

NOTE TO USERS

This reproduction is the best copy available.

UMI[®]

CONCURRENT TRANSACTION LOGIC WITH PRIORITY AND TIMING CONSTRAINTS

by

JIWEN GE

B.Eng

Hangzhou, ZheJiang, P.R.China, 1989

A thesis

presented to Ryerson University

in partial fulfillment of the

requirement for the degree of

Master of Applied Science

in the Program of

Electrical and Computer Engineering

Toronto, Ontario, Canada, 2004

©JIWEN GE 2004

UMI Number: EC52923

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript⁺ and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform EC52923

Copyright 2009 by ProQuest LLC.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest LLC
789 E. Eisenhower Parkway
PO Box 1346
Ann Arbor, MI 48106-1346

Borrower's Page

Ryerson University requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

[illegible]

CONCURRENT TRANSACTION LOGIC WITH PRIORITY AND TIMING CONSTRAINTS

Master of Applied Science 2004

JIWEN GE

Electrical and Computer Engineering

Ryerson University

Abstract

Concurrent Transaction Logic (*CTR*) is a deductive language for programming database transaction applications that integrates queries, updates, and transaction composition in a complete logical framework. The language supports all the properties of classical transactions and the properties found in some new transaction models, e.g., sub-transactions, transaction rollback, and concurrent transactions.

The contributions of this thesis are twofold. First, it extends *CTR* to account for timing-event-based prioritized concurrent systems in which transactions may have priority and timing constraints. This extension of *CTR*, here called *TP-CTR*, provides a high-level logic programming framework for specifying and simulating executions of timed transactions and trigger-events commonly present in real-time concurrent applications. Second, it describes a Prolog implementation of *TP-CTR*. The implemented *TP-CTR* prototype supports the translation from *TP-CTR* to *CTR*. Underlying this protocol, we use a simplified Rate-Monotonic algorithm [10] to schedule the execution of constraint concurrent transactions and built-in timing predicates to handle transaction time-relations.

Acknowledgment

I appreciate the support and encouragement from Dr. Marcus Vinicius Santos , my supervisor. His expertise in Transaction Logic and Logic programming was instrumental for this work.

I would like to thank all the committee members for their participation in my thesis defense.

I would like to thank the Electrical and Computer Engineering Department at Ryerson University for the financial support of my study and research.

I would like to thank all staff of the Electrical and Computer Engineering at Ryerson University for their kind help.

I wish to convey the warmest thanks to my family for their endless support.

Contents

1	Introduction	1
2	An introduction to <i>CTR</i>	5
2.1	Syntax	5
2.2	Database States and Elementary Updates	7
2.2.1	The Relational Oracle	7
2.3	Examples of <i>CTR</i> Formulas	8
2.4	<i>CTR</i> 's Inference System	13
3	The <i>CTR</i> Prototype	16
3.1	System Description of the <i>CTR</i> Prototype	16
3.2	Executing <i>CTR</i> Prototype Formulas	22
3.2.1	Sequential Execution	22
3.2.2	Concurrent Execution	23
3.3	State Updates and Queries	24
4	<i>TP-CTR</i>	26
4.1	Syntax of <i>TP-CTR</i> Priority and Timing Constraints	27
4.2	Informal Semantics of <i>TP-CTR</i>	28
4.3	Inference System	30
4.4	Interpreting Constraint Concurrent Horn Rules	33
4.4.1	Compiling Priority Constraints	33
4.4.2	Using the Inference System to Execute Transactions	34

4.4.3	Compiling Timing Constraints	34
5	The <i>TP-CTR</i> Prototype	36
5.1	The <i>TP-CTR</i> Prototype Compiler	36
5.1.1	Transaction Base Internal Syntax	38
5.1.2	The <i>TP-CTR</i> Prototype Database	41
5.2	The <i>TP-CTR</i> Prototype Interpreter	42
5.2.1	Underlying Scheduling Algorithm in the <i>TP-CTR</i> Interpreter . . .	43
5.2.2	Interpreting and Scheduling Priority Constraints	44
5.2.3	Interpreting Timing Constraints	47
6	Time-based Prioritized <i>TP-CTR</i> Program Examples	49
6.1	A Time-based Financial Transaction Application	49
6.2	A Simplified Elevator Logic Control Application	52
7	Conclusions and Future Works	55
A	<i>TP-CTR</i> Tutorial	57
A.1	<i>TP-CTR</i> Prototype File System	57
A.2	Getting Started	57
A.3	<i>TP-CTR</i> Prototype Commands and Program Files	58
A.3.1	Compiling Commands in <i>TP-CTR</i>	58
A.3.2	Execution Command	59
A.4	The <i>TP-CTR</i> Prototype Syntax	59
A.4.1	Transaction Rules	59
A.4.2	Database Rules	60
A.5	Built-in <i>TP-CTR</i> Predicates	61
A.5.1	Database Declaration	61
A.5.2	Built-in Predicates in Transaction-base	61
A.6	Programming Examples	62

List of Figures

3.1	Use case diagram for the <i>CTR</i> system	17
3.2	The <i>CTR</i> prototype	17
3.3	The <i>CTR</i> Interpreter	18
3.4	The <i>CTR</i> program files	18
3.5	<i>CTR</i> prototype use case scenario	21
5.1	The <i>TP-CTR</i> compiler.	37
5.2	The list <i>priorityList</i> in <i>TP-CTR</i> interpreter.	46
5.3	Timing logic handled in the <i>TP-CTR</i> interpreter.	48

List of Tables

2.1	A deduction for program (2.2)	15
3.1	Corresponding notations used in the <i>CTR</i> prototype and <i>CTR</i>	19
4.1	A deduction for program (4.3)	34

Chapter 1

Introduction

THIS thesis introduces concurrent transaction logic with priority and timing constraints ($TP-CTR$), an extension of Concurrent Transaction Logic (CTR). $TP-CTR$ is designed to provide a high-level logic programming framework for specifying and simulating the control logic of time-related concurrent systems.

Concurrent Transaction Logic is a recently proposed deductive database language, introduced by Bonner and Kifer in [1]. It is based on Transaction Logic (TR) [3], also known as *sequential* Transaction Logic. CTR extends TR with connectives for modeling the concurrent execution of complex processes, in the sense that it integrates concurrency, communication, and database updates in a complete logic framework.

In CTR , concurrency is accomplished by interleaving the execution of concurrent transactions. Such mechanism is implemented in the CTR prototype in terms of a pure Round-Robin algorithm [10]. Hence, in CTR , the underlying scheduling algorithm for the selection of next executed transaction component from a concurrent transaction execution does not follow any explicit priority-based scheme.

