Specification and Verification of Transaction Atomicity, Isolation and Temporal Correctness

Simin Cai*, Barbara Gallina†, Dag Nyström‡, and Cristina Seceleanu§

Mälardalen Real-Time Research Centre

Mälardalen University, Västerås, Sweden

Email: *simin.cai@mdh.se, †barbara.gallina@mdh.se, ‡dag.nystrom@mdh.se, §cristina.seceleanu@mdh.se

Abstract

Although transaction atomicity, isolation and temporal correctness are crucial to the dependability of many real-time database-centric systems, the selected assurance mechanism for one property may lead to violating another. To find an appropriate trade off, one needs to be able to specify and reason about the dependencies between atomicity, isolation and temporal correctness, together with the selected mechanisms such as abort recovery, concurrency control, and scheduling. However, existing design techniques do not take all these properties into account. In this paper, we propose a UML profile, called UTRAN, for specifying atomic concurrent real-time transactions, with explicit support for all three properties and their supporting mechanisms. We also propose a pattern-based modeling framework, called UPPCART, to formalize the transactions and the mechanisms specified in UTRAN, as UPPAAL timed automata. The UTRAN specification can be transformed into a UPPCART model via the proposed reusable patterns, which allow flexible modeling of various abort recovery, concurrency control, and scheduling mechanisms. The desired properties are then formalized using our proposed specification patterns, and verified by the UPPAAL model checker. Our approach facilitates systematic design of desired atomicity, isolation and temporal correctness trade offs with guarantee, thus contributing a dependable real-time database system.

Keywords

Transaction, Atomicity, Isolation, Temporal Correctness, UML, Model Checking

I. Introduction

In many database-centric systems, critical data such as account balance and configuration parameters are stored in databases and managed by DataBase Management Systems (DBMS). To maintain the consistency of data, a DBMS organizes operations as transactions, and manages them with various Abort Recovery (AR) and Concurrency Control (CC) mechanisms [1]. Abort recovery restores the database to a consistent state when a transaction is aborted due to errors, and thus achieves atomicity. Rollback, which undoes all changes of an aborted transaction, is a common AR technique [2]. Concurrency control prevents inconsistency by regulating the concurrent access to data from different transactions. Locks are often applied to avoid arbitrary access of data, as a widely applied CC technique [1]. Together, AR and CC ensure the critical data to be dependable for applications relying on them.

Many database-centric systems are also time-critical, such as industrial control systems [3] and automotive systems [4], whose configurations and states can be stored in databases. Reading an outdated sensor value, or fetching the calibration parameter too late, could result in catastrophic consequences such as loss of lives. In such real-time database systems, therefore, transactions must also be temporally correct, meaning that they must use fresh data, and have to meet specified deadlines [5]. The assurance of atomicity and isolation, however, may jeopardize the assurance of temporal correctness. CC may cause a transaction to be blocked for a long time. AR introduces extra workload when performing recovery. To make matters worse, some CC algorithms may directly abort transactions, while the recovery may again lock data and block other transactions further, which could lead to deadline misses. Therefore, trade offs may need to be considered during the design of Real-Time DBMS (RTDBMS) [6], with respect to the decision of the "variants" [7] of atomicity and isolation, and the selection of the AR and CC mechanisms.

To achieve an appropriate trade off, it is helpful to be able to specify all three properties, as well as the AR and CC mechanisms, explicitly in a high-level language familiar to system designers. To ensure the correctness of the trade off, one should be able to analyze such specifications, and reason about whether the properties can be satisfied with the selected mechanisms. To our knowledge, although the specification and analysis of one or two of these properties have been targeted by the research community, existing techniques do not take all three into consideration. However, since atomicity, isolation and temporal correctness are closely dependent, missing one of them in the analysis may result in the risk in the dependability of the entire RTDBMS.

The contribution of this paper is two-fold. First, we propose a UML (Unified Modeling Language) profile as an extension of the Activity Diagram [8], for the specification of transactions, with explicit support for atomicity, isolation and temporal correctness properties. We choose UML because it has become a well-accepted modeling language for software systems, including real-time systems and database systems, thus can minimize the learning effort for designers. The extension is presented as a UML profile, called UTRAN, which models a transaction as an activity, and includes modeling elements to express abort recovery mechanisms such as rollback and compensation, as well as isolation levels and concurrency control. Time-related properties such as deadlines and periods, which are reused from the UML-MARTE profile [9] for expressing timing information in real-time systems, can be annotated to transactions and operations.

Second, to facilitate the analysis of all three properties, we extend our UPPAAL-based analysis framework [10], which models only transactions with isolation and temporal correctness concerns in UPPAAL Timed Automata (TA) [11]. The new framework, called UPPCART (UPPaal for Concurrent Atomic Real-time Transactions), models transactions with encoded timing information, the selected AR, CC and scheduling mechanisms, and the inconsistency to be avoided, as a network of UPPAAL TA. To reduce the modeling effort, we propose a set of reusable basic modeling units, for modeling various CC algorithms and AR mechanisms, as well as the transactions. Specifications in UTRAN can potentially be automatically transformed into UPPCART models. We also propose patterns for formalizing the atomicity, isolation and temporal correctness properties as UPPAAL specifications [11]. The formalized properties can then be verified rigorously by the tool UPPAAL model checker, which provides a guarantee for the correctness of the designed trade off.

The remainder of the paper is organized as follows. In Section II we present the preliminaries of the paper. In

Section III and Section IV, we introduce our proposed UTRAN profile and UPPCART framework, respectively. The verification of the properties is presented in Section V. We discuss the related work in Section VI, after which we conclude the paper and outline future work in Section VII.

II. Background

In this section, we present the preliminaries of this paper, including the concepts of transactions, atomicity, isolation and temporal correctness (Section II-A), UML profiles (Section II-B), and UPPAAL TA (Section II-C).

A. Transactions: Atomicity, Isolation and Temporal Correctness

In database systems, clients read and write data through a database management system, which guarantees data consistency via transaction management. A *transaction* is a partially-ordered set of logically-related operations, called a Work Unit (WU), that as a whole ensures the *ACID* properties [2]. ACID refers to: *Atomicity* (a transaction either runs completely or makes no changes at all), *Consistency* (a transaction executing alone must not violate logical constraints), *Isolation* (concurrent transactions do not interfere each other), and *Durability* (committed changes are made permanent). The lifecycle of a transaction is managed by the following operations: begin (start a transaction), commit (terminate a transaction while making its changes permanent and visible), and abort (terminate a transaction and recover from its changes). Two types of aborts exist in a database system. System aborts are initiated by the DBMS due to system errors or data contentions. User aborts are issued by clients to stop the transaction deliberately according to the application semantics.

a) Atomicity: Under full atomicity assumption, "commit" means the completion of "all" changes, and "abort" means that "nothing" is changed. In this paper, we are particularly interested in the recovery of transactions terminated by errors. Therefore, our semantics of "commit" remains "all", while "abort" could have various meanings depending on the variants of atomicity.

We refer to the "nothing" semantics of full atomicity as *failure atomicity*. Failure atomicity is achieved by *rollback*, a recovery mechanism to undo all changes and return to the states before the transaction starts when it gets aborted [2]. Let us assume a transaction T_0 that writes data D_0 and D_1 in sequence. Failure atomicity will reverse the changes of D_0 and D_1, in case T_0 is aborted before its commitment. Since failure atomicity may be restricted in terms of performance and functionality, a number of variants of *relaxed atomicity* have also been proposed, which allow changes to be partially undone, or recover inconsistency with extra operations [12]. The following abort recovery mechanisms that support relaxed atomicity are considered in this paper. *Immediate compensation* recovers inconsistency due to abort by immediately executing a sequence of operations, such as to update D_2 that represents the error state. *Deferred compensation*, in contrast, executes a transaction which is scheduled as a normal one to restore consistency. In both variants, the operations are designed flexibly, depending on the application semantics. An atomicity manager, which possesses the knowledge of the atomicity variants, performs the recovery at runtime [13].

b) Isolation: In literature, isolation has been quantified as various levels [14]. An *isolation level* is defined as the avoidance of a particular set of *phenomena*, which are interleaved transaction executions that can lead to

inconsistent data. If we use r_i^j to denote that transaction T_i reads data D_j, w_i^j to denote that T_i writes D_j, the following execution is considered as a phenomenon: <r_0^0, w_1^0, w_1^1, r_0^1>, representing the execution "T_0 reads D_0, T_1 writes D_0, T_1 writes D_1, T_0 reads D_1". In this example, T_0 reads an old version of D_0 before the change of T_1, but a new version of D_1 after the change of T_1. If D_0 and D_1 are a pair of configuration parameters that should be compatible, the consequence of T_0 using these inconsistent parameters may result in unsafe system behaviors. An isolation level precludes a subset of such phenomena, thus avoiding the inconsistency. Isolation levels are also a flexible way to relax isolation, as the precluded phenomena are adjustable according to the particular semantics.

DBMS ensures isolation by associating a concurrency control manager to the managed data, which regulates the interleaved transaction executions according to a selected CC algorithm [1]. *Pessimistic Concurrency Control (PCC)*, a family of CC algorithms commonly applied in DBMS [1], is considered in this paper. PCC exploits locks to prevent unwanted interleavings. Depending on the algorithm, a transaction needs to hold a specific type of lock, before reading or writing the data. Locks are acquired at a certain time point before the operations, and are released at a certain time point afterwards. Upon receiving requests, the CC manager decides which transactions should obtain the lock, wait, or even be aborted, according to the selected algorithm. The atomicity manager, in case a transaction gets aborted by CC, performs the abort and recovery of the transaction.

c) Temporal Correctness: In a real-time database system, *temporal correctness* involves the transaction *timeliness*, and *temporal data consistency* [5]. Transactions need to meet their deadlines, which is referred to as timeliness [5]. Temporal data consistency includes two aspects. *Absolute validity* requires that data read by a transaction must not be older than a specified validity interval. *Relative validity* requires that, if a transaction reads a group of data, these data must be generated within a specified interval so that the results are temporally correct. Temporal correctness is directly influenced by the scheduling policy adopted by the RTDBMS, which schedules the operations issued clients. Commonly applied scheduling policies include First-In-First-Out (FIFO), round-robin, or based on the priorities of the transactions [1]. In addition to deadlines and validity intervals, other important time-related information includes execution times of the operations, and the arrival patterns of transactions (that is, whether a transaction is started with a period, with a bounded inter-arrival interval, or randomly) [5].

The ACID properties often need to be relaxed in order to guarantee temporal correctness [6], for instance, by using compensation rather than rollback for relaxed atomicity [15]. Real-time characteristics of transactions are incorporated in many CC algorithms for better timeliness. A widely applied real-time PCC is *Two Phase locking - High Priority (2PL-HP)* [16]. In this algorithm, a transaction acquires a readlock (write) on data before it performs a read (write) operation, and releases all locks during commitment. If two transactions try to lock the same data, and at least one of them requires a write lock, a CC conflict occurs. The transaction with higher priority will be granted with the lock, while the transaction with lower priority will be aborted by the RTDBMS. As a result, transactions with higher priorities are more likely to meet their deadlines.

B. UML Profiles and MARTE

UML is one of the most widely accepted modeling language in software development [8]. The profile mechanism is designed to extend UML for languages customized for specific domains. A profile defines a package of stereotypes, which are domain-specific concepts derived from existing UML concepts, and constraints to associate them. A

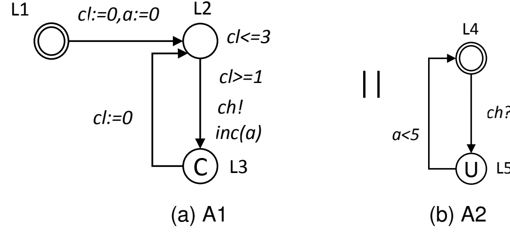

Fig. 1. A network of timed automata

stereotype can have tagged values as attributes. Profiles may be used as specification languages to model systems, or adopted to add supplementary information that is used for analysis or code generation [17]. As timing information is essential for our analysis and thus needs to be supported in the specifications, we reuse the relevant concepts from the MARTE (Modeling and Analysis of Real-Time Embedded systems) [9] profile. MARTE is a profile that defines the basic concepts to support the modeling of real-time and embedded applications, as well as to provide information for performance and schedulability analysis. In this paper, we propose a profile that encodes the information of transaction management and transactional properties for formal analysis. The following MARTE concepts are reused: (i) MARTE::NFP_Duration, a data type for time intervals; (ii) MARTE::ArrivalPattern, a data type for arrival patterns, such as periodic, sporadic and aperiodic patterns.

