

Logic Based Modeling and Analysis of Workflows *

(Extended Abstract)

Hasan Davulcu Michael Kifer C.R. Ramakrishnan I.V. Ramakrishnan
Department of Computer Science
SUNY at Stony Brook, Stony Brook, NY 11794-4400
{davulcu,kifer,cram,ram}@cs.sunysb.edu

To appear in PODS-98
June 1998, Seattle, Washington

*This work is partially supported by a DLA/DARPA contract and by the NSF grants IRI9404629, CCR-9510072, CCR-9705998, CCR-9711386, CCR-9404921, CDA-9303181, CDA-9504275

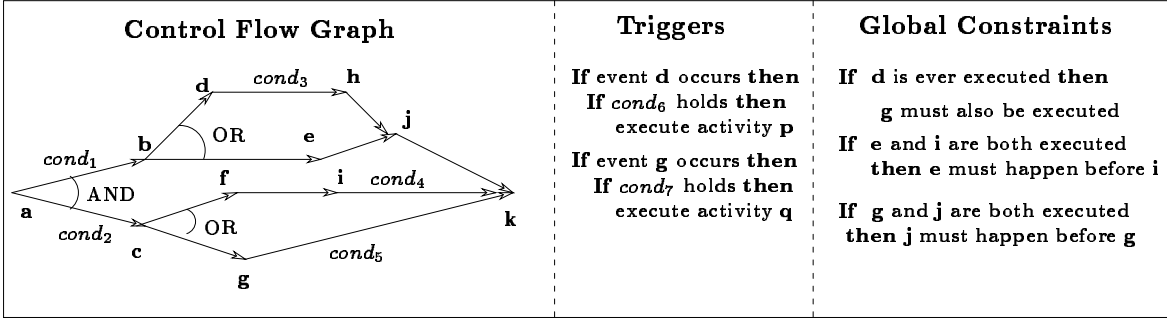


Figure 1: Frameworks for Specifying Workflows

1 Introduction

A *workflow* is a collection of cooperating, coordinated activities designed to carry out a well-defined complex process, such as trip planning, graduate student registration procedure, or a business process in a large enterprise. An activity in a workflow might be performed by a human, a device, or a program. *Workflow management systems* provide a framework for capturing the interaction among the activities in a workflow and are recognized as a new paradigm for integrating disparate systems, including legacy systems [18, 8]. Ideally, they should also help the user in analysis and reasoning about complex business processes.

It has been realized that analysis and reasoning about workflows requires a formal specification model with a well defined semantics [16, 2]. In this paper, we develop a novel framework for *specifying, analyzing* and *executing* workflows based on Transaction Logic [5, 4, 6, 7].

Workflow representation frameworks. Figure 1 depicts three most common frameworks for specifying workflows: *control flow graph*, *triggers* (also known as *event-condition-action rules*), and *temporal constraints*.

The control flow graph is most appropriate for depicting the local execution dependencies of the activities in a workflow; it is a good way to visualize the overall flow of control. Control flow graphs are the primary specification means in most commercial implementations of workflow management systems. A typical graph specifies the initial and the final activity in a workflow, the successor-activities for each activity in the graph, and whether these successors must *all* be executed concurrently, or it suffices to execute just one branch non-deterministically. In Figure 1, all successors of activity **a** must be executed, which is indicated with the “AND”-label. In contrast, “OR” indicates that when **b** is finished, there is a choice of executing **d**, **h**, then **j** *or* **e** then **j**. Successful execution of any one of these branches should suffice for the overall success of the workflow.

Arcs in a control flow graph can be labeled with *transition conditions*. The condition applies to the current state of the workflow (which, in a broad sense, may include the current state of the underlying database, the output of the completed tasks, the current time, etc.). When the task at the tail of an arc completes, the task at the head can begin only if the corresponding transition condition evaluates to true.

The Workflow Management Coalition [9] identifies additional controls, such as loops and sub-workflows. Various researchers have also suggested other types of controls, including alternative execution and compensation for failed activities [11, 15, 13, 22, 25, 1, 12]. However, control flow graphs have one obvious limitation: they cannot be used to specify *global* dependencies between workflow tasks, such as those expressed as global constraints on the right-hand side of Figure 1.

Defining workflows using triggers is yet another possibility [10]. However, this method is not as general as control flow graphs. For instance, like the graphs, triggers cannot be used to specify global task dependencies, and they are not sufficiently expressive when it comes to representing alternatives in workflow execution (depicted as “OR” nodes in Figure 1). In fact, it follows from a result in [7] that triggers with so-called “immediate”

execution semantics can be represented using control flow graphs, and this result can be adapted to triggers with the “eventual” execution semantics as well. Since triggers can be “compiled into” the control flow graph, we shall be treating triggers as part of the control flow graph.

Other researchers proposed the frameworks that rely exclusively on constraints to specify *both* the local and the global properties of workflows. In [23, 24], Singh describes an algebra of temporal constraints, which is believed to cover all useful global dependencies that might arise in workflow systems. For instance, this algebra includes the Klein’s constraints [20], which commonly occur in workflow specifications.¹ This algebra is sufficiently expressive for modeling control flow graphs that have *no* transition conditions attached to arcs. However, such constraints cannot be used to model workflows that query the intermediate state of the workflow and make scheduling decisions based on the outcome. More importantly algebraic approaches are not part of a larger reasoning framework and so it is unclear how they apply to more expressive specification languages (*e.g.*, to represent sub-workflows, failure and compensation, etc.) and how they could be used for reasoning and verifying the correctness of workflows.

Logic-based formalism. In this paper, we base our approach on Concurrent Transaction Logic [6] (abbr., *CTR*). There are many reasons for this choice. First, control flow graphs with transition conditions can be easily and naturally represented in *CTR*. Second, [7] shows that triggers are easy to represent as well. Finally, [5] contains an extensive discussion of the temporal capabilities of *CTR*. In particular, the entire algebra of constraints described in [23] is isomorphic to a small subset of the propositional Transaction Logic. Hence *CTR* provides us with a unifying formalism that subsumes all of the three popular workflow specification frameworks described above. Furthermore, being a full-blown logic, *CTR* allows not only scheduling workflows, but also *reasoning* about their properties. Finally, like in logic programming systems, the proof theory of *CTR* is also a run-time environment for *executing* workflows.

Summary of results. Our approach is based on a transformation procedure, named *Apply*, which accepts a workflow specification that includes a control flow graph \mathcal{G} , triggers (viewed as part of \mathcal{G}), and a set of temporal constraints \mathcal{C} , and constructs an equivalent specification in *CTR* that represents only those executions of \mathcal{G} where \mathcal{C} holds. The resulting specification can be directly used to execute the workflow. Thus, *Apply* can be viewed as a *compilation* process that facilitates verification of workflows and optimizes run-time scheduling.