As for sequential transactions, CTR uses the sequential conjunction operator (\otimes) to denote the order of transaction execution, e.g., $a \otimes b \otimes c$ means: first executes transaction a , then transaction b , then transaction c . Exception handling excluded, in CTR , there is no other mechanism one can use to force a particular execution sequence of transactions. Moreover, CTR does not provide any timing constraint type predicate which is necessary

for embedding scheduled events in a database transaction application system, e.g., a real-time database application in which some events can trigger expected transactions at specified times.

The study on the use of logics in real-time systems has been the subject of substantial research in the past and nowadays. In [4], Belliri and Mattolini reviewed a selection of the most representative temporal logics designed for real-time systems. For the specification of a real-time system behaviour, the time behaviour of a system is naturally described with constraints on event occurrences in [5]. Chen and Tsai present a modification of a pure temporal logic in [6], describing the system behaviours in terms of absolute timing of events as well as their relative ordering which can tell when the state actually occur. Torp and Jensen even assigned data with time properties in [7] .

As well, Ulusoy and Sivasankaran studied the use of logics to specify priority properties in real-time systems [8, 9] respectively. More specifically they focus on the design of scheduling algorithms to improve the protocol efficiency and maximize the number of transactions satisfying their real-time constraints.

However, these works focus either on priority constraints or on timing constraints. They do not combine these two types of constraints in one logic framework. We deem this greatly limits their application as formalism for specifying timing-related systems. In our approach, on the other hand, we are able to handle both types of constraints.

To be able to provide a formalism for specifying and simulating timing-event-based systems with explicit priority, in this thesis we extend *CTR*'s inference system by defining a priority-constraint-based inference system, which allows us to formally execute formulas using a SLD-style refutation mechanism. Such inference system is implemented in the interpreter in terms of a simplified Rate-Monotonic Scheduling algorithm instead of a pure Round Robin scheduling algorithm. Our logical framework also introduces timing constraint formulas for specifying timing properties commonly found in real-time application domains. This improved timing-event-based concurrent transaction logic we call *TP-CTR*. Like *CTR*, *TP-CTR* is a language for programming database transac-

tions and applications. *TP-CTR*'s timing and priority constraints enable programmers simulate real-time features in *TP-CTR*, e.g., time-event-driven and interrupt. The time-event-driven feature can be simulated by the *TP-CTR*'s timing constraints. Also, the key feature of real-time software system, interrupt, can be created by combining these two constraints: timing constraints and priority constraints. These features extend the use of logic programming languages to timing-event-based real-time database transaction applications.

A lot of real-time simulation platforms, current in use or under research, use procedure languages to design the real-time system, such as *c* in VxWorks [11] and RT-Linux [12]. It is unquestionable that the solutions designed by procedural languages usually provide high performance for real-time applications. But this is not always the case. To some real-time systems, e.g., real-time database applications, real-time systems with complicated control logic, simulations designed with procedural languages appear to be difficult and timing-consuming jobs. This is because procedural languages focus on computation result. But complicated logic include numerous computation results corresponding to the large amount of composite scenarios of pre-conditions. In such cases, simulating with procedural language is timing-consuming and low efficient. On the other hand, logic programming shows its advantages on this problem since logic programming focuses on the nature of control logic and explores every possible solution of it. As a type of logic programming, *TP-CTR* not only has the natural advantages of logic programming to handle complicated control logic problems, but also includes real-time functionalities. By combining priority and timing constraints, and the high-level logic programming framework provided by *CTR*-based approaches in a complete logic framework, *TP-CTR* provides a possible better solution to handle these real-time applications with complicated control logic over procedural language *c*.

Along with the introduction of *TP-CTR*, the thesis also presents an implementation of *TP-CTR* in Prolog, based on the *CTR* prototype. It is assumed that the reader is already familiar with Prolog. The prototype runs on the XSB prolog interpreter [15, 16],

currently the most efficient deductive database system.

Thesis Outline

The thesis is organized as follows:

- Chapter 2 briefly presents the essentials of *CTR*'s syntax and semantics. Those readers more acquainted with *CTR* could skip this chapter.
- Chapter 3 describes a system view of the *CTR* prototype software, and gives a detailed description of how program execution and scheduling take place in the prototype.
- Chapter 4 introduces *TP-CTR*, its priority and timing constraints syntax, informal semantics and inference engine, and shows how the inference engine is used to execute constraint concurrent transaction Horn rules.
- Chapter 5 presents the developed *TP-CTR* prototype, focusing on the differences between the *TP-CTR* prototype and the *CTR* prototype.
- Chapter 6 provides *TP-CTR* program examples on two different application areas, namely, time-based database transaction, and real-time control system.
- Chapter 7 concludes the work and elaborates on possible improvements.

Besides, the thesis includes one appendix, which presents a tutorial on how to use the software application developed in this thesis.

Chapter 2

An introduction to CTR

CONCURRENT Transaction Logic is a deductive database language for programming database transactions and applications. The language, an extension of Transaction Logic [3], integrates concurrency, communication, and database updates in a complete logic framework. This chapter outlines the language using the terminology of deductive databases. Details are available in [1, 2].

2.1 Syntax

The syntax of Concurrent Transaction Logic is similar to that of first order logic, except that it extends first-order logic with three new logical connectives: \otimes , called *sequential conjunction*; $|$, called *concurrent conjunction*; and a modality of isolation, \odot , for specifying *atomic* actions that executes atomically and in isolation, i.e., it does not communicate or interact with other programs. These operators are used to specify queries and to combine simple transactions into complex ones. The resulting logical formulas are called *transaction formulas*.

In certain important situations, CTR has an elegant, top-down, SLD-style proof procedure that can be expressed within CTR itself. The definition below, lists the conditions that characterizes these situations. The subset of CTR that satisfies these conditions is called the *Horn fragment* of CTR .

Definition 1 (The Horn fragment of CTR) The syntax of the Horn fragment of CTR is defined recursively as follows:

- An atomic formula is an expression of the form $p(t_1, \dots, t_n)$, where p is a predicate symbol, and t_1, \dots, t_n are terms.
- A *concurrent sequential goal* is in any formula of the form:
 - An atomic formula; or
 - $a_1 \otimes \dots \otimes a_i$; or
 - $a_1 \mid \dots \mid a_i$; or
 - $\odot a_1$

where each a_i is a concurrent sequential goal, and $i \geq 0$.

- If a is a concurrent sequential goal and t is an atomic formula, then $t \leftarrow a$, is a *concurrent Horn rule*.
- A *transaction base* is a set of concurrent Horn rules.

□

If a and b are transaction formulas, then informally:

- $a \otimes b$ means: first execute a , then execute b .
- $a \mid b$ means: execute a and b concurrently.
- $\odot a$ means: execute a “atomically”, i.e., without interleaving with other transactions.
- $t \leftarrow a$ means: to execute t is sufficient to execute a .

