

Scientific Workflow Management and the KEPLER System*

Bertram Ludäscher^{†*} Ilkay Altintas[†] Chad Berkley[‡] Dan Higgins[‡]
Efrat Jaeger[†] Matthew Jones[‡] Edward A. Lee[§] Jing Tao[†] Yang Zhao[§]

September 2004; revised March 2005

Abstract

Many scientific disciplines are now data and information driven, and new scientific knowledge is often gained by scientists putting together data analysis and knowledge discovery “pipelines”. A related trend is that more and more scientific communities realize the benefits of sharing their data and computational services, and are thus contributing to a distributed data and computational community infrastructure (a.k.a. “the Grid”). However, this infrastructure is only a means to an end and scientists ideally should be bothered little with its existence. The goal is for scientists to focus on development and use of what we call *scientific workflows*. These are networks of analytical steps that may involve, e.g., database access and querying steps, data analysis and mining steps, and many other steps including computationally intensive jobs on high performance cluster computers. In this paper we describe characteristics of and requirements for scientific workflows as identified in a number of our application projects. We then elaborate on KEPLER, a particular scientific workflow system, currently under development across a number of scientific data management projects. We describe some key features of KEPLER and its underlying PTOLEMY II system, planned extensions, and areas of future research. KEPLER is a community-driven, open source project, and we always welcome related projects and new contributors to join.

Contents

1	Introduction	2
2	Scientific Workflows	3
2.1	Example Workflows	3
2.1.1	Promoter Identification	3
2.1.2	Mineral Classification	4
2.1.3	Job Scheduling	5
2.2	Requirements and Desiderata	6
2.3	Differences to Business Workflows	8
3	Highlights of KEPLER	8
3.1	Web Service Extensions	8
3.2	Grid and other Extensions	9
3.3	Actor-Oriented Modeling	10
4	Research Issues	13
4.1	Higher-Order Constructs	13
4.2	Third Party Transfers	14
4.3	Other Research Issues	15
4.4	Related Work	15
5	Conclusions	16

*Work supported by NSF/ITR 0225676 (SEEK), DOE SciDAC DE-FC02-01ER25486 (SDM), NSF/ITR CCR-00225610 (Chess), NSF/ITR 0225673 (GEON), NIH/NCRR 1R24 RR019701-01 Biomedical Informatics Research Network Coordinating Center (BIRN-CC), NSF/ITR 0325963 (ROADNet), NSF/DBI-0078296 (Resurgence)

[†]San Diego Supercomputer Center, UC San Diego; ^{*}Dept. of Computer Science & Genome Center, UC Davis; [‡]National Center for Ecological Analysis and Synthesis, UC Santa Barbara; and [§]Department of Electrical Engineering and Computer Sciences, UC Berkeley

“The diversity of the phenomena of nature is so great, and the treasures hidden in the heavens so rich, precisely in order that the human mind shall never be lacking in fresh nourishment.”

— Johannes Kepler, *Mysterium Cosmographicum*

1 Introduction

Information technology is revolutionizing the way many sciences are conducted, as witnessed by new techniques, results, and discoveries from quickly evolving, multi-disciplinary fields such as bioinformatics, biomedical informatics, cheminformatics, ecoinformatics, geoinformatics, etc. To further advance this new data- and information-driven science through advanced IT infrastructure, large investments are made, e.g., in the UK e-Science programme, or in the US through the NSF Cyberinfrastructure initiative and other initiatives from NIH (BIRN: Biomedical Informatics Research Network) and DOE (SciDAC: Scientific Discovery through Advanced Computing, GTL: Genomes to Life), just to mention a few. While many efforts focus on the underlying middleware infrastructure, known as “the Grid”, scientists are ultimately interested in tools that bring the power of distributed databases and other computational Grid resources to the desktop, and allow them to conveniently put together and run their own *scientific workflows*. By these we mean process networks that are typically used as “data analysis pipelines” or for comparing observed and predicted data, and that can include a wide range of components, e.g., for querying databases, for data transformation and data mining steps, for execution of simulation codes on high performance computers, etc. Ideally, the scientist should be able to plug-in almost any scientific data resource and computational service into a scientific workflow, inspect and visualize data on the fly as it is computed, make parameter changes when necessary and re-run only the affected “downstream” components, and capture sufficient metadata in the final products such that the runs of a scientific workflow, when considered as (computational) experiments themselves, help explain the results and make them reproducible by the computational scientist and others. Thus, a scientific workflow system becomes a scientific problem-solving environment, tuned to an increasingly distributed and service-oriented Grid infrastructure.

However, before this grand vision can become reality, a number of significant challenges have to be addressed. For example, current Grid software is still too complex to use for the average scientist, and fast changing versions and evolving standards require that

these details be hidden from the user by the scientific workflow system. Web services seem to provide a simple basis for loosely coupled, distributed systems, but core web service standards such as WSDL [WSD03] only provide simple solutions to simple problems,¹ while harder problems such as web service orchestration, 3rd party transfer (from one service directly to another, circumventing the transfer back to a workflow control engine), and transactional semantics of service-based workflows, remain the subject of emerging or future web service standards. The complexity of the underlying technical issues and the resulting (sometimes overly) complex standards make it less likely that those will be as widely adopted as the core standards such as XML and WSDL.

Another set of challenges arises from the inherent complexity of scientific data itself. For example, how can we capture more of the semantics of scientific data (beyond simple metadata meant for human consumption) and thus inform the system which data sets might be suitable input for a specific analytical pipeline? Similarly, how can we define when it is even potentially meaningful at the conceptual level to compose two independently designed web services, or when an analysis pipeline might be included as a subworkflow in another scientific workflow? Knowledge representation techniques, including formal ontologies, and corresponding Semantic Web standards such as the Web Ontology Language [OWL03] seem promising directions. However, as is the case for Grid middleware, the goal is to hide the underlying complexity as much as possible from the user of a scientific workflow system.

The paper is organized as follows: In Section 2 we introduce scientific workflows by means of several real-world examples from different domains. We use those examples to illustrate some of the characteristic features and requirements of scientific workflows, and compare the latter with business workflows. In Section 3 we present specific features of KEPLER and its underlying PTOLEMY II system. As it turns out, PTOLEMY II provides much more than a user-friendly graphical user-interface (called Vergil) and a ready-to-be-extended open source platform. The main advantage of the system lies in a modeling and design paradigm called *actor-oriented modeling* that has proven to be essential to deal with the complex design issues of scientific workflows. Section 4 presents some ongoing research issues. Finally, in Section 5, we briefly summarize our findings and work.

¹E.g. WSDL mainly provides an XML notation for function signatures, i.e., the types of inputs and outputs of web services.




Figure 1: Conceptual (“napkin drawing”) view of the Promoter Identification Workflow (PIW) [ABB⁺03]

2 Scientific Workflows

There is a growing interest in scientific workflows as can be seen from a number of recent events, e.g., the Scientific Data Management Workshop [SDM03], the e-Science Workflow Services Workshop [eSc03], the e-Science Grid Environments Workshop [eSc04], the Virtual Observatory Service Composition Workshop [GRI04], the e-Science LINK-Up Workshop on Workflow Interoperability and Semantic Extensions [LIN04], and last not least, various activities as part of the Global Grid Forum (e.g. [GGF04]), just to name a few. Scientific workflows also play an important role in a number of ongoing large research projects dealing with scientific data management, including those funded by NSF/ITR (GriPhyN, GEON, LEAD, SCEC, SEEK, ...), NIH (BIRN), DOE (SciDAC/SDM, GTL), and similar efforts funded by the UK e-Science initiative (myGrid, DiscoveryNet, and others). For example, the SEEK project [SEE] is developing an Analysis and Modeling System (AMS) that allows ecologists to design and execute scientific workflows [MBJ⁺04]. The AMS workflow component employs a Semantic Mediation System (SMS) to facilitate workflow design and data discovery via semantic typing [BL04]. Thus SEEK is a good example of a community-driven project in need of a system that allows users to seamlessly access data sources and services, and put them together into reusable workflows. Indeed SEEK is one of the main projects contributing to the cross-project KEPLER initiative and workflow system discussed below.

Aspects and Types of Workflows. Scientific workflows often exhibit particular “traits”, e.g., they can be data-intensive, compute-intensive, analysis-

intensive, visualization-intensive, etc. The workflows in Sections 2.1.1, 2.1.2, and 2.1.3, e.g., exhibit different features, i.e., service-orientation and data analysis, re-engineering and user interaction, and high-performance computing, respectively. Depending on the intended user group, one might want to hide or emphasize particular aspects and technical capabilities of scientific workflows. For example, a “Grid engineer” might be interested in low-level workflow aspects such as data movement and remote job control. Having workflow components (or *actors*) that operate at this level will be beneficial to the Grid engineer. Conversely, a scientific workflow system should hide such aspects from analytical scientists (say an ecologist studying species richness and productivity).