1. *Apply* transformation enables us to determine:
 - (a) Whether every legal execution of a given workflow specification satisfies a particular property. Moreover, if some execution does not satisfy the property, then the verification procedure returns a counter example which is the most general execution of the workflow that violates the property in question.
 - (b) Whether the specification made up of the control flow graph and global constraints is consistent; and
 - (c) Whether some of the specified constraints are redundant.
2. The *Apply* transformation eliminates the parts of the control flow graph that are inconsistent with the constraints, which facilitates scheduling of events at run-time.
3. The separation of control flow graph and global constraints in the workflow specifications leads to tighter complexity results for the verification problem.

Our results also contribute to the theory of Transaction Logic itself. Here, we essentially extend the efficient SLD-style proof procedure of *CTR* from so called concurrent-Horn goals to a larger class of formulas, which incorporates temporal constraints (a more precise formulation appears in Section 2).

¹Klein constraints are of the form: (1) if events **a** and **b** both occur, then **a** occurs earlier than **b**; or (2) if event **a** ever occurs then **b** must occur as well (before or after **a**).

2 An Overview of Concurrent Transaction Logic

This section provides a short summary of the CTR syntax, which is used in this paper to represent workflows. Due to space limitation, we cannot discuss the model theory of the logic or its proof theory. Instead, we rely on the procedural reading of CTR statements.² A thorough treatment of all the main aspects of Transaction Logic appears in [6, 5, 4]. A fairly detailed, yet informal introduction can be found in [19].

CTR is a conservative extension of the classical predicate logic in the sense that both its proof theory and the model theory reduce to those of the classical logic for formulas that do not cause state transitions (but only query the current state).

Basic syntax. The atomic formulas of CTR are identical to those of the classical logic, *i.e.*, they are expressions of the form $p(t_1, \dots, t_n)$, where p is a predicate symbol and the t_i 's are function terms. More complex formulas are built with the help of connectives and quantifiers.

Apart from the classical \vee , \wedge , \neg , \forall , and \exists , CTR has two additional connectives, \otimes (*serial conjunction*) and $|$ (*concurrent conjunction*), and two modal operators, \diamond (*executional possibility*) and \odot (*isolated execution*). For instance, $\odot(p(X) \otimes q(X)) | (\forall Y(r(Y) \vee s(X, Y)))$ is a well-formed formula.

Informal semantics. Underlying the logic and its semantics is a set of database *states* and a collection of *paths*. For the purpose of this paper, the reader can think of the states as just a set of relational databases, but the logic does not rely on the exact nature of the states—it can deal with a wide variety of them.

A *path* is a finite sequence of states. For instance, if s_1, s_2, \dots, s_n are database states, then $\langle s_1 \rangle$, $\langle s_1, s_2 \rangle$, and $\langle s_1, s_2, \dots, s_n \rangle$ are paths of length 1, 2, and n , respectively.

Just as in classical logic, CTR formulas assume truth values. However, unlike classical logic, the truth of CTR formulas is determined over paths, *not* at states. If a formula, ϕ , is true over a path $\langle s_1, \dots, s_n \rangle$, it means that ϕ can *execute* starting at state s_1 . During the execution, the current state will change to s_2, s_3, \dots , etc., and the execution terminates at state s_n .

With this in mind, the intended meaning of the CTR connectives can be summarized as follows:

- $\phi \otimes \psi$ *means*: execute ϕ then execute ψ . Or, model-theoretically, $\phi \otimes \psi$ is true over a path $\langle s_1, \dots, s_n \rangle$ if ϕ is true over a prefix of that path (say, $\langle s_1, \dots, s_i \rangle$) and ψ is true over the suffix (*i.e.*, $\langle s_i, \dots, s_n \rangle$). In terms of control flow graphs (cf. Figure 1), this connective represents arcs connecting adjacent activities.
- $\phi | \psi$ *means*: ϕ and ψ must both execute concurrently, in an interleaved fashion. This connective corresponds to the “AND”-nodes in control flow graphs.
- $\phi \wedge \psi$ *means*: ϕ and ψ must both execute along the *same* path. In practical terms, this is best understood in terms of *constraints* on the execution. For instance, ϕ can be thought of as a transaction and ψ as a constraint on the execution of ϕ . It is this feature of the logic that lets us specify temporal constraints as part of workflow specifications.
- $\phi \vee \psi$ *means*: execute ϕ *or* execute ψ non-deterministically. This connective corresponds to the “OR”-nodes in control flow graphs.
- $\neg\phi$ *means*: execute in any way, provided that this will not be a valid execution of ϕ . There are many uses for this feature. One is that, just as in classical logic, the negation lets us define deductive rules which, in terms of the workflows, correspond to sub-workflow definitions. Negation is also an important component in temporal constraint specifications.
- $\odot\phi$ *means*: execute ϕ in isolation, *i.e.*, without interleaving with other concurrently running activities.

This operator enables us to specify the transactional parts of workflow specifications.

²This is analogous to the procedural reading of Datalog programs.

- $\diamond\phi$ means: check if ϕ is executable at the current state. Section 6 discusses the role of the possibility operator \diamond in workflow modeling.

Concurrent-Horn subset of CTR . Next, we define the implication, $p \leftarrow q$, as $p \vee \neg q$. The form and the purpose of the implication in CTR is similar to that of Datalog: p can be thought of as the name of a procedure and q as the definition of that procedure. However, unlike Datalog, both p and q assume truth values on execution paths, not at states.

More precisely, $p \leftarrow q$ means: if q can execute along a path $\langle s_1, \dots, s_n \rangle$, then so can p . If p is viewed as a subroutine name, then the meaning can be re-phrased as: one way to execute p is to execute q , the definition of p .

Having provided the intuition behind the logical connectives, it is now easy to see how control flow graphs are represented in CTR . For instance, the graph in Figure 1 is represented as:

$$a \otimes \left(\left(\text{cond}_1 \otimes b \otimes ((d \otimes \text{cond}_3 \otimes h) \vee e) \otimes j \right) \mid \left(\text{cond}_2 \otimes c \otimes ((f \otimes i \otimes \text{cond}_4) \vee (g \otimes \text{cond}_5)) \right) \right) \otimes k \quad (2.1)$$

Expressions of the above form are called *concurrent-Horn goals*. Formally:

- any atomic formula is a concurrent-Horn goal;
- $\phi \otimes \psi$, $\phi \mid \psi$, and $\phi \vee \psi$ are concurrent-Horn goals, if so are ϕ and ψ ;
- $\odot\phi$ and $\diamond\phi$ are concurrent-Horn goals, if so is ϕ .

