sequential and light on both compute and data consumption.

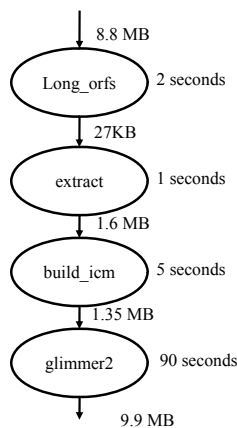


Figure 6: Glimmer workflow. A simple workflow used in educational context to find genes in microbial DNA.

4.2.2 Gene2Life

Let us consider the Gene2Life workflow used for molecular biology analysis through the North Carolina Biportal. This workflow takes an input DNA sequence, discovers genes that match the sequence. It globally aligns the results and attempts to correlate the results based on organism and function. Figure 7 depicts the steps of the workflow and the corresponding output at each stage. In this workflow the user provides a sequence that can be a nucleotide or an amino acid. The input sequence performs two parallel BLAST [Altschul et al. 1990] searches, against the nucleotide and protein databases respectively. The results of the searches are parsed to determine the number of identified sequences that satisfy the selection criteria. The outputs trigger the launch of ClustalW, a bioinformatics application that is used for the global alignment process to identify relationships. These outputs are then passed through parsimony programs for analysis. The two applications that may be available for such analysis are dnapars and protpars. In the last step of the workflow plots are generated to visualize the relationships, using an application called drawgram. This workflow has two parallel sequences.

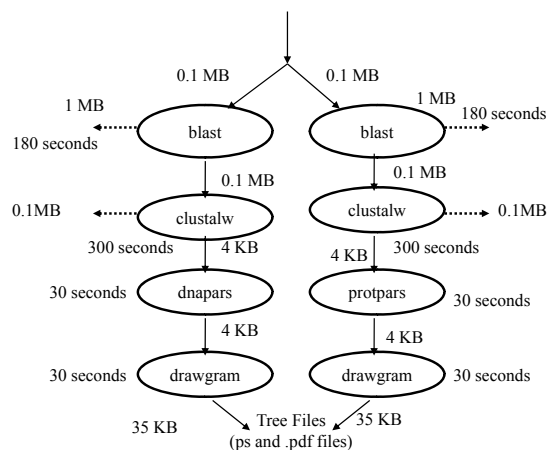


Figure 7: Gene2Life workflow. The workflow is used for molecular biology analysis of input sequences. The dotted arrows show the intermediate products from this workflow that are required by the user and/or might be used to drive other scientific processes.

4.2.3 Motif Network

The MotifNetwork project [Tilson et al. 2007, Tilson et al. 2007], a collaboration between RENCi and NCSA, is a software environment to provide access to domain analysis of genome sized collections of input sequences. The MotifNetwork workflow is computationally intensive. The first stage of the workflow assembles input data and processes the data that is then fed into InterProScan service. The concurrent executions of InterProScan are handled through Taverna and scripts. The results of the domain “scanning” step are passed to a parallelized MPI code for the determination of domain architectures. The motif workflow has a parallel split and merge paradigm where preprocessing spawns a set of parallel tasks that operate on subsets of the data. Finally, the results from the parallel tasks are merged and fed into the

multi-processor application.

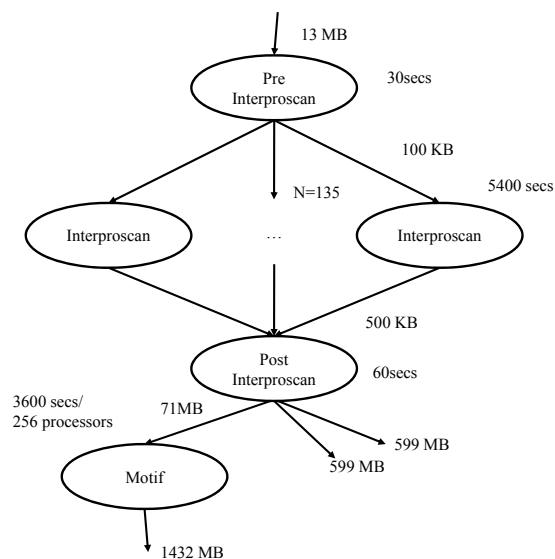


Figure 8: Motif workflow. A workflow used for motif/domain analysis of genome sized collections of input sequences.