The KEPLER system aims at supporting very different kinds of workflows, ranging from low-level “plumbing” workflows of interest to Grid engineers, to analytical knowledge discovery workflows for scientists, and conceptual-level design workflows that might become executable only as a result of subsequent refinement steps [BL05].

In the following we first introduce scientific workflows by means of several examples taken from different projects and implemented using the PTOLEMY II-based KEPLER system [KEP]. We then discuss typical features of scientific workflows and from this derive general requirements and desiderata for scientific workflow systems. We take a closer look at underlying technical issues and challenges in Section 3.

2.1 Example Workflows

2.1.1 Promoter Identification

Figure 1 shows a high-level, conceptual view of a typical scientific *knowledge discovery workflow* that

links genomic biology techniques such as microarrays with bioinformatics tools such as BLAST to identify and characterize eukaryotic promoters² – we call this the *Promoter Identification Workflow* or PIW (see also [Wer01, ABB⁺03, PYN⁺03]): Starting from microarray data, cluster analysis algorithms are used to identify genes that share similar patterns of gene expression profiles that are then predicted to be co-regulated as part of an interactive biochemical pathway. Given the gene-ids, gene sequences are retrieved from a remote database (e.g., GenBank) and fed to a tool (e.g., BLAST) that finds similar sequences. In subsequent steps, transcription factor binding sites and promoters are identified to create a promoter model that can be iteratively refined.

While Figure 1 leaves many details open, some features of scientific workflows can already be identified: There are a number of existing databases (such as GenBank) and computational tools (such as ClustalW and BLAST) that need to be combined in certain ways to create the desired workflow. In the past, accessing remote resources often meant implementing a *wrapper* that mimics a human entering the input of interest, submitting an HTML form, and “screen-scraping” the result from the returned page [LPH01]. Today, more and more tools and databases become accessible via web services, greatly simplifying this task. Another trend are web portals such as NCBI [NCB04] that integrate many tools and databases and sometimes provide the scientist with a “workbench” environment.

Figure 2 depicts snapshots of an early implementation of PIW in KEPLER. KEPLER is an extension of the PTOLEMY II system [PTO04] for scientific workflows. The topmost window includes a loop whose body is expanded below and which performs several steps on each of the given gene-ids: First, an NCBI web service is used to access GenBank data. Subsequently a BLAST step is performed to identify similar sequences to the one retrieved from GenBank. Then a second inner loop is executed (bottom window) for a transcription factor binding site analysis. Using PTOLEMY II terminology, we call the individual steps *actors*, since they act as independent components which communicate with each other only through the channels indicated in the figure. The overall execution of the workflow is orchestrated by a *director* (the green box in Figure 2; see Section 3.3 for details).

This early PIW implementation in KEPLER [ABB⁺03] illustrates a number of features: Actual “wiring” of a scientific workflow can be much more complicated than the conceptual view (Figure 1) sug-

²A promoter is a subsequence of a chromosome that sits close to a gene and regulates its activity.




Figure 2: PIW implemented in KEPLER [ABB⁺03]. Composite actors (subworkflows) expanded below.

gests. A mechanism for collapsing details of a sub-workflow into an abstract component (called *composite actor* in PTOLEMY II) is essential to tame complexity: The windows in Figure 2 have well-defined input and output ports and thus correspond to (sub)-workflows that can be collapsed into a more abstract, composite actor as indicated. Nevertheless, the resulting workflow is fairly complex and we will need to introduce additional mechanisms to simplify the design in particular of loops (see Section 4.1).

2.1.2 Mineral Classification

The second example, from a geoinformatics domain, illustrates the use of a scientific workflow system for *automation* of an otherwise manual procedure, or alternatively, for *reengineering* an existing custom tool in a more generic and extensible environment. The upper left window in Figure 3 shows the top-level workflow: Some samples are selected from a database holding experimentally determined mineral compositions of igneous rocks. This data, together with a set of classification diagrams are fed into a CLASSIFIER subworkflow (bottom left). The manual process of classifying samples involves determining the position of the sample values in a series of diagrams such as the one shown on the right in Figure 3: if the location of a sample point in a non-terminal diagram of order n has been determined (e.g., *diorite gabbro*

anorthosite, bottom right), the corresponding diagram of order $n+1$ is consulted and the point located therein. This process is iterated until the terminal level of diagrams is reached (here shown in the upper right: the classification result is **anorthosite**).

This traditionally manual process has been automated in commercial custom tools, or here in the KEPLER workflow shown in Figure 3. As above, workflows are shown in graphical form using PTOLEMY II’s Vergil user interface [BLL⁺04b]. Note that in Vergil, workflows can be annotated with user comments. Subworkflows (e.g., bottom-left) become visible by right-clicking on a composite actor (such as CLASSIFIER, upper-left) and selecting “Look Inside” from the resulting pop-up menu. Vergil also features simple VCR-like control buttons to *play*, *pause*, *resume*, and *stop* workflow execution (red icons in the top-left toolbar; e.g., right-triangle for *play*).

KEPLER specific features of this workflow include: A searchable library of actors and data sources (*Actor* and *Data* tabs close to the upper-left) with numerous reusable KEPLER actors. For example, the BROWSER actor (used in the bottom-right of the CLASSIFIER subworkflow) launches the user’s default browser and can be used as a powerful generic input/output device in any workflow. In this example, the classification diagrams are generated on the client side as interactive SVG displays in the browser (windows on the right in Figure 3). Moving the mouse over the diagram highlights the specific region and displays the rock name classification(s) for that particular region. The BROWSER actor has proven to be very useful in many other workflows as well, e.g., as a device to display results of a previous step, and as a selection tool that passes user choices (made via HTML forms, check-boxes, etc.) to subsequent workflow steps.

2.1.3 Job Scheduling

The final example workflow, depicted in Figure 4, is from a cheminformatics domain and involves running thousands of jobs of the GAMESS quantum chemical code [SBB⁺93] under the control of the Nimrod/G Grid distribution tool [AGK00]. This is an example of a workflow employing *high-performance computing* (HPC) resources in a coordinated manner to achieve a computationally hard task, in this case a variant of a hybrid quantum mechanics/molecular mechanics (QM/MM) technique; see [GT98] and [SBA⁺04] for details. Interestingly, the workflow in Figure 4 is rather domain-neutral and illustrates some features typical of many high-performance computational experiments:

The main window shows four composite actors,




Figure 4: Workflow for scheduling HPC jobs.

corresponding to the four depicted subworkflows. The first one, PREPAREINPUTS creates a list of input files for the subsequent jobs. These files are then used to create a plan file for Nimrod/G in the PREPAREEXPERIMENT step. The ADDEXPERIMENT subworkflow takes a plan file and generates experiment run files using several COMMANDLINE actors. The latter is shown with a “\$” icon (to indicate a command shell), and has proven to be a very useful rapid-prototyping tool: Existing local applications can be made part of a workflow simply by providing a suitable command line expression and the corresponding command line arguments. The MANAGERESOURCES subworkflow can create new processes (via ADDFORK) to run jobs and subsequently add experiments as new server processes.

This example workflow also highlights the possibility of incremental design and development: At the time of writing, not all components of the overall workflow are operational. Nevertheless, due to the clearly defined input/output interfaces of all subworkflows (a feature inherited from PTOLEMY II), each of them can be designed, implemented, and tested separately. Moreover, the current version of the workflow relies heavily on invoking external applications via the COMMANDLINE actor. Some of these applications might be “promoted” to custom actors with native Java implementations in the future. Such changes are encapsulated by the containing subworkflow and thus do not require changes of other parts of the workflow.




Figure 3: Mineral Classification workflow (left) and generated interactive result displays (right).

2.2 Requirements and Desiderata

In this section we summarize a number of common requirements and desiderata of scientific workflows, as exhibited by the examples above or by other workflows we encountered in various application-oriented research projects including GEON, SEEK and several others [GEO, SEE, SDM, BIR, ROA].

R1: Seamless access to resources and services:

This is a very common requirement (e.g., see the example workflows in Section 2.1), and web services provide a first, simple mechanism for remote service execution and remote database access³ via service calls. However, as mentioned before, web services are a simple solution to a simple problem. Harder problems, e.g., web service orchestration, and 3rd party transfer are not solved by “vanilla” web services alone.

R2: Service composition & reuse and workflow design:

Since web services emerge as the basic building blocks for distributed Grid applications and workflows, the problem of *service composition*, i.e., how to compose simple services to perform complex tasks, has become a

hot research topic [ICA03]. Among the different approaches are those that view service composition as an AI planning problem [BDG03], a query planning problem [LAG03, LN04], or a general design and programming problem. A related issue is how to design components so that they are easily reusable and not geared to only the specific applications that may have driven their original development. As we will see, service composition and reuse are addressed by employing an *actor-oriented* approach at the design level (Section 3.3), but also require flexible means for *data-transformations* at the “plumbing” level (Section 3.2).