It should be clear from the above example how control flow graphs translate into concurrent-Horn goals.

A *concurrent-Horn rule* is a CTR formula of the form $head \leftarrow body$, where $head$ is an atomic formula and $body$ is a concurrent-Horn goal.

In this paper, we limit our attention to non-iterative workflows, which means that we do not allow recursive concurrent rules. Section 6 discusses to what extent our present results apply to recursively defined workflows.

From the workflow point of view, the primary use for the rules is to represent sub-workflows. Indeed, since workflows and sub-workflows can be described using concurrent-Horn goals, we can use the rules of the form $subWorkflowName \leftarrow subWorkflowDefinition$ to define sub-workflows. For instance, $subWorkflowName$ can be used in workflow specifications as if it were a regular activity, thereby completely hiding the underlying structure of the activity from top-level specifications.

Observe that the definition of concurrent-Horn rules and goals *does not* include the connective \wedge . In general, \wedge represents *constrained execution*, which is usually hard to implement, since constraints must be checked at every step of the execution. If a constraint violation is detected, a new execution path must be tried out. In contrast, the concurrent-Horn fragment of CTR is efficiently implementable, and there exist an SLD-style proof procedure that proves concurrent-Horn formulas and *executes* them at the same time [6].

The efficiency gap between concurrent-Horn execution and constrained execution is the main motivation for our results. In logical terms, we show that, for a large class of constraints, formulas of the form $ConcurrentHornGoal \wedge Constraints$ have an equivalent concurrent-Horn form (which, therefore, does not use the connective \wedge). In practical terms, therefore, this means that there is an efficient workflow scheduling strategy and, moreover, this strategy can be determined at “design time” of the workflow (as opposed to run-time scheduling of [24]).

Elementary updates. We complete our informal introduction to CTR by explaining how execution of (some) formulas may actually change the underlying database state. Most of the machinery has already been introduced (albeit very informally). What is missing is the notion of *elementary updates*.

In CTR , elementary updates are represented by ordinary atomic, variable-free formulas. Syntactically, CTR does not distinguish elementary updates in any way, but the user may want to do so by adopting a syntactic

convention (*e.g.*, a convention could be that $ins.p(t)$ represents the act of insertion of tuple t into the relation p).

What distinguishes elementary updates is their semantics. Through some black magic, called *transition oracle*, CTR arranges so that each elementary update is always true along certain arcs, *i.e.*, paths of the form $\langle s_1, s_2 \rangle$. Informally, one can think of an elementary update as a binary relation over states. For instance, if $\langle s_1, s_2 \rangle$ belongs to the relation corresponding to an elementary update u , it means that u can cause a transition from state s_1 to state s_2 . Note that an update can be *non-deterministic* (any one of a number of alternative state transitions might be possible) and it is possible for an update to be inapplicable in certain states (but it is also possible for an update to apply in every state).³

This mechanism is very general. It accounts for a wide variety of elementary state changes: from simple tuple insertions and deletions, to relational assignments, to updates performed by legacy programs, to whatever workflow activities might do. The connectives of CTR are then used to build more complex updates from the elementary ones and then to combine these complex updates into even more complex update programs. This process of building CTR programs from the ground up is very natural and powerful. The reader is referred to [4, 5, 6] for concrete examples.

Now we can explain how the various workflow activities (*e.g.*, the symbols a, b, c , etc., in (2.1)) appear to CTR . Namely, each activity is encoded as a variable-free atomic formula, η , that represents either a sub-workflow defined by a set of concurrent-Horn rules, or it can represent an ordinary activities, in which case η is an elementary update. The latter is appropriate, since individual activities appear to workflow management systems as “black boxes” that perform state changes in ways that are (at best) only partially specified.

3 Events and Temporal Constraints

In workflow systems, tasks are typically modeled in terms of their significant, externally observable events, such as *start*, *commit*, or *abort*. For the purpose of control flow, we can represent these events as regular activities and incorporate them directly into the control flow graph in appropriate spots. The temporal constraints on workflow execution can then be expressed in terms of these events. Without loss of generality (as far as workflow modeling goes), we make the following assumptions:

- *No significant event occurs twice during the execution.*
Indeed, we can always rename different occurrences of the same type of event.
- *Each significant event is represented as an elementary update that applies in every state.* (3.1)
This assumption is appropriate since, typically, a significant event amounts to nothing more than forcing a suitable record into the system log.

The first assumption translates into the following *unique event property*, which limits the kinds of concurrent-Horn goals that we shall consider in this paper:

Definition 3.1 (Unique Event Property). A concurrent-Horn goal \mathcal{G} has the *unique event property* if and only if every significant event occurs at most once in every execution of \mathcal{G} . In such cases, we shall also say that \mathcal{G} is a *unique-event goal*. \square

Unique-event goals can be recognized in linear time in the size of the goal, but we shall not present this algorithm here. Instead, we mention some obvious, yet useful properties of such goals, which suffices for our purposes. Let α be a significant event. Then:

- If $G = E_1 \otimes E_2$ is a unique-event goal and α occurs in E_1 then it cannot occur in E_2 .
- If $G = E_1 \mid E_2$ is a unique-event goal and α occurs in E_1 then it cannot occur in E_2 .
- If $G = E_1 \vee E_2$ then G is a unique-event goal if and only if so are both E_1 and E_2 .

³An example of the first kind is an update that deletes $p(t)$ only iff $p(t)$ is true in the current state. An example of the second update is deletion of $p(t)$ regardless of whether $p(t)$ is true. If $p(t)$ is not true in some state, s , then no state transition takes place, but the update will still be true over the arc $\langle s, s \rangle$.

In the rest of this paper, all concurrent-Horn goals are assumed to have the unique event property.

Transaction Logic can express a wide variety of temporal constraints [5], but here we focus on a relatively simple algebra of constraints, which we denote by \mathcal{C}_{ONSTR} . \mathcal{C}_{ONSTR} is as expressive as Singh’s Event Algebra [24]. Using these constraints we can specify that one task must start before some other task, that the execution of one task causes some other task to be executed or not executed, etc. These constraints are believed to be sufficient for the needs of workflow management systems, and they are far beyond of the capabilities of the currently available commercial systems.

