Scientific Data Management Center (SDM Scidac-2)

FINAL REPORT

to the US Department of Energy SciDAC Program

Grant #DE-FC02-07ER25811

Bertram Ludäscher (UC Davis) & Ilkay Altintas (UC San Diego)

2007 – 2012

1 Introduction
2 Workflow development
3 Generic workflow components and templates
4 Provenance collection and analysis
5 Workflow reliability and fault tolerance
6 Publications and References
Publications
2012
2011
2010
2009
2008
2007
Talks and Presentations
References

1 Introduction

Effectively generating, managing, and analyzing scientific data requires a comprehensive, end-to-end approach that encompasses all stages from the initial data acquisition to the final analysis of the data. As part of the SPA thrust area, we developed a suite of tools and frameworks that integrate into a robust and auditable system for automation of scientific processes to enhance and speed up scientific discovery. Our technologies provide run-time management of the workflow processes, provenance collection, and analysis and display of results. This has led to the deployment of production workflows that allow scientists to a) monitor, in near-real-time, complex tasks such as the execution of large simulation codes, and b) facilitate complex analyses of the process metadata and of the simulation results. This has resulted in significant savings in scientists' time, in more efficient use of resources, and in a more cost-effective scientific discovery process overall.

Workflow technologies have a long history in the database and information systems communities [GHS95]. Similarly, the scientific community has developed a number of problem-solving environments, most of them as integrated solutions [HRG+00]. Component-based solution support systems are also proliferating [CL02, CCA06]. Scientific workflows merge advances in all these areas to automate support for sophisticated scientific problem-solving [LAB+06, LG05, DOE04, ABB+03, BVP00, VS97, SV96]. We use the term scientific workflow as a blanket term describing a series of structured activities and computations (called workflow components or actors) that arise in scientific problem-solving as part of the discovery process. This description includes the actions performed (by actors), the decisions made (control-flow), and the underlying coordination, such as data transfers (dataflow) and scheduling, required to execute the workflow. In its simplest case, a workflow is a linear sequence of tasks, each one implemented by an actor. An example of a scientific workflow is: transfer a configuration file to a large cluster, run a simulation passing this file as an input parameter, transfer the results of the simulation to a secondary system (e.g. a smaller cluster), select a known variable, and generate a movie showing how this variable evolves over time. Scientific workflows can exhibit and exploit data-, task-, and pipeline-parallelism. In science and engineering, process tasks and computations often are large-scale, complex, and structured with intricate dependencies [DOE04, DBN+96, EBV95, Elm66].

Over the past five years, our activities have both established Kepler as a viable scientific workflow environment and demonstrated its value across multiple science applications. We have published numerous peer-reviewed papers on the technologies highlighted in this short paper and have given Kepler tutorials at SC06, SC07, SC08, and SciDAC 2007. Our outreach activities have allowed scientists to learn best practices and better utilize Kepler to address their individual workflow problems.

Our contributions to advancing the state-of-the-art in scientific workflows have focused on the following areas. Progress in each of these areas is described in subsequent sections.

- **Workflow development.** The development of a deeper understanding of scientific workflows "in the wild" and of the requirements for support tools that allow easy construction of complex scientific workflows;
- **Generic workflow components and templates.** The development of generic actors (i.e. workflow components and processes) which can be broadly applied to scientific problems;
- **Provenance collection and analysis**. The design of a flexible provenance collection

and analysis infrastructure within the workflow environment; and

- **Workflow reliability and fault tolerance.** The improvement of the reliability and fault-tolerance of workflow environments.

2 Workflow development

The original base-line contribution of the SPA team has been to co-found the Kepler project – an open source workflow support environment [kepler-project.org]. Kepler is now a widely accepted scientific workflow development and execution environment that powers a number of research and production projects all over the world. Kepler has been the **most cited** (according to recent Web of Science and Google Scholar queries) and a very influential scientific workflow products with hundreds of thousands of downloads. The SPA researchers and engineers continue to regularly contribute to Kepler. We are constantly working with the Kepler Core team [e.g., AJB+04, LAB+06, GBA+07] to enhance Kepler at all levels including the user interface, documentation, and tutorials. Our work has led to a significant reduction in the effort required to generate real workflows [ABC+06, SAC+07]. We also actively partner with science teams to transfer technology to their projects and develop and deploy their workflows. This provides real-life case-studies which are then used to enhance Kepler requirements, to identify Kepler enhancements, generic functionalities, and canonical generic workflow solutions, and to improve user interfaces.

Accomplishments

As active participants, and founding members, in the Kepler research community we **contribute to the development of the Kepler environment**. Specifically, we participate in Kepler and Ptolemy workshops that have produced significant enhancements in the underlying workflow environment, and we have identified and implemented new requirements for scientific workflows, fixed bugs, and improved the overall software development environment, as well as execution-time interfaces and data collection practices [e.g., BMR+08, CA08, NV08, SAC+07, VAB+07]. Some of the specific tools and additions to Kepler are discussed in sections that follow (e.g., provenance recorder, generic actors, ADIOS). An example of using Kepler for simulation monitoring and steering is shown in Figure 1. This particular workflow processes files representing time-steps as they are generated by simulations, processes them for display on a project dashboard, and in parallel invokes a program (ELITE) which can determine if the simulation is stable.

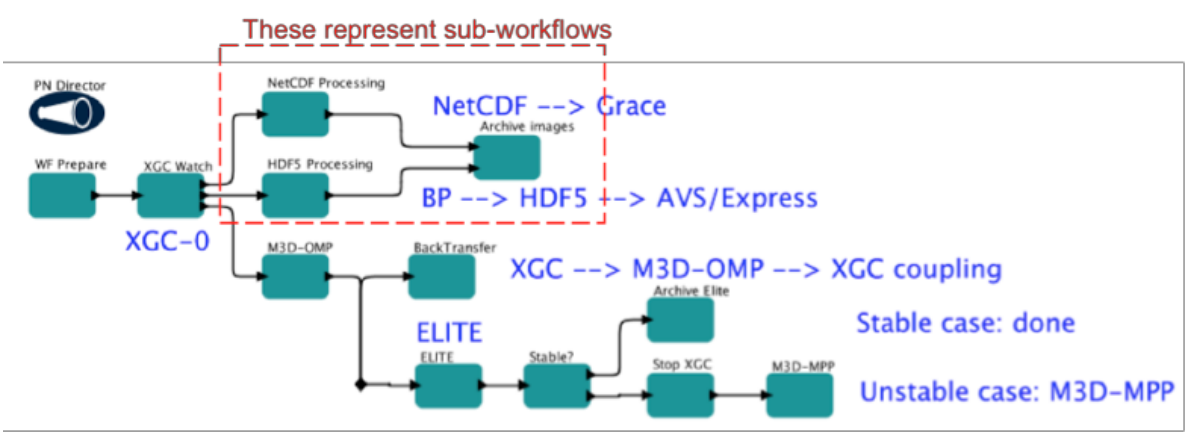

Figure 1: a workflow for monitoring the correctness of a simulation in real-time.