4.2.4 MEME-MAST

The goal of National Biomedical Computation Resource(NBCR) is to facilitate biomedical research by harnessing advanced computational and information technologies. The MEME-MAST (Figure 9) workflow deployed using Kepler [Ludscher et al. 2005, Altintas et al. 2004] allows users to discover signals or motifs in DNA or protein sequences and then search the sequence databases for the recognized motifs. This is a simple workflow often used for demonstration purposes. The workflow is a sequential workflow similar to Glimmer.

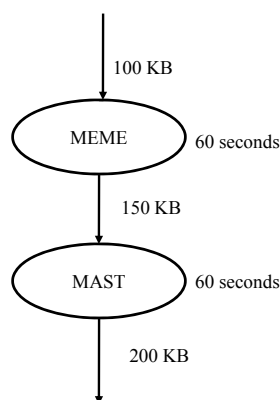


Figure 9: MEME-MAST workflow. A simple demonstration workflow used to discover signals in DNA sequences.

4.2.5 Molecular Sciences

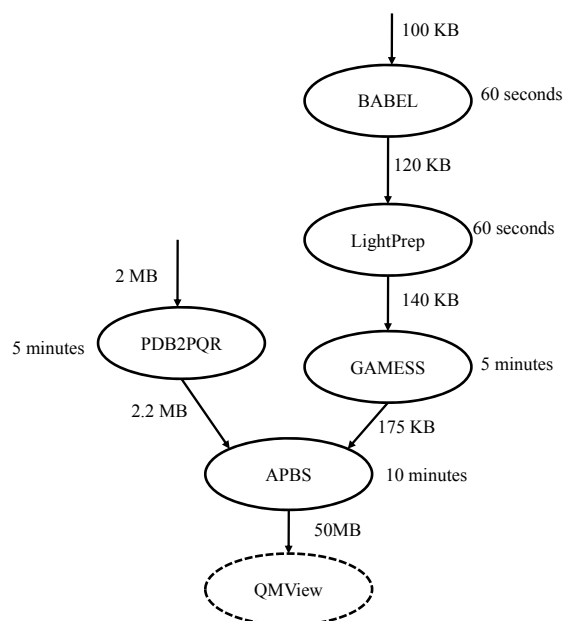


Figure 10: Molecular Sciences workflow is used in study of atomic structures of proteins and ligands.

An important step in the drug-design process is understanding the three-dimensional atomic structures of proteins and ligands. Gemstone, a client interface to a set of computational chemistry and biochemistry tools, provides the biomedical community access to a set of tools that allows users to analyze and visualize atomic structures. Figure 10 shows an example molecular science workflow. The workflow runs in an interactive mode where each step of the workflow is manually launched by the user once the previous workflow task has finished. The first few steps of the workflow involve downloading the desired protein and ligand from the Protein Data Bank (PDB) database and converting it to a desired format. Concurrent preprocessing is done on the ligand using the Babel and LigPrep services. Finally GAMESS and APBS are used to analyze the ligand and protein. The results are finally visualized using the QMView which is done as an offline process. First few steps have small data and small compute and finally produce megabytes of data.

4.2.6 Avian Flu

The Avian Flu Grid project is developing a global infrastructure for the study of Avian Flu as an infectious agent and as a pandemic threat. Figure 11 shows a workflow that is used in drug design to understand the mechanism of host selectivity and drug resistance. The workflow has a number of small preprocessing steps followed by a final step where upto 1000 parallel tasks are spawned. The data products from this workflow are small.

4.2.7 caDSR

The cancer Biomedical Informatics Grid(caBIG) is a virtual infrastructure that connects scientists with data and tools towards a federated cancer research environment. Figure 12 shows a workflow using the caDSR (Cancer Data Standards Repository) and EVS (Enterprise Vocabulary Ser-