We specify all significant events in the system as propositions drawn from a set, denoted by \mathcal{E}_{VENT} . In addition, we introduce one special proposition, *path*, which is defined as $\phi \vee \neg\phi$, for any \mathcal{CTR} formula. This means that *path* is true on *all* possible execution paths.⁴

Definition 3.2 (Constraints). The basic building blocks of \mathcal{C}_{ONSTR} are formulas of the form $\text{path} \otimes e \otimes \text{path}$, where $e \in \mathcal{E}_{VENT}$. To save space, we shall use a shortcut for such formulas: $\nabla\phi \equiv \text{path} \otimes \phi \otimes \text{path}$, by definition. Then the following constraints form the constraint algebra \mathcal{C}_{ONSTR} :

1. **Primitive constraints:** If $e \in \mathcal{E}_{VENT}$ then ∇e (event e must happen) and $\neg\nabla e$ (event e must *not* happen) are *primitive* constraints in \mathcal{C}_{ONSTR} . The constraint ∇e is a *positive* primitive constraint and $\neg\nabla e$ is a *negative* primitive constraint.
2. **Serial constraints:** If $s_1, \dots, s_n \in \mathcal{E}_{VENT}$ are *positive* primitive constraints, then $s_1 \otimes \dots \otimes s_n \in \mathcal{C}_{ONSTR}$ is a *serial* constraint. For convenience, primitive constraints are also viewed as serial constraints.
3. **Complex constraints:** If $C_1, C_2 \in \mathcal{C}_{ONSTR}$ then so are $C_1 \wedge C_2$, and $C_1 \vee C_2$.

Nothing else is in \mathcal{C}_{ONSTR} . □

To get a better grasp of the capabilities of \mathcal{C}_{ONSTR} , here are a few typical constraints and their real-world interpretation:

- $\nabla e \wedge \nabla f$ — events e and f must both occur (in some order);
- $\neg\nabla e \vee \neg\nabla f$ — it is not possible for e and f to happen together.
- $\neg\nabla e \vee (\nabla e \otimes \nabla f)$ — if event e occurs, then f must occur some time later;
- $\neg\nabla e \vee \neg\nabla f \vee (\nabla e \otimes \nabla f)$ — if both e and f occur, then e must come before f . This is known as Klein’s order constraint [20].
- $\neg\nabla f \vee (\nabla e \otimes \nabla f)$ — if event f has occurred, then event e must have occurred some time prior to that;
- $\neg\nabla e \vee \nabla f$ — if event e occurs, then f must also occur (before or after e). This is known as Klein’s existence constraint [20].

Note that Definition 3.2 does not explicitly state that \mathcal{C}_{ONSTR} is closed under negation. Nevertheless, we can show that it is.

Claim 3.3 (Splitting Serial Constraints). *Under the assumptions (3.1), any serial constraint is equivalent to a \wedge -conjunction of serial constraints, each composed of no more than two primitive constraints.*

Proof. Consider a positive serial constraint composed of more than two primitive constraints: $\nabla e_1 \otimes \nabla e_2 \otimes s$, where s is a serial constraint. We can show that this is equivalent to $(\nabla e_1 \otimes \nabla e_2) \wedge (\nabla e_2 \otimes s)$. □

A serial constraint of the form $\nabla\alpha \otimes \nabla\beta$ is called an *order constraint*; it says that α and β must both occur and α must occur before β . (Note that this is somewhat stronger than Klein’s order constraint mentioned earlier.)

⁴This is one of the counterparts of “true” in classical logic. In \mathcal{CTR} , one can define other propositions that express various truths. For instance, we can express the proposition *state*, which is true precisely on paths of length 1, i.e., at states. It is also possible to express formulas that are true precisely on arcs, etc.

Lemma 3.4 (Constraint Negation). *Let $C \in \mathcal{C}_{ONSTR}$. Then \mathcal{C}_{ONSTR} has a constraint that is equivalent to $\neg C$ under the assumptions (3.1).*

Proof. We can push negation down to the serial constraints in C using the classical De Morgan’s laws for \wedge , \vee , and \neg , which are valid also in \mathcal{CTR} :

$$\begin{aligned} \neg(\phi \vee \psi) &\equiv \neg\phi \wedge \neg\psi \\ \neg(\phi \wedge \psi) &\equiv \neg\phi \vee \neg\psi \\ \neg\neg\phi &\equiv \phi \end{aligned}$$

Since $\neg\neg\forall e$ is equivalent to $\forall e$, we only need to show that $s = \neg(\forall e_1 \otimes \dots \otimes \forall e_n)$ is equivalent to some constraint in \mathcal{C}_{ONSTR} .

By Claim 3.3, we can assume that $n < 3$. If $n = 1$, then $s = \neg\forall e_1$ is a negative primitive constraint. If $n = 2$, then $s = \neg(\forall e_1 \otimes \forall e_2)$, which is equivalent (under the assumptions (3.1)) to $\neg\forall e_1 \vee \neg\forall e_2 \vee (\forall e_2 \otimes \forall e_1)$. \square

The previous results lead to the following normal form for the members of \mathcal{C}_{ONSTR} :

Corollary 3.5 (Normal Form for Constraints). *Every constraint in \mathcal{C}_{ONSTR} is equivalent to a constraint of the form $\vee_i(\wedge_j \text{serialConstr}_{i,j})$ where each $\text{serialConstr}_{i,j}$ is either a primitive constraint or a serial constraint composed of two positive primitive constraints.*

Proof. Follows from Claim 3.3, Lemma 3.4, and the fact that, as in classical logic, \vee distributes through \wedge and vice versa. \square

Lemma 3.4 helps express certain constraints much more easily. For instance:

- $\neg(\forall e \otimes \forall f)$ — it is not possible for f to occur after e (and for e before f).
- $\neg(\forall e \otimes \forall f \otimes \forall g)$ — if e happens and then f does, the event g cannot happen later.

4 Consistency, Verification, and Scheduling Problems for Workflows

This section and the next assumes that the control flow graphs do not have transition conditions on the arcs and that the specification does not contain concurrent-Horn rules that define sub-workflows. In Section 6, we discuss how our results apply to graphs that include these features.

Let \mathcal{G} be a concurrent-Horn goal with unique event property (Definition 3.1), which represents a control flow graph, and let $\mathcal{C} \subset \mathcal{C}_{ONSTR}$ be a set of constraints. The three central problems in workflow management systems can be formulated as follows:

Consistency: *Determine whether \mathcal{G} is consistent with \mathcal{C} .*

Verification: *Determine whether every legal execution of the workflow satisfies some property $\Phi \in \mathcal{C}_{ONSTR}$.*

Scheduling: *Find an execution path (or all paths) in \mathcal{G} where \mathcal{C} holds.*

In \mathcal{CTR} , the consistency problem is tantamount to the existence of an execution for the formula $\mathcal{G} \wedge \mathcal{C}$.