If the simulation is considered unstable, it stops the run (see box "stop XCG") thus saving compute resources, and calls another program (M3D-MPP) to recomputed new parameters for the next simulation run. This example shows the power of using workflows for near real-time simulation monitoring, generating summaries, graphs and images "on-the-fly" which can greatly help scientists follow up and understand their simulations behavior quickly.

Another critical component of the SPA activities is **transfer of workflow technology to science teams**. We have deployed workflows to a number of science teams to a) help introduce Kepler into their domains, and b) feedback lessons learned into Kepler to improve the Kepler environment. This includes:

- **ScalaBLAST** (Bioinformatics): This workflow performs a comparison of N submitted genomes against M genomes stored in a library. This requires generating the NxM individual comparisons, submitting the results on a cluster, and aggregating the results.
- **Groundwater Modeling and Analysis:** A set of Kepler-based scientific workflows have been constructed to support subsurface flow and transport modeling using the STOMP (*Subsurface* Transport Over Multiple Phases) simulation. The high-level groundwater modeling workflow involved specific computational tasks including clustering, multivariate interpolation, subsurface flow and transport simulation, and data visualization. Additional low-level workflows were developed to support data staging and simulation job submission. Furthermore, an iterative workflow was designed to collect several input variable ranges and perform a parameter study using fixed code and combinations of input variables values.
- **Center for Plasma Edge Simulation (CPES)**: This fusion project has been a main user of our technologies providing both requirements and feedback. As the first project using Kepler for monitoring and post-processing of High Performance Computing applications, our actors for SSH and job management found immediate use. eSiMon has been used as the front-end to fusion scientists and the content was created by Kepler workflows. Provenance tracking was first used in the fusion monitoring workflow and by fusion users of eSiMon. The above tools comprise EFFIS, the End-to-end Framework for Fusion Integrated Simulation [CKP+10]. The integrated simulation is driven by a Kepler "coupling" workflow. The plasma state in the simulation by the XGC0 code on a supercomputer is constantly monitored, which involves a data conversion step using another fusion code (M3D) and a parameter-study using yet another code (Elite) on an analysis cluster. If the plasma state is found to be unstable, the XGC0 simulation is stopped and a combined XGC0 and M3D magneto-hydrodynamic simulation is started to go through the turbulent phase of the fusion reaction. After that, XGC0 is continued again alone. Additionally, the workflow creates plots and 2D visualizations from each code's output at every timestep for eSiMon, which can be used for on-line monitoring of the coupled simulation as well as for post-processing runs.
- **S3D** (Jackie Chen, Combustion) – Three types of workflows were identified: Preparation, Production, and Post-Production workflows. The preparation workflows are designed to save unnecessary faulty runs on the supercomputer. Before accepting a particular build of S3S for production use, it is necessary to run a series of test cases. The production workflows are similar to CPES runs, but, of course, the diagnostic plots to be generated during the simulation are different. The post-processing workflows are designed to use the dashboard as an index to archived runs in order to facilitate collaboration with their visualization partners.
- **CCSM** (climate): This recent effort with the National Center for Atmospheric Research (NCAR) is helping develop Kepler-driven provenance for the Community Climate System

Model (CCSM[1]). CCSM belongs to an elite category of computer-based simulation models known as general-circulation models. The automatic capturing of provenance is an essential part of this activity, especially as the number and volume of simulations is expected to significantly grow.

We expect, and have seen, that over time, the science teams will assume responsibility for the maintenance of existing workflows and creation of new ones. Thus, we do not view our development of specific workflows as an ongoing obligation (or as a limit on the utilization of the tools, since other teams are developing workflows without our direct support) but rather a proven approach for deploying new technology through a win-win scenario.

In addition to our direct engagement with certain science teams, we have been **disseminating information on workflow technologies to the broader community** through a series of publications (see Publications), tutorials and presentations. In particular, we have given well attended tutorials at the last three Supercomputing conferences (SC06, SC07, and SC08), the SciDAC 08 conference, and the 2007 meeting of the NCCS. We have also provided tutorials for interested groups, upon request.

3 Generic workflow components and templates

Many workflows contain sections with very similar functionality but subtle differences in how that functionality is obtained. For example, transferring a file between machines, submitting a job to a batch processing system, monitoring the execution of a running job, and remotely executing a command are found in almost all of the scientific workflows we have developed. However, the actual implementation of these capabilities varies dramatically depending on features such as the specific machine configurations (e.g., which batch processor is used), the security requirements (e.g. ssh or rsh, certificates or one-time-passwords), and the workflow requirements (e.g. failover options, fault tolerance requirements, validation options). Tools that provide the instantiation of a case-specific workflow or workflow component from more general templates or components would substantially improve the ability to reuse existing solutions, leading to greater productivity when developing workflows.

Accomplishments

We have developed a large number of **actors and actor packages** to support SDM-type workflows, including the following: A *ProcessFile* actor that allows to execute external commands on a remote machine and supports checkpointing of successful commands (useful for "smart resume" in case a restart of the workflow is necessary) and logging of errors; an *SSH-FileWatcher* that allows to list and watch remote directories with a set of "file-masks;" this is a frequently used step when monitoring loosely-coupled, remote processes. These and various actors have been bundled into an ssh and job management package for remote execution, file transfer, and job submissions to various job management systems (PBS, Condor, SGE, and LoadLeveler). Another set of actors has been developed to deal with web service-oriented workflow steps; these include: *WSWithComplexTypes* (permits sending and receiving of structured web service parameters), various "shim actors" that allow conversion between strings, records, and XML structures, and a number of array manipulation actors.