2.2 Database States and Elementary Updates

In $CT\mathcal{R}$, a pair of oracles, called *state data oracle* and *state transition oracle*, specify elementary database operations. The state data oracle specifies a set of database state queries, and the state transition oracle specifies a set of database elementary updates. These oracles are not fixed because any pair of oracles can be *plugged into* a $CT\mathcal{R}$ theory. For ease of reference, below we present the definition of these oracles, introduced in [3].

Definition 2 (state data oracle) A *state data oracle* O^d , is a mapping from sets of state identifiers to sets of first-order formulas. \square

Intuitively, if D_i is a state identifier, then $O^d(D_i)$ is the set of formulas considered to be all the truths known about the state D_i . In practice, it is not necessary to materialize all these truths. Because, given a logical formula ϕ and a state identifier D_i , the proof theory for $CT\mathcal{R}$ only needs to know whether $\phi \in O^d(D_i)$. Thus, to do inference in $CT\mathcal{R}$, an enumeration of $O^d(D_i)$ is all that is needed.

Definition 3 (state transition oracle) A *state transition oracle* O^t , is a mapping from pairs of state identifiers to sets of ground atomic formulas. These ground atoms are referred to as *elementary transitions*. \square

Intuitively, if D_1 and D_2 are two state identifiers, and $b \in O^t(D_1, D_2)$, then b is the set of elementary updates that change state D_1 into state D_2 . An elementary update can thus be non-deterministic, since for each update, the transition oracle defines a binary relation on states. In practice, this relation does not have to be materialized. Instead, for a given update u , and a given state D_1 , the proof theory of $CT\mathcal{R}$ only needs an enumeration of the possible successor states, D_2 .

2.2.1 The Relational Oracle

The examples in this thesis use the notion of relational databases, in which a state is a set of tuples, and elementary transactions consist of the insertion and deletion of individual

tuples from the database.

In [1, 2, 3], Bonner and Kifer represent relational databases in the usual way as sets of ground atomic formulas. Moreover, they use two predicates, *ins* and *del*, to insert and delete atoms from the database. The definition of *Relational Oracles* formalizes this idea.

Definition 4 (*Relational Oracles*) : A state D is a set of ground atomic formulas. The data oracle simply returns all these formulas. Thus $O^d(D) = D$.

Moreover, for each p in D , the transition oracle defines two new predicates, $ins(p)$ and $del(p)$, representing the insertion and deletion of single atom p in D respectively as follows:

$$ins(p) \in O^t(D_1, D_2) \text{ iff } D_2 = D_1 + \{p\}$$

$$del(p) \in O^t(D_1, D_2) \text{ iff } D_2 = D_1 - \{p\}$$

□

It should be noted, however, that CTR is restricted neither to relational databases, nor to update operations based on single tuple. For instance, databases could be deductive, object-oriented, disjunctive, or a collection of scientific objects, such as matrices or DNA sequences. Likewise, database operations could include SQL-style bulk updates, or the insertion and deletion of rules, or complex scientific calculations, such as the Fourier transformation and matrix inversion. In CTR , the set of states is determined by the data oracle. Changing the oracles can change the set of states, and thus the set of semantic structures. This is one way in which different oracles give rise to different versions of CTR . In this thesis, we used the relational oracle for the simplicity.

2.3 Examples of CTR Formulas

This section gives some examples to illustrate CTR 's syntax and informal semantics. The full description of CTR semantics is available in [1, 3]. We start with simple examples of sequential goals and concurrent sequential goals.

Sequential conjunction and database update predicates: The formula below illustrates the sequential conjunction (\otimes) combined with the elementary update predicates *del* and *ins*:

$$del(r(a)) \otimes ins(s(a))$$

“Delete $r(a)$ from the database, and then insert $s(a)$ ”.

Concurrent sequential goal: The formula below illustrates the use of the concurrent and sequential conjunctions, $|$ and \otimes , respectively, in the specification of concurrent transactions:

$$(\phi_1 \otimes \phi_2) | (\varphi_1 \otimes \varphi_2)$$

“Execute concurrently the transactions $\phi_1 \otimes \phi_2$ and $\varphi_1 \otimes \varphi_2$. To execute $\phi_1 \otimes \phi_2$, first do ϕ_1 then ϕ_2 , and similarly for $\varphi_1 \otimes \varphi_2$ ”.

Horn rules: Like classical logic, *CTR* has a Horn-like fragment with both a procedural and a declarative semantics. The formula:

$$q(X) \leftarrow r(X) \otimes del(r(X)) \otimes ins(s(X))$$

defines a subroutine with name q and parameter X . Given the parameter value a , $q(a)$ commits if the atom $r(a)$ exists in the databases before the updates execute and assuming $r(X)$ is a updatable database tuple.

In the following examples, we show how *CTR* can be used to combine elementary operations into complex transactions.

Example 1 (Database transactions for operating an online store) Assume a relation *inventory*(*Name*, *Amt*) represents the amount of available goods in a store’s inventory, where *Name* is the goods’ name and *Amt* is an integer representing the amount available. To simplify, we ignore the price and other factors. The rules below define three transactions in the transaction base:

```

process_order(CustomInfo, Name, Amt) ←
    ⊙ ( inventory(Name, OldAmt)
      ⊗ OldAmt ≥ Amt
      ⊗ NewAmt is OldAmt − Amt
      ⊗ change_inventory(Name, OldAmt, NewAmt)
      ⊗ ins(record(CustomInfo, Name, Amt)))

supply(SupplierInfo, Name, Amt) ←
    ⊙ ( not inventory(Name, OldAmt)
      ⊗ ins(inventory(Name, Amt))
      ⊗ ins(supplier(SupplierInfo, Name, Amt)))

supply(SupplierInfo, Name, Amt) ←
    ⊙ ( inventory(Name, OldAmt)
      ⊗ NewAmt is OldAmt + Amt
      ⊗ change_inventory(Name, OldAmt, NewAmt)
      ⊗ ins(supplier(SupplierInfo, Name, Amt)))

change_inventory(Name, OldAmt, NewAmt) ←
    ⊙ ( del(inventory(Name, OldAmt))
      ⊗ ins(inventory(Name, NewAmt)))

```

The first rule specifies: to sell an amount *Amt* of goods *Name* to a customer *Customer-Info*, first check how many goods are available in the store's inventory database. If the available amount, *OldAmt*, is no less than the requested amount, *Amt*, then deduct *Amt* from the available amount *OldAmt*, and update the latest inventory amount, and then record this transaction in the database; otherwise the *process_order* transaction should fail since the available amount of goods can not meet the demand of the order. The second rule specifies the supplying transaction: to supply an amount *Amt* of goods into the inventory, if the goods' name are new to the inventory, add a new inventory record with the goods' name and amount, and then add another supplying record to the supplier's account; Otherwise, if the goods have already a record in the inventory, add the amount into that inventory record, and add another record to the supplier's account in the same way. The last rule is in charge of changing the amount of a goods in the inventory record: first delete the old record with the *OldAmt*, then insert a new record with *Name* and

NewAmt.