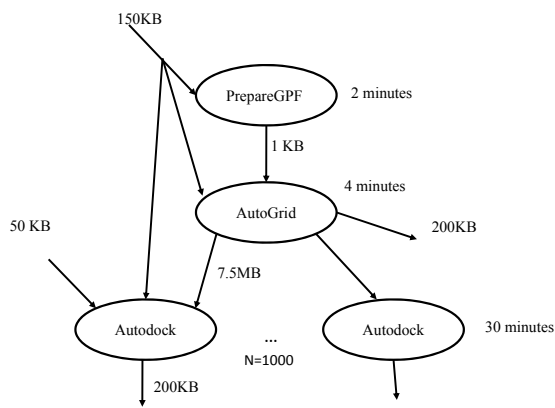


Figure 11: Avian Flu workflow. A workflow used in drug design to study the interaction of drugs with the environment has a large fan-out at the end.

vices) services [CaGrid Taverna Workflows] to find all the concepts related to a given context. The caDSR service is used to define and manage standardized metadata descriptors for cancer research data. EVS in turn facilitates terminology standardization across the biomedical community. This workflow is predominantly a query type workflow and the compute time is very small in the order of seconds.

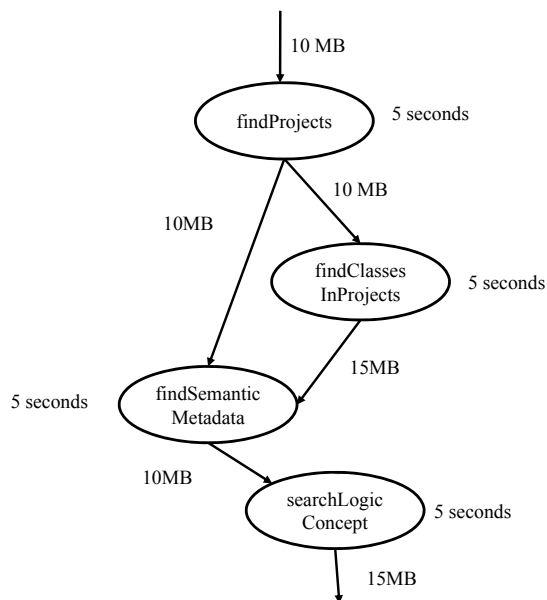


Figure 12: Cancer Data Standards Repository workflow. A workflow primarily queries concepts related to an input context.

4.3 Astronomy and Neutron Science

In this subsection we consider scientific workflow examples from the astronomy and neutron science community.

4.3.1 Astronomy workflow

The goal of the Pan-STARRS's (Panoramic Survey Telescope And Rapid Response System) project [et al. 2005] is a continuous survey of the entire sky. The data collected by the currently deployed prototype telescope 'PS1' will be used to detect hazardous objects in the Solar System, and other astronomical studies including cosmology and Solar System astronomy. The astronomy data from Pan-STARRS is managed by the teams at John Hopkins University and Microsoft Research through two workflows. The first PSLoad workflow (Figure 13) stages incoming data files from the telescope pipeline and loads them into individual relational databases each night. Periodically the online production databases that can be queried by the scientists, are updated with the databases collected over the week by the PSMerge workflow (Figure 14). The infrastructure to support the PS1 telescope data is still under development. Both the Pan-STARRS workflows are data intensive but require coordination and orchestration of resources to ensure reliability and integrity of the data products. The workflows have a high degree of parallelism through substantial partitioning of data into small subsets.

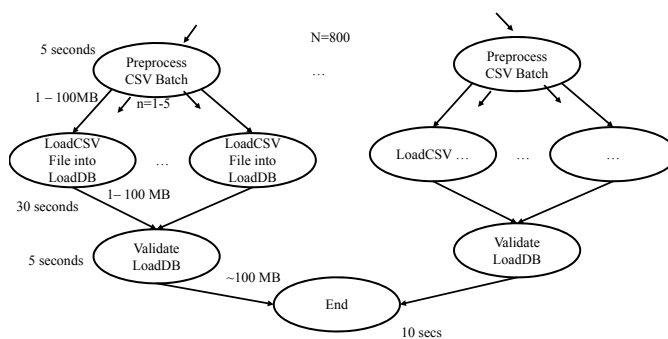


Figure 13: PSLoad workflow. Data arriving from the PS1 telescope is processed and staged in relational databases each night.