C. UPPAAL Timed Automata (TA)

Timed Automata are finite-state automata extended with real-valued clock variables [18]. UPPAAL TA extends TA with discrete data variables, synchronization channels, user-defined functions, among other modeling features [11]. Multiple UPPAAL TA can form a network of TA via parallel composition, in which individual automaton can carry out its internal actions, while pairs of automata can perform hand-shake synchronization.

We use an example to illustrate UPPAAL TA. Fig. 1 shows a network of TA modeling a simple concurrent real-time system, in which automaton A1 sporadically increments a variable *a* and synchronizes with automaton A2. A1 consists of a set of locations (L1, L2 and L3), and edges connecting them. A clock variable *cl* is defined in A1 to measure the elapse of time, and progresses continuously. A discrete variable *a* is defined globally, and shared by A1 and A2. Semantically, a state of the network of TA consists of the current locations of the automata, together with the values of the clock and discreet variable. At each location, an automaton may stay at the location, as long as the **invariant**, which is a conjunction of clock constraints associated with the location, is satisfied. Alternatively and non-deterministically, the automaton may take a transition along an edge, if the **guard**, which is a conjunction of constraints on discrete or clock variables associated with the edge, is satisfied. In Fig. 1, A1 may delay in L2 as long as cl $\leq$ 3, or follow the edge to L3 when cl $\geq$ 1. Each edge may have an associated action, which is the synchronization with other automata via a **channel**. Binary channels are used to synchronize one sender (indicated by a mark "!") with a single receiver (indicated by a mark "?"). In Fig.1, A1 sends a synchronization message ti A2 via binary channel *ch*, while it takes the edge from L2 to L3. The synchronization can take place only if both the sender and the receiver are ready to fire the edge. A broadcast channel is used to pass messages between one sender and an arbitrary number of receivers. Using broadcast channels, the sender does not block even if some of the receivers are not ready. An edge may have an *assignment*, which resets the clocks or updates discreet variables when the edge is fired. In UPPAAL TA, both guards and assignments can be encoded as user-defined function in

a subset of C language, which brings high flexibility and expressiveness for modeling. In our example, when A1 moves from L2 to L3, a is incremented using the function *inc(a)*.

A location marked as "U" is an urgent location, meaning that the automaton must leave the location without delay in time. Another automaton may take transitions as long as the time does not progress. A location marked as "C" is a committed location, which indicates no delay in time, and immediate transition. In this case, another automaton may NOT take any transitions, unless it is also at a committed location. L3 is a committed location.

The UPPAAL model checker can verify properties specified in the UPPAAL language [11], called UPPAAL queries, which is a decidable subset of Computation Tree Logic (CTL) [19]. For instance, the invariance property "A1 never reaches location L3" can be specified as "$A[\,]\,not\,A1.L3$", in which "A" is a path operator and reads "for all paths", whereas "$[\,]$" is the "always" temporal operator. If a safety property is not satisfied, the model checker will provide a counterexample. The liveness property "If A1 reaches L2, it will eventually reach L3" can be specified, using the "lead-to ($\rightarrow$)" operator, as "$A1.L2 \rightarrow A1.L3$".

Our previous work has proposed to model a concurrent transaction system as a network of UPPAAL TA, denoted as N', which is defined as follows [10], [20]:

$$N' ::= A_0 \parallel ... \parallel A_{n-1} \parallel A_{CCManager} \parallel O_0 \parallel ... \parallel O_{k-1} \parallel D_0 \parallel ... \parallel D_{l-1},$$

where A_0, ..., A_{n-1} are the TA of work units of transactions T_0, ..., T_{n-1}, respectively. $A_{CCManager}$ is the CCManager automaton that models the CC algorithm. O_0, ..., O_{k-1} are the TA of IsolationObservers that observe the phenomena to be precluded by isolation, respectively. D_0, ..., D_{l-1} are the TA that monitors the time of data. Isolation and temporal correctness can then be verified by UPPAAL model checker. We extend this framework in this paper to include atomicity and abort recovery.

III. UTRAN Profile for Specification of Atomic Concurrent Real-time Transactions

In this section, we present our proposed UML profile for real-time transactions, called UTRAN. We first present the domain model of real-time transactions in Section III-A, after which we introduce the UML profile diagram in Section III-B, and illustrate its usage using an example in Section III-C.

A. Domain View

The domain model of real-time transactions is presented in Fig. 2. A transaction can be conceptually modeled as an activity in the UML activity diagram. In order for the RTDBMS to manage the life cycle of a transaction, a unique *id* is assigned to each transaction when it is started. A transaction may be assigned with a *TemporalCorrectnessSpecification* for time-related properties. This specification may specify a *priority* for the transaction, and a *relative deadline* that defines the maximum allowed time interval between the start and the termination of the transaction. The arrival *pattern* can be specified for the transaction, such as periodic, sporadic and aperiodic, as well as the value of the *period* (or minimum inter-arrival time) if applicable. A transaction may also have a specified *relative validity interval*, for the validity of a group of data read by the transaction.

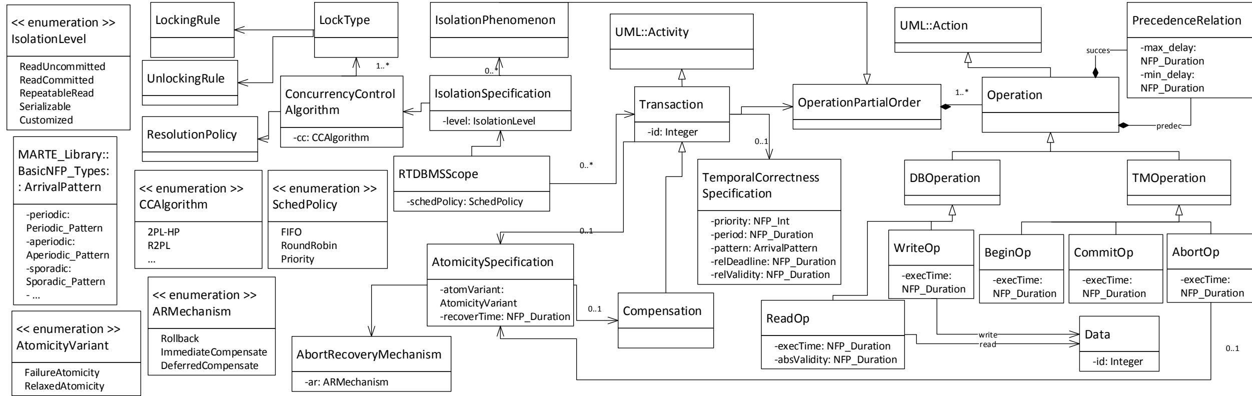

Fig. 2. Domain model of real-time transactions

A transaction consists of a set of partially-ordered operations, which are represented as actions in an activity. Two types of operations are considered explicitly in a transaction: *DBOperations* and *TMOperations*. DBOperations directly performs read and write access to the data. Such read and write operations, denoted as *ReadOP* and *WriteOP* respectively, are atomic, whose *worst-case execution times* are also known. A ReadOP may be assigned with an *absolute validity interval* for the data it reads.

Atomicity and isolation of transactions are also included in the domain model. Multiple transactions managed by the same RTDBMS are related to an *RTDBMSScope*. A *scheduling policy*, which can be FIFO, RoundRobin and Priority-based, can be specified for the RTDBMSScope. An *isolation level* is specified for the RTDBMSScope, indicating the degree of isolation that should be provided for the set of transactions. Essentially, an isolation level defines a set of *IsolationPhenoma* to be precluded, which are illegal sequence of operations that lead to logical inconsistency. Therefore, an IsolationPhenomenon is basically an OperationPartialOrder. While some isolation levels have been defined in the SQL-92 standards (ReadUncommitted, ReadCommitted, RepeatableRead, Serializable) which can be selected from, customized isolation levels can also be defined with specified IsolationPhenoma.

A *ConcurrencyControlAlgorithm* specifies the lock-based concurrency control algorithm selected for the specified isolation. Such an algorithm defines a set of *lock types*, each having its rules about, not only to which data it should apply, but also when a lock should be acquired and released. These rules are specified as *LockingRule* and *UnlockingRule*. A ConcurrencyControlAlgorithm also needs to specify a *resolution policy*, which describes how the conflicts are resolved when two transactions try to lock the same data.

An *AtomicitySpecification* specifies the atomicity variant to conduct in order to restore consistency when it gets terminated by error, as well as the desired *recovery time*. An AtomicitySpecification can either be attached to a transaction, which specifies the atomicity handling when the latter is aborted by the transaction management system, or to an abort operation, specifying the handling of abort issued by the clients. An AtomicitySpecification contains an *AtomicityVariant*, which is an enumeration of the supported atomicity variants, including *FailureAtomicity* and *RelaxedAtomicity*. An *AbortRecoveryMechanism* is associated with the atomicity variant. For FailureAtomicity, *Rollback* is the AR mechanism, by which the RTDBMS will undo all the changes in the database that have been done by aborted transaction. *ImmediateCompensate* and *DeferredCompensate* are the AR mechanisms for RelaxedAtomicity by compensation to restore the consistency of the database. The difference is that, the former

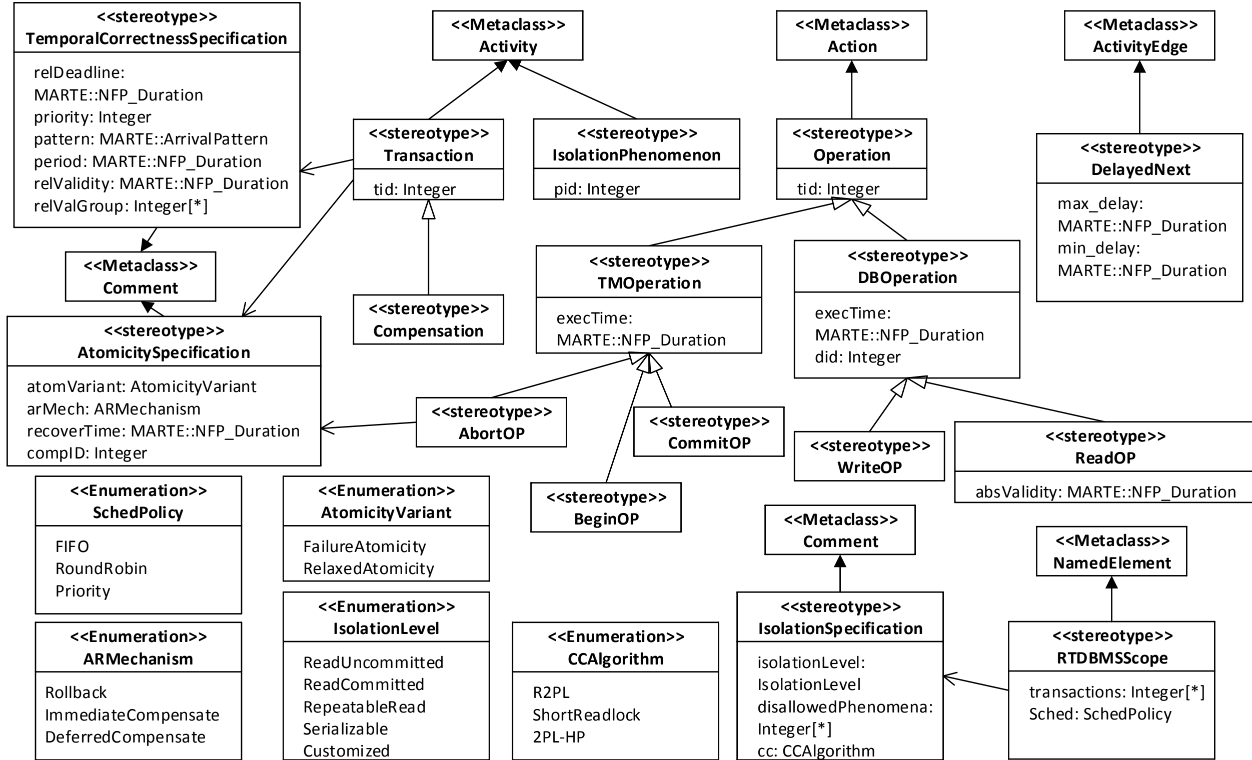

Fig. 3. UTRAN profile for real-time transactions

allows the compensation to be executed immediately with highest priority, while in the latter case the compensation is scheduled as a separate transaction with the same priority as the aborted one. If no AtomicitySpecification is specified, atomicity is totally relaxed, and the partially changed data will not be recovered or compensated at all.

