



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Structured Composition of Dataflow and Control-Flow for Reusable and Robust Scientific Workflows

S. Bowers, B. Ludaescher, A. Ngu,
T. Critchlow

September 9, 2005

Symposium on Applied Computing
Dijon , France
April 23, 2006 through April 27, 2006

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

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Structured Composition of Dataflow and Control-Flow for Reusable and Robust Scientific Workflows

ABSTRACT

Data-centric scientific workflows are often modeled as dataflow process networks. The simplicity of the dataflow framework facilitates workflow design, analysis, and optimization. However, some workflow tasks are particularly “control-flow intensive”, *e.g.*, procedures to make workflows more fault-tolerant and adaptive in an unreliable, distributed computing environment. Modeling complex control-flow directly within a dataflow framework often leads to overly complicated workflows that are hard to comprehend, reuse, schedule, and maintain. In this paper, we develop a framework that allows a *structured* embedding of control-flow intensive subtasks within dataflow process networks. In this way, we can seamlessly handle complex control-flows without sacrificing the benefits of dataflow. We build upon a flexible actor-oriented modeling and design approach and extend it with (actor) frames and (workflow) templates. A *frame* is a placeholder for an (existing or planned) collection of components with similar function and signature. A *template* partially specifies the behavior of a sub-workflow by leaving “holes” (*i.e.*, frames) in the subworkflow definition. Taken together, these abstraction mechanisms facilitate the separation and structured re-combination of control-flow and dataflow in scientific workflow applications. We illustrate our approach with a real-world scientific workflow from the astrophysics domain. This data-intensive workflow requires remote execution and file transfer in a semi-reliable environment. For such workflows, we propose a 3-layered architecture: The top-level, typically a dataflow process network, includes Generic Data Transfer (GDT) frames and Generic remote eXecution (GX) frames. At the second level, the user can specialize the behavior of these generic components by embedding a suitable template (here: transducer templates for control-flow intensive tasks). At the third level, frames inside the transducer template are specialized by embedding the desired implementation. Our approach yields workflows that are more robust (fault-tolerance strategies can be define by control-flow driven transducer templates) and at the same time more reuseable, since the embedding of frames and templates yields more structured and modular workflows.

1. INTRODUCTION

Scientific workflow systems [19, 17, 2, 16] are increasingly being used by scientists to construct and execute complex scientific analyses. Such analyses are typically data-centric and involve “gluing” together data retrieval, computation, and visualization components into a single executable analysis pipeline. The components may be part of the workflow system, part of another application (invoked through system calls, executing R or MATLAB scripts, etc.), or even external, accessed via web or grid services. In addition to providing scientists with a mechanism to compose and configure other-

wise heterogeneous components, scientific workflow systems aim to support end-to-end workflow management, *e.g.*, through tools for access and querying of external data sources, archival of intermediate results, and monitoring of workflow execution. Most scientific workflow systems treat workflows as *dataflow process networks* [13], a model of computation that comes with “built-in” support for stream-based and concurrent execution. Thus, dataflow is a natural paradigm for data-driven and data-intensive scientific workflows such as, *e.g.*, the terabyte-sized Fusion Plasma Simulation [3] and the Terascale Supernova Initiative [21]. It can be efficiently analysed and scheduled, and is also a simple and intuitive model for workflow designers [4]. While dataflow has become the standard model of scientific workflows, some amount of control-flow modeling is often necessary for engineering fault-tolerant, robust, and adaptive workflows.

In this paper, we address the problem of combining dataflow and control-flow for scientific workflows. It has been noted [14] that modeling control-flow using *only* dataflow constructs can quickly lead to overly complex workflows that are hard to understand, reuse, reconfigure, maintain, and schedule [11]. In particular, modeling control-flow using dataflow involves inserting and linking various low-level and specialized control components alongside dataflow components, thus making it difficult to distinguish control-flow from dataflow aspects (since they are “entangled”) and often requiring complex component connections including loops [21].

The organization and contributions of the paper are as follows: We describe a framework that “untangles” dataflow and control-flow aspects and instead supports a structured embedding of control-intensive subtasks within dataflow process networks (Section 2). Our approach is to encapsulate generic behavioral specifications (*i.e.*, control-flow) in workflow *templates*. Templates are distinct and separate components and thus can be easily reused in other workflows. Templates are partial specifications and contain “holes”, so-called *frames*, that act as placeholders for independently defined subcomponents. Composing templates with existing dataflow components results in applying the associated behavior to the component in such a way that the separation between control-flow and dataflow is maintained, thus allowing the underlying dataflow component to be easily changed (typically through a configuration parameter of the template). This approach allows workflow designers to change complex control-flow behavior, by simply using different templates. Our approach was inspired by the notion of hierarchical finite state machines [8] and can also be seen as an extension of *actor-oriented modeling* [12] with *frames* and *templates* (Figure 2).