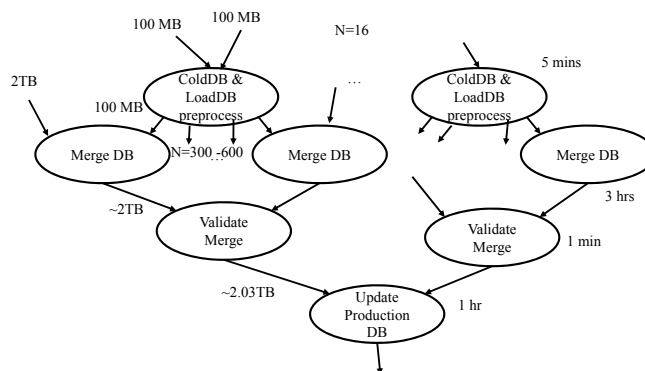


Figure 14: PSMerge workflow. Each week, the production databases that astronomers query are updated with the new data staged during the week.

4.3.2 McStas workflow

Neutron science research enables study of structure and dynamics of molecules that constitute materials. Neutron Source SNS at Oak Ridge National Laboratory connect large

neutron science facilities that contain instruments with computational resources such as the TeraGrid [Lynch et al. 2008]. The Neutron Science TeraGrid Gateway enables virtual neutron scattering experiments. These experiments simulate a beam line and enables experiment planning and experimental analysis. Figure 15 shows a virtual neutron scattering workflow using McStas, VASP, and nMoldyn. VASP and nMoldyn are used for molecular dynamics calculations and McStas is used for neutron ray-trace simulations. The workflow is computationally intensive and currently runs on ORNL supercomputing resources and TeraGrid resources. The workflow’s initial steps run for a number of days and are followed by additional compute intensive steps. The workflow is sequential and data set sizes remarkably small.

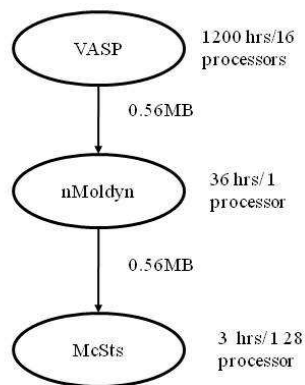


Figure 15: McStats Workflow. This workflow is used for Neutron ray-trace simulations. The data set sizes are remarkably small relative to compute time.

5. DISCUSSION

This paper presents workflows gathered through surveying of six scientific domains. The survey demonstrates similarities and differences of workflows along axes of structural pattern, data pattern, usage, compute time and data sizes. We summarize the results of the survey in Tables 2 and 3. The total number of tasks, length of longest chain and widest degree of parallelization describe the structure of the workflow. The workflows in our survey vary from a handful of tasks to thousands of components. The maximum processor width of a task gives an indication of computational needs and can be useful in resource planning. A number of the workflows are simple, requiring a single processor per task. However others, including motif, flood-plain mapping and McStats often require multiple processors for parallel data processing either in a message passing (MPI) style application or tightly coupled parallel application. The computation and data sizes convey a median time and size magnitude. The majority of workflows work with megabytes to gigabytes of data. However a few workflows such as PanSTARRS Merge can yield gigabyte to terabyte databases as outputs.

The classification we propose can be used to study workflow types in greater detail. For example, Table 4 is an attempt understand the relationship between computational and data sizes. Our workflow survey consists of a large number of workflows that have small data sizes and these tend to vary from taking a few seconds to minutes (e.g. LEAD

Data Mining) to hours (e.g., Storm surge) to days (e.g., McStats). Workflows with medium to large data sets (i.e., from megabytes to terabytes) tend to take longer time to process from hours to days as seen by the second and third rows of the table.

6. CONCLUSIONS

This paper reports on a survey of workflows from physical and natural sciences that vary in structure, and computational and data requirements. The workflows vary significantly in their structure, user constraints associated with them and environments in which they might run. Our proposed workflow classification model helps us understand the characteristics of the workflows and can serve as a foundation for design of next-generation workflow technologies.

7. ACKNOWLEDGEMENTS

This work is funded in part by the National Science Foundation under grant OCI-0721674 and the Linked Environments for Atmospheric Discovery funded by National Science Foundation under Cooperative Agreement ATM-0331480 (IU). This work was supported by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

The authors would like to thank the people who contributed to workflow survey including Suresh Marru and the entire LEAD team, Brian Blanton, Howard Lander, Steve Thorpe, Jeffrey Tilson, Sriram Krishnan, Luca Clementi, Ravi Madduri, Wei Tan, Cem Onyuksel, Yogesh Simmhan, Sudharshan Vazhkudai, Vickie Lynch. The authors would also like to thank Kelvin K. Droegemeier for explaining various concepts of mesoscale meteorology in great detail and Jaidev Patwardhan for feedback on the paper.