B. *Profile*

This subsection describe our UTRAN profile that contains the extensions to model the concepts in the previous domain model. The profile diagram is presented in Fig. 3.

The «Transaction» stereotype, extending the UML Activity metaclass, maps the Transaction domain element. Each activity stereotyped with «Transaction» may have a «TemporalCorrectnessSpecification» and an «AtomicitySpecification», which are associated comments that extends the Comment metaclass. A «TemporalCorrectnessSpecification» contains the information about the deadline, priority, arrival pattern, period, and relative validity of the transaction. An «AtomicitySpecification» specifies the selected AtomicityVariant and ARMechanism, from an enumeration of supported variants, as well as the recovery time, and the id of the compensation transaction which is a special transaction specified by the stereotype «Compensation».

The actions in a «Transaction» are stereotyped as «Operation», each having the transaction id it belongs to. «DBOperation», «TMOperation» and «ClientOperation» map the DBOperation, TMOperation and ClientOperation, respectively. A «DBOperation» specifies the execution time to execute such an operation, and the id of the data it accesses. «ReadOP» and «WriteOP» extends «DBOperation», to map the ReadOP and WriteOP, respectively. A

«TMOperation» specifies the execution time for the transaction management operation, which can be «BeginOP», «CommitOP», or «AbortOP».

The stereotype «RTDBMSScope» maps the RTDBMSScope concept, which contains the transactions managed by the RTDBMS. The stereotype also specifies the CC algorithm and scheduling policy, selected from the enumerations CCAlgorithm and SchedPolicy. The stereotype «IsolationSpecification» maps IsolationSpecation in the domain view, which specifies the isolation level, as well as the disallowed phenomena explicitly. Each phenomenon is modeled as an activity stereotyped as «IsolationPhenomenon», which contains a sequence of actions stereotyped as «Operation».

C. Example

Let us take two autonomous wheel loaders and their controller as an example. All data, including the positions of the wheel loaders, their working conditions, the work plan, and the speed configurations, are stored in the controller's database. Assume that loader A is patrolling on the location with its predefined speed, and periodically update its position in the database. A human operator can update the configuration data, that is, the work plan, and the speed settings of the wheel loader A. These data should be updated at the same time. If the update fails, the data should be rolled back to the previous values. The controller may get a command to start a job with loader B. It updates the status of the job to "start", reads the current location of loader A, read the work plan in the database, read the speed configuration of loader A, and calculate the estimated speed and direction of the wheel loader B. Loader B then moves to the position and informs the controller, which updates the job to "finish". If this job fails, due to some reason, the estimated position should be updated as "unknown", and the job status as "failed".

We consider three transactions in this scenario. The first transaction (UpdateConfTrans) updates the configuration data. The second (JobTrans) controls the loader B to do the job. The third transaction (UpdateLocA) updates the location of A periodically. UpdateConfTrans and JobTrans, as well as the specifications of their atomicity, isolation and temporal correctness properties, are specified in UTRAN and shown in Fig. 4, while UpdateLocA is omitted. The temporal correctness properties are specified in their respective attached «TemporalCorrectnessSpecification», with their deadlines and validity intervals. The atomicity variant of UpdateConfTrans, which is rollback, is specified in its «AtomicitySpecification». On the contrary, JobTrans selects ImmediateCompensate, as specified in its «AtomicitySpecification», and compensates its failure with the compensation transaction LogError, which updates the estimated position and logs the error. The transactions are in the scope of the RTDBMS, stereotyped with «RTDBMSScope», whose isolation level is set to be RepeatableRead in its «IsolationSpecification», which disallow the phenomena InconsistencyConfigs1 and InconsistencyConfigs2, both stereotyped as «IsolationPhenomenon».

IV. UPPCART Framework for Modeling Atomic Concurrent Real-time Transactions

In this section, we propose a pattern-based framework, called UPPCART (UPPaal for Concurrent Atomic Real-time Transactions), for modeling real-time transactions with concurrency control and abort recovery in UPPAAL TA. The proposed framework is based on our previous work [10], which focuses on isolation and timeliness only. In this work, we extend the previous framework in several substantial aspects. First, we extend the framework to incorporate the modeling of atomicity, via a set of reusable patterns for modeling various abort recovery mechanisms, especially their interplays with concurrency control. We also extend existing patterns for time-related behaviors, such

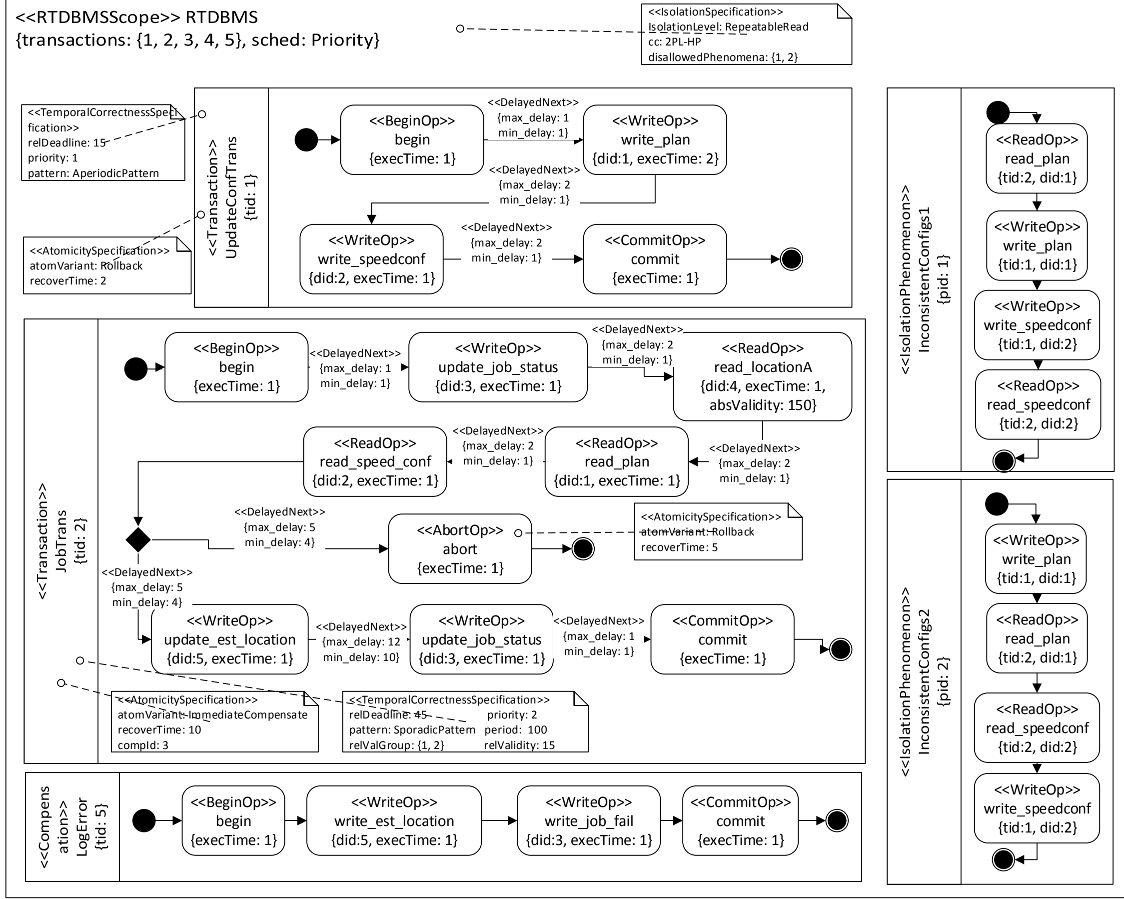

Fig. 4. Activity diagram of the wheel loader example using our UTRAN profile

as periodicity and delays. In addition, the proposed framework improves reusability of existing modeling units for transactions and CC, by proposing a unified reusable modeling unit for all read, write, begin and commit operations, as well as a more generic unit for CCManager that is suitable for a wider range of concurrency control algorithms. In addition, we propose an algorithm to construct the UPPAAL TA models from a UTRAN specification.

We model the transactions, together with the CC algorithm and the AR mechanisms, as a network of UPPAAL TA, denoted as N, which is defined as follows:

$$N ::= A_0 \parallel ... \parallel A_{n-1} \parallel A_{CCManager} \parallel O_0 \parallel ... \parallel O_{k-1} \parallel A_{ATManager},$$

where A_0, ..., A_{n-1} are the TA of work units of transactions T_0, ..., T_{n-1}, respectively. They also model the WU's interaction with the transaction manager with respect to concurrency control and abort recovery. $A_{CCManager}$ is the CCManager automaton that models the CC algorithm, and interacts with the work unit TA. O_0, ..., O_{k-1} are the TA of IsolationObservers that observe the phenomena to be precluded by isolation, respectively. $A_{ATManager}$ is the ATManager automaton that models the atomicity controller of recovery mechanisms upon abort of transactions.

We define two types of reusable structures for constructing the TA models. A *pattern*, consisting of a set of variables, locations, edges and even other patterns, are parametrized structures representing the repetitive modeling

TABLE I. MODELING ELEMENTS OF THE WORK UNIT SKELETON

Element	Type	Explanation
ti	parameter	transaction id
p	parameter	transaction priority
PERIOD	parameter	period/minimal inter-arrival time of the transaction
DEADLINE	parameter	deadline of transaction commitment
RECOVERY_DEADLINE	parameter	deadline of transaction recovery
tc	clock variable	tracking the elapsed time of the transaction
tr	clock variable	tracking the elapsed time of abort recovery
abort_trans[ti]	channel	communication with ATManager
initialize(ti, p)	function	initialization of the transaction

units in our framework. A pattern can be composed with the rest of the automaton after instantiation. A *skeleton* is a special type of pattern that defines the basic structure of a type of constituent automata of N, that is, a work unit skeleton, a CCManager skeleton, an ATManager skeleton, and an IsolationObserver skeleton.

In the following texts, we first introduce the details of UPPCART in Section IV-A, followed by an algorithm to construct UPPCART models from UTRAN in Section IV-B, and an illustrative example in Section IV-C.

A. The Proposed Modeling Framework

The skeletons and patterns for the various automata in UPPCART are presented as follows. We first introduce the skeletons and patterns for work units and concurrency control, as well as the skeleton of IsolationObserver. After this, we present the skeletons and patterns for atomicity and abort recovery mechanisms, and show how they are integrated with the work units and the CCManager.