Due to the \odot operator, *process_order*, *supply*, and *change_inventory* are all executed “atomically”. That is, they execute either entirely or not at all. The transaction below specifies the concurrent execution of *supply* and *processing_order*:

$$\textit{supply}(\textit{supplier1}, \textit{tape1}, \textit{Amt1}) \mid \textit{processing_order}(\textit{customer1}, \textit{tape1}, \textit{Amt2})$$

That is, an amount *Amt1* of goods *tape1* is supplied by *supplier1* while a customer places an order for *tape1*. Because this is intended to be a transaction, if one sub-transaction fails, then both sub-transactions are rolled back whatever the other sub-transaction succeeds or fails.

Notice that, the *CTR* program behaves correctly while an equivalent Prolog program could not: updates in Prolog are not logical. If the execution fails after an update is performed, the update cannot be undone. So although execution in Prolog can be backtracked, the database does not roll back to its initial state if updates are involved. Thus transaction fail in a Prolog program will lead to database inconsistency. \square

The above example illustrates the combination of concurrency and updates that *CTR* supports. The next example shows how *CTR* also supports communication, where two processes synchronize themselves by exchanging messages via communication signals.

Example 2 (Synchronization between two elevator control processes) Assume *stopRequest_control* and *moving_control* are two processes specified to control the stop request control system and the moving control system of an elevator model, respectively. Here, our goal is to illustrate the communication between the two concurrent processes. Hence, we simplify our model to enable one round processing of these two control processes, without considering the loop condition.

$$\begin{aligned}
\text{stopRequest_control} &\leftarrow \text{task}_{11} \\
&\otimes \text{send}(\text{requestSignal}, \text{FloorNum}) \\
&\otimes \text{receive}(\text{stopSignal}, \text{Flr}) \\
&\otimes \text{task}_{12} \\
\\
\text{moving_control} &\leftarrow \text{task}_{21} \\
&\otimes \text{receive}(\text{requestSignal}, \text{FloorNum}) \\
&\otimes \text{moveTo}(\text{FloorNum}) \\
&\otimes \text{send}(\text{stopSignal}, \text{Flr}) \\
&\otimes \text{task}_{22}
\end{aligned}$$

where task_{11} and task_{21} are initialization tasks of $\text{stopRequest_control}$ and moving_control , respectively; task_{12} and task_{22} are the post-handling tasks of $\text{stopRequest_control}$ and moving_control , respectively.

The next transaction specifies the concurrent execution of $\text{stopRequest_control}$ and moving_control :

$$\text{stopRequest_control} \mid \text{moving_control}$$

During execution, $\text{stopRequest_control}$ and moving_control can communicate and synchronize the execution of their tasks effectively by sending and receiving messages along channel requestSignal and stopSignal . Note: $\text{send}(Ch, Msg)$ and $\text{receive}(Ch, Msg)$ are two communication predicates, which are used to send and receive a message Msg along a channel Ch , as shown in [1]. moveTo cannot start moving the elevator to a specific floor until task_{11} finishes and a FloorNum message is received along the requestSignal channel from the stop request control process. Likewise, $\text{stopRequest_control}$ process cannot continue until moveTo finishes and a stop message Flr is received along the stopSignal channel. In such a way, the two processes communicate with each other and synchronize their execution steps.

□

2.4 $CT\mathcal{R}$'s Inference System

This section first introduces the SLD-style resolution of the $CT\mathcal{R}$ inference system. This inference system is used to formally execute transactions. In Chapter 4, we will introduce an extension of this inference system, which allow us to formally execute transactions involving constraints.

SLD-style resolution

The inference system manipulates expressions called *sequents*, which have the form

$$\mathbf{P}, \mathbf{D} \vdash (\exists) \phi$$

where \mathbf{P} is a program, \mathbf{D} is any legal database state, and ϕ is a concurrent sequential goal. The informal meaning of such a sequent is that, based on program \mathbf{P} the formula $(\exists) \phi$ can be proved from state \mathbf{D} .

Let the concurrent sequential goal clause be the expression

$$\begin{array}{c} \leftarrow G_0 \\ \text{where } G_0 \text{ is the sequent } \mathbf{P}, \mathbf{D}_1 \vdash (\exists) \phi \end{array} \quad (2.1)$$

A SLD-style refutation of $\leftarrow G_0$ is a sequence of goal clauses $\leftarrow G_0 \cdots \leftarrow G_n$ where G_n is the *empty clause*, i.e., the sequent $\mathbf{P}, \mathbf{D}_n \vdash ()$, where \mathbf{D}_n is a database state, and $()$ denotes the empty formula. This sequent is an axiom of the inference system, and this axiom states that the empty formula is true on any database state. Each $\leftarrow G_{i+1}$ is obtained from $\leftarrow G_i$ by using the inference system later presented in this section.

Before presenting $CT\mathcal{R}$'s inference system, we deem relevant to also present the notion of *hot components* introduced in [1]. Summarily, hot components of a transaction ϕ , denoted $hot(\phi)$, is the set of transactions ready for execution in ϕ . Formally:

Definition 5 (Hot components) Let ϕ be a concurrent sequential goal. Its set of hot components, $hot(\phi)$, is defined recursively as follows:

- $hot(()) = \{ \}$, where $()$ denotes the empty goal;
- $hot(b) = \{b\}$, if b is an atomic formula;
- $hot(\psi_1 \otimes \cdots \otimes \psi_n) = hot(\psi_1)$;
- $hot(\psi_1 \mid \cdots \mid \psi_n) = hot(\psi_1) \cup \cdots \cup hot(\psi_n)$;
- $hot(\odot\psi) = \{\odot\psi\}$

□

Definition 6 (CTR's Inference system) The inference system consists of one axiom and four inference rules.

Axiom: $P, D \vdash ()$, for any state D

Inference rules: In rules 1-4, σ is a substitution, ψ and ψ' are concurrent sequential goals, and $hot(\psi) = a$.

1. *Applying rule definitions:* Suppose $b \leftarrow \beta$ is a rule in P whose variables have been renamed so that the rule shares no variables with ψ . If a and b unify with mgu σ , then

$$\frac{P, D \vdash (\exists) \psi' \sigma}{P, D \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by replacing an element of a by β .

2. *Querying the database:* If $\mathcal{O}^d(D_i) \models^c (\exists) a\sigma$, and $a\sigma$ and $\psi'\sigma$ share no variables, then

$$\frac{P, D \vdash (\exists) \psi' \sigma}{P, D \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting an element of a .