R3: Scalability: Some workflows involve large volumes of data and/or require high-end computational resources, e.g., running a large number of parallel jobs on a cluster computer (such as workflow in Section 2.1.3). To support such data-intensive and compute-intensive workflows, suitable interfaces to Grid middleware components (sometimes called *Compute-Grid* and *Data-Grid*, respectively) are necessary.

R4: Detached execution: Long running workflows require an execution mode that allows the workflow control engine to run in the background on a remote server, without necessarily staying connected to a user’s client application that

³We do not elaborate on the important challenges of *data integration* [She98]; see, e.g., [Hal01] for a survey of query rewriting techniques, and [NL04] and [LGM03, BLL04a] for related issues of query capabilities and semantics, respectively.

has started and is controlling workflow execution (such as the Vergil GUI of KEPLER).

R5: Reliability and fault-tolerance: Some computational environments are less reliable than others. For example, a workflow that incorporates a new web service can easily “break”, as the latter can often fail, change its interface, or just become unacceptably slow (as it becomes more popular). To make a workflow more resilient in an inherently unreliable environment, contingency actions must be specifiable, e.g., fail-over strategies with alternate web services.

R6: User-interaction: Many scientific workflows require user decisions and interactions at various steps.⁴ For example, an improved version of PIW (Section 2.1.1) allows the user to inspect intermediate results and select and re-rank them before feeding them to subsequent steps. An interesting challenge is the need for user interaction in a detached execution. Using a notification mechanism the user might be asked to reconnect to the running instance and make a decision before the paused (sub-)workflow can resume.

R7: “Smart” re-runs: A special kind of user interaction is the change of a parameter of a workflow or actor. For example, in a visualization pipeline or a long running workflow, the user might decide to change some parameters after inspecting intermediate or even final results. A “smart” re-run would not execute the workflow from scratch, but only those parts that are affected by the parameter change. In dataflow-oriented systems (e.g., visualization pipeline systems such as AVS, OpenDX, SCIRun, or the KEPLER system) this is easier to realize than in more control-oriented systems (e.g., business workflow systems), since data and actor dependencies are already explicit in the system. Another useful technique in this context is *checkpointing*, which allows to backtrack (in the case of a parameter change or even a system failure; cf. (R5)) to a previously saved state without starting over from scratch.

R8: “Smart” (semantic) links: A scientific workflow system should assist workflow design and data binding phases by suggesting which actor components might possibly fit together (this is also an aspect of (R2), service composition), or

by indicating which data sets might be fed to which actors or workflows. To do so, some of the semantics of data and actors has to be captured. However, capturing data semantics is a hard problem in many scientific disciplines: e.g., measurement contexts, experimental protocols, and assumptions made are often not adequately represented. Even if corresponding metadata is available, it is often not clear how to best make it useable by the system. It seems clear though that ontologies provide a very useful *semantic type system* for scientific workflows, in addition to the current (structural) type systems [BL04].

R9: Data provenance: Just as the results of a conventional wet lab experiment should be reproducible, computational experiments and runs of scientific workflows should be reproducible and indicate which specific data products and tools have been used to create a derived data product. Beyond the conventional capture of metadata, a scientific workflow system should be able to automatically log the sequence of applied steps, parameter settings and (persistent identifiers of) intermediate data products. A related desiderata is *automatic report generation*: The system should allow the user to generate reports with all relevant provenance and runtime information, e.g., in XML format for archival and exchange purposes and in HTML (generated from the former, e.g., via an XSLT script) for human consumption.

Data provenance can be seen as a prerequisite to (R8): In order to provide semantic information about a derived data product, suitable provenance information is needed.

While the above list of requirements and desiderata for scientific workflow systems is by no means complete, it should be sufficient to capture many of the core characteristics. Other requirements include the use of an intuitive GUI to allow the user to compose a workflow visually from smaller components, or to “drill-down” into subworkflows, to animate workflow execution, to inspect intermediate results, etc.

A scientific workflow system should also support the combination of different *workflow granularities*. For example, coarse-grained workflows, akin to Unix pipelines or web service-based workflows, consist mainly of “black box” actors whose contents are unknown to the system. Scientific workflows may also be very fine-grained, or include fine-grained subworkflows. In that case, components are “white boxes” containing, e.g., the visual programming equivalent

⁴In fact, when workflow management was still called “office automation”, humans were the main processors of tasks – the workflow system was just used for book-keeping; cf. Section 2.3.

of an algorithm, or a system of differential equations to be solved, in other words, a detailed specification known to the system.

2.3 Differences to Business Workflows

The characteristics and requirements of scientific workflows are partially overlapping those of *business workflows*. Indeed, the term ‘scientific workflows’ seems to indicate a very close relationship with the latter, while a more detailed comparison reveals a number of significant differences. Historically, business workflows have roots going back to *office automation* systems of the 1970’s and 80’s, and gained momentum in the 90’s under different names including *business process modeling* and *business process engineering*; see, e.g., [AM97, vdAvH02, zM04].

Today we see some influence of business workflow standards in the web services arena, specifically standards for *web service choreography*.⁵ For example, the Business Process Execution Language for Web Services (BPEL4WS) [CGK⁺02], a merger of two earlier standards, IBM’s WSFL and Microsoft’s XLANG, has received some attention recently.

When analyzing the underlying design principles and execution models of business workflow approaches, a focus on *control-flow* patterns and *events* becomes apparent, whereas dataflow is often a secondary issue. For example, [vdAtHKB03] describe a large number of workflow design patterns that can be used to analyze and compare business workflow standards and products in terms of their control features and expressiveness.

Scientific workflow systems, on the other hand, tend to have execution models that are much more *dataflow-oriented*. This is true, e.g., for academic systems including KEPLER, TAVERNA [TAV], and TRIANA [TRI], and for commercial systems such as Inforsense’s DISCOVERYNET or Scitegic’s PIPELINE-PILOT. With respect to their modeling paradigm and execution models, these systems seem closer to an “AVS for scientific data and services” than to the more control-flow and task-oriented business workflow systems, or to their early scientific workflow predecessors [CM95, MVW96, AIL98].

The difference between dataflow-orientation and control-flow orientation can also be observed in the underlying formalisms. For example, visualizations of business workflows often resemble flowcharts, state transition diagrams, or UML activity diagrams, all of which emphasize events and control-flow over

dataflow. Formal analysis of workflows usually involves studying their control-flow patterns [Kie02], and is often conducted using Petri nets.

Conversely, the underlying execution model of current scientific workflow systems usually resembles or is even directly implemented as a *dataflow process network* [KM77, LP95], having traditional application areas, e.g., in digital signal processing. Dataflow-oriented approaches are applicable at very different levels of granularity, from low-level CPU operations found in certain processor architectures, to high-level programming paradigms such as *flow-based programming* [Mor94]. Scientific workflow systems and visualization pipeline systems can also be seen as dataflow-oriented *problem solving environments* [WBB96] that scientists use to analyze and visualize their data. Last not least, there is also a close relationship between dataflow-oriented approaches and (pure) functional languages, including non-strict variants such as Haskell (cf. Section 4.1).

3 Highlights of KEPLER

In this section, we discuss some highlights of the current KEPLER system as well as some upcoming extensions. Many features directly address the requirements and desiderata from Section 2. More research-oriented extensions are described in Section 4.

3.1 Web Service Extensions

A basic requirement for scientific workflows is seamless access to remote resources and services (see (R1) in Section 2.2 and the examples in Section 2.1). Since web services are emerging as the standard means for remote service execution of loosely coupled systems, we extended KEPLER early on to handle web services. Given the URL of a web service description [WSD03], the generic WEBSERVICE actor of KEPLER can be *instantiated* to any particular operation specified in the service description. After instantiation, the WEBSERVICE actor can be incorporated into a scientific workflow as if it were a local component. In particular, the WSDL-defined inputs and outputs of the service are made explicit via the instantiated actor’s input and output ports.

Figure 5 shows screenshots of an extended web service harvesting feature, implemented by a special web service HARVESTER component.⁶ As in the case of the generic WEBSERVICE actor, a URL is first provided (see (1) in Figure 5), however this time not to an individual WSDL description of a web service,

⁵Despite the long history of business workflows, it is surprising how short-lived some of the so-called standards are, as “most of them die before becoming mature” [vdA03].

⁶Inspiration came from a similar feature in TAVERNA.




Figure 5: KEPLER web service HARVESTER in action: repository access (1-2), harvesting (3), and use (4).

but to a web service repository. The repository URL might point to a UDDI repository, or simply to a web page listing multiple WSDL URLs as shown in (2). The HARVESTER then retrieves and analyzes all WSDL files of the repository, creating instantiations of web service actors in the user’s local actor library; see (3). For example, one of the harvested services, the BLAST web service, comprises five service operations which are imported into a corresponding subdirectory. The user can then drag-and-drop any of these service operations on the workflow canvas for use in a scientific workflow (4). The HARVESTER feature facilitates rapid prototyping and development of web service-based applications and workflows in a matter of minutes – that is, provided

- (i) the web services are alive when needed, and
- (ii) they can be wired together more or less directly to perform the desired complex task.