1) Modeling Work Units:

a) ***Work Unit Skeleton****:* A WU automaton models the work unit of a transaction and its interaction with the CC and atomicity managers. A WU skeleton, as shown in Fig. 5, is a parametrized structure that consists of the common variables, locations and edges of a WU automaton. The parameters, as well as other modeling elements, are listed in Table I. Starting from the *initial* location, the automaton immediately initializes the transaction with the specified id *ti* and priority *p* using function *initialize(ti, p)*, and moves to the location *ready*. After an arbitrary delay, it moves to the location *begin*, indicating the begin of the transaction, and sets clock variable *tc* to 0. The location *trans_committed* represents the committed state of the transaction. Between *begin* and *trans_committed* are a set of connected instantiated operation patterns that model the database and transaction management operations, and delays between the operations. If the value of *tc* is greater than the specified *DEADLINE*, the automaton moves to the location *miss_deadline*, indicating a deadline miss. Otherwise, it waits until the specified *PERIOD* has reached, and moves to *begin* for the next activation. During the operations, the WUA may receive a message from the atomicity manager ATManager via channel *abort_trans[ti]*, and moves to the instantiated abort recovery pattern, which models the AR mechanism. The location *trans_aborted* represents the aborted state of the transaction. Similarly, if the value of *tr* is greater than a specified *RECOVERY_DEADLINE*, timeliness is breached, and the WUA moves to *miss_deadline*.

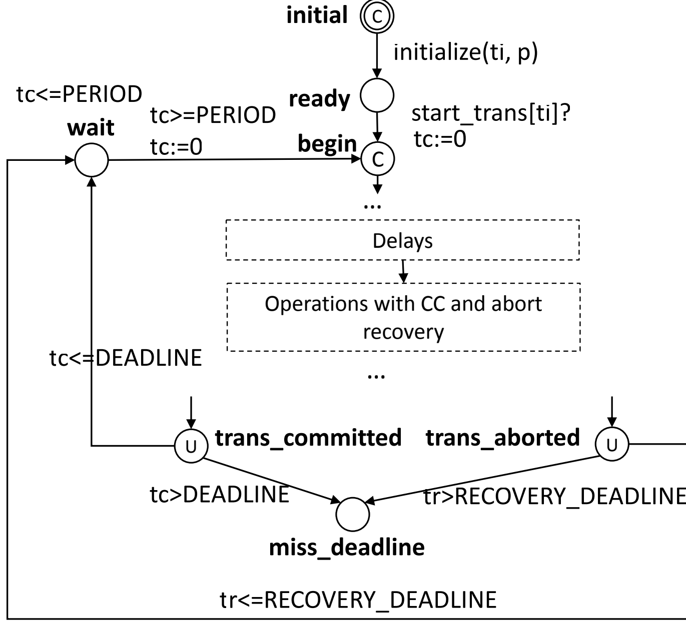

Fig. 5. TA skeleton for a work unit

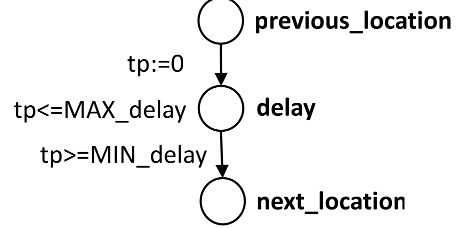

Fig. 6. Delay pattern

*b) **Delay Pattern**:* The pattern in Fig 6 models the delays between operations. The automaton may stay at location *delay* for at most (least) *MAX_delay* (*MIN_delay*) time units, which are provided as parameters.

2) Modeling CC and Isolation Phenomena:

*a) **Operation-CC, Locking and Unlocking Patterns**:* We define a pattern to model the begin, commit, read and write operations in each work unit. Since within each operation, the work unit may interact with the CC manager according to the specific CC algorithm, our operation pattern also comprises CC-related activities such as locking and unlocking as sub-patterns. Our Operation-CC pattern is presented in Fig. 7, whose modeling elements are listed in Table II. Scheduling is modeled by three functions, *enq_shc(ti)*, *deq_shc(ti)* and *shc()*, which specify the selected scheduling policy. After the *start_operation* location, the *enq_shc(ti)* function pushes the transaction id into the scheduling queue. On the edges from the location *check_sched*, the function *shc()* checks if the transaction is the next to be scheduled. If it is the case, the automaton moves to *do_operation*; otherwise, the automaton waits at location *wait*, until some transaction or the CCManager releases CPU via the *cpu_free* channel. The automaton may stay at *do_operation* for at most *WCRT_op* time units, and at least *BCRT_op* time units, which represent the longest and shortest time to complete the operation. Upon the completion of the operation, a signal is sent to the IsolationObservers via channel *notify_op[ti]*. Before reaching *finish_operation*, the CPU is set to be free, and the transaction is moved from the scheduling queue by the function *deq_shc(ti)*.

According to the selected CC algorithm, the transaction needs to lock and unlock data, before or after the operations. This is modeled by the Locking and Unlocking patterns inserted into the operation patterns, as illustrated in Fig. 7. The Locking and Unlocking patterns are presented in Fig. 8. In the Locking pattern, the automaton sends a request to the CCManager via channel *locktype[ti][di]*, in which “locktype” is parametrized for the particular type of lock, such as a readlock, specified by the CC algorithm. The automaton then either moves to location *finish_locking*,

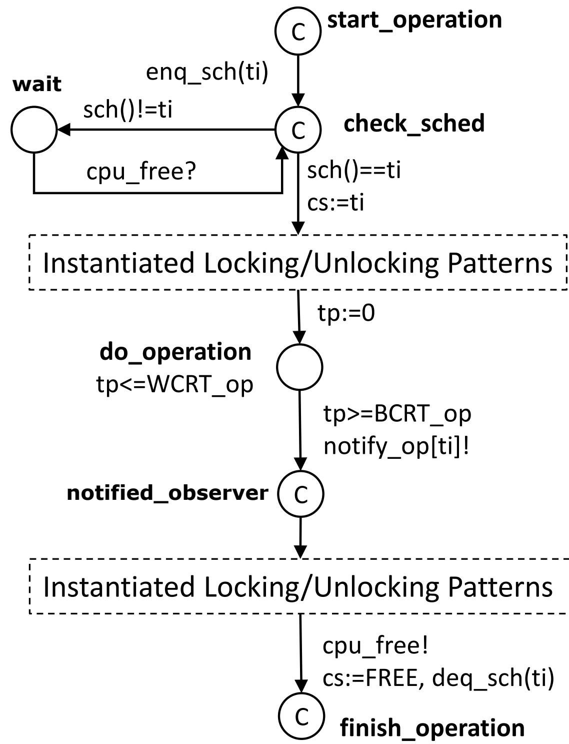

Fig. 7. Operation-CC pattern

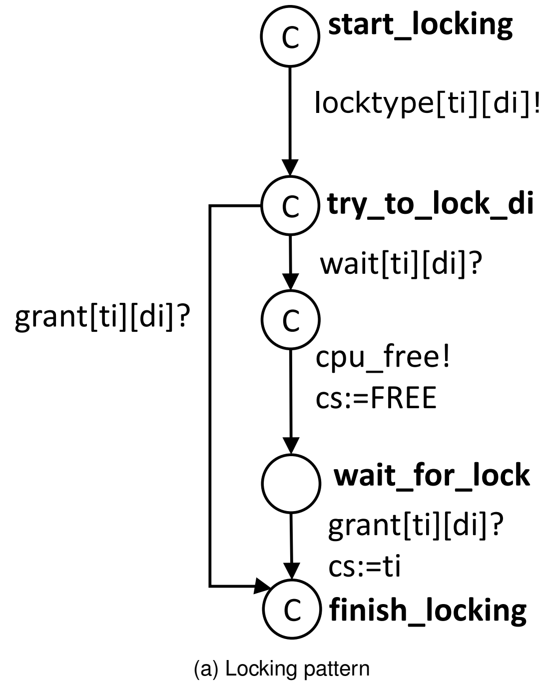

(a) Locking pattern

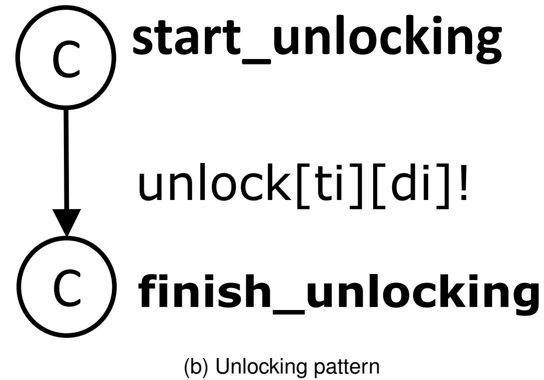

(b) Unlocking pattern

Fig. 8. Locking and unlocking patterns

if it is granted by CCManager via channel *grant[ti][di]*; or releases CPU and gets blocked at location *wait_for_lock*, until CCManager grants it later. In the Unlocking pattern, the automaton sends the request via channel *unlock[ti][di]*, which is received and processed by the CCManager.

b) ***CCManager Skeleton:*** The CCManager skeleton, presented in Fig. 9, provides a common structure for modeling various CC algorithms, and the interaction with the transactions and the atomicity manager. The particular resolution policy of a CC algorithm is encoded in the functions listed in Table III. When CCManager receives a locking request, it updates the status of the transaction and the data by calling *updateRequest()*, and judges whether the requester can obtain the lock by calling *satisfyPolicy()*. The satisfying requester is granted with the lock, if the algorithm does not abort any transactions in order to resolve conflicts. If any transactions need to be aborted due to concurrency conflicts, as suggested by *needAbort()*, CCManager sends a signal to ATManager via channel *cc_conf*, and waits until all abort and recovery are handled, before it grants the lock to the requester. On the other hand, if the requester does not satisfy the policy, it is either aborted, decided *needAbort()* according to the CC algorithm,

TABLE II. MODELING ELEMENTS OF THE OPERATION, LOCKING AND UNLOCKING PATTERNS

Element	Type	Explanation
di	parameter	id of data to be accessed
BERT_op (WERT_op)	parameter	best (worst) case response time of the operation
tp	clock variable	temporary variable for tracking the time of individual operations
cs	integer variable	indicating the possession of the CPU
FREE	constant	indicating that the CPU is free
cpu_free	broadcast channel	release of CPU
locktype[ti][di]	channel	request CCManager for a "locktype" of lock on data di
grant[ti][di]	channel	grant of lock on data di from CCManager
wait[ti][di]	channel	reject of lock on data di from CCManager
unlock[ti][di]	channel	unlocking data di
notify_op[ti]	broadcast channel	notification of completion of operation
enq_sch(ti)	function	adding transaction ti in the scheduling queue
sch()	function	returning the next transaction from the scheduling queue according to the selected policy
deq_sch()	function	removing transaction ti from the scheduling queue

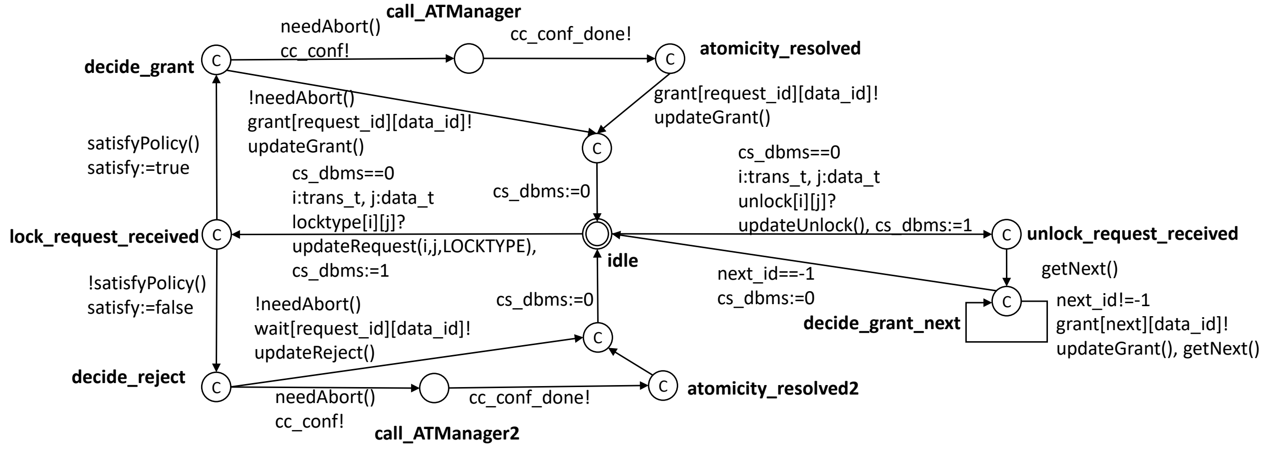

Fig. 9. TA skeleton for the CCManager

or blocked. When the CCManager receives an unlocking request, it updates the status of the transaction and the locks, and grants locks to all legitimated blocked transactions. The next transaction to be granted with a lock is obtained using the *getNext()* function.

c) ***IsolationObserver Skeleton****:* The skeleton for an IsolationObserver is shown in Fig. 10. Each IsolationObserver observes a specified sequence of operations, by accepting the corresponding notification messages from the work unit automata via the *notify_op[ti][di]* channel when an operation is completed. If the monitored sequence indicating the phenomeon occurs, the automaton moves to the *isolation_phenomenon* location.