3. *Executing elementary updates:* If $\mathcal{O}^t(D_1, D_2) \models^c (\exists) a\sigma$, and $a\sigma$ and ψ' share no variables, then

$$\frac{\mathbf{P}, \mathbf{D}_2 \vdash (\exists) \psi' \sigma}{\mathbf{P}, \mathbf{D}_1 \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting an element of a .

4. *Executing atomic transactions:* If $\odot \alpha$ is the hot component in ψ , then

$$\frac{\mathbf{P}, \mathbf{D} \vdash (\exists) (\alpha \otimes \psi')}{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting an element of $\odot \alpha$.

□

Based on these inference rules, one can prove the execution sequence of both sequential executions and concurrent executions. Let us consider the following program.

$$\begin{aligned} p &\leftarrow a_1 \mid a_2 \\ a_1 &\leftarrow \text{ins}(a) \\ a_2 &\leftarrow \text{ins}(b) \end{aligned} \tag{2.2}$$

The deduction of the transaction p is illustrated in Table 2.1.

Table 2.1: A deduction for program (2.2)

Sequents	Inference rule	Hot components
$\mathbf{P}, \{\} \vdash p$	1	$\{p\}$
$\mathbf{P}, \{\} \vdash a_1 \mid a_2$	1	$\{a_1, a_2\}$
$\mathbf{P}, \{\} \vdash \text{ins}(a) \mid a_2$	3	$\{\text{ins}(a), a_2\}$
$\mathbf{P}, \{a\} \vdash a_2$	1	$\{a_2\}$
$\mathbf{P}, \{a\} \vdash \text{ins}(b)$	3	$\{\text{ins}(b)\}$
$\mathbf{P}, \{a, b\} \vdash ()$	<i>Axiom</i>	$\{\}$

Each sequent in the table is derived from the one below by an inference rule. The deduction succeeds because the bottom-most sequent is an axiom. Carried out top-down, the deduction corresponds to an execution of the transaction p in which atoms a and b are inserted into the empty database in the order, first a , then b .

Chapter 3

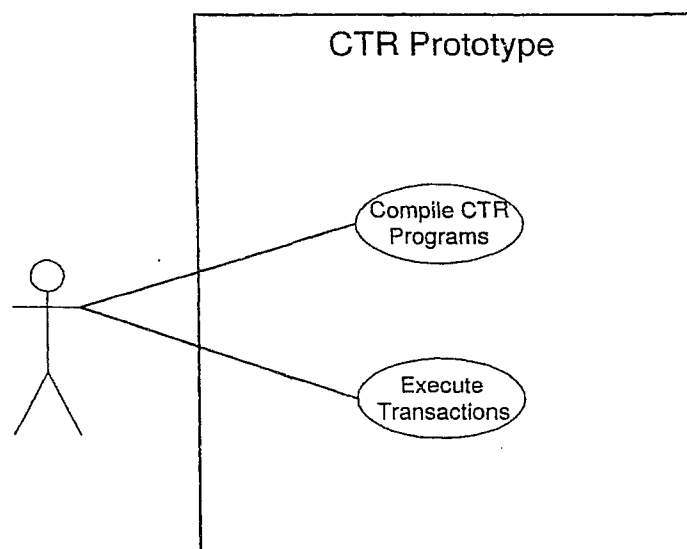
The *CTR* Prototype

THE *CTR* prototype [13, 17] is an implementation of the Horn fragment of *CTR* and *CTR*'s inference system.

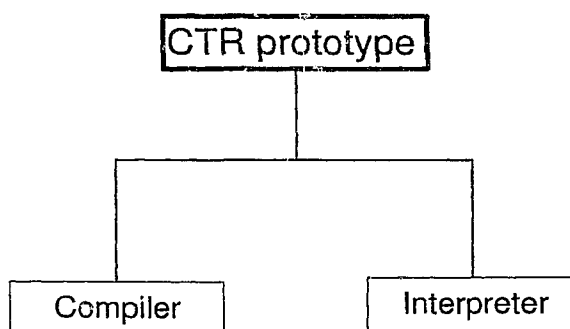
To facilitate the understanding of the *TP-CTR* prototype later introduced in Chapter 5, in this chapter we outline the *CTR* prototype by presenting a system view of the software application available in [17]. This chapter is organized as follows. First it presents a system description of the *CTR* prototype, describing its major components. Then it shows how the prototype handles sequential execution, concurrent execution, and database updates and queries.

3.1 System Description of the *CTR* Prototype

Figure 3.1 shows a *use case* diagram of the *CTR* prototype. As such, it shows the system functionalities seen from the user's viewpoint. These functionalities are two: to *compile* a *CTR* program into an internal program; and to *execute* transactions based on the translated internal program.

Fig. 3.1: Use case diagram for the *CTR* system

To handle these functions, the *CTR* prototype system can be seen as consisting of a compiler and an interpreter, as shown in Figure 3.2 .


Figure 3.2: The *CTR* prototype

The role of the *CTR* compiler introduced in [17] is to translate a *CTR* program into an internal format recognizable to the *CTR* interpreter. The *CTR* interpreter is an

implementation of the *CTR* inference engine, which is used to prove transactions based on the relational oracle and the translated *CTR* programs.

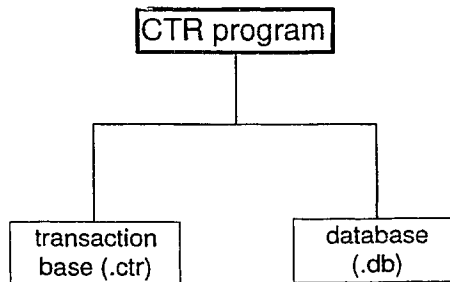
Figure 3.3 shows the two most important components of the *CTR* interpreter: the inference engine and the relational oracle.

Figure 3.3: The *CTR* Interpreter



In the *CTR* prototype, a *CTR* program consists of a *transaction base* file and a *database* file, as shown in figure 3.4.

Figure 3.4: The *CTR* program files



Transaction-base files have the extension *.ctr* and database files have the extension *.db*. A transaction base is a program consisting of a set of concurrent Horn rules. Table

3.1 illustrates the notations used in the *CTR* prototype and the corresponding notations used in *CTR*.

Table 3.1: Corresponding notations used in the *CTR* prototype and *CTR*

<i>CTR</i> Prototype	<i>CTR</i>	Meaning
:-	\leftarrow	<i>Logical implication</i>
$*$	\otimes	<i>Sequential conjunction</i>
$\#$	$ $	<i>Concurrent conjunction</i>
o	\odot	<i>Modality of isolation</i>

A *CTR* prototype database is a Prolog rulebase, consisting of *database rules* and *database tuple declarations*. A database rule is a normal Prolog rule, which has one of the following forms:

- an atomic formula; or
- $\phi \text{ :- } \alpha$, where ϕ is an atomic formula and α is a clause.