The problem with (i) is that, while harvested web services look like local components, their runtime failure can easily “break” a scientific workflow, reminding the user that the service *interface* has been harvested, not the actual code.⁷ We are currently extending KEPLER to make workflows with web services more reliable. One simple approach is to avoid the association of a service operation with a fixed URL. Instead, a list of alternate services can be provided when the workflow is launched, and service failure can then be compensated by invocation of one of the alternate services. Another option is to insert special control tokens into the data stream, indicating

⁷Which is of course the whole point of web services.

to downstream actors the absence of certain results. Long running workflows may thus more gracefully react to web service failures and produce at least partial results. This idea has been further developed for “collection-oriented” (in the functional programming sense) workflows: via so-called “exception-catching actors”, invalid (due to failures) data collections can be filtered out of the data stream, while valid subcollections pass through unaffected [McP05]. An interesting research question is how to extend PTOLEMY II’s pause-resume model to a full-fledge transaction model that can handle service failures.

The problem (ii) is even more fundamental and has different aspects: At the *design level* the challenge is how to devise actors that can be reused easily. In Section 3.3 we give a brief introduction to *actor-oriented modeling*, the underlying paradigm of PTOLEMY II, and discuss how it facilitates component composition and reuse. At the “*plumbing level*” it is often necessary to apply data transformations between two consecutive web services (called “shims” in TAVERNA). Such data transformations are supported through various actors in KEPLER, e.g., XSLT and XQuery actors to apply transformations to XML data, or Perl and Python actors for text-based transformations.

3.2 Grid and other Extensions

Figure 6 depicts a number of KEPLER actors that facilitate scientific workflows, including workflows that make use of “the Grid”. In the upper left, the previously discussed generic WEBSERVICE actor and some

instantiations are shown. Note how the latter specialize their actor interface via their input/output ports: e.g., `BLAST_SEARCHSIMPLE` has three input ports and one output port, for the search arguments and result, respectively. The naming scheme used is `WSN_OP`, where `WSN` is the name of the web service and `OP` is a specific web service operation.




Figure 6: Grid actors and other KEPLER extensions.

The upper right shows two Grid actors, called `FILEFETCHER` and `FILESTAGER`, respectively. These actors make use of GridFTP [Pro00] to retrieve files from, or put files to, remote locations on the Grid. The `GLOBUSJOB` actor below is another Grid actor, in this case for running a Globus job [Glo]. At the bottom of Figure 6 a small workflow is shown that takes a Globus proxy and some input files, staging the files to where the job is run, then fetching the results from the remote location and displaying them on the client side. The green box specifies that this workflow is executed using an SDF (Synchronous Data-Flow) director. This director analyzes the dataflow dependencies and token consumption and production rates of actors (here: token = file), and schedules the execution of actors accordingly.

On the right, a number of actors that use the SDSC Storage Resource Broker [SRB] are shown, e.g., to connect and disconnect from SRB and to get and put files from and to SRB space, respectively. We are currently in the process of providing all commonly used SRB commands as actors. This will allow the KEPLER user to design and execute Grid workflows involving a number of different tools, e.g., SRB for

data handling aspects, and Globus, Nimrod and other tools for computational aspects and job scheduling.

In the center and left of Figure 6, various other KEPLER actors are shown: The `COMMANDLINE` actor can be used to incorporate any application into a workflow, provided it can be accessed from the command line.⁸ The “\$” icon is reminiscent of a shell prompt. The actor is parameterized with the arguments of the shell command, making it easy to create generic or specialized command line invocations. A `BROWSER` actor is shown directly below (cf. Section 2.1.2). It takes as input an HTML file or URL and displays it in the user’s default browser. This makes the actor an ideal output device for displaying intermediate or final workflow results in ways that are well-known to users. Another extremely useful application of this actor is as an input device for user interactions. The result file of an upstream actor might have been transformed to an HTML file (e.g., using the XSLT actor) and augmented with HTML forms, check boxes, or other input forms that are displayable to the user in a standard web browser. Upon executing the desired user interaction, an `http-post` request is sent to a special KEPLER web server, acting as a listener, and from there the workflow is resumed.

The `EMAIL` actor in the center of the figure provides a simple notification mechanism to inform the user of specific situations in the workflow. Together, the `EMAIL` and `BROWSER` actors address core issues of requirement (R6) in Section 2.2. The `PAUSE` actor (red down-triangle) pauses workflow execution at specific points, allowing the user to inspect intermediate results, possibly changing parameter values, and resuming the workflow subsequently (addressing (R7) in Section 2.2).

Finally, actors for accessing *real-time data streams* from ROADNet sensor networks [ROA] have recently been added. These actors (e.g., `ORBWAVEFORM-SOURCE`) can be integrated easily into KEPLER, since many of the underlying PTOLEMY II directors support streaming execution.⁹

3.3 Actor-Oriented Modeling

Arguably the most unique feature of KEPLER comes from the underlying PTOLEMY II system:

*“The focus [of the Ptolemy project] is on assembly of concurrent components. The key underlying principle ... is the use of well-defined **mod-***

⁸E.g., KEPLER workflows can include data analysis steps via calls to R [R].

⁹This should come as no surprise, since dataflow process networks are defined on token streams in the first place.

*els of computation that govern the interaction between components.”*¹⁰

This focus together with the *actor-oriented modeling* paradigm make PTOLEMY II an ideal starting point for tackling the breadth of challenges in scientific workflow design and execution. In Ptolemy, a system or model thereof (in our case, a scientific workflow) is viewed as a composition of independent components called *actors*. Communication between actors happens through interfaces called *ports*. We distinguish between *input ports* and *output ports*. In addition to the ports, actors have *parameters*, which configure and customize the behavior.¹¹ For example, a generic filter actor might consume a stream of input tokens via an input port, letting through to the output port only those tokens that satisfy a condition specified by a parameter.




Figure 7: The semantics of component interaction is determined by a director, which controls execution and supplies the objects (called receivers) that implement communication.

Actors, or more precisely their ports, are connected to one another via *channels*. Given an interconnection of actors, however, there are many possible execution semantics that one could assign to the diagram. For example, actors might have their own thread of control, or their execution might be triggered by the availability of new inputs.

A key property of PTOLEMY II is that the execution semantics is specified in the diagram by an object called a *director* (see Figure 7). The director defines how actors are executed and how they communicate with one another. Consequently, the execution model is less an emergent side-effect of the various interconnected actors and their (possibly ad-hoc) orchestration, and more a prescribed concurrent semantics as one might find in a well-defined concurrent programming language. The execution model defined by the director is called the *model of computation*. Patterns of concurrent interaction are factored out into the design of the directors, rather than being individually

constructed by the designer of the workflow. Figure 7 depicts a producer and a consumer actor whose ports are connected by a unidirectional channel. The diagram is annotated by a director, which might, for example, execute the producer prior to the consumer so as to respect data precedences. The communication between the actors is mediated by an object called a *receiver*, which is provided by the director, not by the actors. Thus, for example, whether the communication is buffered or synchronous is determined by the designer of the director, not by the designer of the actor. This hugely improves the reusability of actor designs.

Process Networks. The Process Network (PN) director is a popular choice for designers of scientific workflows. It gives a diagram the semantics of (dataflow) process networks [KM77, LP95]. In this semantics, actors are independent processes that execute concurrently, each with its own thread of control, and communicate by sending *tokens* through unidirectional channels with (in principle) unbounded buffering capacity. Writing to a channel is a non-blocking operation, while reading from a channel can block until sufficient input data are available. This model of computation is similar to that provided by Unix pipes, as in the following example of a Unix command-line composition of processes:

```
cat foo.txt | bar | baz
```

This example shows three independently executing processes (`cat`, `bar`, and `baz`) that are connected to one another through unidirectional pipes. The stream of tokens flowing between the processes also synchronizes them if necessary. For example if `bar` and `baz` are filter operations working on a single line of text at a time (e.g., `grep xyz`), then a Unix process executing `bar` will block until a line of text is provided by the process executing `cat foo.txt`. Unlike Unix pipes, however, the PN director in PTOLEMY II tolerates feedback loops and forking and merging of data streams. It performs deadlock detection, and manages buffers to keep memory requirements bounded (if possible).

The PN director is only one example of a large number directors available in PTOLEMY II. There is also, for example, the SDF (Synchronous Data-Flow) director, which can be used for specialized process networks with fixed token production and consumption rates per firing (see below). The SDF director performs static analysis on a workflow that guarantees absence of deadlocks, determines required buffer sizes, and optimizes the scheduling of actor execution. Other directors have been constructed for modeling

¹⁰<http://ptolemy.eecs.berkeley.edu/objectives.htm>.