8. REFERENCES

- [Altintas et al. 2004] ALTINTAS, I., BERKLEY, C., JAEGER, E., JONES, M., LUDSCHER, B., AND MOCK, S. 2004. Kepler: An Extensible System for Design and Execution of Scientific Workflows.
- [Altschul et al. 1990] ALTSCHUL, S. F., GISH, W., MILLER, E. M., AND LIPMAN, D. 1990. Basic Local Alignment Search Tool. *Journal of Molecular Biology* 214, 1-8.
- [Atkins 2002] ATKINS, D. 2002. A Report from the U.S. National Science Foundation Blue Ribbon Panel on Cyberinfrastructure. In *CCGRID '02: Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid*. IEEE Computer Society, Washington, DC, USA, 16.
- [Blanton et al. 2008] BLANTON, B., LANDER, H., LUETTICH, R. A., REED, M., GAMIEL, K., AND GALLUPPI, K. 2008. Computational Aspects of Storm Surge Simulation. *Ocean Sciences Meeting*.
- [CaGrid Taverna Workflows] CaGrid Taverna Workflows. CaGrid Taverna Workflows. http://www.cagrid.org/wiki/CaGrid:How-To:Create_CaGrid_Workflow_Using_Taverna.
- [Churches et al. 2006] CHURCHES, D., GOMBAS, G., HARRISON, A., MAASSEN, J., ROBINSON, C., SHIELDS, M., TAYLOR, I., AND WANG, I. 2006. Programming Scientific and Distributed Workflow with Triana

Workflow Name	Structural Pattern	Data Pattern	Usage
LEAD Weather Forecasting	Sequential	Data production	Interactive, Event-driven, User Constrained
LEAD Data Mining	Sequential	Data reduction	Event-driven
Storm Surge	Parallel-merge	Data reduction	Interactive, Event-driven, User Constrained
Flood-plain mapping	Mesh	Data production	Event-driven
Glimmer	Sequential	Data processing	Interactive
Gene2Life	Parallel	Data reduction	Interactive
Motif	Parallel-split	Data production	Interactive
MEME-MAST	Sequential	Data processing	Interactive
Molecular Sciences	Parallel-merge	Data production	Interactive
Avian Flu	Parallel-split	Data processing	Interactive
caDSR	Sequential	Data processing	Interactive
PanSTARRS Load	Parallel-split-merge	Data processing	Event-driven
PanSTARRS Merge	Parallel-split-merge	Data processing	Event-driven
McStats	Sequential	Data processing	Interactive

Table 2: Workflow Survey Pattern and Usage Summary. The table captures the structural, data and usage patterns.

Workflow Name	Total no. of tasks	Longest Chain	Max Width	Max task processor width	Comp.	Data sizes
LEAD Weather Forecasting	6	4	3	16	hours	megabytes to gigabytes
LEAD Data Mining	3	3	1	1	minutes	kilobytes
Storm Surge	6	2	5	16	minutes- hours	megabytes
Flood-plain mapping	7	4	2	256	days	gigabytes
Glimmer	4	4	1	1	minutes	megabytes
Gene2Life	8	4	2	1	minutes	kilobytes to megabytes
Motif	138	4	135	256	hours	megabytes to gigabytes
MEME-MAST	2	2	1	1	minutes	kilobytes
Molecular Sciences	6	5	2	1	minutes	megabytes
Avian Flu	~ 1000	3	1000	1	minutes	kilobytes to megabytes
caDSR	4	4	1	1	seconds	megabytes
PanSTARRS Load	~ 1600 - 41000	4	800 - 40000	1	minutes	megabytes
PanSTARRS Merge	~ 4900 - 9700	4	4800 - 9600	1	hours	gigabytes to terabytes
McStats	3	3	1	128	days	kilobytes to megabytes

Table 3: Workflow Survey Summary. The table captures the structural, computational and data characteristics of the workflows.