3) Modeling Data: Fig. 11 presents the skeleton of data. The clock variable *age*, which is reset every time a write operation is performed on the data, represents how old the data is since the last update.

4) Modeling Atomicity and Abort Recovery: We separate the atomicity control model into an ATManager automaton, and the abort recovery parts in work unit automata. The ATManager models the behavior of deciding the transactions to be aborted upon errors, conflicts or user's instructions. The work unit automata include the

TABLE III. MODELING ELEMENTS OF THE CCMANAGER SKELETON

Element	Type	Explanation
LOCKTYPE	parameter	type of the lock
request_id	integer variable	id of the requesting transaction
data_id	integer variable	id of the requested data
next_id	integer variable	id of the next transaction to be granted with locks
cs_dbms	integer variable	indicating critical section for handling request atomically
satisfy	boolean variable	indicating whether the requester should be granted with the lock
cc_conf	channel	notification of CC conflict to ATManager
cc_conf_done	channel	resolution of CC conflict by ATManager
satisfyPolicy()	function	checking if the requester should be granted with the lock according to the selected CC algorithm
needAbort()	function	checking if any transaction should be aborted due to CC
getNext()	function	getting the next transaction to be granted with locks
updateRequest()	function	updating status of transaction and data on request
updateGrant()	function	updating status of transaction and data after grant
updateReject()	function	updating status of transaction and data after reject
updateUnlock()	function	updating status of transaction and data after unlock

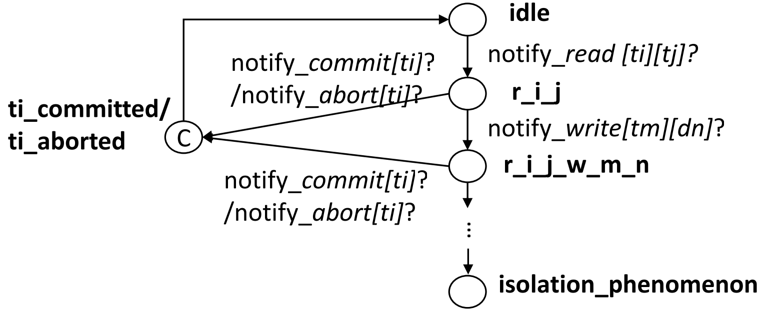

Fig. 10. IsolationObserver skeleton

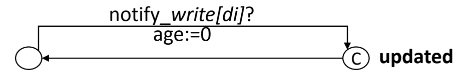

Fig. 11. TA skeleton for data

instantiated abort recovery patterns that model the selected mechanisms for the specific transactions. We distinguish two types of abort, which are user abort that is issued by a client using an abort operation deliberately, and system abort that occurs due to internal conflicts and system failures, such as CC conflicts.

*a) **ATManager Skeleton**:* Our ATManager skeleton provides a common structure for modeling the atomicity manager. The proposed skeleton, as shown in Fig. 12, the ATManager may receive user abort requests via *user_abort[i]* channel, or system abort due to CC via *cc_conf* channel from CCManager. Other types of errors, such as communication errors, can be modeled in a similar way. The function *getAbort()* specifies the logic to decide the transaction to be aborted. The automaton then sends the abort signal to the corresponding work unit automaton via channel *abort_trans[abort_id]*, and waits until the abort is done by the work unit automaton. ATManager then updates the status and locks of transactions and data using the function *updateAbort()*, and checks if more transactions need to be aborted.

*b) **Abort Recovery Patterns: RollbackImComp and DeferredComp**:* The abort recovery mechanisms are modeled by the RollbackImComp pattern (Fig. 13), and the DeferredComp pattern (Fig. 14), respectively, which are composed into the work unit automata. The former models the rollback and immediate compensation mecha-

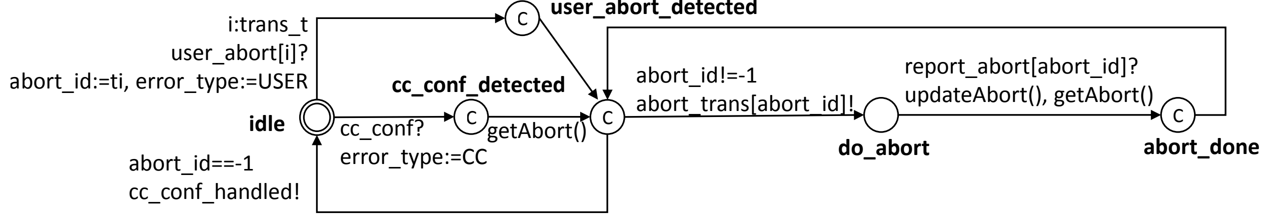

Fig. 12. TA skeleton for the ATManager

nisms, while the latter models the deferred compensation mechanism. The Rollback-ImmeComp pattern models the execution of a series of operations by the DBMS. In case of rollback, the operations are the ones completed before the abort of the transaction. In case of immediate compensation, the operations are specified for the transaction.

In the RollbackImComp pattern, each operation is represented by a location *op_n*, at which the automaton may stay for at most (least) *WCRT_opn* (*BCRT_opn*) time units. When all operations are completed, the work unit reports the completion of recovery to the ATManager via channel *report_abort[ti]*, removes the transaction from the scheduling queue by function *deq_sch(ti)*, and notifies the IsolationObserver via channel *notify_abort[ti]*.

In case of deferred compensation, the compensating transaction is modeled as a separate work unit, using the work unit skeleton and the operation patterns. The DeferredComp pattern starts the compensation transaction via channel *start_trans[ci]*, where *ci* is the id of the compensating transaction. The work unit automaton then immediately reports to ATManager and removes the transaction from the scheduling queue. When the compensating transaction *ci* commits, the work unit automaton receives the notification of *ci*, and notifies that transaction is aborted and recovered via channel *notify_abort[ti]*.

*c) **SystemAbort Pattern**:* We distinguish two types of abort, which are user abort that is issued by a client using an abort operation deliberately, and system abort that occurs due to internal conflicts and system failures, such as CC conflicts. System abort is modeled as a composition of an instantiated operation pattern with a Rollback-ImmComp pattern or a DefComp pattern, as shown in Fig. 15. We refer to this compensation as the SystemAbort pattern. In this pattern, when the WU automaton receives an *abort_trans[ti]* signal from the ATManager, it moves to the corresponding abort recovery patterns.

*d) **UserAbort Pattern**:* The UserAbort pattern is defined in Fig. 16. When the work unit is scheduled as the next one to execute by function *sch(ti)*, it issues the abort request to ATManager via channel *user_abort[ti]*. After it gets the permission from ATManager, the automaton moves to the corresponding abort recovery pattern. When the recovery is completed, the automaton sets the CPU to be free, and moves to location *trans_aborted*.

B. Model Construction Algorithm from UTRAN

The pseudo code listed in Fig. 17 accepts a UTRAN activity diagram specifying the RTDBMS and transactions as input, and returns a UPPCART model. The aforementioned skeletons and patterns are used to construct the model. The main procedure modelNTA traverses each «Transaction» in the diagram, and constructs a work unit automaton for each of them (line 5). The construction of work unit automata is done by the procedure modelWUA (line 18 to 26), which calls procedure modelOP (line 28 to 33) to construct instantiated patterns for operations with CC,

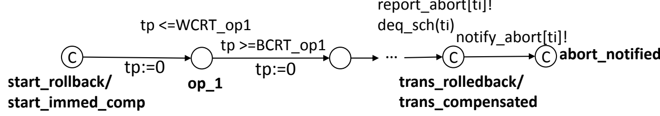

Fig. 13. RollbackImComp pattern

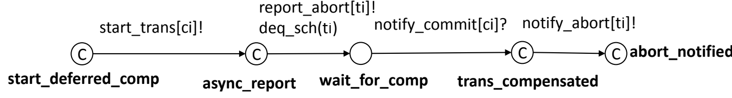

Fig. 14. DeferredComp pattern

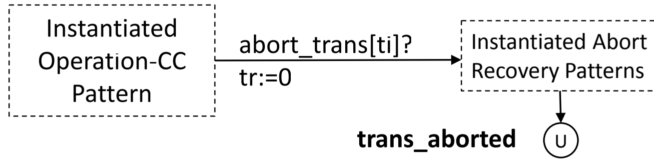

Fig. 15. SystemAbort pattern

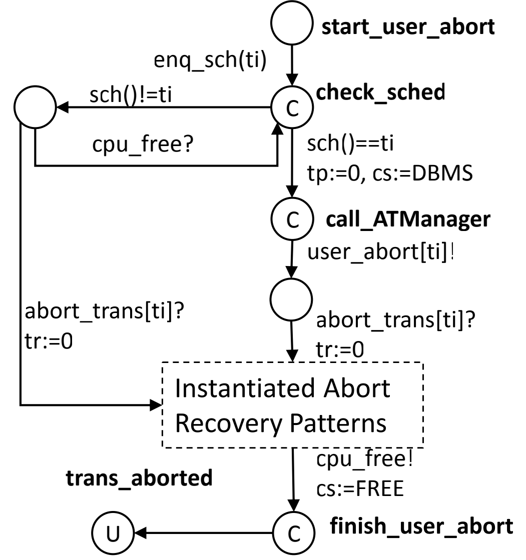

Fig. 16. UserAbort pattern

and procedure modelATOMICITY (line 34 to 51) to model the «AtomicitySpecification» of user abort operations. The main procedure then models the atomicity aspects of system abort for each «Transaction» (line 7), by calling modelATOMICITY, followed by creating CCManager (line 11), ATManager (line 12), and IsolationObservers (line 13). Before ending, the main procedure instantiates the functions in the automata models with the specific code from the selected CCAlgorithm and SchedPolicy (line 14).

C. Example

We illustrate the modeling framework using the example in Section III-C. The work unit automaton of transaction UpdateConfTrans is shown in Fig. 18. Its operations, as shown in the figure, are modeled by instantiating the Operation-CC pattern, in which the CC-related interaction described by the 2PL-HP algorithm is modeled by the instantiated Locking and Unlocking patterns. The rollback mechanism is modeled with the SystemAbort pattern, consisting of the instantiated RollbackImComp pattern. The other transactions in the system are modeled in the similar way, as shown in Fig. 19 and Fig. 20, respectively. The functions that models the priority-based scheduling, which is specified in the UML diagram in Fig. 18, are listed briefly in Program 1 in Appendix.