There are two main reasons for putting rules in a *CTR* database:

1. to update them (only those rules in the *CTR* database can be changed at runtime).
2. to express queries with negation-as-failure.

To illustrate, below we show an excerpt of a *CTR* database:

```
finished :- not(unfinished).
unfinished :- a(X).
a(100). a(99). ... a(1). a(0).
```

The first and second rules specify a query with negation as failure. The atomic formulas $a(_)$ specify which atoms are true in the initial database state. To update a database atom during execution, it must be declared as an updatable database tuple. *Database tuple declarations* are of the form

updatable R/n.

For instance, considering the database formulas above, the database tuple declaration below specifies that the tuple *a* with arity one, is updatable.

updatable a/1.

The example below illustrates the functionality of *declarations* and *ground atomic formulas*.

Example 3 (A simple book order application) The *CTR* program below exemplifies a simple book order system. The book inventory is represented by the tuple *book/1* in the database.

- The transaction-base file *book.ctr*:

$$\text{order}(X) \text{ :- } o(\text{book}(X) * \text{del}(\text{book}(X))) * \text{writeln}(\text{succeeded}).$$

$$\text{order}(X) \text{ :- } \text{writeln}(\text{failed}).$$
- The database file *book.db*:

$$\text{updatable book/1.}$$

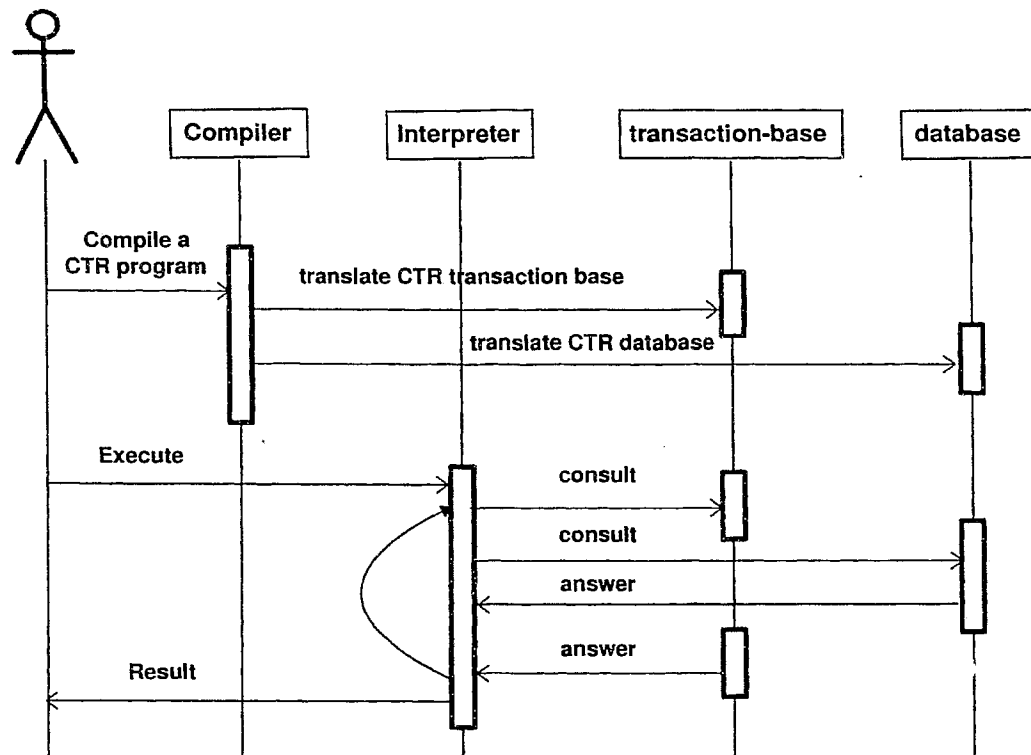
$$\text{book}(1).$$

$$\text{book}(2).$$

The transaction base specifies how orders of books are placed. The database defines an updatable tuple *book/1*, and includes two initial instantiations of the tuple *book*, *book(1)* and *book(2)*. □

Figure 3.5 shows a use case scenario describing how the user interacts with the *CTR* prototype to compile and interpret a *CTR* program:

- The user compiles a *CTR* program.
 - The *CTR* prototype compiler translates the Horn rules in the *CTR* program's transaction base.

Figure 3.5: *CTR* prototype use case scenario

- The *CTR* prototype compiler translates the rules in the *CTR* program's database.
- When the user executes a transaction, the *CTR* prototype interpreter executes the transaction with inference engine. In each resolution step, the *CTR* prototype interpreter communicates with the transaction base and the database according to the executed transaction component in the following ways:
 - Consulting the transaction base.
 - Consulting the database.
 - Updating the database.
 - The *CTR* interpreter outputs the result to the user.

3.2 Executing *CTR* Prototype Formulas

The *CTR* prototype interpreter uses a *Round-Robin scheduling algorithm* to schedule multi-processes in concurrent transaction executions. This algorithm assumes that each transaction process in a concurrent execution has the same priority level. Thus, in the *CTR* prototype, these concurrent processes have the same opportunity to execute.

In the *CTR* prototype, the *hot component set* introduced in Chapter 2 is implemented as a Prolog list. The transaction in the head of the list is selected for execution. If this transaction is defined as a rule in the transaction base, then the transaction formula in the body of the rule replaces the transaction, and then is inserted at the end of the list. When a transaction is selected, the *CTR* prototype interpreter executes one element of the transaction, and then feeds the remaining elements of the transaction back into the end of the list. This is consistent with the *CTR* inference system presented in Chapter 2. In this way, the concurrent sequential transactions in the queue are able to execute in equal turn.

In the rest of this section, we analyze sequential executions and concurrent executions in the *CTR* prototype respectively.

3.2.1 Sequential Execution

Sequential execution is the simpler type of execution, where transactions are connected with sequential conjunction operators. These transactions are executed in sequential order from left to right.

When executing a sequential conjunction, such as $p_1 * p_2 * \dots * p_n$, the *CTR* prototype interpreter starts the execution from p_1 , then p_2 , and repeats until p_n is completed. If any given transaction p_i fails, the whole transaction will roll back.

Also, any transaction p_i could have been defined by rules. For instance, p_1 could have been defined as another sequential conjunction, i.e., $p_i :- a * b$. During execution, the *CTR* prototype interpreter would decompose any composite operation, e.g., replacing p_1

with $a * b$.