Comp. Time Data Size	seconds to minutes	hours	days
Kilobytes to Megabytes	LEAD Data Mining, Glimmer, Gene2Life, MEME-MAST, Molecular Sciences, Avian Flu, caDSR, PanSTARRS Load	Storm surge	McStats
Megabytes to Gigabytes		LEAD Weather Forecasting, Motif	Flood-plain mapping
Gigabytes to Terabytes		PanSTARRS Merge	

Table 4: Comparison of Data Sizes of Workflows with Computation Time.

- Services. *Concurrency and Computation: Practice and Experience (Special Issue: Workflow in Grid Systems)* 18, 10, 1021–1037.
- [Condor DAGMan] CONDOR DAGMAN. Condor DAGMan. <http://www.cs.wisc.edu/condor/dagman/>.
- [Deelman et al. 2003] DEELMAN, E., BLYTHE, J., GIL, Y., AND KESSELMAN, C. 2003. Workflow Management in GriPhyN. Grid Resource Management, J. Nabrzyski, J. Schopf, and J. Weglarz editors, Kluwer.
- [Deelman et al. 2003] DEELMAN, E., BLYTHE, J., GIL, Y., KESSELMAN, C., MEHTA, G., VAHI, K., LAZZARINI, A., ARBREE, A., CAVANAUGH, R., AND KORANDA, S. 2003. Mapping Abstract Complex Workflows onto Grid Environments. *Journal of Grid Computing*, Vol. 1, No. 1,.
- [Deelman and Gil 2006] DEELMAN, E. AND GIL, Y. 2006. Report from the NSF Workshop on the Challenges of Scientific Workflows. *Workflow Workshop*.
- [Droegemeier et al. 2005] DROEGEMEIER, K. K., GANNON, D., REED, D., PLALE, B., ALAMEDA, J., BALTZER, T., BREWSTER, K., CLARK, R., DOMENICO, B., GRAVES, S., JOSEPH, E., MURRAY, D., RAMACHANDRAN, R., RAMAMURTHY, M., RAMAKRISHNAN, L., RUSHING, J. A., WEBER, D., WILHELMSON, R., WILSON, A., XUE, M., AND YALDA, S. 2005. Service-Oriented Environments for Dynamically Interacting with Mesoscale Weather. *Computing in Science and Engg.* 7, 6, 12–29.
- [et al. 2005] ET AL., N. K. 2005. Pan-STARRS Collaboration. *American Astronomical Society Meeting* 206.
- [Kramer and Ryan 2003] KRAMER, W. AND RYAN, C. 2003. Performance Variability of Highly Parallel Architectures. International Conference on Computational Science.
- [Lander et al. 2008] LANDER, H. M., FOWLER, R. J., RAMAKRISHNAN, L., AND THORPE, S. R. 2008. Stateful Grid Resource Selection for Related Asynchronous Tasks. Tech. Rep. TR-08-02, RENCI, North Carolina. April.
- [Li et al. 2008] LI, X., PLALE, B., VIJAYAKUMAR, N., RAMACHANDRAN, R., GRAVES, S., AND CONOVER, H. 2008. Real-time Storm Detection and Weather Forecast Activation through Data Mining and Events Processing. *Earth Science Informatics*.
- [Ludscher et al. 2005] LUDSCHER, B., ALTINTAS, I., BERKLEY, C., HIGGINS, D., JAEGER, E., JONES, M., LEE, E., TAO, J., AND ZHAO, Y. 2005. Scientific Workflow Management and the Kepler System.
- [Luettich et al. 1992] LUETTICH, R., WESTERINK, J. J., AND W. SCHEFFNER, N. 1992. ADCIRC: An Advanced Three-dimensional Circulation Model for Shelves, Coasts and Estuaries; Report 1: Theory and Methodology of ADCIRC-2DDI and ADCIRC-3DL. *Technical Report DRP-92-6, Coastal Engineering Research Center, U.S. Army Engineer Waterways Experiment Station, Vicksburg, MS*.
- [Lynch et al. 2008] LYNCH, V., COBB, J., FARHI, E., MILLER, S., AND TAYLOR, M. 2008. Virtual Experiments on the Neutron Science TeraGrid Gateway. *TeraGrid*.
- [North Carolina Floodplain Mapping Program] North Carolina Floodplain Mapping Program. North Carolina Floodplain Mapping Program. <http://www.ncfloodmaps.com/>.
- [Oinn et al. 2006] OINN, T., GREENWOOD, M., ADDIS, M., ALPDEMIR, M. N., FERRIS, J., GLOVER, K., GOBLE, C., GODERIS, A., HULL, D., MARVIN, D., LI, P., LORD, P., POCOCK, M. R., SENGER, M., STEVENS, R., WIPAT, A., AND WROE, C. 2006. Taverna: Lessons in Creating a Workflow Environment for the Life Sciences: Research Articles. *Concurr. Comput. : Pract. Exper.* 18, 10, 1067–1100.
- [Plale et al. 2006] PLALE, B., GANNON, D., BROTZGE, J., DROEGEMEIER, K. K., KUROSE, J., MCLAUGHLIN, D., WILHELMSON, R., GRAVES, S., RAMAMURTHY, M., CLARK, R. D., YALDA, S., REED, D. A., JOSEPH, E., AND CHANDRASHEKAR, V. 2006. CASA and LEAD: Adaptive Cyberinfrastructure for Real-time Multiscale Weather Forecasting. *IEEE Computer* 39, 66–74.
- [Ramakrishnan et al. 2006] RAMAKRISHNAN, L., BLANTON, B. O., LANDER, H. M., LUETTICH, R. A., JR, REED, D. A., AND THORPE, S. R. 2006. Real-time Storm Surge Ensemble Modeling in a Grid Environment. In *Second International Workshop on Grid Computing Environments (GCE), Held in conjunction ACM/IEEE*