We have also investigated several alternative approaches to development of **generic actors**, including the separation of the control flow and data flow aspects of the workflow. Our

[1] http://www.ccsm.ucar.edu/, http://www.esmf.ucar.edu/

current approach is to use data-driven code to flexibly adapt to the current environment by intelligently selecting the most appropriate execution path based on a combination of user input and available resources. We are also investigating the development of a broader workflow context resource, which would allow effective sharing of information (such as login certificates) between actors within a single, executing workflow. One of the key capabilities that have been used in several of our workflows is the movement of files between machines; for example, transferring information from the LCF machine to Ewok at ORNL on the CPES workflow. As a result, our first generic actor capability has been the deployment of a **GenericFileTransfer** actor. This actor can utilize a variety of protocols, including ftp, scp, and SRM-lite, to transfer a file between two computers. It also automatically confirms that the transfer completed properly (i.e. there were no errors and the resulting files are the same size) and if an error was encountered, retransmits the file up to a specified number of times.

Currently we have a **genericSSH** actor that can establish connection with SSH servers using a password, phassphrase, passcode, or an SSH certificate. We are currently working towards extending this functionality to support GSI-OpenSSH server. This requires implementing GSI authentication in addition to the existing authentication methods. We have also deployed a **genericJobLaunch** actor that encapsulates all the tasks involved in job launching into a single actor. This will enable users to use a single actor and configure its parameters instead of replicating the sub workflow and configuring its individual components.

We have developed a set of pilot **Kepler workflow templates** that we are evaluating. Similar to generic actors in their advantages, workflow templates can then be used, with only minor modifications, to implement new workflows. The current set is concerned with file creation monitoring and data transfers during long-term workflow runs with fusion codes. The intent is to create sufficiently generic components that can be used in long-term workflows in other disciplines. Templates have been presented and discussed as part of the Kepler tutorial at the Supercomputing 2008. We have also developed a new approach to workflow design called Collection-Oriented Modeling and Design (COMAD) and implemented a new COMAD director. We have demonstrated that workflows in this modeling paradigm tend to be easier to understand and maintain [MBZ+08].

4 Provenance collection and analysis

In scientific applications, effectively managing data provenance is extremely important [MGBM05, GMF+05, CFS+05, SPG05]. Data provenance [MGBM05] can be thought of as the complete processing history of a data product, for example, actor identification and invocation parameters (or application codes launched by those actors), properties such as the time, location, and userid of invocation, relevant environment and configuration parameters. This information needs to be persistent and permanently associated with a data product so that its provenance is readily available. It should also be searchable, so that data with certain provenance can be easily identified – for example, if a bug is identified and corrected, the provenance can help identify which runs should be repeated.

Accomplishments

We have **extended Kepler to support obtaining and recording provenance information** through a facility called "provenance recorder." This extension allows a workflow to automatically collect provenance information at a user-defined level of granularity [VAB07, CA08]. This information is recorded in a database allowing easy searching and retrieval. However, in the distributed, heterogeneous environment utilized by

scientific computing, simply recording provenance information from the workflow is insufficient. In a typical scientific workflow, while the workflow engine directs the activity, significant data manipulation occurs on derived products which are never fully incorporated into the workflow environment. For example, the workflow engine may request the creation of a movie from a simulation output, but only the file names, not the underlying data, move through the workflow environment. In order to ensure the recording of the full data provenance, we have enabled the provenance capture environment to allow recording of events external to the workflow – including the compilation of the underlying simulation codes. This provides a more complete record of the workflow execution, increasing the ability of scientists to reliably replicate the results.

We have developed and released several versions of the **provenance recorder** within Kepler. This has provided an entirely new set of capabilities and made Kepler even more valuable. For example, one of the primary reasons CCSM developers are interested in Kepler is to enable provenance capture in their workflows. We distinguish four types of provenance: data, process, workflow and system.

The provenance recorder is located within the Kepler framework and tracks all events of interest generated by workflow execution. These events include actor executions, token transfers between actors, and any additional information that an actor may choose to provide. Because a large, long-running workflow may create thousands of events, the most recent version of the recorder allows the workflow developer to identify which information is recorded. This significantly reduced the information flow load and improves the performance during both workflow execution and during evaluation. Since the release of Kepler v2.1.0 (30-Sep-2010, kepler-project.org), Kepler includes the provenance recorder.

Provenance information has multiple purposes, and thus requires multiple levels of granularity. For example, while a workflow is executing, significant provenance information may be collected to restart the workflow in case of error (see Fault Tolerance section below). However, once the workflow is complete, only a small fraction of that provenance may be required for validation purposes. In order to effectively reduce the amount of provenance information we store, we are actively working on approaches for archiving and managing provenance information based on its usage requirements. In addition to the workflow engine, external processes must have efficient access to the provenance store. This ensures that relevant information generated outside of the workflow engine, such as from a process spawned by Kepler, can be effectively recorded. It also allows tools such as the Dashboard to utilize this information without having to execute a workflow. However, that raises the matter of security and privacy of the provenance information. That information may sometimes be at least as valuable as the raw data. Our provenance framework solves that problem in part through its access mechanisms [NV08], but more work is needed. Provenance recording capabilities of Kepler and some open issues are described in [VAB07, NV08].

5 Workflow reliability and fault tolerance

There are two basic forms of run-time fault-tolerance: forward-recovery (e.g., failure masking and redundancy based failover), and backward-recovery (e.g. check-pointing) [LB90, MV06, Vou05]. Exception handling is a very traditional way of managing run-time problems [MV96, Vou05]. It is also used in the workflow-oriented environments [HA96, CCPP99]. Exception handling can involve forward-recovery, backward-recovery, or graceful termination. More recently the web services community has recognized the need

for some form of standardized fault-tolerance in the service provisioning through replication [SPPJ06]. An important component is collection of sufficient amount of meta-data (provenance information) about the workflow to enable fault-tolerance actions. It is the provenance information that is being collected through our provenance recorder that has the capability of providing meta-data needed to detect and locate workflow run-time issues.

In another line of research, we observed that scientific workflow systems frequently are used to execute a variety of long-running computational pipelines prone to premature termination due to network failures, server outages, and other faults. Researchers have presented approaches for providing fault tolerance for portions of specific workflows, but no solution handles faults that terminate the workflow engine itself when executing a mix of stateless and stateful workflow components. In [KRZ+11] we presented a general framework for efficiently resuming workflow execution using information commonly captured by workflow systems to record data provenance. Our approach facilitates fast workflow replay using only such commonly recorded provenance data. We also proposed a checkpoint extension to standard provenance models to significantly reduce the computation needed to reset the workflow to a consistent state, thus resulting in much shorter re-execution times. Our work generalizes the rescue-DAG approach used by DAGMan to richer workflow models that may contain stateless and stateful multi-invocation actors as well as workflow loops.