3.2.2 Concurrent Execution

The prototype simulates concurrency by interleaving the execution of concurrent processes with a Round-Robin scheduling algorithm. Consider the concurrent transaction execution:

$$task :- a_1 * a_2 * \dots * a_n \# b_1 * b_2 * b_3 * \dots * b_n.$$

To run this transaction, first, the Round-Robin algorithm picks up the first element from the first concurrent sub-transaction, a_1 , and executes it, then followed by the first element from the other concurrent sub-transaction, b_1 . After that, the Round-robin algorithm continues to pick up elements from the two concurrent sub-transactions in the same fashion and execute them until all the elements in the two sub-transactions are executed. The resulting execution sequence is: $a_1, b_1, a_2, b_2, \dots, a_n, b_n$.

This applies for transactions consisting of more than two concurrent transactions, e.g.,

$$task :- a_1 * a_2 \# b_1 * b_2 \# c_1 * c_2.$$

The resulting execution sequence is $a_1, b_1, c_1, a_2, b_2, c_2$. And for transactions defined as concurrent rules, e.g.,

$$task :- a_1 * \alpha \# \beta * b_3.$$

$$\alpha :- a_2 * a_3.$$

$$\beta :- b_1 * b_2.$$

After replacing transactions α and β with their respective definitions, the resulting execution sequence is :

$$task :- a_1 * a_2 * a_3 \# b_1 * b_2 * b_3.$$

In this way, a composite transaction can be decomposed to an execution sequence consisting of elementary operations.

3.3 State Updates and Queries

In *CTR*, the semantics of database states queries and updates are defined by a pair of relational oracles, called, state data oracle and state transition oracle. The *CTR* prototype uses the following predicates to implement the state queries and updates defined in the oracles.

- $db(p)$ determines whether atom p is true in the current database state.
- $ins(p)$ inserts atom p in the database.
- $del(p)$ deletes atom p in the database.

To be able to preserve the database consistency when a transaction that updates the database fails, the *CTR* prototype uses a backtracking mechanism for undoing updates, i.e., to rollback the failed execution and enable re-execution. It performs a *relative commit* of the updates, as opposed to an *absolute commit*, which commits only after the entire transaction succeeds. Considering the following transaction:

$$ins(a) * db(b) * ins(c)$$

Assume atom a and c are not in the initial database. The first update $ins(a)$ inserts the atom a into the database. Then, the second operation queries the database. If the atom b exists in the database, then the third operation, $ins(c)$, will take place and thus the three atoms are all in the final database. However, if the atom b is not in the database, then the transaction execution fails, and the update $ins(a)$ is undone. In this case, although the update performed by $ins(a)$ has been already executed, the database is able to restore to its initial state and the database consistency is thus guaranteed. The undo is possible because the commit was not absolute, but was relative to the overall transaction.

To undo updates in a failed transaction, the *CTR* prototype actually introduces two underlying atom tags to represent the existence of an updatable atom p , namely, $inserted(p)$ and $deleted(p)$.

The atom p is true in the database if $inserted(p)$ is true and $deleted(p)$ is not true in the database. The atom p is not true in the database if $deleted(p)$ is true or $inserted(p)$ is not true.

By implementing these two atom tags, the technique in [17] leads to transaction rollback or partial rollback, and re-execution when a failure occurs, and thus prevents the Prolog database from being corrupted by backtracking through updates.

Chapter 4

TP-CTR

ALTHOUGH *CTR* is a logic for specifying and executing database processes involving queries and updates concurrently as well as serially, *CTR* does not support explicit priority constraints in its logic framework. In *CTR*, the *programmer*¹ uses the sequential and concurrent connectives to specify the ordering of a transaction execution. In a concurrent conjunction, there is no preference over which transaction will be selected for execution. Hence, *CTR*'s orientation is towards logic *eventuality* rather than real-time *immediacy*. This particular feature of the logic has limited its application extended to real-time domains. Also, *CTR* can not handle timing-related applications due to lack of timing constraints, e.g., a timing relation to specify when to execute a process.

To overcome this limitation and extend concurrent transaction logic to the real-time domain, we extend the *CTR* logic framework to a so-called timing-event-based prioritized concurrent transaction logic (*TP-CTR*), which provides a high-level formalism for specifying priority constraints and timing constraints in timing-event-based transaction applications. Summarily, *TP-CTR* extends *CTR* in the sense that it introduces *constraints* in concurrent Horn rules, a *translation mechanism* to translate such rules, and an *inference system* to handle prioritized transactions.

In this chapter, we describe the *TP-CTR*'s syntax, informal semantics, the inference system and the interpretation of constraint concurrent Horn rules.

¹Like *CTR*, *TP-CTR* is a logic for programming transactions. Hence in this thesis, *users* and *programmers* are considered synonyms

4.1 Syntax of $TP-CTR$ Priority and Timing Constraints ²⁷

CTR does not include priority constraints in its logic framework. In CTR , the programmer uses the sequential and concurrent connectives to specify the transaction execution order in programs. For instance, in a concurrent conjunction, all the concurrent processes have the same opportunity to be selected for execution. This feature of the logic has limited its application in real-time domains.

$TP-CTR$, an extension of CTR , is designed to provide a high-level framework for specifying priority constraints and timing constraints in concurrent transaction logic systems. To allow the specification of priority constraints and timing constraints that may be used to trigger or interrupt other transactions, we use *constraint concurrent Horn rules*. In $TP-CTR$, $b \leftarrow \phi : \psi$ is a constraint concurrent Horn rule if $b \leftarrow \phi$ is a concurrent Horn rule and ψ is a *constraint formula*. A $TP-CTR$ program is composed of a set of constraint concurrent Horn rules. The definition of the concurrent Horn rule $b \leftarrow \phi$ here is the same as that in CTR , as shown in definition 1. Below is the definition of the constraint formula ψ .

Definition 7 (Constraint formula) A Constraint formula is any formula of the form:

- $priority(\alpha, p)$, where α is an atomic formula in ϕ and p is an integer number, representing the priority level of the atomic formula α .
- $timeElapsed(t) \rightarrow triggerEvent(\alpha)$, where t is a positive integer number, and α is an atom in ϕ .
- $\psi_1 \wedge \dots \wedge \psi_k$, where each ψ_i is a constraint formula, and $K \geq 0$.

□

With the above extended syntax definition, we can use *constraint concurrent Horn rules* to specify priority and timing constraints, which may be used to trigger or interrupt other transactions.

When a constraint concurrent Horn rule has no constraint formula, it becomes a concurrent Horn rule. A *TP-CTR program* consists of both a set of concurrent Horn rules and a set of constraint concurrent Horn rules.

4.2 Informal Semantics of *TP-CTR*

Next we introduce the informal semantics of constraint concurrent Horn rules from the users' viewpoint.

Priority constraints

In *TP-CTR*, priority constraints specify the execution priority of transactions occurring in a concurrent goal formula. For instance, the constraint concurrent Horn rule below specifies the execution priorities of the concurrent transactions p , q , and r .

$$s \leftarrow p \mid q \mid r : \text{priority}(p, 2) \wedge \text{priority}(q, 1) \wedge \text{priority}(r, 1) \quad (4.1)$$

“To execute s , execute p , q , and r concurrently, observing that p has higher priority than q and r , and q has the same priority as r .”