- Conference for High Performance Computing, Networking, Storage and Analysis.*
- [Ramakrishnan et al. 2006] RAMAKRISHNAN, L., REED, M. S., TILSON, J. L., AND REED, D. A. 2006. Grid Portals for Bioinformatics. In *Second International Workshop on Grid Computing Environments (GCE), Held in conjunction with ACM/IEEE Conference for High Performance Computing, Networking, Storage and Analysis.*
- [Ramakrishnan et al. 2007] RAMAKRISHNAN, L., SIMMHAN, Y., AND PLALE, B. 2007. Realization of Dynamically Adaptive Weather Analysis and Forecasting in LEAD. In *In Dynamic Data Driven Applications Systems Workshop (DDDAS) in conjunction with ICCS (Invited).*
- [SCOOP Website] SCOOP Website. SCOOP Website. <http://scoop.sura.org>.
- [Taylor et al. 2004] TAYLOR, I., SHIELDS, M., AND WANG, I. 2004. Resource Management for the Triana Peer-to-Peer Services. In *Grid Resource Management*, J. Nabrzyski, J. M. Schopf, and J. Węglarz, Eds. Kluwer Academic Publishers, 451–462.
- [Taylor et al. 2006] TAYLOR, I. J., DEELMAN, E., GANNON, D. B., AND SHIELDS, M. 2006. *Workflows for e-Science: Scientific Workflows for Grids*. Springer.
- [Tilson et al. 2007] TILSON, J., BLATECKY, A., RENDON, G., MAO-FENG, G., AND JAKOBSSON, E. 2007. Genome-Wide Domain Analysis using Grid-enabled Flows. *Proceedings 7th International Symposium on BioInformatics and BioEngineering, IEEE*.
- [Tilson et al. 2007] TILSON, J., RENDON, G., MAO-FENG, G., AND JAKOBSSON, E. 2007. MotifNetwork: A Grid-enabled Workflow for High-throughput Domain Analysis of Biological Sequences: Implications for Study of Phylogeny, Protein Interactions, and Intraspecies Variation. *Proceedings 7th International Symposium on BioInformatics and BioEngineering, IEEE*.
- [Unidata Local Data Manager (LDM)] Unidata Local Data Manager (LDM). Unidata Local Data Manager (LDM). <http://www.unidata.ucar.edu/software/ldm/>.
- [van der Aalst et al. 2003] VAN DER AALST, W. M. P., TER, KIEPUSZEWSKI, B., AND BARROS, A. P. 2003. Workflow Patterns. *Distributed and Parallel Databases 14*, 1 (July), 5–51.
- [Yu and Buyya 2005] YU, J. AND BUYYA, R. 2005. A Taxonomy of Scientific Workflow Systems for Grid Computing. *SIGMOD Rec. 34*, 3 (September), 44–49.