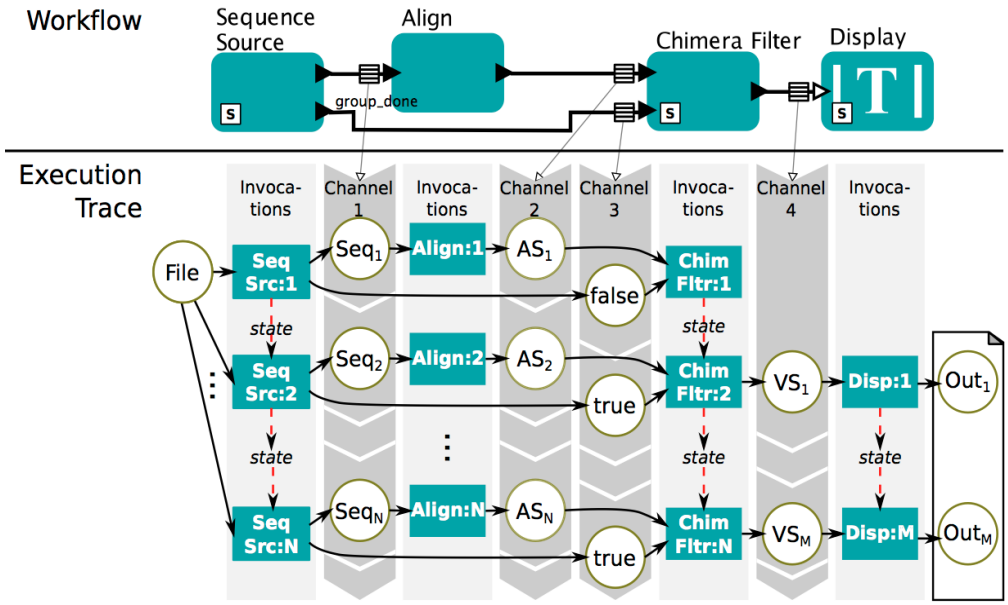

Figure 2: Example workflow with stateful actors. To recover the workflow execution after a fault, unconsumed tokens inside workflow channels and internal states of all actors except the stateless Align have to be restored [KRZ+11].

Accomplishments

We have implemented a Kepler-based fault-tolerance (FT) framework that leverages collected provenance information to provide options for forward recovery as well as

backward recovery of the workflows [MCA+10, YAC+10]. The FT framework addresses the majority of the known issues and faults in current production workflows and is composed of 3 major components:

- An error Handling Layer, which monitors the components beyond the workflow engine's direct control, such as visualization services, and reacts accordingly when an error is detected
- A Contingency actor, which provides a recovery block mechanism within the workflow.
- Checkpointing and Smart resume capabilities for when the above 2 mechanisms fail are not sufficient to prevent the workflow from terminating abnormally.

The Contingency actor handles faults at the Workflow Layer. This actor contains a primary sub-workflow, and possibly additional sub-workflows. During execution, data read by the actor is transferred first to the input of the primary sub-workflow. The sub-workflow contains the primary task to be executed, and if an error occurs, the Contingency actor may re-execute the primary sub-workflow or transfer the original input data to a different sub-workflow. The choice of which sub-workflow to execute is governed by a finite state machine. Additionally, the Contingency actor may be configured to pause between re-executions of a sub-workflow.

The extension of the model beyond the workflow control plane allows us to catch and process problem signals from environments that Kepler has no direct control over [MCA+10, YAC+10]. The key concept in that context is operational profiles – the frequency of usage of different Kepler and other operations, and their relationship to run-time failures. Operational profiles are an essential part of software reliability engineering. Typically they are created from the software requirements, and through customer reviews. Creation of operational profiles often is laborious and requires human intervention. Our approach builds an operational profile based on the actual usage from execution logs. The difficulty in using execution logs is that the amount of data to be analyzed is extremely large (more than a million records per day in many applications). Our solution constructs operational profiles by identifying all the possible clustered sequences of events (patterns) that exist in the logs. This is done very efficiently using suffix arrays data structure [NWV09, NVW+08].

6 Publications and References

Publications are papers published by the SDM center staff since the beginning of SciDAC-2. *References* are documents by others or selected publications by SDM center staff during SciDAC-1, or before the SDM center's inception in 2001.

Publications

2012

- [KGC+12] Sven Köhler, Supriya Gulati, Gongjing Cao, Quinn Hart, Bertram Ludäscher: Sliding Window Calculations on Streaming Data using the Kepler Scientific Workflow System. ICCS 2012: 1639-1646

2011

- [AAV+11] Ilkay Altintas, Manish Kumar Anand, Trung N. Vuong, Shawn Bowers, Bertram Ludäscher, Peter M. A. Sloot: A Data Model for Analyzing User Collaborations in Workflow-Driven e-Science. I. J. Comput. Appl. 18(3): 160-179 (2011)
- [DZM+11] Lei Dou, Daniel Zinn, Timothy M. McPhillips, Sven Köhler, Sean Riddle, Shawn Bowers, Bertram Ludäscher: Scientific workflow design 2.0: Demonstrating streaming data collections in Kepler. ICDE 2011: 1296-1299
- [KRZ+11] Sven Köhler, Sean Riddle, Daniel Zinn, Timothy M. McPhillips, Bertram Ludäscher: Improving Workflow Fault Tolerance through Provenance-Based Recovery. SSDBM 2011: 207-224
- [CSC+11] G. Chin Jr., C. Sivaramakrishnan, T. Critchlow, K. Schuchardt, A.H.H. Ngu, "Scientist-Centered Workflow Abstractions via Generic Actors, Workflow Templates, and Context-Awareness for Groundwater Modeling and Analysis", The Fifth IEEE International Workshop on Scientific Workflows (SWF 2011) published within the Proceedings of the Seventh IEEE World Conference on Services (Services 2011). Washington, DC, July 2011.
- [NJC+11] A. H. Ngu, A. Jamnagarwala, G. Chin Jr., C. Sivaramakrishnan, T. Critchlow, "Kepler Scientific Workflow Design and Execution with Contexts", International Journal of Computers and Their Applications (IJCA) Special Issue on Scientific Workflows, Provenance and Their Applications, In Press.
- [ZHM+11] Daniel Zinn, Quinn Hart, Timothy McPhillips, Bertram Ludäscher, Yogesh Simmhan, Michail Giakkoupis and Viktor K. Prasanna. Towards Reliable, Performant Workflows for Streaming-Applications on Cloud Platforms, 11th International Symposium on Cluster, Cloud and Grid Computing (ccGrid), May, 2011.
- [WKA+11] J. Wang, Korambath, P., and Altintas, I., "A Physical and Virtual Compute Cluster Resource Load Balancing Approach to Data-Parallel Scientific Workflow Scheduling", in *IEEE 2011 Fifth International Workshop on Scientific Workflows (SWF 2011), at 2011 Congress on Services (Services 2011)*, Washington, DC, USA, 2011.
- [CWA+11] Daniel Crawl, Jianwu Wang, Ilkay Altintas. Provenance for MapReduce-based Data-Intensive Workflows. In Proceedings of the 6th Workshop on Workflows in Support of Large-Scale Science (WORKS11) at Supercomputing 2011 (SC2011) Conference. ACM 2011, pages 21-29.