In Section 3 we first present a specialized 3-layered architecture of our framework. It allows the designer to select and reuse different control-flow intensive behaviors for generic top-level components, by embedding inside of them suitable transducer templates

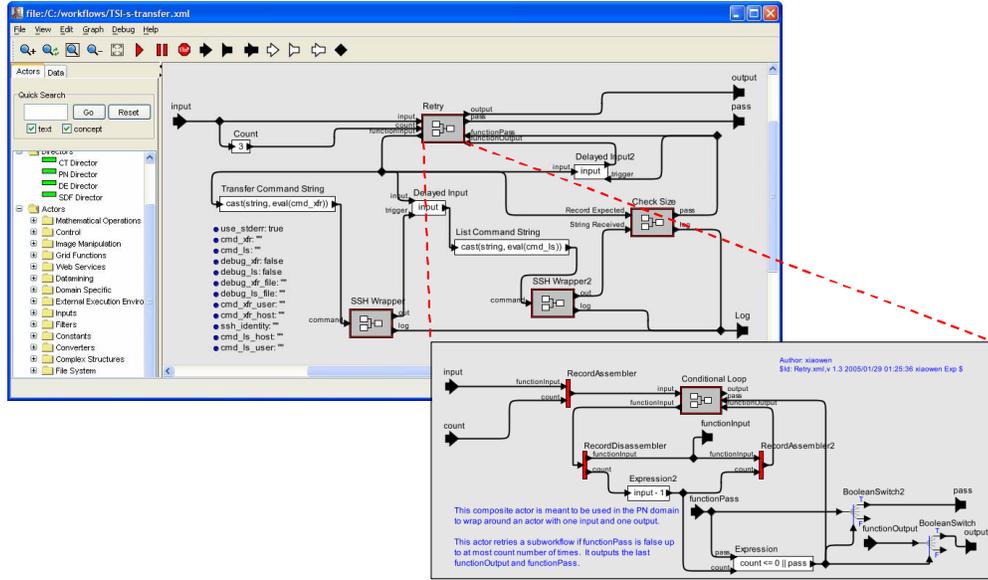


Figure 1: Control-flow intensive astrophysics workflow in KEPLER [21]. “Retry”, a composite actor for fault-tolerant data transfer (top), contains a subworkflow (bottom), which itself contains a “ConditionalLoop” subworkflow (inside not shown). Complex feedback loops and the use of boolean switches illustrate the complexity of modeling control-flow directly in a dataflow process network.

as the middle-layer. The concrete implementation of frames inside of transducer templates is independently selected via the bottom-layer. We then describe a Generic Data Transfer (GDT) component, which is part of our prototypical implementation on top of the open-source KEPLER system [2, 15], an extension of PTOLEMY II [5] for scientific workflows. GDT was motivated by earlier work on a control-intensive astrophysics workflow [21]. As shown in Figure 1, this workflow uses dataflow constructs to implement a fault-tolerance scheme (involving “retry”) for transferring files, resulting in a very complex process network. In the new approach, the GDT component encapsulates this and other transfer behaviors as templates in which workflow designers can select from a set of behaviors as well as the desired underlying transfer protocols (e.g., scp or ftp). Given a particular behavior and protocol, the GDT automatically composes these into the desired executable component. At any time, the behavior and the underlying protocols can be easily changed by simply *reconfiguring* GDT. In the original workflow this would be a complex and error-prone *programming task*, involving the insertion, deletion, and re-wiring of various control-flow and dataflow components. We also describe a Generic remote eXecution (GX) component, whose middle-layer employs exactly the same control-intensive behaviors (via transducer templates) as GDT to support fault-tolerance, demonstrating the versatility of our approach and the improved component reusability it creates.