¹¹Parameters are usually not shown in the figures.

Discrete Event systems (DE), Continuous-Time models (CT, which solve ordinary differential equations), and Communication Sequential Processes (CSP), to mention just a few [BLL⁺04b].

By relieving actors from the details of component interaction, the actors themselves become much more reusable (cf. (R2) in Section 2.2). The behavior of an actor adapts to the execution and communication semantics provided by the director. This feature of actor-oriented modeling is called *behavioral polymorphism*. For example, a single PTOLEMY II actor implementation of an arithmetic operation, say PLUS, can be connected to any number of input operands and reused within different models of computation and under the control of different directors. An SDF director, e.g., schedules the actor invocation (or “firing”) as soon as *all* inputs have data, which it knows since actors declare their fixed token consumption and production rates in the SDF domain. In contrast, when the PLUS actor is governed by a DE director, additions happen when *any* input has data, corresponding to the different overall execution model in the Discrete Event domain. In addition to behavioral polymorphism, the PTOLEMY II type system also supports *data polymorphism*, again increasing the reusability of actors. For example, our PLUS actor can be implemented in such a way that it dynamically chooses the correct numeric addition (integer, float, double, complex), depending on the types of inputs it receives. Moreover, on other data types, e.g., strings, vectors, matrices, or user-defined types, the PLUS actor¹² can execute appropriate actions, e.g., string concatenation, vector or matrix addition, etc.

Actor-Oriented Programming Interface.

Actor-oriented modeling addresses several challenges in the design of complex systems [EJL⁺03]. We have already mentioned improved component reusability due to behavioral and data polymorphism. Another aspect is *hierarchical modeling*. As illustrated by the examples in Section 2.1, subworkflows can be abstracted into (composite) actors themselves (e.g., see the CLASSIFIER actor/subworkflow in Figure 3) and thus arbitrarily nested. In the following, we give a simplified introduction on some implementation aspects of PTOLEMY II’s actor-oriented approach. These can be adapted to the context of scientific workflows and distributed, service-oriented environments, leading to a more structured approach to service composition and workflow design.

The structure we propose is based on various phases and methods in PTOLEMY II’s actor-oriented

<i>execution</i>	→ preinitialize , <i>type-check</i> , <i>run</i> [*] , wrapup
<i>run</i>	→ initialize , <i>iteration</i> [*]
<i>iteration</i>	→ prefire , fire [*] , postfire

Figure 8: AOPI execution *phases* and actor **methods**.

programming interface (AOPI), see Figure 8. These AOPI methods are used by a director to orchestrate overall execution. Symbols in boldface denote actual methods that actor implementations have to provide; the remaining symbols describe other phases¹³ of the overall execution.

When a director starts a workflow *execution*, it invokes the **preinitialize** method of all actors. Since this method is invoked only once per lifetime of an execution (even if there are multiple runs), and prior to all other activities, this is a good time to put in place the receiver components of actors, and for actors to “advertise” their supported port data types, transport protocols, etc.

Next the director *type-checks* all connections and ports. This includes checking each port’s data types, all (previously advertised) type constraints, and the validity of port types being connected through channels. A type inference algorithm is used to determine the most general types satisfying the given constraints. For scientific workflows, we can modify directors to also type-check which transport protocol to use, or to check whether producer and consumer actors exchange data directly or via *handles*:¹⁴ For example, if an actor A declares its output port to be of *handle type* “**http** | **ftp**” and a connected actor B declares its input port to be of handle type “**http**”, then type-checking can establish that the connection is valid, provided A’s output port is subtyped to use **http** handles only. Indeed such information can and should be passed to the actor with the invocation of the **initialize** method.

Other possible actions during execution of **initialize** are: Web service actors can “ping” the web services they represent and signal failure-to-initialize if the corresponding service is not alive. A “fail-over-aware” director can use this information to replace the defective web service with an equivalent one that is alive (see (R5) in Section 2.2). A workflow *execution* will often consist only of one *run*, but if a workflow is re-run, **initialize** is called again. A *run* usually includes multiple *iterations*, each of which includes a call to **prefire**, **fire** (possibly called repeat-

¹²This actor is called ADDSUBTRACT in PTOLEMY II.

¹³Some correspond to methods of other PTOLEMY II entities, e.g., *director methods* or *manager methods* [BLL⁺04b].

¹⁴By *handle* we mean a unique identifier that can also be used to retrieve data, e.g., a URL.

edly by some special directors), and a call to **postfire**. The main actor operation finally happens in the **fire** method, e.g., a web service actor will make the actual remote service call here.

Towards Actor-Oriented Scientific Workflows.

The idea of *actor-oriented scientific workflows* is to apply the principles of actor-orientation and hierarchical modeling, underlying the Ptolemy approach [EJL⁺03, BLL⁺04b], to the modeling and design of scientific workflows. In particular, web service operations, which provide the building blocks of many loosely coupled workflows, should be structured into different phases, corresponding to the different phases and methods used in actor-oriented modeling. For example to implement a web service w_A , the service developer should think of specific web service operations such as w_A .**initialize** and w_A .**prefire** in addition to the main “worker” method w_A .**fire**. As in the case of Ptolemy actors, this will lead to more generic and reusable components and even facilitate more complex extensions such as stateful web services.¹⁵

4 Research Issues

In this section we briefly discuss some technical issues that we have begun addressing for KEPLER, but that are less mature and require some additional research.

4.1 Higher-Order Constructs

The early implementation of the Promoter Identification Workflow (PIW) depicted in Figure 2 demonstrated the feasibility and some advantages of implementing scientific workflows in the KEPLER extension of PTOLEMY II [ABB⁺03]. However, it also highlighted some inherent challenges of the dataflow-oriented programming paradigm [LA03]. We have argued in Section 2.3 that many current scientific workflow systems are more dataflow-oriented than business workflow systems and approaches, which tend to emphasize event-based control-flow rather than dataflow. When designing real-world scientific workflows it is necessary, however, to handle complex control-flows within a dataflow-oriented setting as well. It is well-known that control-flow constructs require some thought in order to handle them properly. The fairly intricate network topology in Figure 2 includes backward-directed “dataflow” channels, having the sole purpose of sending control tokens that

¹⁵Statefulness is an established concept in actor-oriented modeling and dataflow networks; e.g., it can be represented explicitly via feedback loops.




Figure 9: PIW variant with map iterator.

initiate another iteration of a subworkflow. While such complicated structures achieve the desired effect (here, a special kind of loop), they are hard to understand, design, and maintain. Such ad-hoc constructions also increase the complexity of workflow design while diminishing the overall reusability of workflow components (see (R2) in Section 2.2). Fortunately, there are better ways to incorporate *structured control* into a dataflow-oriented system, thereby directly supporting workflow design as required by (R2).

In [LA03] we have illustrated how higher-order functional programming constructs can be used to improve the design of PIW. In particular, the higher-order function $\text{map} :: (\alpha \rightarrow \beta) \rightarrow [\alpha] \rightarrow [\beta]$ has proven to be very useful to implement a certain type of iteration. It takes a function f (from α to β) and a list of elements of type α , and applies f to each list element, returning the list of result elements (each of type β). Thus map is defined as

$$\text{map } f [x_1, x_2, \dots, x_n] = [f(x_1), f(x_2), \dots, f(x_n)]$$

For example, $\text{map } f [1, 2, 3] = [1, 4, 9]$ for $f(x) = x^2$.

Figure 9 shows an improved version of the PIW workflow from Section 2.1.1 and Figure 2, now using the higher-order map function. Note how backward-directed flows of control-tokens are avoided. Instead, iterations are realized as nested subworkflows inside a higher-order MAP actor. For example, to implement a look-up of a list of gene sequences via a GenBank web service that can only accept one gene at a time, we simply create the higher-order construct $\text{MAP}(\text{GENBANKWS})$ as shown in Figure 9 (the “stack” icon indicates that the contained workflow is applied multiple times).

Other higher-order functional programming constructs, e.g., foldr (for “fold right”) can be similarly used to provide more abstract and modular iteration and control constructs in a dataflow setting, and we plan to add those to KEPLER in the future. The

utility of declarative functional programming methods for dataflow-oriented systems is no coincidence; see, e.g., [Ree95] for more on the close links between dataflow, functional, and visual programming, and [NA01] for interesting applications in implicit parallel programming. Here we only give a simple illustration using a core subworkflow of PIW in a Haskell specification; see [LA03] for details:

```


d0 = $Gid                % input: some gene-id
d1 = genBankG in         % get its gene sequence
d2 = blastP d1          % find candidates from similar seqs
d3 = map genBankP d2    % get promoter sequences
d4 = map promoterRegion d3 % compute regions
d5 = map transfac d4    % compute transcr. factor sites
d6 = zip d2 d4 % create list of (promoter-id,region) pairs
d7 = map gpr2str d6     % accumulate into string list
d8 = concat d7          % create a single file
d9 = putStr d8          % output to subsequent steps

```

The input and output (ports) of this workflow are given by d_0 and d_9 , respectively. Note the use of `map` to iterate over lists where the available services (e.g. `genBankP`) can only handle one item at a time. Also note that these ten equations establish a simple forward-only dataflow process network with the d_i representing named channels, and the expressions on the right of the equation representing processes (i.e., actors). A merge of two parallel branches happens, e.g., through the function `zip` that creates a single stream of pairs (*promoter-id, promoter-region*) in channel d_6 from the two streams in d_2 and d_4 .