The verification problem is a special case of the consistency problem. Indeed, every legal execution of the workflow satisfies Φ iff $\mathcal{G} \wedge \mathcal{C} \wedge \neg\Phi$ cannot execute (*i.e.*, \mathcal{G} is inconsistent with $\mathcal{C} \wedge \neg\Phi$).

The verification problem also subsumes the *redundancy* problem: $\Phi \in \mathcal{C}$ is redundant iff every legal execution of $\mathcal{G} \wedge (\mathcal{C} - \{\Phi\})$ satisfies Φ .

In this paper, we solve the verification problem constructively by transforming the formula $\mathcal{G} \wedge \mathcal{C}$ into an *equivalent concurrent-Horn* formula \mathcal{G}' , which is always executable; or if this is impossible, $\mathcal{G} \wedge \mathcal{C}$ reduces to $\neg\text{path}$ — a non-executable transaction, which is the \mathcal{CTR} analog of the classical *false*. Our algorithm is exponential in the size of \mathcal{C} (in the worst case), which turns out to be inherent to the verification problem:

Proposition 4.1 (Complexity of Verification). *Let \mathcal{G} be a concurrent goal and $\mathcal{C} \subset \mathcal{C}_{ONSTR}$ be a set of constraints. Then determining whether $\mathcal{G} \wedge \mathcal{C}$ is executable in CTR is NP-complete in the size of \mathcal{C} .*

The proof is by reduction to satisfiability of propositional logic [14]. A similar result has been previously obtained in [21]. However, their NP-completeness result is based on *synchronizer*-constraints. Each synchronizer corresponds to a combination of an *existence constraint*⁵ and an *order constraint*⁶ in our formalism. We tighten their complexity result by showing that synchronization per se is not a culprit: the problem is NP-complete even in the presence of just the existence constraints. In fact, it follows from our solution to the consistency problem that for order constraints the verification problem can be solved in polynomial time.

The scheduling problem needs more explanation. Workflow literature distinguishes two approaches to the problem: *passive* and *pro-active*.

Passive schedulers receive sequences of events from an external source, such as a workflow or a transaction manager, and validate that these sequences satisfy all global constraints (possibly after reordering some events in the sequences). Several such schedulers are described in [23, 3, 17]. To validate a particular sequence of events, each of these schedulers takes at least quadratic time in the number of events.⁷ However, in passive scheduling environments, it is left to an unspecified external system to do consistency checking, to ensure the liveness of the scheduling strategy and to select the event sequences for validation. The known algorithms for these tasks are worst-case exponential.

In contrast to passive scheduling, our approach is pro-active. In particular, we do not rely on any external system. Instead, we construct a “compressed” explicit representation of all allowed executions (*i.e.*, executions that are known to satisfy all constraints). This representation can be used to enumerate all allowed executions at linear time per execution path (linear in the size of the path). In this way, at each stage in the execution of a workflow, the scheduler knows all events that are eligible to start. There is no need to validate constraints at run time, since the constraints are “compiled into” the structure.

More precisely, our solution to the scheduling problem capitalizes on the solution to the consistency problem. First, we verify that the specifications are consistent by transforming $\mathcal{G} \wedge \mathcal{C}$ into an equivalent concurrent-Horn goal \mathcal{G}' , as explained above. The formula \mathcal{G}' plays the role of the aforesaid explicit representation for the set of all allowed executions of $\mathcal{G} \wedge \mathcal{C}$.

If the transformation succeeds (*i.e.*, the specifications are consistent), enumerating all execution paths of \mathcal{G}' takes time linear in \mathcal{G} per path (note: linear in the original graph, not in the much larger graph \mathcal{G}' !). This means that after the compilation, we can pick a legal schedule for workflow activities in time *linear* in the size of the original control flow graph. In contrast, the event scheduler of [24] has quadratic complexity.

Thus, while expanding the effort on consistency checking (which needs to be done anyway), we compile the original specifications into a form that lets us find allowable schedules much more efficiently than with the passive approaches of [24, 3, 17] (and, furthermore, these latter algorithms do not do consistency checks).

5 Compiling Constraints into the Control Flow Graph

We define the process of compiling the constraints in \mathcal{C}_{ONSTR} into unique-event goals by starting with simple events and extending the transformation to more complex ones. The unique-event property assumption is crucial for the correctness of the results in this section.

Compiling primitive constraints. The following transformation takes a primitive constraint of the form $\forall \alpha$ or $\neg \forall \alpha$ and a control flow graph (expressed as a concurrent unique-event goal) and returns a concurrent-Horn goal whose executions are precisely those executions of the original graph that satisfy the constraint. Intuitively, this means that the constraint is compiled into the graph.

⁵Existence constraints form the subset of \mathcal{C}_{ONSTR} obtained from primitive constraints by combining them with \wedge and \vee only.

⁶Order constraints form the subset of \mathcal{C}_{ONSTR} that does not use \vee .

⁷Two of these schedulers are actually exponential in the size of the largest global constraint, but it is reasonable to assume that in practice this size is bound by a small constant.

Definition 5.1 (Primitive Constraint Compilation). Let $\alpha, \beta \in \mathcal{E}_{\text{VENT}}$. Then:

$$\begin{aligned} \text{Apply}(\nabla\alpha, \alpha) &= \alpha \\ \text{Apply}(\nabla\alpha, \beta) &= \neg\text{path} && \text{if } \alpha \neq \beta \\ \text{Apply}(\neg\nabla\alpha, \alpha) &= \neg\text{path} \\ \text{Apply}(\neg\nabla\alpha, \beta) &= \beta && \text{if } \alpha \neq \beta \end{aligned}$$

Let T and K be concurrent-Horn goals and let σ stand for $\nabla\alpha$ or $\neg\nabla\alpha$. Then:

$$\begin{aligned} \text{Apply}(\nabla\alpha, T \otimes K) &= \begin{cases} (\text{Apply}(\alpha, T) \otimes K) & \text{if } \alpha \text{ occurs in } T \\ (T \otimes \text{Apply}(\alpha, K)) & \text{if } \alpha \text{ occurs in } K \end{cases} \\ \text{Apply}(\neg\nabla\alpha, T \otimes K) &= \text{Apply}(\neg\nabla\alpha, T) \otimes \text{Apply}(\neg\nabla\alpha, K) \\ \text{Apply}(\alpha, T \mid K) &= \begin{cases} (\text{Apply}(\alpha, T) \mid K) & \text{if } \alpha \text{ occurs in } T \\ (T \mid \text{Apply}(\alpha, K)) & \text{if } \alpha \text{ occurs in } K \end{cases} \\ \text{Apply}(\neg\nabla\alpha, T \mid K) &= \text{Apply}(\neg\nabla\alpha, T) \mid \text{Apply}(\neg\nabla\alpha, K) \\ \text{Apply}(\sigma, \odot T) &= \odot(\text{Apply}(\sigma, T)) \\ \text{Apply}(\sigma, T \vee K) &= \text{Apply}(\sigma, T) \vee \text{Apply}(\sigma, K) \quad \square \end{aligned}$$