2. ACTOR-ORIENTED DESIGN EXTENSIONS

In KEPLER, users develop workflows by selecting appropriate components called *actors* and placing them on a design canvas, after which they can be “wired” together to form the desired workflow graph (cf. Figure 1). Actors have input and output ports which provide the communication interface to other actors. Workflows can be hierarchically structured, yielding *composite actors* that encapsulate subworkflows (e.g., see the bottom-right in Figure 1). A novel feature of KEPLER, inherited from PTOLEMY II, is that the overall execution and component interaction semantics of a workflow

is *not* defined by the components, but is factored out into a separate component called a *director* (not shown here). Taken together, workflows, actors, ports, connections, and directors represent the basic building blocks of *actor-oriented modeling and design* [12].

In this section we define scientific workflows as dataflow process networks and describe two extensions to actor-oriented modeling, i.e., frames and templates. Frames form the basis of our approach for embedding control-flow intensive behaviors (via *workflow templates*) inside of dataflow process networks.

2.1 Scientific Workflows as Process Networks

An actor-oriented *workflow graph* $W = \langle \mathbf{A}, \mathbf{D} \rangle$ consists of a set \mathbf{A} of *actors* representing components or tasks and a set of directed dataflow connections \mathbf{D} (see below), representing communication channels that connect actors via ports, and along which actors communicate by passing *tokens*.

Let $\text{ports}(A)$ denote the set of *ports* of actor A . Each port $p \in \text{ports}(A)$ is designated as either *input* or *output*. Some input ports may be distinguished as *parameters* $\text{pars}(A) \subseteq \text{in}(A)$ which can be used for configuring A ’s behavior. For convenience, we write $A.p$ to emphasize that port p belongs to actor A . The *signature* Σ_A of an actor is given by its ports; we write $\Sigma_A = \text{in}(A) \rightarrow \text{out}(A)$.

Actors are wired together through their ports via dataflow connections. A *dataflow connection* $d \in \mathbf{D}$ is a directed hyperedge $d = \langle \mathbf{o}, \mathbf{i} \rangle$, connecting n output ports $\mathbf{o} = \{o_1, \dots, o_n\}$ with m input ports $\mathbf{i} = \{i_1, \dots, i_m\}$. A dataflow connection $d = \langle \mathbf{o}, \mathbf{i} \rangle$ corresponds to a *merge* step of output tokens from \mathbf{o} , followed by a *copy* step, delivering all tokens to the input ports \mathbf{i} .

A *composite actor* A_W encapsulates a *sub-workflow* W . The (external) ports of A_W consist of a distinguished set of ports from W , i.e., A_W might not expose all of its subworkflow’s ports. A *hierarchical workflow* is a workflow graph that contains at least one composite actor. Since subworkflows can themselves be hierarchical, any level of nesting can be modeled.

A port p may have a *structural data type* constraining the al-

lowed set of values accepted by p (if p is an input port) or produced by p (if p is an output port). The PTOLEMY II type system includes simple types (e.g., string and int) and complex types (such as nested record and list structures). A port p may also have a *semantic type*, denoting a concept from a description logic ontology [4]. For example, a semantic type is:

MEASUREMENT \sqcap \forall ITEM.SPECIESOCCURRENCE

indicating that the corresponding port accepts (or produces) data tokens that are measurements where the measured item is a species occurrence (as opposed to, e.g., a temperature). In addition to port semantic types, an actor A itself may also be associated with a semantic type, describing the overall function or purpose of A . While structural (i.e., data) type safety ensures that actors can “work with” incoming data tokens at runtime, semantic type safety avoids actor connections at design time that are not meaningful in terms of their concept annotations (e.g., occurrence data cannot be used where temperature data is expected).

So far, the execution semantics of a workflow graph W has not been specified. Indeed, in PTOLEMY II and thus in KEPLER, the workflow designer can choose among different models of computation, each one being represented by a so-called *director*. A director specifies and mediates all inter-actor communication, separating workflow orchestration and scheduling (the director’s concern) from individual actor execution (the actor’s concern). This separation achieves a form of *behavioral polymorphism* [12], resulting in more reusable actor components and subworkflows. KEPLER (through PTOLEMY II) provides a variety of directors that implement process network (PN and SDF), discrete event (DE), continuous time (CT), and finite state transducer (FST) semantics.

2.2 Frames

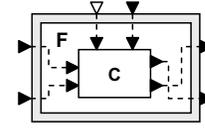
Actors in actor-oriented modeling and design are always *concrete*: they correspond to particular implementations and can be directly executed in a workflow. We extend actor-oriented modeling with a new entity called *frame*, which is an abstraction that denotes a set of alternative actor implementations (or templates) with similar, but not necessarily identical functionality. For workflow designers, frames are placeholders for components that will be instantiated and specialized later. Thus, a designer can place a frame F on the design canvas, and connect it with other workflow components, without prematurely specifying which component C is to be used. For component developers, frames can be used as abstractions for a family of components (actors or templates) with similar function.