- [CWA+11] Daniel Crawl, Jianwu Wang, Ilkay Altintas. Provenance for MapReduce-based Data-Intensive Workflows. In Proceedings of the 6th Workshop on Workflows in Support of Large-Scale Science (WORKS11) at Supercomputing 2011 (SC2011) Conference. ACM 2011, pages 21-29.

2010

- [AAC+10] Ilkay Altintas, Manish K. Anand, Daniel Crawl, Adam Belloum, Paolo Missier, Carole A. Goble, Peter M.A. Sloot. "Understanding Collaborative Studies Through Interoperable Workflow Provenance ," in *Provenance and Annotation of Data and Processes*, 2010, vol. 6378, pp. 42–58..
- [ABA+10] Manish Kumar Anand, Shawn Bowers, Ilkay Altintas, Bertram Ludäscher. "Approaches for Exploring and Querying Scientific Workflow Provenance Graphs," in *Provenance and Annotation of Data and Processes*, 2010, vol. 6378, pp. 17–26.
- [ABL10] Manish Kumar Anand, Shawn Bowers, Bertram Ludäscher: Techniques for efficiently querying scientific workflow provenance graphs. EDBT 2010: 287-298
- [MCA+10] Pierre Mouallem, Daniel Crawl, Ilkay Altintas, Mladen Vouk and Ustun Yildiz, A Fault-Tolerance Architecture for Kepler-based Distributed Scientific Workflows, In 22nd International Conference on Scientific and Statistical Database Management, (SSDBM'2010).
- [NJC+10] A. Ngu, A. Jamnagarwala, G. Chin Jr. , C. Sivaramakrishnan, T. Critchlow, "Context-Aware Scientific Workflow Systems Using KEPLER", To appear in the International Journal of Business Process Integration. 2010.
- [WKK+10] Jianwu Wang, Prakashan Korambath, Seonah Kim, Scott Johnson, Kejian Jin, Daniel Crawl, Ilkay Altintas, Shava Smallen, Bill Labate, Kendall N. Houk. Theoretical Enzyme Design Using the Kepler Scientific Workflows on the Grid. Accepted by 5th Workshop on Computational Chemistry and Its Applications (5th CCA) at International Conference on Computational Science (ICCS 2010).
- [YPM_10] Ustun Yildiz, Pierre Mouallem, Mladen Vouk, Daniel Crawl and Ilkay Altintas, "Fault-Tolerance in Dataflow-based Scientific Workflow Management," in *SERVICES*, 2010, pp. 336–343.
- [ZSSL10] Zinn Daniel, Shawn Bowers, Sven Köhler, Bertram Ludäscher: Parallelizing XML data-streaming workflows via MapReduce . J. Comput. Syst. Sci. 76(6): 447-463 (2010)

2009

- [ABL09] Manish Kumar Anand, Shawn Bowers, Bertram Ludäscher, A navigation model for exploring scientific workflow provenance graphs. SC-WORKS 2009
- [ABML09a] M. Anand, S. Bowers, T. McPhillips, and B. Ludäscher, Exploring Scientific Workflow Provenance Using Hybrid Queries over Nested Data and Lineage Graphs, In 21st Intl. Conf. on Scientific and Statistical Database Management (SSDBM), New Orleans, 2009, 237 – 254.
- [ABML09b] Manish Kumar Anand, Shawn Bowers, Timothy M. McPhillips, Bertram Ludäscher: Efficient provenance storage over nested data collections. EDBT 2009: 958-969.
- [GBA+09] Antoon Goderis, Christopher Brooks, Ilkay Altintas, Edward A. Lee, Carole A. Goble: Heterogeneous composition of models of computation. Future Generation Comp. Syst. 25(5): 552-560 (2009).
- [LAB+09] B. Ludäscher, I. Altintas, S. Bowers, J. Cummings, T. Critchlow, E. Deelman,

D. D. Roure, J. Freire, C. Goble, M. Jones, S. Klasky, T. McPhillips, N. Podhorszki, C. Silva, I. Taylor, and M. Vouk. In A. Shoshani and D. Rotem, editors Scientific Process Automation and Workflow Management, Scientific Data Management: Challenges, Existing Technology, and Deployment, Computational Science Series, chapter 13. Chapman & Hall/CRC, 2009.

- [LBM09] B. Ludäscher, S. Bowers, and T. McPhillips. M. T. Ozsu and L. Liu, editors, Scientific Workflows, Encyclopedia of Database Systems. Springer, 2009.
- [LWM+09] Scientific Workflows: Business as Usual? , B. Ludäscher, M. Weske, T. McPhillips, S. Bowers. In 7th Intl. Conf. on Business Process Management (BPM), Ulm, Germany, 2009.
- [LZK+09] Jay Lofstead, Fang Zheng, Scott Klasky, Karsten Schwan, Adaptable, Metadata Rich IO Methods for Portable High Performance IO, IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2009.
- [MBZL09] Timothy M. McPhillips, Shawn Bowers, Daniel Zinn, Bertram Ludäscher: Scientific workflow design for mere mortals. Future Generation Comp. Syst. 25(5): 541-551 (2009).
- [MVK+09] Pierre Mouallem, Mladen Vouk, Scott Klasky, Norbert Podhorszki and Roselyne Barreto: "Tracking Files Using the Kepler Provenance Framework", Proceedings of 21st International Conference on Scientific and Statistical Database Management, SSDBM'09, LNCS 5566, pp. 273-282, New Orleans, LA, USA, June 2009.
- [WCA+09] Jianwu Wang, Daniel Crawl, Ilkay Altintas. Kepler + Hadoop – A General Architecture Facilitating Data-Intensive Applications in Scientific Workflow Systems. In Proceedings of the 4th Workshop on Workflows in Support of Large-Scale Science (WORKS09) at Supercomputing 2009 (SC2009) Conference. ACM 2009, ISBN 978-1-60558-717-2.
- [YGN09a] U. Yildiz, A. Guabtni, and A. H. H. Ngu, Towards scientific workflow patterns," in Proceedings of the 4th Workshop on Workflows in Support of Large-Scale Science, In conjunction with Supercomputing, SC (USA), ACM Press, 2009.
- [YGN09b] U. Yildiz, A. Guabtni, and A. H. H. Ngu, Business versus scientific workflows: A comparative study," in Proceedings of the IEEE Third International Workshop on Scientific Workflows, SWF (In conjunction with 7th IEEE International Conference on Web Services (ICWS 2009)), (USA), IEEE Computer Society Press, 2009.
- [ZBML09] Daniel Zinn, Shawn Bowers, Timothy M. McPhillips, Bertram Ludäscher: X-CSR: Dataflow Optimization for Distributed XML Process Pipelines. ICDE 2009: 577-580
- [ZSML09b] Daniel Zinn, Shawn Bowers, Timothy M. McPhillips, Bertram Ludäscher: Scientific workflow design with data assembly lines. SC-WORKS 2009