Observe that, due to the properties given in (3.2), the above transformation preserves the unique-event property of concurrent-Horn goals. For example, if T is $\gamma \otimes (\alpha \vee \beta \vee \eta) \otimes \delta$, then:

$$\begin{aligned} \text{Apply}(\nabla\alpha, T) &= \gamma \otimes \alpha \otimes \delta \\ \text{Apply}(\neg\nabla\alpha, T) &= \gamma \otimes (\beta \vee \eta) \otimes \delta \end{aligned}$$

Proposition 5.2 (Primitive Constraint Compilation). If T is a concurrent-Horn goal and σ is a primitive constraint, then $\text{Apply}(\sigma, T) \equiv T \wedge \sigma$.

Compiling order constraints. Next we extend Apply to work with order constraints, *i.e.*, constraints of the form $\nabla\alpha \otimes \nabla\beta$.

Definition 5.3 (Order Compilation). Let $\alpha, \beta \in \mathcal{E}_{\text{VENT}}$ and let T be a concurrent-Horn goal. Then:

$$\text{Apply}(\nabla\alpha \otimes \nabla\beta, T) = \text{sync}(\alpha < \beta, \text{Apply}(\nabla\alpha, \text{Apply}(\nabla\beta, T)))$$

The transformation sync is designed to synchronize events in the desired order. It is defined as follows:

$$\text{sync}(\alpha < \beta, T) = T'$$

where T' is like T , except that every occurrence of event α is replaced with $\alpha \otimes \text{send}(\xi)$ and every occurrence of event β is replaced with $\text{receive}(\xi) \otimes \beta$, where ξ is a new constant.

The actions send and receive are easily expressed in \mathcal{CTR} (see [6]) and their semantics is what one would expect of such synchronization primitives: $\text{receive}(\xi)$ is true if and only if $\text{send}(\xi)$ has been previously executed. In this way, β cannot start before α is done. \square

It is easy to verify that, due to (3.2), the above transformation preserves the unique-event property of concurrent-Horn goals. The following examples illustrate the transformation:

$$\begin{aligned} \text{Apply}(\nabla\alpha \otimes \nabla\beta, \gamma \vee (\beta \otimes \alpha)) &= \text{receive}(\xi) \otimes \beta \otimes \alpha \otimes \text{send}(\xi) \\ \text{Apply}(\nabla\alpha \otimes \nabla\beta, \alpha \mid \beta \mid \rho_1 \mid \dots \mid \rho_n) &= (\alpha \otimes \text{send}(\xi)) \mid (\text{receive}(\xi) \otimes \beta) \mid \rho_1 \mid \dots \mid \rho_n \end{aligned}$$

Proposition 5.4 (Order Compilation). Let T be a concurrent-Horn goal and $\alpha, \beta \in \mathcal{E}_{\text{VENT}}$. Then $\text{Apply}(\nabla\alpha \otimes \nabla\beta, T) \equiv T \wedge (\nabla\alpha \otimes \nabla\beta)$.

Compiling general constraints. We are now ready to extend `Apply` to handle the general constraints in `CONSTRAINT`.

Definition 5.5 (Compiling General Constraints). Let T be a concurrent-Horn goal. We assume that workflows are specified by a set of constraints \mathcal{C} and each individual constraint is represented in the normal form of Corollary 3.5. Therefore, \mathcal{C} can be written as a single dependency of the form

$$\delta_1 \wedge \delta_2 \wedge \dots \wedge \delta_n \quad (5.1)$$

where each δ_i is in the normal form. In particular, all serial constraints are assumed to have been split into simpler order constraints. To extend `Apply` to such constraints, we only need to define:

$$\begin{aligned} \text{Apply}(\mathcal{C}_1 \vee \mathcal{C}_2, T) &\equiv \text{Apply}(\mathcal{C}_1, T) \vee \text{Apply}(\mathcal{C}_2, T) \\ \text{Apply}(\mathcal{C}_1 \wedge \mathcal{C}_2, T) &\equiv \text{Apply}(\mathcal{C}_1, \text{Apply}(\mathcal{C}_2, T)) \quad \square \end{aligned}$$

As before, it is easy to see that the above transformation preserves the unique event property.

Proposition 5.6 (Compiling General Constraints). Let T be a concurrent-Horn goal and let δ be a constraint of the form (5.1). Then $\text{Apply}(\delta, T) \equiv T \wedge \delta$.

Knots. After compiling the constraints \mathcal{C} into the graph \mathcal{G} , several things still need to be done. First, the result of the compilation, $\mathcal{G}_{\mathcal{C}}$, can have literals of the form $\neg\text{path}$ so, strictly speaking, $\mathcal{G}_{\mathcal{C}}$ is not a concurrent-Horn goal. However, we can use the following `CTR` tautologies to simplify $\mathcal{G}_{\mathcal{C}}$:

$$\begin{aligned} \neg\text{path} \otimes \phi &\equiv \phi \otimes \neg\text{path} \equiv \neg\text{path} \\ \neg\text{path} \mid \phi &\equiv \phi \mid \neg\text{path} \equiv \neg\text{path} \\ \neg\text{path} \vee \phi &\equiv \phi \vee \neg\text{path} \equiv \phi \end{aligned}$$

The result would be either a concurrent-Horn goal or $\neg\text{path}$.

If the result is not $\neg\text{path}$, this still does not mean that we have a directly executable workflow specification. The reason is that the *send/receive* synchronization primitives may cause a “deadlock”. In model-theoretic terms, this means that such a formula is `CTR`-equivalent to $\neg\text{path}$, and in proof-theoretic terms this means that the proof procedure would halt and declare that no execution exists. In this case, we rewrite $\mathcal{G}_{\mathcal{C}}$ into $\neg\text{path}$.

Also, when the proof procedure declares a failure, it produces a concurrent Horn goal, $\mathcal{G}_{\text{fail}}$, which in a sense is the smallest subpart of the original workflow that is inconsistent with the constraints. In this way, the workflow designers can be given a feedback that might help them find the bug in their specifications.

Even if the proof procedure of `CTR` does find a proof and thus $\mathcal{G}_{\mathcal{C}}$ is an executable workflow specification, $\mathcal{G}_{\mathcal{C}}$ may have sub-formulas where the *send/receive* primitives cause a cyclic wait, which we call *knots*.