4.2 Third Party Transfers

Scientific workflows can involve large volumes of data (see (R3) in Section 2.2). In a web service setting, this creates a problem since so-called 3rd party transfers are not currently supported by web services: Let us consider two web services w_A and w_B , located at two sites s_1 and s_2 , respectively. w_A takes some input x and produces some data d that we would like to pass on to w_B , which produces the final output data y . We can depict this as follows:



Assume that the overall execution of this workflow WF is coordinated and controlled by a workflow engine E (e.g., KEPLER) running at some site s_3 . Current web service implementations do not allow the engine E to call $w_A@s_1$, telling it to route d directly to $w_B@s_2$. Instead, web service invocations and the input/output dataflows that go with them, all go

through $E@s_3$. In pseudo-code this means:

$$WF@s_3(\mathbf{in} \ x, \mathbf{out} \ y) = \{ \\ d@s_3 := w_A@s_1(x@s_3); \\ y@s_3 := w_B@s_2(d@s_3) \}$$

How do we execute the “remote assignments” shown here? To execute $WF@s_3$, the workflow engine E first sends a request message containing x to $w_A@s_1$. Upon completion, w_A replies back to $E@s_3$ with the result d . Now $WF@s_3$ can proceed and E forwards d to s_2 where w_B can work on it. The final result y is then sent from s_2 back to s_3 . This simple call/return execution is quite desirable from a modeling and design point of view since control-flow and dataflow go hand in hand, and since the control engine E does not have to worry about the status of direct (i.e., 3rd party) transfers of data d from w_A to w_B . The downside, however, is that data is moved around more often than necessary. Let us trace the “data shipments” of x , d , and y :

1. **ship** $x@s_3 \rightsquigarrow x@s_1$ % part of request to w_A
2. s_1 **execute** $d := w_A(x)$ % execute w_A
3. **ship** $d@s_1 \rightsquigarrow d@s_3$ % part of reply from w_A
4. **ship** $d@s_3 \rightsquigarrow d@s_2$ % part of request to w_B
5. s_2 **execute** $y := w_B(d)$ % execute w_B
6. **ship** $y@s_2 \rightsquigarrow y@s_1$ % part of reply from w_B

If d is very large, executing both steps (3) and (4) is wasteful: first d is sent from s_1 to s_3 where the workflow engine E runs, only to be sent to s_2 in the next step. Instead of sending d over the wire twice, the more direct 3rd party transfer $w_A@s_1 \xrightarrow{d} w_B@s_2$ moves d only once, but as mentioned before, is not currently supported by web services.¹⁶ The question becomes: How can we avoid unnecessary transfers and achieve the efficiency of 3rd party transfer, while retaining the above simple call/return execution model?

A Handle-Oriented Approach. A simple solution to the above problem is that w_A does not send the actual data d but a *handle* h_d to it. Such a handle corresponds to a “logic pointer” and can be represented by a globally unique URI, but may also be a URL and indicate the protocol by which d is to be accessed, e.g., `http`, `ftp`, `GridFTP00`, `scp`, or `SRB`. If we replace all data occurrences x , d , and y by handles h_x , h_d , and h_y , respectively, we obtain:

1. **ship** $h_x@s_3 \rightsquigarrow h_x@s_1$ % request to w_A

¹⁶And even if it were, “divorces” control-flow and dataflow, resulting in more complex execution models.

2. **@s₁ execute** $h_d := w_A(h_x)$ % execute w_A
3. **ship** $h_d@s_1 \rightsquigarrow h_d@s_3$ % reply from w_A
4. **ship** $h_d@s_3 \rightsquigarrow h_d@s_2$ % request to w_B
5. **@s₂ execute** $h_y := w_B(h_d)$ % execute w_B
6. **ship** $h_y@s_2 \rightsquigarrow h_y@s_1$ % reply from w_B

Now, instead of sending (the possibly very large) d over the wire twice in (3) and (4), we only do so for the (constant size) handle h_d . We cannot hope to further reduce this since a reply message from w_A to E and a new request from E to w_B are necessary for the overall control of workflow execution.

In order to implement the above handle-solution, we need to slightly extend our web services: in steps (2) and (5), w_A and w_B need to process handles by dereferencing them or by creating new ones. The former happens when a web service acts as a *consumer* of data (w_A consumes x), while the latter is needed in the role of a data *producer* (w_A produces d).

Consider, e.g., the case where handles are represented as URLs with **http** as the transport protocol. In step (2) above, w_A needs to dereference h_x before it can execute its function. h_x might be, e.g., **http://foobar.com/f17**. When dereferenced via **http-get** it yields the actual data x .¹⁷ To properly process handles as a data consumer, the operation “receive x ” has to be replaced by “receive h_x ”, followed by a “dereference and get” operation $x := \mathbf{http-get}(h_x)$. All subsequent read operations can then operate on x as before.

In the role of a data producer, we have the reverse situation. We want to avoid shipping of the actual result data d and instead send a handle h_d . Thus, we need to first create this handle, e.g., by creating a new file **f18** that can be accessed via $h_d = \mathbf{http://baz.edu/f18}$. All subsequent write access to d will proceed unchanged, provided the file name **f18** is used for d . Finally, we need to replace “send d ” with “send h_d ”.

We are currently working on extensions of KEPLER that make the system “handle-aware” [Lud04]. For example, during the type-checking phase (Figure 8) a handle-aware director could determine whether two web service actors A and B that invoke the web services w_A and w_B , respectively, support compatible handle types. For this to work seamlessly, web services themselves should offer an actor-oriented programming interface as presented in Section 3.3.

¹⁷Note that while the handle h_x is sent from s_3 to s_1 in step (1), x might actually *not* reside at s_3 .

4.3 Other Research Issues

Higher-order constructs and the handle-approach to 3rd party transfers are only two of a number of pressing research issues in scientific workflows.¹⁸ For example, detached execution (R4), reliability and fault-tolerance (R5), semantic links (R8), and data provenance (R9) are all scientific workflow requirements that need further attention in the future. For example, [BL04] presents some initial work on the use of ontologies as semantic types to help generate data transformation mappings between consecutive workflow steps. These kinds of semantic extensions can help at both levels, at the “plumbing” level to create data transformations as in [BL04], and at the design level to create more reusable components (R2) and to support “smart” links in workflows (R8).

4.4 Related Work

In Section 3 we have described some of the features of KEPLER and the underlying PTOLEMY II system on which KEPLER is based. PTOLEMY II aims at modeling and design of heterogeneous, concurrent systems. In contrast, KEPLER aims at the design and execution of scientific workflows. Consequently, KEPLER extensions to PTOLEMY II include numerous actors and capabilities that facilitate scientific workflows (e.g., web service actors and harvester, GridFTP, SRB and database actors, command-line and secure shell actors, etc.) Additional components are constantly added, e.g., to support statistics packages (such as R), GIS functionality (e.g., Grass and ArcIMS couplings), and other scientific data analysis and visualization capabilities [WPS⁺05].

The research and development on KEPLER also benefits from interactions and collaborations with other groups. On one hand, development is driven by application scientists, the ultimate “customers” of scientific workflow system, on the other hand, work in related projects also influences KEPLER developments. For example, TAVERNA [TAV, OAF⁺04] is a system that focuses on web service-based bioinformatics workflows. In contrast, TRIANA [TRI, CGH⁺05] provides mechanisms for coupling workflows more tightly with Grid middleware tools. Cross-fertilization between these and other projects has happened, e.g., through e-Science LINK-UP workshops [LIN04], meetings and workshops at GGF [GGF04], etc. Other scientific workflow tools include Pegasus [DBG⁺03], Chimera, and job scheduling tools such as Condor/G [DTL04] and Nimrod/G [AGK00]. For a taxonomy of workflow management

¹⁸Addressing (R2) and (R3), respectively.

systems for Grid computing and a comparison of systems see [YB05]. Future work will address the various outstanding research issues and workflows requirements that have not yet been (fully) met. For example, some projects contributing to KEPLER plan to provide couplings to highly-interactive visualization tools such as SCIRun [WPS⁺05] and GeoVista [TG02].