The CCManager for 2PL-HP shares the same structure with the CCManager skeleton in Fig. 9, as presented in Fig. 21. The instantiation in particular for 2PL-HP is done through the core functions of the algorithm, which are listed in Program 2 in Appendix. Similarly, the ATManager is constructed using the ATManager skeleton in Fig. 12, and instantiated with the functions listed in Program 3 in Appendix. Fig. 23 presents the example of ATManager. Using the IsolationObserver skeleton, we construct an IsolationObserver for an «IsolationPhenomenon» InconsistentConfig1, which specifies a sequence of operations. The two isolation phenomena in the example are presented in Fig. 24 and Fig. 25, respectively.

```
1: Input: a UTRAN diagram, denoted as U
2: Output: a Network of Timed Automata, NTA
3: procedure MODELNTA(U)
4:     for  each «Transaction» T in U do
5:         call modelWUA(T)
6:         if T has an AtomicityVariant as VARIANT then
7:             call modelAtomicity(SYSABORT, VARIANT)
8:         end if
9:         instantiate a Delay pattern for each «DelayedNext»
10:    end for
11:    create CCManager using CCManager Skeleton
12:    create ATManager using ATManager Skeleton
13:    create an IsolationObserver using IsolationObserver Skeleton for each «IsolationPhenomenon» in U
14:    instantiate the functions in the NTA according to the selected CCAlgorithm and SchedPolicy
15: end procedure
16:
17: procedure MODELWUA(T)
18:    create a WU automaton W using the Work Unit Skeleton
19:    for  each «BeginOp», «CommitOp», «ReadOp» and «WriteOp» in T, denoted as P do
20:        call modelOP(W, P)
21:    end for
22:    for for each «AbortOp» P in T, whose AtomicityVariant as VARIANT do
23:        call modelAtomicity(USERABORT, VARIANT)
24:    end for
25: end procedure
26:
27: procedure MODELOP(A, P)
28:    instantiate Operation-CC pattern for P in A
29:    if transaction needs to lock/unlock data D before/after P according to CCAlgorithm then
30:        insert instantiated Locking/Unlocking patterns before/after P
31:    end if
32: end procedure
33:
34: procedure MODELATOMICITY(TYPE, VARIANT)
35:    if TYPE is USERABORT then
36:        instantiate a UserAbort pattern
37:    else if TYPE is SYSABORT then
38:        instantiate a SystemAbort pattern
39:    end if
40:    call modelAbortRecovery(VARIANT)
41: end procedure
42:
43: procedure MODELABORTRECOVERY(VARIANT)
44:    if VARIANT is Rollback then
45:        instantiate RollbackImComp pattern, rollback operations are the operations completed before abort
46:    else if VARIANT is ImmediateCompensation then
47:        instantiate RollbackImComp pattern, compensation operations are listed in the associated «Compensation»
48:    else if VARIANT is DeferredCompensation then
49:        instantiate DeferredComp pattern, call constructWUA(C) for the associated «Compensation» C
50:    end if
51: end procedure
```

Fig. 17. Model Construction Algorithm

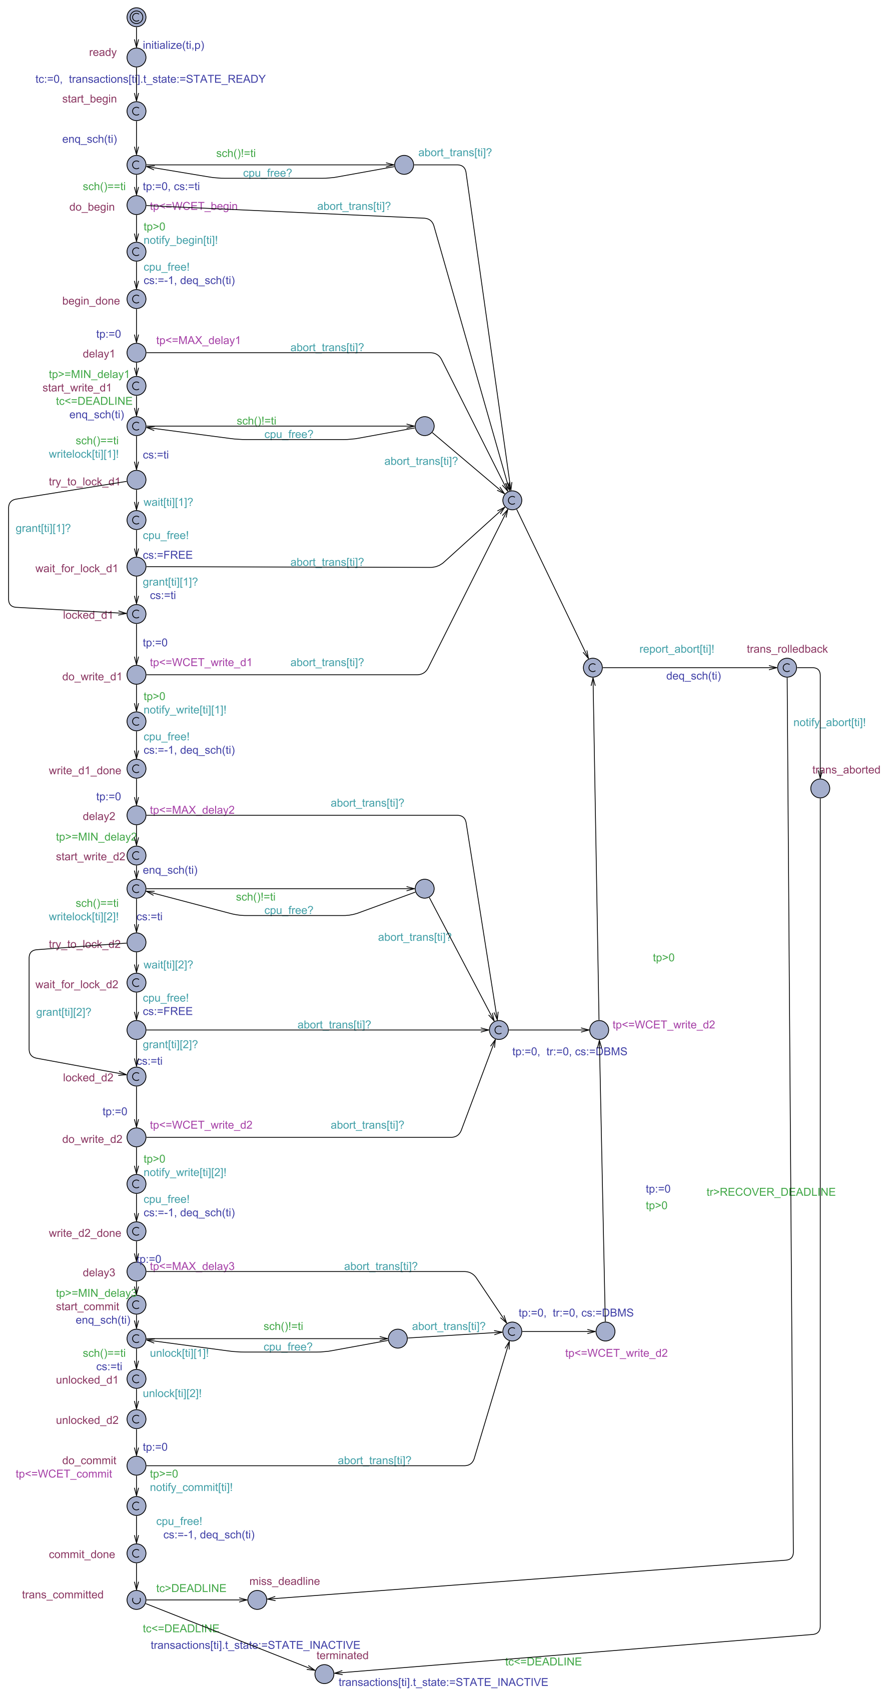

Fig. 18. TA of UpdateConfTrans

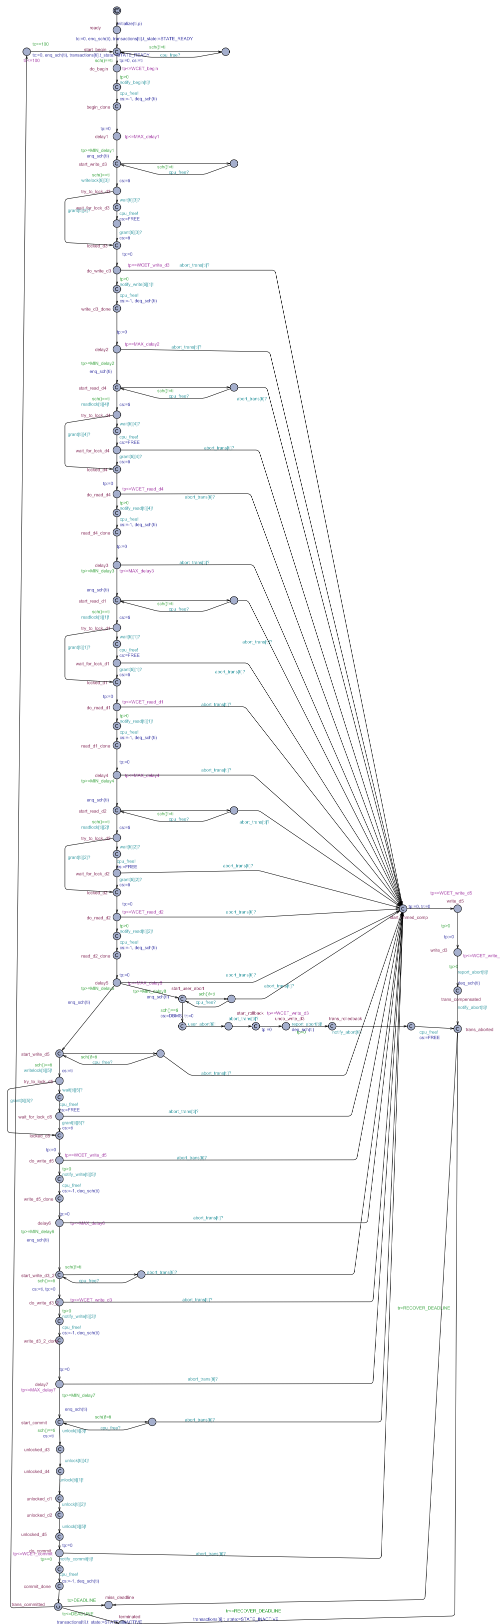

Fig. 19. TA of JobTrans

initialize(ti,p)
tc:=0, enq_sch(ti), transactions[ti].t_state:=STATE_READY
tc<=MAXT
tc>=MINT
start_read_d4
tc:=0, enq_sch(ti), transactions[ti].t_state:=STATE_READY
sch()!=ti
cpu_free?
sch()==ti
writelock[ti][4]!
cs:=3
try_to_lock_d4
wait[ti][4]?
grant[ti][4]?
cpu_free!
cs:=FREE
wait_for_lock_d4
grant[ti][4]?
cs:=ti
locked_d4
tp:=0
tp<=WCET_write_d4
do_write_d4
tp>0
notify_write[ti][4]!
cpu_free!
cs:=-1, deq_sch(ti)
read_d4_done
unlock[ti][4]!
tc<=DEADLINE
tc>DEADLINE
miss_deadline
transactions[ti].t_state:=STATE_INACTIVE
unlocked_d4