Formally, a frame is a named entity F that acts as a placeholder for a component C to be “plugged into” F (see Figure 2 a). When devising a frame F , a family of components \mathbf{C}_F is envisioned, with each $C \in \mathbf{C}_F$ being a possible alternative for embedding into F . Like an actor, a frame has input, output, and parameter ports, structural types, and semantic types; taken together they form the *frame signature* Σ_F . This signature represents the common API of the family \mathbf{C}_F of components that F abstracts.

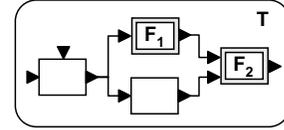
An *embedding* $F[C]$ of a component C into a frame F is a set of pairs associating (or “wiring”) ports of C with ports of F , i.e., $F[C] \subseteq \text{ports}(F) \times \text{ports}(C)$. We indicate the wiring type of a pair $(x, y) \in F[C]$ as follows:

- $F.x \blacktriangleright C.y$; if $x \in \text{in}(F), y \in \text{in}(C)$ (input)
- $F.x \blacktriangleleft C.y$; if $x \in \text{out}(F), y \in \text{out}(C)$ (output)
- $F.x \blacktriangledown C.y$; if $x \in \text{pars}(F), y \in \text{pars}(C)$ (parameter)

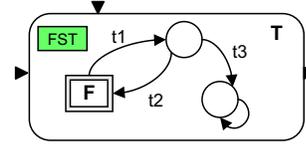
The embedded component C may also introduce new ports not in $\text{ports}(F)$. We denote these ports as $\triangleright C.y$, $\triangleleft C.y$, and $\triangledown C.y$ for input, output, and parameter ports y , respectively.



a). Embedding $F[C]$



b). Workflow Template $T(F_1, F_2)$



c). Transducer Template $T(F)$

Figure 2: a) Embedding of component C in frame F ; b) workflow template $T(F_1, F_2)$; c) finite state transducer template $T(F)$.

Similarly, an embedding $F[C]$ may not use all the ports of C . We denote these unused ports as $F.x \triangleleft$, $F.x \triangleright$, and $F.x \triangledown$ for input, output, and parameter ports x , respectively. We note that parameter ports $F.x$ can also be connected to input ports $C.y$ and vice versa. However, other connection types $(x, y) \in F[C]$ are not allowed. More precisely, an embedding $F[C]$ is *well-formed* if the input and output port directions are observed, i.e., F ’s inputs (outputs) are wired only to inputs (outputs) of C (Figure 2 a). A well-formed embedding $F[C]$ is *structurally well-typed* if the structural types align, and *semantically well-typed* if the semantic types align:

More precisely, we require for each connection between a port of F and a port of C having (structural or semantic) types τ_F and τ_C that: (1) $\tau_F \preceq \tau_C$ for input structural types; (2) $\tau_C \preceq \tau_F$ for output structural types; (3) $\tau_F \sqsubseteq \tau_C$ for input semantic types; and (4) $\tau_C \sqsubseteq \tau_F$ for output semantic types.¹ Thus, we use contravariant subtyping for both structural and semantic types: when embedding a component C in a frame F , C should be able to handle F ’s inputs. Conversely, F should be able to handle outputs of C (or, equivalently, C should not produce output that is more general than what F anticipates). The signature $\Sigma_{F[C]}$ after embedding C in F includes (unless specified otherwise by the designer) the unused ports of F plus the new ports introduced by C .

When a workflow designer chooses a component C to embed within a frame F , we can use the port types of C and F to semi-automatically compute the appropriate connections for $F[C]$. In addition, component types can be used to help workflow designers search repositories for plausible components to be embedded within a given frame.

2.3 Workflow Templates

A frame F imposes some constraints on the set \mathbf{C}_F of components for which it stands. In particular, embeddings $F[C]$ should be well-formed and well-typed for any $C \in \mathbf{C}_F$ as explained above. However, no assumptions can be made about the “inner workings” of C . A *workflow template* T provides a similar level of abstraction

¹“ \preceq ” denotes the standard subtyping relation between data types, while “ \sqsubseteq ” denotes concept subsumption in description logics.

for a set of workflows \mathbf{W}_T . Unlike a frame, however, a template T (partially) specifies the behavior of the workflows it represents.