The problem with knots is that, when they exist, finding an execution path in $\mathcal{G}_{\mathcal{C}}$ may not be a linear task (despite what we have promised in Section 4). Fortunately, it is easy to show that the proof theory of `CTR` can be used to excise all knots from $\mathcal{G}_{\mathcal{C}}$ in time linear in the size of $\mathcal{G}_{\mathcal{C}}$. This procedure, which we call `Excise`, yields either a knot-free concurrent-Horn goal equivalent to $\mathcal{G}_{\mathcal{C}}$, or $\neg\text{path}$, if $\mathcal{G}_{\mathcal{C}}$ is inconsistent.

Due to space limitation, we cannot fully describe `Excise` here, but we illustrate this process with the following example.

Example 5.7 (Knots). Let the graph \mathcal{G} be $\gamma \otimes (\eta \vee (\alpha \mid \beta \mid \eta))$ and let the constraints be as follows: $c_1 \equiv \neg\nabla\alpha \vee (\nabla\alpha \otimes \nabla\beta)$, $c_2 \equiv \neg\nabla\beta \vee (\nabla\beta \otimes \nabla\eta)$, $c_3 \equiv \neg\nabla\alpha \vee (\nabla\eta \otimes \nabla\alpha)$. The constraint c_1 says, *If α takes place, then β must also happen afterwards*. The other constraints have similar interpretation. Omitting some intermediate steps, we have:

$$\begin{aligned} \text{Apply}(c_1, \mathcal{G}) &= \text{Apply}(\neg\nabla\alpha, \mathcal{G}) \vee \text{Apply}(\nabla\alpha \otimes \nabla\beta, \mathcal{G}) = \mathcal{G}_1 \vee \mathcal{G}_2, \\ &\quad \text{where } \mathcal{G}_1 \equiv \gamma \otimes (\alpha \otimes \text{send}(\xi_1) \mid \text{receive}(\xi_1) \otimes \beta \mid \eta) \text{ and } \mathcal{G}_2 \equiv \gamma \otimes \eta \\ \text{Apply}(c_2, \mathcal{G}_1 \vee \mathcal{G}_2) &= \text{Apply}(c_2, \mathcal{G}_1) \vee \text{Apply}(c_2, \mathcal{G}_2) = \mathcal{G}_3 \vee \mathcal{G}_2, \\ &\quad \text{where } \mathcal{G}_3 \equiv \gamma \otimes (\alpha \otimes \text{send}(\xi_1) \mid \text{receive}(\xi_1) \otimes \beta \otimes \text{send}(\xi_2) \mid \text{receive}(\xi_2) \otimes \eta) \\ \text{Apply}(c_3, \mathcal{G}_3 \vee \mathcal{G}_2) &= \mathcal{G}_4 \vee \mathcal{G}_2, \\ &\quad \text{where } \mathcal{G}_4 \equiv \gamma \otimes (\text{receive}(\xi_3) \otimes \alpha \otimes \text{send}(\xi_1) \mid \text{receive}(\xi_1) \otimes \beta \otimes \text{send}(\xi_2) \mid \text{receive}(\xi_2) \otimes \eta \otimes \text{send}(\xi_3)) \end{aligned}$$

Finally, $\text{Excise}(\mathcal{G}_4 \vee \mathcal{G}_2) = \text{Excise}(\mathcal{G}_4) \vee \text{Excise}(\mathcal{G}_2)$. The proof procedure of CTR finds no knots in \mathcal{G}_2 , so $\text{Excise}(\mathcal{G}_2) = \mathcal{G}_2$. On the other hand, it detects a knot in \mathcal{G}_4 as follows.

First, the proof procedure “executes” γ and deletes it from \mathcal{G}_4 . This results in a goal where each concurrent conjunct starts with a *receive* and the corresponding *send*’s are slated to occur only later in the execution. Therefore, the proof procedure halts and we declare a knot in \mathcal{G}_4 . Thus $\text{Excise}(\mathcal{G}_4) = \neg\text{path}$. Hence, $\text{Excise}(\text{Apply}(c_1 \wedge c_2 \wedge c_3, \mathcal{G})) \equiv \mathcal{G}_2$. \square

Main results. We are now ready to summarize how the `Apply` and `Excise` transformations help solve the consistency, verification, and related problems. Theorems 5.8 through 5.10 assume that every activity in the workflow (except for the *receive* primitive) always succeeds. Without this assumption, only the “if”-part of Theorem 5.8 holds and its corollaries, Theorems 5.9 and 5.10, must be adjusted accordingly.

Theorem 5.8 (Consistency Checking). *Given a workflow specification $\mathcal{G} \wedge \mathcal{C}$, it is inconsistent iff $\text{Excise}(\text{Apply}(\mathcal{C}, \mathcal{G})) = \neg\text{path}$.*

Proof. Follows from Proposition 5.6 and the soundness and completeness of CTR proof theory. \square

Theorem 5.9 (Property Verification). *Given a workflow specification $\mathcal{G} \wedge \mathcal{C}$ and a property $\Phi \in \text{CONSTRAINT}$, there is a constructive way of verifying whether every execution of the workflow satisfies Φ .*

Proof. Φ is satisfied by every execution of the workflow if and only if $\text{Excise}(\text{Apply}(\neg\Phi \wedge \mathcal{C}, \mathcal{G})) = \neg\text{path}$. Otherwise, $\text{Excise}(\text{Apply}(\neg\Phi \wedge \mathcal{C}, \mathcal{G}))$ rewrites to the most general counter example where Φ fails to hold. \square

Theorem 5.10 (Redundancy Elimination). *Given a workflow specification $\mathcal{G} \wedge \mathcal{C}$ and a constraint $\Phi \in \mathcal{C}$, we can verify whether Φ is redundant.*

Theorem 5.11 (Complexity). *Let \mathcal{G} be a control flow graph and $\mathcal{C} \subset \text{CONSTRAINT}$ be a set of global constraints in the normal form of Corollary 3.5. Let $|\mathcal{G}|$ denote the size of \mathcal{G} , N be the number of constraints in \mathcal{C} , and d be the largest number of disjuncts in a constraint in \mathcal{C} . Then*

- *The worst-case size of $\text{Apply}(\mathcal{C}, \mathcal{G})$ is $O(d^N \times |\mathcal{G}|)$.*
- *The worst-case time complexity of applying `Excise` is proportional to the size of $\text{Apply}(\mathcal{C}, \mathcal{G})$.*

A simple corollary of Theorem 5.11 is: If \mathcal{C} consists of serial constraints only, then $d = 1$ and the size of $\text{Apply}(\mathcal{C}, \mathcal{G})$ is proportional to $|\mathcal{G}|$.