Fig. 20. TA of UpdateLocA

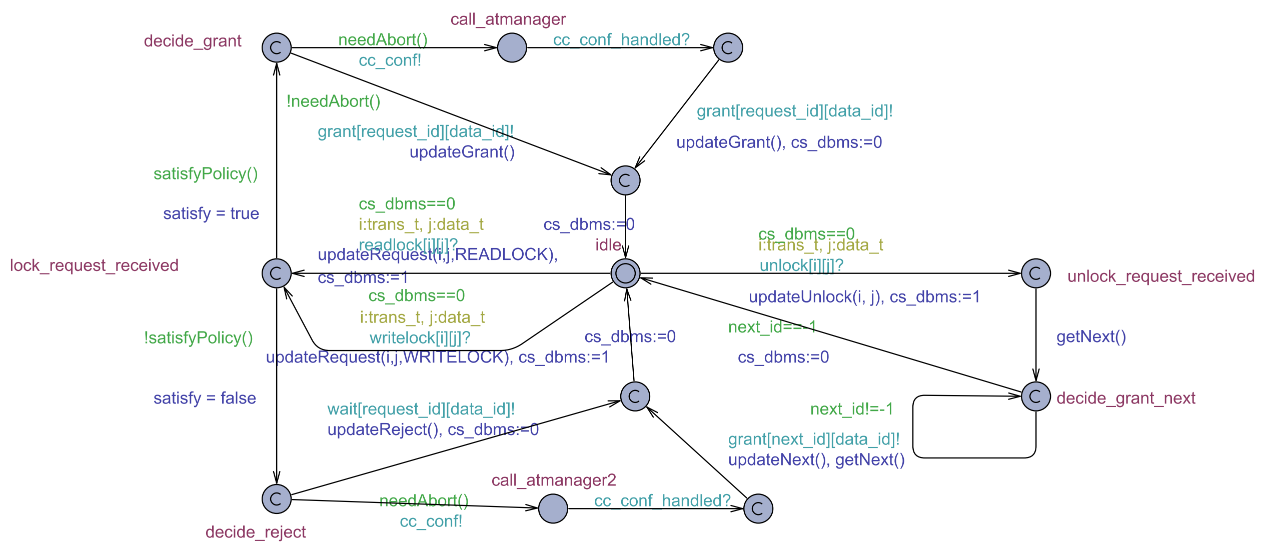

Fig. 21. TA of CCManager

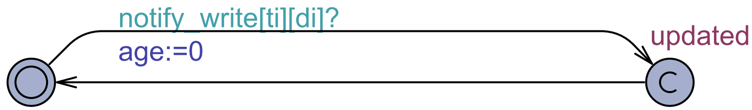

Fig. 22. TA of Data

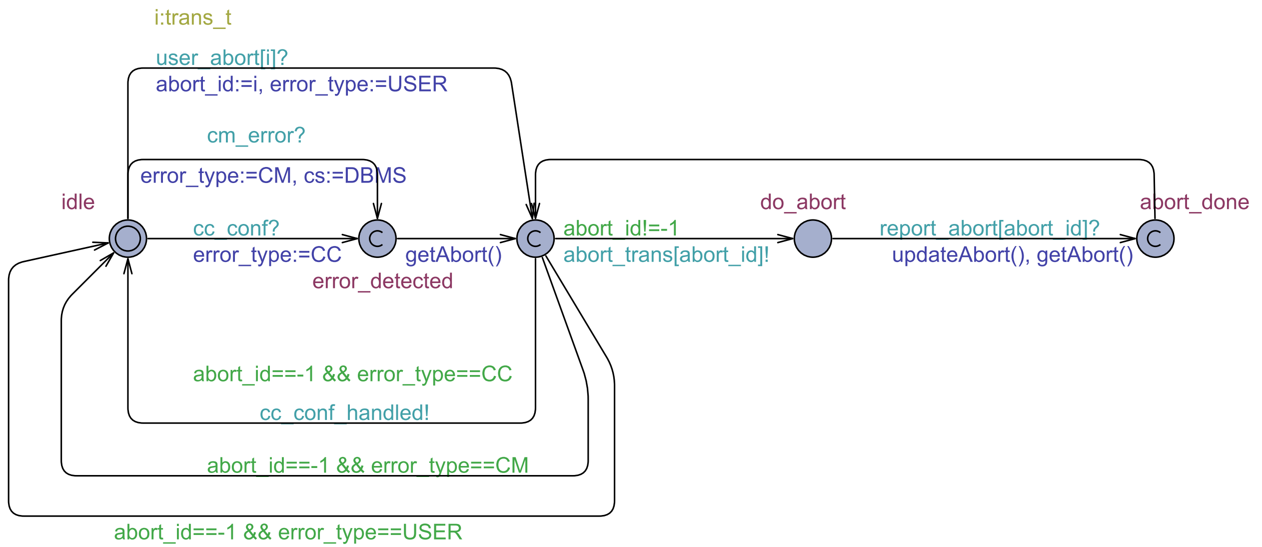

Fig. 23. TA of ATManager

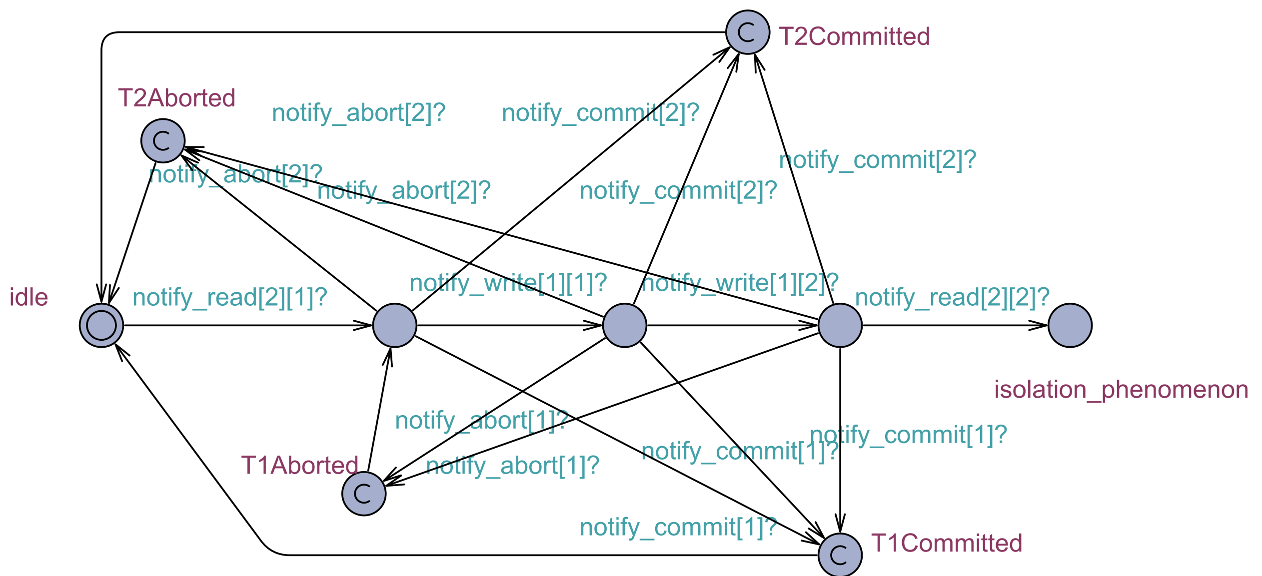

Fig. 24. TA of InconsistentConfig1

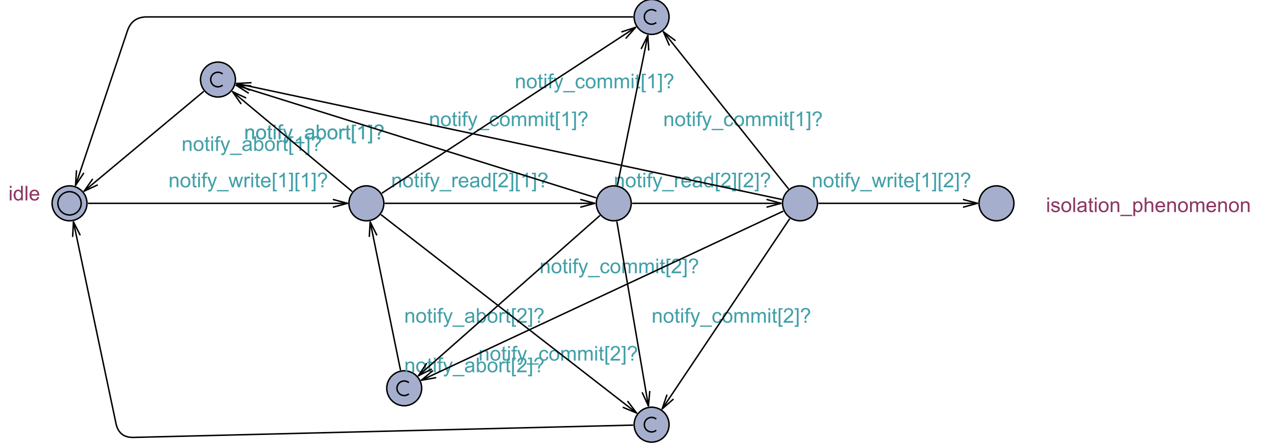

Fig. 25. TA of InconsistentConfig2

TABLE IV. UPPAAL QUERY PATTERNS FOR VERIFYING TRANSACTIONAL PROPERTIES

Property Type	Property Description	UPPAAL Query Pattern
Atomicity	T_i aborted due to ERRORTYPE is eventually rolled back (compensated)	$(ATManager.abort_id == i$ $\&\& \; ATManager.error_type == ERRORTYPE)$ $\rightarrow Ai.trans_rolledback \; (Ai.trans_compensated)$
Isolation	The specified isolation phenomena never occur	$A\,[\,]\; not\; (O_1.isolation_phenomenon \,\|\|\, ... \,\|\|\, O_n.isolation_phenomenon)$
Timeliness	T_i never misses its deadline	$A\,[\,]\; not\; Ai.miss_deadline$
Absolute Validity	When read by T_i, D_j is never older than the absolute validity interval AVI(j)	$A\,[\,]\,(Ai.read_di_done\; imply\; Dj.age <= AVI(j))$
Relative Validity	Whenever T_i reads D_j or D_l, the age differences of D_j and D_l is smaller than or equal to the relative validity interval RVI(j,l)	$A[\,]\,((Ai.read_dj_done \,\|\|\, Ai.read_dl_done)\; imply$ $((Dj.age - Dl.age <= RVI(j,l))\; \&\&$ $(Dl.age - Dj.age <= RVI(j,l))))$

V. VERIFICATION

With the transactions as well as the atomicity and concurrency control mechanisms modeled in UPPAAL TA, we are able to formally verify the atomicity, isolation and temporal correctness properties. Table IV lists the patterns to formalize the properties in UPPAAL queries. Among them, atomicity is formalized as a liveness property, while isolation and temporal correctness are formalized as invariance properties.

We use the aforementioned patterns to formalize the properties for the system in Section III-C, and verify them using the UPPAAL model checker. The verification results are listed in Table V, which shows that all properties are satisfied with the current design of CC and atomicity mechanisms.

VI. RELATED WORK

A number of high-level description languages have been proposed for transaction-based systems. Some of them, like ours, extend UML or existing profiles. Marouane et al. [21] propose a MATRE-based profile for real-time database systems. Timing properties can be specified using their profile, while atomicity and isolation are not considered. Unified Transaction Modeling Language (UTML) [22] and its extension [23] are UML-based languages

TABLE V. VERIFICATION RESULTS OF THE EXAMPLE SYSTEM

Property Type	UPPAAL Query Pattern	Verification Time	Explored States	Verification Result
Atomicity	$(ATManager.abort_id == 1 \,\&\&\, ATManager.error_type == CC)$ $\rightarrow A1.trans_rolledback$	5.99s	344107	Satisfied
Atomicity	$(ATManager.abort_id == 2 \,\&\&\, ATManager.error_type == USER)$ $\rightarrow A2.trans_rolledback$	6.13s	344475	Satisfied
Atomicity	$(ATManager.abort_id == 2 \,\&\&\, ATManager.error_type == CC)$ $\rightarrow A2.trans_compensated$	6.13s	355050	Satisfied
Isolation	$A\,[\,]\ not\,(InconsistentConfig1.isolation_phenomenon$ $\|\|\, InconsistentConfig2.isolation_phenomenon)$	5.60s	336405	Satisfied
Timeliness	$A\,[\,]\ not\,(A1.miss_deadline \,\|\|\, A2.miss_deadline \,\|\|\, A3.miss_deadline)$	5.62s	336405	Satisfied
Absolute Validity	$A\,[\,]\,(A2.read_d4_done\, imply\; D4.age <= 150)$	9.45s	423960	Satisfied
Relative Validity	$A[\,]\,((A2.read_d1_done \,\|\|\, A2.read_d2_done)\; imply$ $((D1.age - D2.age <= 15) \,\&\&\, (D2.age - D1.age <= 15)))$	18.35s	547479	Satisfied

for transactions that enables selection of the ACID properties. Atomicity and isolation are treated as monolithic properties respectively, rather than a spectrum of variants. Timeliness is not the authors' focus. The Business Process Execution Language (BPEL) [24] and the Business Process Model and Notation (BPMN) [25] are XML-based, high-level description languages for specifying business processes, which is a flexible transaction model with various atomicity. Rollback and compensation can be specified at transaction level and for internal activities. Charfi et al. [26] and Sun et al. [27] introduce extra concepts for transactions to BPEL, which allow transaction policies for atomicity to be specified explicitly. Compared with their work, our proposed profile can specify variants of isolation as well as timing properties. Watahiki et al. [28] introduce temporal constraints to BPMN, which can be verified by UPPAAL. Isolation and CC are not part of this framework. ReflecTS [29] allows specification of various ACID properties of flexible transaction models. Both ASSET [30] and KALA [31] use procedural languages for flexible transaction models, in which operations and AR mechanisms are specified using primitives provided by the languages. Compared to these works, our supports specification of timeliness, and the selection of CC algorithms.