Like actors and frames, a template T has the usual port signature $\Sigma_T : \text{in}(T) \rightarrow \text{out}(T)$. In addition, a template includes an “inner” workflow graph W_T , where some of the components of W_T are not concrete actors, but frames (Figure 2 b). Let F_1, \dots, F_n be the frames that occur in W_T , either directly, or indirectly through nested templates. Then we can view T as a partial workflow specification $T(F_1, \dots, F_n)$, whose frames F_i can be independently specialized by embedded components (actors or templates) C_i . The resulting embedding $T(F_1[C_1], \dots, F_n[C_n])$ is a concrete, executable workflow if no C_i has itself a frame; otherwise the embedding is a (more refined) template.

In addition to providing input/output constraints through the port signature Σ_T and behavioral constraints through the workflow graph structure W_T (with frames acting as placeholders), a template T can also constrain the intended model(s) of computation by providing one or more directors: In Figure 2 c, a *transducer template* $T(F)$ is shown. This template includes a workflow graph W_T with a frame F . Moreover, an FST director is inscribed in T , meaning that the workflow graph is to be executed as a finite state transducer². A director dictates the execution model of a workflow graph W_T (e.g., SDF or PN for synchronous dataflow and process network execution, respectively; or here: FST), and may also impose constraints on the graph structure. In the case of FST, nodes (components) are not called actors but *states* (depicted as circles in Figure 2 c); connections are called *state transitions* (depicted as curved arcs). In response to a state transition, the FST director calls a *state implementation* if one has been associated with the state [8, 12]. In our case, we can create a more generic behavior for the finite state transducer by delaying the specification of a concrete actor to implement a state, and instead introducing a frame. In this way the same control-flow driven behavior can be reused with different underlying state implementations.

3. GENERIC DATA TRANSFER

Here we consider a particular design pattern³ for structuring frames and templates into generic workflow components that can be executed using alternative control behaviors and alternative task implementations. We define the Generic Data Transfer (GDT) and Generic eXecution (GX) components using this pattern. We also describe how the GDT and GX components are implemented within KEPLER.

3.1 A Generic Control-Flow Component Pattern

The generic control-flow component pattern consists of three levels, as shown in Figure 3. The top level is represented as a frame within a dataflow graph and denotes a particular task (e.g., data transfer or remote execution). This top-level frame can be embedded with one of many finite state transducer templates (the middle level), each of which defines a control-flow behavior for the task. A transducer template has one or more *state frames* that can be embedded with a particular task implementation (e.g., scp or ssh). The various frame implementations form the bottom-level of the pattern.

We use finite state transducers for modeling embedded control-flow because they offer a more natural, intuitive, and typically more succinct language for specifying control behavior, compared to dataflow process networks. Finite state machines (or transducers) are

²a kind of finite state machine that not only consumes input tokens but that also produces output tokens

³similar in spirit to software design patterns [7]

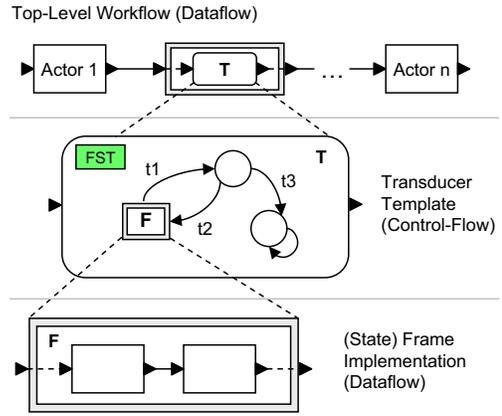


Figure 3: A pattern for modeling generic control-flow components that consists of an outer frame (top), a nested transducer template (middle), and state-frame embeddings (bottom)

often used to model business workflows [1], which are primarily control-flow oriented (as opposed to dataflow oriented), and underpin many of the web-service orchestration languages [6].

We define a finite state transducer (FST) in the normal way: An FST is a tuple $M = \langle I, O, Q, q_0, T \rangle$, where I and O are sets of input and output events, respectively, Q is a finite set of states, q_0 is the initial state, and T is a finite set of transitions, each of which has the form $t : q \xrightarrow{c/a} q'$. Here, c is an optional *condition* that guards the transition t (i.e., t can only be executed if c is true), and a is an optional *action*. The FST M starts in the initial state. When M is “called” from the outside, it transitions from the current state q to the next state q' , based on the current input events I and the conditions of transitions emanating from q . In addition, we consider FST states that can be associated with a subworkflow (called *state refinements* in PTOLEMY II [8]), where the subworkflow is executed upon entry into the state.