6 Conclusion

We presented a logic-based formalism for *specifying*, *verifying*, and *executing* workflows. We developed an algorithm for consistency checking of workflows and for their property verification. The algorithm compiles global constraints on workflow execution into the control flow graph. In addition to solving the consistency and verification problems, this compilation technique helps optimize the run-time scheduling of workflow events.

Our work can be extended to handle additional features, such as the following:

Transition conditions. As a compilation technique, our algorithm cannot fully account for control flow graphs with transition conditions that query the database. For such graphs, our algorithm is sound but not complete for consistency and verification problems. However, if additional semantic information about the database is available, it can be incorporated into our framework to yield more accurate results.

Sub-workflows. It is straightforward to augment our compilation technique to handle sub-workflows defined via concurrent-Horn rules. Furthermore, when global dependencies do not span sub-workflow boundaries, the complexity reported in Theorem 5.11 can be brought down somewhat. Indeed, it can be shown that, if M is the largest number of dependencies in a sub-workflow, then the size of $\text{Apply}(\mathcal{C}, \mathcal{G})$ is $O(d^M \times |\mathcal{G}|)$.

Loops and iteration. Loops in control flow graph can be expressed using recursive CTR rules. Our techniques can be extended to handle workflows with loops. However, in this case, we can only obtain conservative algorithms for testing inconsistency of workflow specifications.

Failure semantics. Failure atomicity is built into CTR semantics. However, more complex workflows require more advanced failure semantics, such as *compensation* [13]. Some such semantics can be expressed using the possibility operator of CTR , \diamond . Work is in progress on extending our framework to handle other failure semantics.

Acknowledgements. The authors would like to thank Tony Bonner for the helpful comments on a draft of this paper.

References

- [1] G. Alonso, D. Agrawal, A. El Abbadi, M. Kamath, R. Günthör, and C. Mohan. Advanced transaction models in workflow contexts. In *International Conference on Data Engineering*, New Orleans, Louisiana, February 1996.
- [2] G. Alonso, D. Agrawal, A. El Abbadi, and C. Mohan. Functionality and limitations of current workflow management systems. In *IEEE-Expert (to appear in a special issue on Cooperative Information Systems)*, 1997.
- [3] P. Attie, M. Singh, A. Sheth, and M. Rusinkiewicz. Specifying and enforcing intertask dependencies. In *Intl. Conference on Very Large Data Bases*, 1993.
- [4] A.J. Bonner and M. Kifer. An overview of transaction logic. *Theoretical Computer Science*, 133:205–265, October 1994.
- [5] A.J. Bonner and M. Kifer. Transaction logic programming (or a logic of declarative and procedural knowledge). Technical Report CSRI-323, University of Toronto, November 1995. <http://www.cs.toronto.edu/~bonner/transaction-logic.html>.
- [6] A.J. Bonner and M. Kifer. Concurrency and communication in transaction logic. In *Joint Intl. Conference and Symposium on Logic Programming*, pages 142–156, Bonn, Germany, September 1996. MIT Press.
- [7] A.J. Bonner, M. Kifer, and M. Consens. Database programming in transaction logic. In A. Ohori C. Beeri and D.E. Shasha, editors, *Proceedings of the International Workshop on Database Programming Languages, Workshops in Computing*, pages 309–337. Springer-Verlag, February 1994. Workshop held on Aug 30–Sept 1, 1993, New York City, NY.
- [8] O. Bukhres and E. Kueshn, Eds. Special issue on software support for work flow management. *Distributed and Parallel Databases—An International Journal*, 3(2), April 1995.
- [9] Workflow Management Coalition. Terminology and glossary. Technical Report (WFMC-TC-1011), Workflow Management Coalition, Brussels, 1996.
- [10] U. Dayal, M. Hsu, and R. Ladin. Organizing long-running activities with triggers and transactions. In *ACM SIGMOD Conference on Management of Data*, 1990.
- [11] A. Elmagarmid, Y. Leu, W. Litwin, and M. Rusinkiewicz. A multidatabase transaction model for interbase. In *Intl. Conference on Very Large Data Bases*, 1990.

- [12] A.K. Elmagarmid, editor. *Database Transaction Models for Advanced Applications*. Morgan-Kaufmann, San Mateo, CA, 1992.
- [13] H. Garcia-Molina and K. Salem. Sagas. In *Intl. Conference on Very Large Data Bases*, pages 249–259, May 1987.
- [14] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman and Company, San Francisco, CA, 1978.
- [15] D. Georgakopoulos, M. Hornick, P. Krychniak, and F. Manola. Specification and management of extended transactions in a programmable transaction environment. In *International Conference on Data Engineering*, Houston, TX, February 1994.
- [16] D. Georgakopoulos, M. Hornick, and A. Sheth. An overview of workflow management: From process modeling to infrastructure for automation. *Journal on Distributed and Parallel Database Systems*, 3(2):119–153, April 1995.
- [17] R. Gunthor. Extended transaction processing based on dependency rules. In *Proceedings of the RIDE-IMS Workshop*, 1993.
- [18] M. Hsu, Ed. Special issue on workflow systems. *Bulletin of the Technical Committee on Data Engineering (IEEE Computer Society)*, 18(1), March 1995.
- [19] M. Kifer. Transaction logic for the busy workflow professional. Unpublished manuscript. <ftp://ftp.cs.sunysb.edu/pub/TechReports/kifer/tr-for-wf.ps.Z>, August 1996.
- [20] J. Klein. Advanced rule-driven transaction management. In *IEEE COMPCON*. IEEE, 1991.
- [21] M.E. Orlowska, J. Rajapakse, and A.H.M. ter Hofstede. Verification problems in conceptual workflow specifications. In *Intl. Conference on Conceptual Modelling*, volume 1157 of *Lecture Notes in Computer Science*, Cottbus, Germany, 1996. Springer-Verlag.
- [22] M. Rusinkiewicz and A. Sheth. Specification and execution of transactional workflows. In W. Kim, editor, *In Modern Database Systems: The Object Model, Interoperability, and Beyond*. ACM Press, 1994.
- [23] M.P. Singh. Semantical considerations on workflows: An algebra for intertask dependencies. In *Proceedings of the International Workshop on Database Programming Languages*, Gubbio, Umbria, Italy, September 6–8 1995.
- [24] M.P. Singh. Synthesizing distributed constrained events from transactional workflow specifications. In *Proceedings of 12-th IEEE Intl. Conference on Data Engineering*, pages 616–623, New Orleans, LA, February 1996.
- [25] H. Wachter and A. Reuter. The ConTract model. In [12], chapter 7, pages 220–263. 1992.