5 Conclusions

We have provided an overview of scientific workflow management issues, motivated by real-world examples that we encountered in a number of application-oriented projects. The spectrum of what can be called a scientific workflow is wide and includes scientific discovery workflows (e.g., Section 2.1.1), workflows that automate manual procedures or reengineer custom tools (e.g., Section 2.1.2), and data and compute-intensive workflows (e.g., Section 2.1.3). Scientific workflow support is needed for practically all information-oriented scientific disciplines, including bioinformatics, cheminformatics, ecoinformatics, geoinformatics, physics, etc. We identified a number of common requirements and desiderata of scientific workflows (Section 2.2) and contrasted them with business workflows.

The KEPLER system addresses many of the core requirements (Section 3) and provides support for web service-based workflows and Grid extensions. The source code of KEPLER is freely available [KEP] and a first alpha-release has been distributed earlier this year. A unique feature of KEPLER is inherited from the underlying PTOLEMY II system: the actor-oriented modeling approach. This approach facilitates modeling and design of complex systems and thus provides also a very promising direction for pressing problems such as web service composition and orchestration. The way data polymorphism and behavioral polymorphism are supported by an actor-oriented approach that “concentrates” component interaction in a separate *director* entity, can also shed light on other efforts to create reusable component architectures such as CCA [AGG⁺99]. Areas of research include modeling issues such as the use of higher-order functional constructs for workflow design (Section 4.1), and optimization issues such as the use of virtual data references (handles) to facilitate data-intensive, web service-based workflows (Section 4.2).

Acknowledgements. KEPLER is an open source, cross-project collaboration that would not exist with-

out the contributions of the many team members. We thank all current and past contributors to PTOLEMY II – the KEPLER systems would not be possible without them. We also thank all KEPLER members for their contributions, in particular, Tobin Fricke for implementing actors that access the wonderful world of ROADNet real-time data streams, Steve Neuendorffer and Christopher Brooks for sharing their insights into PTOLEMY II, Rod Spears for QBE facilities, Xiaowen Xin for many contributions including to the PIW workflow, Zhengang Cheng for providing some of the first web service actors, Werner Krebs for EOL extensions, Steve Mock for Globus actors, Shawn Bowers for his work on semantic types for KEPLER, and last not least the many scientists and PIs that provide direct or indirect support to this effort, among them Bill Michener, Chaitan Baru, Kim Baldrige, Mark Miller, Arie Shoshani, Terence Critchlow, and Mladen Vouk.

References

- [ABB⁺03] I. Altintas, S. Bhagwanani, D. Butler, S. Chandra, Z. Cheng, M. Coleman, T. Critchlow, A. Gupta, W. Han, L. Liu, B. Ludäscher, C. Pu, R. Moore, A. Shoshani, and M. Vouk. A Modeling and Execution Environment for Distributed Scientific Workflows. In *15th Intl. Conf. on Scientific and Statistical Database Management (SSDBM)*, Boston, Massachusetts, 2003.
- [AGG⁺99] R. Armstrong, D. Gannon, A. Geist, K. Keahey, S. Kohn, L. McInnes, S. Parker, and B. Smolinski. Toward a Common Component Architecture for High-Performance Scientific Computing. In *8th IEEE Intl. Symposium on High Performance Distributed Computation*, August 1999.
- [AGK00] D. Abramson, J. Giddy, and L. Kotler. High Performance Parametric Modeling with Nimrod/G: Killer Application for the Global Grid. In *Intl. Parallel and Distributed Processing Symposium (IPDPS)*, Cancun, Mexico, May 2000. <http://www.csse.monash.edu.au/~david/nimrod/>.
- [AIL98] A. Ailamaki, Y. E. Ioannidis, and M. Livny. Scientific Workflow Management by Database Management. In *10th Intl. Conf. on Scientific and Statistical Database Management (SSDBM)*, Capri, Italy, 1998.
- [AM97] G. Alonso and C. Mohan. Workflow Management Systems: The Next Generation of Distributed Processing Tools. In

- S. Jajodia and L. Kerschberg, editors, *Advanced Transaction Models and Architectures*. 1997. [DTL04]
- [BDG03] J. Blythe, E. Deelman, and Y. Gil. Planning for workflow construction and maintenance on the Grid. In ICAPS [ICA03]. [EJL⁺03]
- [BIR] Biomedical Informatics Research Network Coordinating Center (BIRN-CC), University of California, San Diego. <http://nbirn.net/>.
- [BL04] S. Bowers and B. Ludäscher. An Ontology Driven Framework for Data Transformation in Scientific Workflows. In *International Workshop on Data Integration in the Life Sciences (DILS)*, LNCS 2994, Leipzig, Germany, March 2004. . [eSc03]
- [BL05] S. Bowers and B. Ludäscher. Actor-Oriented Design of Scientific Workflows. submitted for publication, 2005. [eSc04]
- [BLL04a] S. Bowers, K. Lin, and B. Ludäscher. On Integrating Scientific Resources through Semantic Registration. In *16th Intl. Conf. on Scientific and Statistical Database Management (SSDBM)*, Santorini Island, Greece, 2004. [GEO]
- [BLL⁺04b] C. Brooks, E. A. Lee, X. Liu, S. Neuenendorffer, Y. Zhao, and H. Zheng. Heterogeneous Concurrent Modeling and Design in Java (Volumes 1-3). Technical report, Dept. of EECS, University of California, Berkeley, 2004. Technical Memoranda UCB/ERL M04/27, M04/16, M04/17. [GGF04]
- [CGH⁺05] D. Churches, G. Gombas, A. Harrison, J. Maassen, C. Robinson, M. Shields, I. Taylor, and I. Wang. Programming Scientific and Distributed Workflow with Triana Services. *Concurrency and Computation: Practice and Experience. Special Issue on Scientific Workflows*, 2005. [Glo]
- [CGK⁺02] F. Curbera, Y. Golland, J. Klein, F. Leyman, D. Roller, S. Thatte, and S. Weerawarana. Business Process Execution Language for Web Services (BPEL4WS), Version 1.0, 2002. <http://www.ibm.com/developerworks/library/ws-bpel/>. [GRI04]
- [CM95] I. Chen and V. Markowitz. The Object-Protocol Model: Design, Implementation, and Scientific Applications. *ACM Transactions on Information Systems*, 20(5), 1995. [GT98]
- [DBG⁺03] E. Deelman, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, K. Vahi, K. Blackburn, A. Lazzarini, A. Arbree, R. Cavanaugh, and S. Koranda. Mapping Abstract Complex Workflows onto Grid Environments. *Journal of Grid Computing*, 1(1):25–39, 2003. [Hal01]
- [ICA03] Proceedings of the ICAPS Workshop on Planning for Web Services, Trento, Italy, June 2003. [KEP]
- [Kie02] B. Kiepuszewski. *Expressiveness and Suitability of Languages for Control Flow Modelling in Workflows*. PhD thesis, Queensland University of Technology, 2002. [KM77]
- [LA03] B. Ludäscher and I. Altintas. On Providing Declarative Design and Programming Constructs for Scientific Workflows. T. T. Douglas Thain and M. Livny. Distributed Computing in Practice: The Condor Experience. *Concurrency and Computation: Practice and Experience*, 2004. J. Eker, J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuenendorffer, S. Sachs, and Y. Xiong. Taming Heterogeneity – the Ptolemy Approach. In *Proceedings of the IEEE*, volume 91(1), January 2003. *e-Science Workflow Services Workshop*, e-Science Institute, Edinburgh, Scotland, December 2003. <http://www.nesc.ac.uk/esi/events/303/index.html>. *e-Science Grid Environments Workshop*, e-Science Institute, Edinburgh, Scotland, May 2004. <http://www.nesc.ac.uk/esi/events/>. NSF/ITR: GEON: A Research Project to Create Cyberinfrastructure for the Geosciences. www.geogrid.org. *Workflow in Grid Systems Workshop*, GGF10, Berlin, Germany, March 2004. <http://www.extreme.indiana.edu/groc/Worflow-call.html>. The Globus Alliance. www.globus.org. *GRIST Workshop on Service Composition for Data Exploration in the Virtual Observatory*, California Institute of Technology, July 2004. <http://grist.caltech.edu/sc4devo/>. J. Gao and M. A. Thompson, editors. *Combined Quantum Mechanical and Molecular Mechanical Methods*. American Chemical Society, 1998. A. Halevy. Answering Queries Using Views: A Survey. *VLDB Journal*, 10(4):270–294, 2001. Proceedings of the ICAPS Workshop on Planning for Web Services, Trento, Italy, June 2003. KEPLER: A System for Scientific Workflows. <http://kepler-project.org>. B. Kiepuszewski. *Expressiveness and Suitability of Languages for Control Flow Modelling in Workflows*. PhD thesis, Queensland University of Technology, 2002. G. Kahn and D. B. MacQueen. Coroutines and Networks of Parallel Processes. In B. Gilchrist, editor, *Proc. of the IFIP Congress 77*, pp. 993–998, 1977. B. Ludäscher and I. Altintas. On Providing Declarative Design and Programming Constructs for Scientific Workflows