Components that implement this generic control-flow pattern enable workflow designers to easily configure both the behavior and underlying implementation of the component. A workflow designer can (i) insert into a workflow the generic component (as shown at the top of Figure 3), (ii) select a behavior from the available transducer templates associated with the component, and (iii) select task implementations from those available for the state frames of the template. The behaviors and implementations that a workflow designer selects from may originally be specified by the component developer or can potentially be reused and repurposed from other generic components.

3.2 The Generic Data Transfer Component

A common task in scientific workflows is data transfer between hosts. Current solutions “hardwire” into the overall workflow both the underlying transport protocol (e.g., scp) and the dynamic behavior used to operate the protocol (e.g., reactions to exceptions and number of retries). For example, the astrophysics workflow previously discussed (see Figure 1) hardcodes the transfer of local simulation data from one host to the host in which a particular analysis is performed. Data transfer is also commonly performed in scientific workflows to store and archive the results of analytical processes.

Data transfer using our framework can instead be specified as follows. The designer first selects the GDT component whose sig-

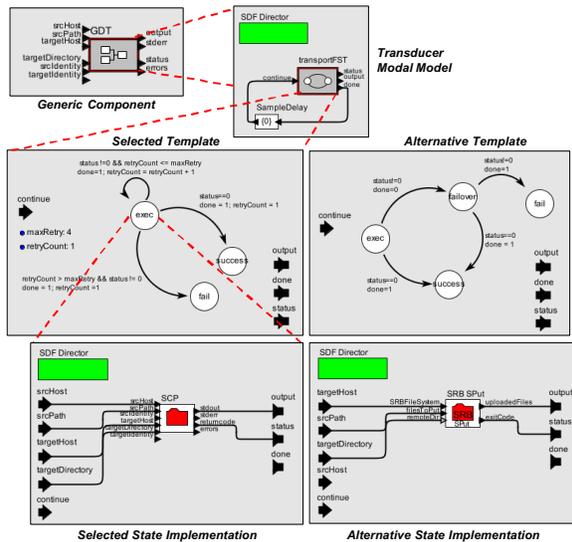


Figure 4: The Generic Data Transfer component

nature specifies the common inputs and outputs such as source and target hosts, file names and locations, and user information. Using the GDT, the designer can then select a transducer template with the desired data-transfer behavior (e.g., from a library of predefined transducer templates). The designer may choose, e.g., a `retry-failover` template that: (1) attempts a transfer protocol p_1 up to n times, (2) if p_1 is not successful, attempts an alternative transfer protocol p_2 up to m times, and (3) if p_2 also fails, reports a failure condition. Note that n , m , p_1 , and p_2 are configuration parameters of the template where p_1 and p_2 denote state-frame implementations. The designer can then select appropriate state-frame implementations through the GDT. The designer might select, e.g., an `scp` state implementation for p_1 and an `ftp` state implementation for p_2 . Finally, based on the signatures and configured parameters of the GDT component, `retry-failover` template, and the state implementations (`scp` and `ftp`), the proper embeddings are performed resulting in a fully instantiated (i.e., “ground”) GDT component that can then be executed from within the overall workflow.

Using KEPLER we have implemented an initial version of the GDT component, as shown in Figure 4. In this implementation, the GDT component (top, left) is a special actor (more precisely, an extension of a composite actor) that provides necessary frame functions for supporting the generic control-flow pattern. The GDT component contains an intermediate subworkflow (upper-right). This subworkflow contains a “modal model” actor, which is required in PTOLEMY II for nesting FSTs within dataflow process networks. This subworkflow also permits multiple executions of the transducer template on each firing of the GDT component.

Two transducer templates are shown in Figure 4 for the GDT component. The selected template (middle, left) is a simple retry loop, which executes the desired protocol `maxRetry` times before entering a fail state. Note that this template performs an equivalent function as the control components of Figure 1. The other template (middle, right) of Figure 4 provides a simple fail-over behavior in which an initial protocol is attempted, and if it fails, a fail-over protocol is used. Finally, `scp` and `SRB sput` state implementations are shown at the bottom of Figure 4. In the figure, the `scp` im-

plementation has been selected, and its signature propagated to the GDT component.