2008

- [BMR+08] Shawn Bowers, Timothy McPhillips, Sean Riddle, Manish Anand, Bertram Ludäscher, Kepler/pPOD: Scientific Workflow and Provenance Support for Assembling the Tree of Life, in the proceedings of the International Provenance and Annotation Workshop (IPAW) June 17-18, 2008, Salt Lake City, published in "Provenance and Annotation of Data and Processes, editors J. Freire, D. Koop and L. Moreau, Springer-Verlag, Berlin Heidelberg, 2008, 70-77.
- [BPL08] Shawn Bowers, Timothy M. McPhillips, Bertram Ludäscher: Provenance in collection-oriented scientific workflows. Concurrency and Computation: Practice and Experience 20(5): 519-529 (2008)

- [CA08] Daniel Crawl and Ilkay Altintas, A Provenance-Based Fault Tolerance Mechanism for Scientific Workflows, in the proceedings of the International Provenance and Annotation Workshop (IPAW) June 17-18, 2008, Salt Lake City, published in "Provenance and Annotation of Data and Processes, editors J. Freire, D. Koop and L. Moreau, Springer-Verlag, Berlin Heidelberg, 2008, pp. 152-159
- [GBA+08] Antoon Goderis, Christopher Brooks, Ilkay Altintas, Edward A. Lee and Carole Goble, Composing Heterogeneous Models of Computation in Kepler and Ptolemy II, accepted for publication in Future Generation Computer Systems (FGCS) in 2008. Early version published as a Technical Report, No. UCB/EECS-2007-139 EECS Department University of California Berkeley, November 27, 2007
- [LPA+08] Bertram Ludäscher, Norbert Podhorszki, Ilkay Altintas, Shawn Bowers, Timothy M. McPhillips: From computation models to models of provenance: the RWS approach. Concurrency and Computation: Practice and Experience 20(5): 507-518 (2008)
- [MBZ+08] T. McPhillips, S. Bowers, D. Zinn, B. Ludäscher, Scientific Workflow Design for Mere Mortals, . Future Generation Computer Systems, 2008
- [NBH+08] Anne H. H. Ngu, Shawn Bowers, Nicholas Haasch, Timothy M. McPhillips, Terence Critchlow, "Flexible Scientific Workflow Modeling Using Frames, Templates, and Dynamic Embedding," Statistical and Scientific Database Management, 20th International Conference, SSDBM 2008, Bertram Ludäscher, Nikos Mamoulis (Eds.), Hong Kong, China, July 9-11, 2008, Proceedings. Lecture Notes in Computer Science 5069 Springer 2008, ISBN 978-3-540-69476-2, pp. 566-572
- [Zin08] Modeling and Optimization of Scientific Workflows, Daniel Zinn, EDBT PhD Workshop, Nantes, France, 2008

2007

- [ACC+07] Ilkay Altintas, George Chin, Daniel Crawl, Terence Critchlow, David Koop, Jeff Ligon, Bertram Ludäscher, Pierre Mouallem1, Meiyappan Nagappan, Norbert Podhorszki, Claudio Silva, MladenVouk, "Provenance in Kepler-based Scientific Workflow Systems," Poster # 41, at Microsoft eScience Workshop Friday Center, University of North Carolina, Chapell Hill, NC, October 13 - 15, 2007, pp. 82
- [DEL+07] Provenance in scientific workflow systems, S. Davidson, S. Boulakia, A. Eyal, B. Ludäscher, T. McPhillips, S. Bowers, M. Anand, and J. Freire, IEEE Data Eng. Bull 30(4):44-50, 2007.
- [GBA+07] Antoon Goderis, Christopher Brooks, Ilkay Altintas, Edward A. Lee and Carole Goble. Composing Different Models of Computation in Kepler and Ptolemy II, Proc. of the 2nd Int. Workshop on Workflow Systems in e-Science (WSES 07) in conjunction with the Int. Conference on Computational Science (ICCS) 2007, Beijing, China, May 27-30, 2007
- [PLK07] Workflow automation for processing plasma fusion simulation data, N. Podhorszki, B. Ludäscher, S. Klasky, 2nd Workshop on Workflows in Support of Large-Scale Science (WORKS 07).
- [PLK07a] Archive Migration through Workflow Automation, N. Podhorszki, B. Ludäscher, S. Klasky, Intl. Conf. on Parallel and Distributed Computing and Systems (PDCS), November 1921, 2007, Cambridge, Massachusetts.
- [SAC+07] Arie Shoshani, Ilkay Altintas, Alok Choudhary, Terence Critchlow, Chandrika Kamath, Bertram Ludäscher, Jarek Nieplocha, Steve Parker, Rob Ross, Nagiza Samatova, Mladen Vouk , :SDM Center Technologies for Accelerating Scientific Discoveries," SciDAC 2007 Proceedings, Dec 2007, Journal of Physics, Conference Series, Vol. 78, paper

#012068, 5 pages.

- [SAC+07a] Arie Shoshani, Ilkay Altintas, Alok Choudhary, Terence Critchlow, Chandrika Kamath, Bertram Ludäscher, Jarek Nieplocha, Steve Parker, Rob Ross, Nagiza Samatova, Mladen Vouk, Scientific Data Management: Essential Technology for Accelerating Scientific Discoveries, CTWatch Quarterly, Volume 3, Number 4, November 2007.
- [VAB+07] M. A. Vouk, I. Altintas, R. Barreto, J. Blondin, Z.Cheng, T. Critchlow, A. Khan, S. Klasky, J. Ligon, B. Ludäscher, P. A. Mouallem, S. Parker, N. Podhorszki, A. Shoshani, C. Silva: "Automation of Network-Based Scientific Worflows" published “Grid-based Problem Solving Environments, IFIP, Volume 239, eds. Gaffney. PW, Pool JCT (Boston: Springer), pp 35-61, 2007.