- based on Process Networks. Technical Report SciDAC-SPA-TN-2003-01, San Diego Supercomputer Center, 2003. <http://kbi.sdsc.edu/SciDAC-SDM/scidac-tn-map-constructs.pdf>.
- [LAG03] B. Ludäscher, I. Altintas, and A. Gupta. Compiling Abstract Scientific Workflows into Web Service Workflows. In *15th Intl. Conf. on Scientific and Statistical Database Management (SS-DBM)*, Boston, Massachusetts, 2003. <http://kbis.sdsc.edu/SciDAC-SDM/ludaescher-compiling.pdf>.
- [LGM03] B. Ludäscher, A. Gupta, and M. E. Martone. A Model-Based Mediator System for Scientific Data Management. In Z. Lacroix and T. Critchlow, editors, *Bioinformatics: Managing Scientific Data*. Morgan Kaufmann, 2003.
- [LIN04] *LINK-Up Workshop on Scientific Workflows*, San Diego Supercomputer Center, October 2004. <http://kbis.sdsc.edu/events/link-up-11-04/>.
- [LN04] B. Ludäscher and A. Nash. Web Service Composition Through Declarative Queries: The Case of Conjunctive Queries with Union and Negation. In *20th Intl. Conf. on Data Engineering (ICDE)*, 2004.
- [LP95] E. A. Lee and T. Parks. Dataflow Process Networks. *Proceedings of the IEEE*, 83(5):773–799, May 1995. <http://citeseer.nj.nec.com/455847.html>.
- [LPH01] L. Liu, C. Pu, and W. Han. An XML-Enabled Data Extraction Tool for Web Sources. *Intl. Journal of Information Systems, Special Issue on Data Extraction, Cleaning, and Reconciliation*, 2001.
- [Lud04] B. Ludäscher. Towards Actor-Oriented Web Service-Based Scientific Workflows (or: How to Handle Handles). Technical report, San Diego Supercomputer Center, September 2004.
- [MBJ⁺04] W. K. Michener, J. H. Beach, M. B. Jones, B. Ludäscher, D. D. Pennington, R. S. Pereira, A. Rajasekar, and M. Schildhauer. A Knowledge Environment for the Biodiversity and Ecological Sciences. *Journal of Intelligent Information Systems*, 2004. to appear.
- [McP05] T. M. McPhillips. Pipelined scientific workflows for inferring evolutionary relationships. Natural Diversity Discovery Project, 2005. manuscript.
- [Mor94] J. P. Morrison. *Flow-Based Programming – A New Approach to Application Development*. Van Nostrand Reinhold, 1994.
- [MVW96] J. Meidanis, G. Vossen, and M. Weske. Using Workflow Management in DNA Sequencing. In *Intl. Conf. on Cooperative Information Systems (CoopIS)*, 1996.
- [NA01] R. S. Nikhil and Arvind. *Implicit Parallel Programming in pH*. Morgan Kaufmann, 2001.
- [NCB04] National Center for Biotechnology Information (NCBI). <http://www.ncbi.nlm.nih.gov/>, 2004.
- [NL04] A. Nash and B. Ludäscher. Processing Unions of Conjunctive Queries with Negation under Limited Access Patterns. In *9th Intl. Conf. on Extending Database Technology (EDBT)*, LNCS 2992, pp. 422–440, Heraklion, Crete, Greece, 2004.
- [OAF⁺04] T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M. R. Pocock, A. Wipat, and P. Li. Taverna: A tool for the composition and enactment of bioinformatics workflows. *Bioinformatics Journal*, 20(17):3045–3054, 2004.
- [OWL03] OWL Web Ontology Language Reference, W3C Proposed Recommendation, December 2003. www.w3.org/TR/owl-ref/.
- [Pro00] G. Project. GridFTP – Universal Data Transfer for the Grid, 2000. see <http://www.globus.org/datagrid/gridftp.html>.
- [PTO04] PTOLEMY II project and system. Department of EECS, UC Berkeley, 2004. <http://ptolemy.eecs.berkeley.edu/ptolemyII/>.
- [PYN⁺03] L. Peterson, E. Yin, D. Nelson, I. Altintas, B. Ludäscher, T. Critchlow, A. J. Wyrobek, and M. A. Coleman. Mining the Frequency Distribution of Transcription Factor Binding Sites of Ionizing Radiation Responsive Genes. In *New Horizons in Genomics, DOE/SC-0071*, Santa Fe, New Mexico., March 30–April 1 2003.
- [R] R – Statistical Data Analysis. <http://www.r-project.org>.
- [Ree95] H. J. Reekie. *Realtime Signal Processing: Dataflow, Visual, and Functional Programming*. PhD thesis, School of Electrical Engineering, University of Technology, Sydney, 1995.
- [ROA] ROADNet: Real-time Observatories, Applications and Data management Network. roadnet.ucsd.edu.
- [SBA⁺04] W. Sudholt, K. Baldrige, D. Abramson, C. Enticott, and S. Garic. Parameter Scan of an Effective Group Difference

- Pseudopotential Using Grid Computing. *New Generation Computing*, 22:137–146, 2004. [WBB96]
- [SBB⁺93] M. Schmidt, K. Baldrige, J. Boatz, S. Elbert, M. Gordon, J. Jensen, S. Koseki, N. Matsunaga, K. Nguyen, S. Su, T. Windus, M. Dupuis, and J. Montgomery. The General Atomic and Molecular Electronic Structure System. *Journal of Computational Chemistry*, 14:1347–1363, 1993. cf. <http://www.msg.ameslab.gov/GAMESS/GAMESS.html>.
- [SDM] Scientific Data Management Center (SDM). <http://sdm.lbl.gov/sdmcenter/>, see also <http://www.npaci.edu/online/v5.17/scidac.html>.
- [SDM03] *Scientific Data Management Framework Workshop*, Argonne National Labs, August 2003. <http://sdm.lbl.gov/~arie/sdm/SDM.Framework.wshp.htm>.
- [SEE] NSF/ITR: Enabling the Science Environment for Ecological Knowledge (SEEK). seek.ecoinformatics.org.
- [She98] A. Sheth. Changing Focus on Interoperability in Information Systems: From System, Syntax, Structure to Semantics. In M. Goodchild, M. Egenhofer, R. Fegeas, and C. Kottman, editors, *Interoperating Geographic Information Systems*, pp. 5–30. Kluwer, 1998.
- [SRB] SDSC Storage Resource Broker. <http://www.npaci.edu/DICE/SRB/>.
- [TAV] The Taverna Project. <http://taverna.sf.net/>.
- [TG02] M. Takatuska and M. Gahegan. Geovista Studio: A codeless visual programming environment for geoscientific data analysis and visualization. *Computers and Geosciences*, 28(2):1131–1144, 2002.
- [TRI] The Triana Project. <http://www.trianacode.org/>.
- [vdA03] W. van der Aalst. Don't go with the flow: Web services composition standards exposed. *IEEE Intelligent Systems. Web Services – Been there done that? Trends & Controversies*, Jan/Feb 2003. <http://tmitwww.tm.tue.nl/research/patterns/download/ieeewebflow.pdf>.
- [vdAtHKB03] W. van der Aalst, A. ter Hofstede, B. Kiepuszewski, and A. Barros. Workflow Patterns. *Distributed and Parallel Databases*, 14(3):5–51, July 2003.
- [vdAvH02] W. van der Aalst and K. van Hee. *Workflow Management: Models, Methods, and Systems (Cooperative Information Systems)*. MIT Press, 2002.
- H. Wright, K. Brodlie, and M. Brown. The Dataflow Visualization Pipeline as a Problem Solving Environment. In M. Göbel, J. David, P. Slavik, and J. J. van Wijk, editors, *Virtual Environments and Scientific Visualization*, pp. 267–276. Springer, 1996.
- [Wer01] T. Werner. Target gene identification from expression array data by promoter analysis. *Biomolecular Engineering*, 17:87–94, 2001.
- [WPS⁺05] D. Weinstein, S. Parker, J. Simpson, K. Zimmerman, and G. Jones. Visualization in the SCIRun Problem-Solving Environment. In C. Hansen and C. Johnson, editors, *Visualization Handbook*, pp. 615–632. Elsevier, 2005.
- [WSD03] Web Services Description Language (WSDL) Version 1.2. <http://www.w3.org/TR/wsdl12>, June 2003.
- [YB05] J. Yu and R. Buyya. A Taxonomy of Workflow Management Systems for Grid Computing. Technical Report GRIDS-TR-2005-1, Grid Computing and Distributed Systems Laboratory, University of Melbourne, 2005. <http://www.gridbus.org/reports/GridWorkflowTaxonomy.pdf>.
- [zM04] M. zur Muehlen. *Workflow-based Process Controlling*. Logos Verlag, Berlin, 2004.