In our current implementation, all configuration including the selection of templates and state implementations is performed by assigning specific attributes of the GDT component.⁴ For example, when a workflow designer configures the GDT component, a dialog box is presented that contains a drop-down list of available templates. Once a template is selected, the user can also select an associated state implementation (note that at anytime after selecting a template it can be navigated to within the KEPLER GUI). The GDT actor reacts to these attribute changes dynamically, assigning the transducer to the modal model, refining the appropriate transducer states, and making the appropriate port connections. A workflow designer can also re-configure the GDT component by assigning different templates or state implementations.

Finally, in addition to the GDT component, we have also implemented the Generic eXecution (GX) component, shown in Figure 5. The GX component is similar to the GDT component, but is tailored for remote execution (e.g., `ssh`) as opposed to file transfer. The GX component, however, can directly reuse the control-flow templates defined for the GDT component (middle of Figure 5). The ability to reuse control-flow in this way is a significant advantage of this approach. Indeed, the ability to reuse both control-flow and dataflow components via templates and frames can lead to more robust, intuitive, and ultimately reusable scientific workflows.

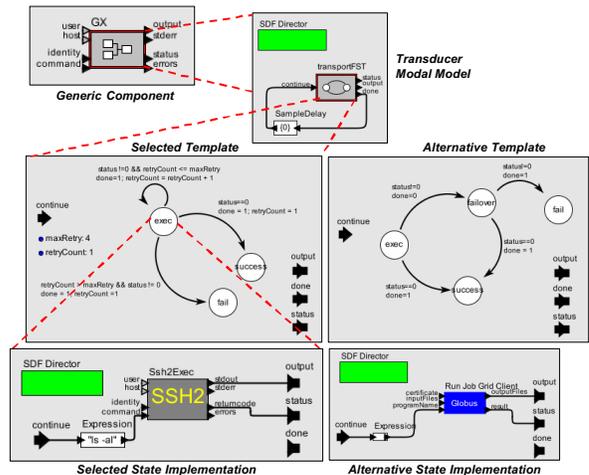


Figure 5: The Generic eXecution component

4. RELATED WORK

Scientific workflow systems [19, 17, 2, 16] are often based on a dataflow model, due to the data-centric and data-driven nature of many scientific workflows. In contrast, business workflow systems [1] and systems for web-service composition (e.g., BPEL4WS [6] and OWL-S) often use control-based models such as finite state machines or Petri nets. Few systems seamlessly integrate control-flow and dataflow within a single model. Our approach for embedding control-flow into dataflow was inspired by hierarchical finite state machines [8] in PTOLEMY II and more generally, PTOLEMY’s ability to nest heterogeneous computation models [12, 5].

A number of systems, e.g., for business workflows and more recently for web-service composition [10], support separate interface

⁴In PTOLEMY II, and *attribute* is a static property.

declarations from underlying control-flow models. For example, in the Collaboration Management Infrastructure (CMI) [20], each interface declaration has an associated state machine that can be selected at design time. This approach has been shown to enable improved workflow constructs [18]. Our approach goes further by combining dataflow and control-flow in a structured way using the frame and template abstractions. This enables complex “interface” definitions not possible in other scientific workflow systems.

The development of “rigid” workflow modeling and design frameworks have recently been identified as a major bottleneck for scientific workflow reuse and repurposing (*i.e.*, reconfiguring existing workflows for new purposes) [9]. We have shown that our approach can significantly enhance reusability in scientific workflows. Moreover, the use of frames and templates together with semantic port types also yields improved search mechanisms for scientific workflow repositories.

5. CONCLUDING REMARKS

While scientific workflows are primarily dataflow-oriented, certain workflow tasks can be control-intensive, *e.g.*, procedures for providing fault-tolerant and adaptive distributed data transfer. Modeling these tasks directly using dataflow constructs can lead to workflows that are overly complex and difficult to maintain and reuse. We have described a framework for structured embedding of generic control-flow components within dataflow process networks. In particular, we have introduced (actor) *frames* and (workflow) *templates* and shown how they can be used to develop robust workflows via reusable control-intensive subtasks. This framework has been prototypically implemented on top of the KEPLER scientific workflow system.

As future work, we intend to further implement frames and templates as separate modeling constructs within KEPLER, and leverage them to create more complex generic control-flow components. We also want to extend our current data transfer and remote execution components with added support for automated signature matching. One of our goals is to populate KEPLER with a number of generic components, including rich libraries of supporting task implementations and transducer templates, to support a wide range of scientific workflows. We are also interested in investigating FST composition operations so that designers can select multiple transducer templates and combine them to dynamically create and reuse complex control-intensive behavior.