Talks and Presentations

Invited talk: T. Critchlow, I. Altintas, S. Klasky, M. Vouk, S. Parker, B. Ludaescher, "Accelerating Scientific Exploration Using Workflow Automation Systems", SIAM Computational Science & Engineering (CSE'07), February 2007.

Invited talk: Ilkay Altintas, et al, Accelerating the scientific exploration process with scientific workflows, SciDAC 2006 conference, http://www.iop.org/EJ/abstract/1742-6596/46/1/065.

Tutorial: Ilkay Altintas, Bertram Ludaescher, Scott Klasky, Mladen A. Vouk. "Introduction to scientific workflow management and the Kepler system,", Proceedings of the 2006 ACM/IEEE conference on Supercomputing, Tampa, Florida, Nov 06, Article No. 205, 2006, ISBN:0-7695-2700-0

Tutorial: Scott Klasky, Mladen A. Vouk, Ilkay Altintas, Jeff Ligon, Pierre Moualem, Mei Nagappan, "Scientific Data Workflows," National Center for Computational Sciences (NCCS) 2007 Users Meeting at Oak Ridge National Laboratory, March 27-29, 2007

Tutorial: Altintas; Ilkay, Mladen Vouk, Scott Klasky, Norbert Podhorszki, Daniel Crawl, "Introduction to scientific workflow management and the Kepler system," presented at Supercomputing 2007, Reno, NV, 11 Nov 07, Complete materials are available on-line at: http://sdm.ncsu.edu/tutorial/ or http://groups.google.com/group/spa-tutorial

Tutorial: Timothy McPhillips, Sean Riddle, Manish Anand, Scientific Workflow Support for AToL Data Analysis & Management, Bertram Ludaescher, Shawn Bowers, NSF/AToL meeting, New Orleans, March 8-9, 2008

Tutorial: Ilkay Altintas, Mladen Vouk, Scott Klasky, Norbert Podhorszki, "Introduction to scientific workflow management and the Kepler system," SciDaC 2008, Seattle, WA, 18-Jul-08.

Tutorial: Ilkay Altintas, Mladen Vouk, Scott Klasky, Norbert Podhorszki, Daniel Crawl, "Introduction to scientific workflow management and the Kepler system," Supercomputing 2008, Austing, TX, Nov 08

References

- [ABB+03] Altintas I., S. Bhagwanani, D. Buttler, S. Chandra, Z. Cheng, M. Coleman, T. Critchlow, A. Gupta, W. Han, L. Liu, B. Ludäscher, C. Pu, R. Moore, A. Shoshani, M. Vouk, "A Modeling and Execution Environment for Distributed Scientific Workflows," Proc. 15th IEEE International Conference on Scientific and Statisical Database Management (SSDBM 2003).
- [ABD+99] LAPACK: E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen., LAPACK Users' Guide, Third Edition. ed. Philadelphia: SIAM, 1999.
- [AJB+04] Altintas, E. Jaeger, C. Berkley, M. Jones, B. Ludascher, and S. Mock, "Kepler: An Extensible System for Design and Execution of Scientific Workflows, 16th Intl. Conf. on Scientific and Statistical Database Management (SSDBM), Santorini, Greece, 2004 http://users.sdsc.edu/~ludaesch/Paper/ssdbm04-kepler.pdf
- [BCC+97] ScaLAPACK: L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley, ScaLAPACK User's Guide, 1997.
- [BLAST09] mpiBLAST: Open-Source Parallel BLAST, http://www.mpiblast.org/
- [BVP00] Balay R.I, Vouk M.A., Perros H., "Performance of Network-Based Problem-Solving Environments," Chapter 18, in Enabling Technologies for Computational Science Frameworks, Middleware and Environments, editors Elias N. Houstis, John R. Rice, Efstratios Gallopoulos, Randall Bramley, Hardbound, ISBN 0-7923-7809-1, 2000
- [CAL+06] Avery Ching, Alok Choudhary, Wei-keng Liao, Lee Ward, and Neil Pundit. Evaluating I/O characteristics and methods for storing structured scientific data. In Proceedings of the International Parallel and Distributed Processing Symposium, April 2006.
- [CCA06] http://www.cca-forum.org/, accessed February 2006
- [CCPP99] Fabio Casati, Stefano Ceri, Stefano Paraboschi, and Giuseppe Pozzi, "Specification and Implementation of Exceptions in Workflow Management Systems," ACM Transactions on Database Systems 24(3), Sept. 1999
- [CFL+06] Avery Ching, Wu-chun Feng, Heshan Lin, Xiaosong Ma, and Alok Choudhary. Exploring I/O strategies for parallel sequence database search tools with S3aSim. In Proceedings of the International Symposium on High Performance Distributed Computing, June 2006.
- [CL02] I. Crnkovic and M. Larsson (editors), Building Reliable Component-Based Software Systems, Artech House Publishers, ISBN 1-58053-327-2, 2002, http://www.idt.mdh.se/cbse-book/
- [CLR+00] P. Carns, W. Ligon III, R. Ross, and R. Thakur, "PVFS: A Parallel File System For Linux Clusters," Proceedings of the 4th Annual Linux Showcase and Conference, Atlanta, GA, October 2000, pp. 317–327.
- [DBN+96] R.L. Dennis, D.W. Byun, J.H. Novak, K.J. Galluppi, C.C. Coats, and M.A. Vouk, "The Next Generation of Integrated Air Quality Modeling: EPA's Models-3," Atmospheric Environment, accepted, in print, expected 1996.
- [DOE04] R. Mount et al., Department of Energy, Office of Science report, "Data Management Challenge". Nov 2004, http://www.er.doe.gov/ascr/Final-report-v26.pdf
- [EBV95] Elmaghraby S.E., Baxter E.I., and Vouk M.A., "An Approach to the Modeling and Analysis of Software Production Processes," Intl. Trans. Operational Res., Vol. 2(1), pp.