Many researchers have also made efforts to formally model and analyze transaction properties. The ACTA framework [32] specifies transaction models in first order logic and allows for formal reasoning. Gallina [7] uses higher-order logic to specify transaction properties, which can be formally analyzed by the Alloy tool. Both frameworks are restricted in the formal specification and analysis of ACID, while timeliness, especially the impact of CC and abort recovery mechanisms, are not included. Derks et al. [12] propose to model and verify transactions with atomicity variants in Petri nets. Liu et al. [33] model and analyze a transaction model using Maude. Lanotte et al. [34] propose a timed-automata-based language for long running transactions with timing constraints. Committing protocols for atomicity variants can be modeled and analyzed. In contrast to these works, our work provides a formal framework for modeling transactions together with abort recovery and CC mechanisms, in which atomicity, isolation, timeliness, as well as their impacts on each other, can be analyzed in a unified framework.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed a UML profile called UTRAN for specifying atomic concurrent real-time transactions. UTRAN supports specification of transaction atomicity, isolation and temporal correctness, as well as the selection of AR, CC and scheduling mechanisms, in high level. Specified as UML activities with UTRAN, transactional properties can be explicitly specified, and be extracted and further analyzed by tools.

We have also proposed a framework based on UPPAAL TA to formally model the UTRAN specification, which allows the specified properties to be rigorously verified by UPPAAL model checker. We provide a set of parametrized automata skeletons and patterns to model the transaction system. Via instantiation and composition, these skeletons and patterns enable flexible modeling of a wide range of abort recovery mechanisms and CC algorithms. Properties are formalized as UPPAAL queries for verification. We have also proposed an algorithm to construct the TA model from an UTRAN specification, which can potentially be automated by a tool.

Our furture work will focus on a tool chain that facilitates the entire process, from high-level specification, to automatic model generation and verification. Another future work is to improve the scalability of our formal framework. In case of large systems, exhaustive model checking may not converge due to state explosion. Other formal techniques, such as statistical model checking, could be integrated in order to achieve better scalability.

Acknowledgment: The Swedish Research Council (VR) is gratefully acknowledged for supporting this research by the project "Adequacy-based Testing of Extra-Functional Properties of Embedded Systems".

REFERENCES

[1] R. A. Elmasri and S. B. Navathe, *Fundamentals of Database Systems*. Addison-Wesley Longman Publishing Co., Inc., 2004.

[2] J. Gray and A. Reuter, *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., 1992.

[3] S. Han, K.-Y. Lam, J. Wang, K. Ramamritham, and A. K. Mok, "On co-scheduling of update and control transactions in real-time sensing and control systems: Algorithms, analysis, and performance," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 10, pp. 2325–2342, 2013.

[4] S. Cai, B. Gallina, D. Nyström, and C. Seceleanu, "Customized real-time data management for automotive systems: A case study," in *Industrial Electronics Society, IECON 2017-43rd Annual Conference of the IEEE*. IEEE, 2017, pp. 8397–8404.

[5] K. Ramamritham, "Real-time databases," *Distributed and Parallel Databases*, vol. 1, no. 2, pp. 199–226, 1993.

[6] J. A. Stankovic, S. H. Son, and J. Hansson, "Misconceptions about real-time databases," *Computer*, vol. 32, no. 6, pp. 29–36, 1999.

[7] B. Gallina, "Prisma: a software product line-oriented process for the requirements engineering of flexible transaction models," Ph.D. dissertation, University of Luxembourg, 2010.

[8] "The unified modeling language specification version 2.5.1," Object Management Group, Standard. [Online]. Available: https://www.omg.org/spec/UML/2.5.1/

[9] "Uml profile for marte specification version 1.1," Object Management Group, Standard. [Online]. Available: https://www.omg.org/spec/MARTE/1.1/

[10] S. Cai, B. Gallina, D. Nyström, and C. Seceleanu, "A formal approach for flexible modeling and analysis of transaction timeliness and isolation," in *Proceedings of the 24th International Conference on Real-Time Networks and Systems*. ACM, 2016, pp. 3–12.

[11] K. G. Larsen, P. Pettersson, and Y. Wang, "Uppaal in a nutshell," *International Journal on Software Tools for Technology Transfer*, vol. 1, no. 1, pp. 134–152, 1997.

[12] W. Derks, J. Dehnert, P. Grefen, and W. Jonker, "Customized atomicity specification for transactional workflows," in *The Proceedings of the 3rd International Symposium on Cooperative Database Systems for Advanced Applications*. IEEE, 2001, pp. 140–147.

[13] J. Kienzle, E. Duala-Ekoko, and S. Gélineau, "Aspectoptima: A case study on aspect dependencies and interactions," in *Transactions on Aspect-Oriented Software Development V*. Springer, 2009, pp. 187–234.

[14] A. Adya, B. Liskov, and P. O'Neil, "Generalized isolation level definitions," in *Proceedings of the 16th ICDE*, 2000, pp. 67–78.

[15] N. Soparkar, E. Levy, H. F. Korth, and A. Silberschatz, "Adaptive commitment for distributed real-time transactions," in *Proceedings of the third international conference on Information and knowledge management*. ACM, 1994, pp. 187–194.

[16] R. K. Abbott and H. Garcia-Molina, "Scheduling real-time transactions: A performance evaluation," *ACM Trans. Database Syst.*, vol. 17, no. 3, pp. 513–560, Sep. 1992.

[17] B. Selic, "A systematic approach to domain-specific language design using uml," in *Proceedings of the 10th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing*. IEEE, 2007, pp. 2–9.

[18] R. Alur and D. L. Dill, "A theory of timed automata," *Theoretical computer science*, vol. 126, no. 2, pp. 183–235, 1994.

[19] E. M. Clarke, E. A. Emerson, and A. P. Sistla, "Automatic verification of finite-state concurrent systems using temporal logic specifications," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 8, no. 2, pp. 244–263, 1986.

[20] S. Cai, B. Gallina, D. Nyström, and C. Seceleanu, "Towards the verification of temporal data consistency in real-time data management," in *Modelling, Analysis, and Control of Complex CPS (CPS Data), 2016 2nd International Workshop on*. IEEE, 2016, pp. 1–6.

[21] H. Marouane, C. Duvallet, A. Makni, R. Bouaziz, and B. Sadeg, "An uml profile for representing real-time design patterns," *Journal of King Saud University-Computer and Information Sciences*, 2017.

[22] G. Nektarios and S. Christodoulakis, "Utml: Unified transaction modeling language," in *Proceedings of the 3rd International Conference on Web Information Systems Engineering*. IEEE, 2002, pp. 115–126.

[23] D. Distante, G. Rossi, G. Canfora, and S. Tilley, "A comprehensive design model for integrating business processes in web applications," *International Journal of Web Engineering and Technology*, vol. 3, no. 1, pp. 43–72, 2006.

[24] "Web services business process execution language version 2.0," Organization for the Advancement of Structured Information Standards, Standard. [Online]. Available: http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html

[25] "Business process model and notation specification version 2.0," Object Management Group, Standard. [Online]. Available: https://www.omg.org/spec/BPMN/2.0/

[26] A. Charfi, B. Schmeling, and M. Mezini, "Transactional bpel processes with ao4bpel aspects," in *Fifth European Conference on Web Services*. IEEE, 2007, pp. 149–158.

[27] C.-a. Sun, E. el Khoury, and M. Aiello, "Transaction management in service-oriented systems: Requirements and a proposal," *IEEE Transactions on Services Computing*, vol. 4, no. 2, pp. 167–180, 2011.

[28] K. Watahiki, F. Ishikawa, and K. Hiraishi, "Formal verification of business processes with temporal and resource constraints," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2011, pp. 1173–1180.

[29] A.-B. Arntsen and R. Karlsen, "Reflects: a flexible transaction service framework," in *Proceedings of the 4th workshop on Reflective and adaptive middleware systems*. ACM, 2005, p. 4.

[30] A. Biliris, S. Dar, N. Gehani, H. Jagadish, and K. Ramamritham, "Asset: A system for supporting extended transactions," in *ACM SIGMOD Record*, vol. 23, no. 2. ACM, 1994, pp. 44–54.

[31] J. Fabry and T. D'Hondt, "Kala: Kernel aspect language for advanced transactions," in *Proceedings of the 2006 ACM symposium on Applied computing*. ACM, 2006, pp. 1615–1620.

[32] P. K. Chrysanthis and K. Ramamritham, "Synthesis of extended transaction models using acta," *ACM Trans. Database Syst.*, vol. 19, no. 3, pp. 450–491, 1994.

[33] S. Liu, P. C. Ölveczky, M. R. Rahman, J. Ganhotra, I. Gupta, and J. Meseguer, "Formal modeling and analysis of ramp transaction systems," in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*. ACM, 2016, pp. 1700–1707.

[34] R. Lanotte, A. Maggiolo-Schettini, P. Milazzo, and A. Troina, "Modeling long-running transactions with communicating hierarchical timed automata," in *Formal Methods for Open Object-Based Distributed Systems*. Springer, 2006, pp. 108–122.

[35] Ï. B. Arpinar, U. Halici, S. Arpinar, and A. Doğaç, "Formalization of workflows and correctness issues in the presence of concurrency," *Distributed and Parallel Databases*, vol. 7, no. 2, pp. 199–248, 1999.

[36] A. Zarras and V. Issarny, "A framework for systematic synthesis of transactional middleware," in *Proceedings of the IFIP International Conference on Distributed Systems Platforms and Open Distributed Processing*, 1998, pp. 257–272.

APPENDIX A
USER-DEFINED FUNCTIONS IN THE EXAMPLE MODELS

Program 1. Functions for priority-based scheduling

```
//Push ti to the queue, sorted by priority
void enq_sch(ti) {
...
for(i=0;i<queue.size;i++) {
if(ti.priority < queue[i].priority) {
queue[i+1] = queue[i];
queue[i] = ti;
...} }}

//Delete ti from the queue, and sort the rest
void deq_sch(ti) {
...
for(i=0;i<queue.size;i++) {
if(ti == queue[i]) {
queue[i] = queue[i+1];
...} }}

//Return the first ready transaction in the queue,
//and the CPU is not occupied by others
int sch() {
...
for(i=0;i<queue.size;i++) {
if((cs==i||cs==FREE) && queue[i].state==READY) {
return i;  }}}
```

Program 2. Functions for 2PL-HP CCManager

```
//Check if the requester should be granted with the lock
bool satisfyPolicy() {
        ...
        if(data_id not locked) return true;
        else if(data_id is readlocked) {
                if(locktype == readlock) return true;
                if(locker has lower priority) return true;
                else return false;
        } else {
                if(locker has lower priority) return true;
                else return false;
        }
}

//If the requester is allowed to lock the data that is
//already locked, the locker needs to be aborted
bool needAbort() {
        ...
        if(requester.data is locked by another transaction)
                return true;
}

//Get the next transaction that waits to lock the data
void getNext() {
  ...
  for(i=0;i<queue.size;i++) {
    if(queue[i].state==BLOCKED && queue[i].data==data_id) {
      next_id = i;
  ...} }}

//Update the status of the transactions and the data
updateRequest();
updateGrant();
updateNext();
updateUnlock();
updateReject();
```

Program 3. Functions for ATManager

```
//Get the next transaction to abort
bool getAbort() {
  ...
  if(error_type==CC) {
        for (each transaction t) {
      if(t has locked data_id)
            abort_id=t; return; }
        } else if(error_type==USER) {
          abort_id=-1; return; }
        else //other errors
        ... }

//Update the status of the transactions and the data
updateAbort();
```