6. REFERENCES

- [1] G. Alonso and C. Mohan. Workflow Management Systems: The Next Generation of Distributed Processing Tools. In *Advanced Transaction Models and Architectures*. 1997.
- [2] I. Altintas, C. Berkley, E. Jaeger, M. Jones, B. Ludäscher, and S. Mock. Kepler: An Extensible System for Design and Execution of Scientific Workflows. In *SSDBM*, 2004.
- [3] V. Bhat, S. Klasky, S. Atchley, M. Beck, D. McCune, and M. Parashar. High Performance Threaded Data Streaming for Large Scale Simulations. In *GRID '04: Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing (GRID'04)*, pages 243–250, Washington, DC, USA, 2004. IEEE Computer Society.
- [4] S. Bowers and B. Ludäscher. Actor-Oriented Design of Scientific Workflows. In *24th Intl. Conf. on Conceptual Modeling (ER)*, 2005.
- [5] C. Brooks, E. A. Lee, X. Liu, S. Neuendorffer, Y. Zhao, and H. Zheng. Heterogeneous Concurrent Modeling and Design in Java. Technical Report Technical Memorandum UCB/ERL M05/21, Univ. of California, Berkeley, 2005.
- [6] F. Curbera, Y. Goland, J. Klein, F. Leyman, D. Roller, S. Thatte, and S. Weer-awarana. *Business Process Execution Language for Web Services (BPEL4WS), Version 1.0*. 2002.
- [7] E. Gamma, R. Helm, R. Johnson, and J. J. Vlissides. *Design Patterns*. Addison-Wesley Professional, 1995.
- [8] A. Girault, B. Lee, and E. A. Lee. Hierarchical Finite State Machines with Multiple Concurrency Models. *IEEE Transactions on CAD*, 18(6), 1999.
- [9] A. Goderis, C. Goble, U. Sattler, and P. Lord. Seven bottlenecks to workflow reuse and repurposing. In *International Semantic Web Conference (ISWC2005)*, 2005. to appear.
- [10] R. Hull and J. Su. Tools for Composite Web Services: A Short Overview. *SIGMOD Record*, 34(2), 2005.
- [11] E. A. Lee and D. G. Messerschmitt. Static scheduling of synchronous data flow programs for digital signal processing. *IEEE Trans. Comput.*, C-36, 1987.
- [12] E. A. Lee and S. Neuendorffer. Actor-oriented Models for Codesign: Balancing Re-Use and Performance. In *Formal Methods and Models for System Design*. Kluwer, 2004.
- [13] E. A. Lee and T. M. Parks. Dataflow Process Networks. *Proc. of the IEEE*, 83(5), 1995.
- [14] B. Ludäscher and I. Altintas. On Simplifying Collection Handling and Control-Flow Issues in SPA/Ptolemy-II. Technical Report SciDAC-SPA-TN-2003-01, San Diego Supercomputer Center, 2003.
- [15] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, and Y. Zhao. Scientific Workflow Management and the Kepler System. *Concurrency and Computation: Practice & Experience*, 2005. to appear.
- [16] R. S. MacLeod, D. M. Weinstein, J. Davison de St. Germain, C. R. Johnson, S. G. Parker, and D. . Brooks. SCIRun/BioPSE: Integrated Problem Solving Environment for Bioelectric Field Problems and Visualization. In *Proc. of the IEEE Intl. Symposium on Biomedical Imaging (ISBI): From Nano to Macro*. IEEE, 2004.
- [17] S. Majithia, M. S. Shields, I. J. Taylor, and I. Wang. Triana: A Graphical Web Service Composition and Execution Toolkit. In *Proc. of the IEEE Intl. Conf. on Web Services (ICWS)*. IEEE Computer Society, 2004.
- [18] A. H. H. Ngu, D. Georgakopoulos, D. Baker, A. Cichocki, J. Desmarais, and P. Bates. Advanced Process-based Component Integration in Telcordia’s Cable OSS. In *ICDE*, 2002.
- [19] T. M. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, R. 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*, 20(17), 2004.
- [20] M. Rusinkiewicz and D. Georgakopoulos. From Coordination of Workflow and Group Activities to Composition and Management of Virtual Enterprises. In *International Symposium on Database Applications in Non-Traditional Environments*, 1999.
- [21] X. Xin. Case Study: Terascale Supernova Initiative Workflow (TSI-Swesty). LLNL Technical Note, 2004. <http://www-casc.llnl.gov/sdm/documentation/casestudy-tsi-s.doc>.