117-135, 1995.
- [Elm66] Elmaghraby S.E., "On generalized activity networks," J. Ind. Eng., Vol. 17, 621-631, 1966
- [FFG+05] Evan J. Felix, Kevin Fox, Kevin Regimbal, Jarek Nieplocha. "Active Storage Processing in a Parallel File System". 6th LCI International Conference on Linux Clusters: The HPC Revolution. Chapel Hill, North Carolina, on April 26, 2005.
- [GHS95] D. Georgakopoulos, M. Hornick, and A. Sheth, "An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure,"Distributed and Parallel Databases, Vol. 3(2), April 1995.
- [HA98] Claus Hagen, Gustavo Alonso, "Flexible Exception Handling in the OPERA Process Support System," ICDCS 1998, pp. 526-533
- [HRG+00] Elias N. Houstis, John R. Rice, Efstratios Gallopoulos, Randall Bramley (editors), Enabling Technologies for Computational Science Frameworks, Middleware and Environments, Kluwer-Academic Publishers, Hardbound, ISBN 0-7923-7809-1, 2000.
- [IG96] R. Ihaka and R. Gentleman. R: A Ranguage For Data Analysis And Graphics. Journal of Computational and Graphical Statistics, 5:299-314, 1996.
- [JPR06] ProRata TOOLbox, Journal of Proteome Research, Vol. 5, No. 11, 2006
- [LAB+06] Scientific Workflow Management and the Kepler System, B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, and Y. Zhao, Concurrency and Computation: Practice & Experience, 2006.
- [LB90] J.C. Laprie, and C. Beounes, "Definition and Analysis of Hardware- and Software-Fault-Tolerant Architectures", IEEE Computer Society Press, Volume 23, Issue 7, Pages: 39 – 51, July 1990.
- [LCC+06] Wei-keng Liao, Kenin Coloma, Alok Choudhary, Lee Ward, Eric Russell, and Neil Pundit. Scalable Design and Implementations for MPI Parallel Overlapping I/O. In the IEEE Transactions on Parallel and Distributed Systems, vol. 17, no. 11, pp. 1264-1276, Nov., 2006.
- [LG05] B. Ludäscher and C. A. Goble, editors. ACM SIGMOD Record, Special Section on Scientific Workflows, volume 34(3), September 2005.
- [LLC+03] J. Li, W. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, and M. Zingale, "Parallel netCDF: A High-Performance Scientific I/O Interface," Proceedings of SC2003, Phoenix, AZ, November 2003.
- [LMC+05] Lin H, Ma X, Chandramohan P, Geist A, Samatova NF, Efficient Data Access for Parallel BLAST, Proceedings of 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2005), pp. 72, 04-08 April 2005, doi: 10.1109/IPDPS.2005.190.
- [Mou04] Pierre Mouallem, "Fault Tolerance and Reliability in Scientific Workflows," MS Thesis, NC State University, 2004
- [Mou04] Richard Mount, editor, The Office of Science Data-Management Challenge,Report from the DOE Office of Science Data-Management Workshops, March–May 2004, http://www-user.slac.stanford.edu/rmount/dm-workshop-04/Final-report.pdf.
- [MV96] D.F. McAllister, and M.A. Vouk, "Software Fault-Tolerance Engineering," Chapter 14 in Handbook of Software Reliability Engineering, McGraw Hill, pp. 567-614, January 1996.
- [PGL+07] Swapnil V. Patil, Garth A. Gibson, Sam Lang, and Milo Polte, "GIGA+: scalable directories for shared file systems," Parallel Data Storage Workshop in conjunction with Supercomputing, 2007.

- [PKM+06] Pan C, Kora G, McDonald WH, Tabb DL, VerBerkmoes NC, Hurst GB, Pellletier DA, Samatova NF, Hettich RL., ProRata: A quantitative proteomics program for accurate protein abundance ratio estimation with confidence interval evaluation, Anal Chem., vol. 78, no. 20, pp. 7121-31, Oct 2006, PMID: 17037911.
- [PKT+06] Pan C, Kora G, Tabb DL, Pelletier DA, McDonald WH, Hurst GB, Hettich RL, Samatova NF., Robust estimation of peptide abundance ratios and rigorous scoring of their variability and bias in quantitative shotgun proteomics, Anal Chem., vol. 78, no. 20, pp. 7110-20, Oct 2006, PMID: 17037910
- [PST+08] Milo Polte, Jiri Simsa, Wittawat Tantisiriroj, Garth Gibson, "Fast Log-Based Concurrent Writing of Checkpoints," Petascale Data Storage Workshop, November 2008.
- [RPM09] RPM RScaLAPACK Download, http://rpm.pbone.net/index.php3/stat/4/idpl/4766700/com/R-RScaLAPACK-0.5.1-9.fc7.i386.rpm.html
- [SBY05] N. Samatova, D. Bauer, and S. Yoginath, The taskPR Package, pp. 6, May 2005, http://cran.r-project.org/doc/packages/taskPR.pdf
- [SPPJ06] J. Salas, F. Perez, M. Patia-Martinez, R. Jiminez-Peris, "WS-Replication: A Framework for Highly Available Web Services," WWW Conf., Edinburgh, Scotland, May 2006.
- [SS07] H. Shan and J. Shalf, "Using IOR to Analyze the I/O Performance of HPC Platforms," 2007 Cray User Group Conference, Seattle WA, May 2007.
- [SV96] Singh M.P., Vouk M.A., "Scientific workflows: scientific computing meets transactional workflows," Proceedings of the NSF Workshop on Workflow and Process Automation in Information Systems: State-of-the-Art and Future Directions, Univ. Georgia, Athens, GA, USA; 1996, pp.SUPL28-34.
- [TGL99] Rajeev Thakur, William Gropp, and Ewing Lusk, "On Implementing MPI-IO Portably and with High Performance," in Proc. of the Sixth Workshop on I/O in Parallel and Distributed Systems, May 1999, pp. 23-32.
- [Vou05] Mladen A Vouk, "Software Reliability Engineering of Numerical Systems," Chapter 13, in Accuracy and Reliability in Scientific Computing, Editor: Bo Einarsson, ISBN 0-89871-584-9, SIAM, 2005, pp 265-300.
- [VS97] Vouk M.A., and M.P. Singh, "Quality of Service and Scientific Workflows," in The Quality of Numerical Software: Assessment and Enhancements, editor: R. Boisvert, Chapman & Hall, pp.77-89 , 1997.
- [YSB+05] Yoginath S, Samatova NF, Bauer D, Kora G, Fann G, Geist A, RScaLAPACK: High-performance parallel statistical computing with R and ScaLAPACK, Proceedings of the 18th International Conference on Parallel and Distributed Computing Systems (PDCS-2005), Sep 12-14, 2005, Las Vegas, Nevada.
- [Yu09] H. Yu. The Rmpi Package. http://cran.r-project.org/web/packages/Rmpi/index.html, 2009.