Notice that the second argument of the priority predicate represents the execution priority level of the predicate occurring in its first argument. The bigger the number representing the priority level, the higher its execution priority is. Priority levels ranges from 1 to 100. The default priority level of a transaction predicate is 1, i.e., if a constraint concurrent Horn rule does not specify the priority of a transaction predicate explicitly, then it is given the lowest execution priority.

But then one might ask: why use the complicated formula above when one can more succinctly write the equivalent and more straightforward formula below?

$$s \leftarrow p \otimes (q \mid r)$$

This is because priority constraints like the ones in Expression (4.1) are useful specially when combined with timing constraints. The examples in the rest of this section illustrate the point.

Timing Constraints

Timing properties are usually specified by the timing relations among events. In [5], the timing relations refer to terms of time that specify constraints among events. We are interested in these timing properties and adopt some of them into our system. Below we illustrate some examples involving timing constraints.

Let b , c , and d be transactions. The rule below specifies a timing constraint for b :

$$a \leftarrow b \mid c \mid d : \text{timeElapsed}(100) \rightarrow \text{triggerEvent}(b) \quad (4.2)$$

The rule above specifies the following: to execute a , execute b , c , and d concurrently, but delaying the execution of b 100 seconds. That is, for the first 100 seconds, only c and d should run concurrently, and then after 100 seconds, transaction b should be triggered, and thus added to the concurrent execution.

The example below illustrates how priority and timing constraints can be combined to provide interesting real-time *interrupt* behavior.

Let again b , c , and d denote transactions. The constraint concurrent Horn rule below, specifies timing and priority constraints for transaction b .

$$a \leftarrow b \mid c \mid d : \text{timeElapsed}(100) \rightarrow \text{triggerEvent}(b) \wedge \text{priority}(b, 2)$$

The priority constraint specifies that b has execution priority over c and d . Notice that no constraint is specified for c or d , i.e., they run as a flat concurrency (with the default priority level 1). The timing constraint specifies that the execution of b should start 100 seconds after a is activated.

Let's assume that c and d represent complex transactions whose respective execution times last for more than 200 seconds time. Given these assumptions, since b has high

execution priority, the activation of b should interrupt the execution of c and d . That is to say, after 100 seconds, b starts execution, and c and d are suspended because of their lower priority level. Until b finishes its execution, the execution of c and d can be resumed. This behavior of concurrent transaction execution, created by the combined use of timing and priority constraints, illustrates the interrupt functionality. This is what often happens in a real-time domain scenarios, and the reason why we introduce the priority constraint into the extended system.

Some comments on the use of constraints in $TP-CTR$: one may have noticed that priority constraints are not suitable for all transactions in a concurrent sequential goal. For example, suppose a , b , and c are transactions, and r is defined by the rule below:

$$r \leftarrow a \otimes b \otimes c$$

It does not make sense to specify a priority for the transactions occurring in the body of the rule above, since the semantics of the \otimes connective specifies the execution order of the goal $a \otimes b \otimes c$. Priority constraints only make sense when they refer to transactions occurring in a concurrent conjunction, since it can change the execution sequence of the transactions involved in the conjunction.

4.3 Inference System

The inference system introduced here differs from the inference system introduced in CTR in the sense that in CTR , there is no criteria for selecting which transaction in a concurrent goal will be executed. This means in CTR , the execution order of concurrent transactions is non-deterministic. In $TP-CTR$, on the other hand, a simplified Rate Monotonic scheduling algorithm based on the priority level of a predicate, is employed as the selection criteria. Thus, the execution order of concurrent transactions with different priority levels is deterministic.

The next definition formalizes this idea. It defines the transaction to be executed amongst the candidates for execution. We refer to it as the “hottest” component of a

concurrent sequential goal.

Definition 8 (Hottest component) Let ϕ be a concurrent sequential goal. Its hottest component, denoted $hottest(\phi)$, is defined recursively as follows:

- $hottest(()) = \{\}$, where $()$ is the empty goal;
- $hottest(b) = \{b\}$, if b is a atomic formula;
- $hottest(\psi_1 \otimes \dots \otimes \psi_k) = hottest(\psi_1)$;
- $hottest(\psi_1 \mid \psi_2 \mid \dots \mid \psi_k) = \begin{cases} hottest(\psi_1) & , \text{ if } pLevel(hottest(\psi_1)) \geq \\ & pLevel(hottest(\psi_2 \mid \dots \mid \psi_k)) \\ hottest(\psi_2 \mid \dots \mid \psi_k) & , \text{ otherwise} \end{cases}$
where $pLevel(\phi)$ denotes the priority level of ϕ .
- $hottest(\odot\psi) = \odot\psi$.

□

Since the compiler adds an extra first parameter to all prioritized predicates and encodes in this parameter the priority level of the predicate, one can obtain the priority level of a sequential goal as follows:

- $pLevel(p(priority(l), \dots)) = l$, if $p(\dots)$ is an atomic formula;
- $pLevel(\psi_1 \otimes \dots \otimes \psi_k) = pLevel(\psi_1)$

Like CTR , $TP-CTR$ also uses the SLD-style resolution introduced in Chapter 2, the only difference between them being the inference rules and the notion of hot component used in CTR , and hottest component used in $TP-CTR$. In CTR , hot component is a set representing the sub-transactions ready for execution. In $TP-CTR$, hottest component is the first highest priority-level component of the CTR hot component set.

Axiom: $P, D \vdash ()$, for any state D

Inference rules: In rules 1-3, σ is a substitution, ψ and ψ' are concurrent sequential goals, and $hottest(\psi) = a$.

1. *Applying rule definitions:* Suppose $b \leftarrow \beta$ is a rule in \mathbf{P} whose variables have been renamed so that the rule shares no variables with ψ . If a and b unify with mgu σ , then

$$\frac{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi' \sigma}{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by replacing a by β .

2. *Querying the database:* If $\mathcal{O}^d(\mathbf{D}_i) \models^c (\exists) a \sigma$, and $a \sigma$ and $\psi' \sigma$ share no variables, then

$$\frac{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi' \sigma}{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting a .

3. *Executing elementary updates:* If $\mathcal{O}^t(\mathbf{D}_1, \mathbf{D}_2) \models^c (\exists) a \sigma$, and $a \sigma$ and ψ' share no variables, then

$$\frac{\mathbf{P}, \mathbf{D}_2 \vdash (\exists) \psi' \sigma}{\mathbf{P}, \mathbf{D}_1 \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting a .

4. *Executing atomic transactions:* If $\odot \alpha$ is the hottest component in ψ , then

$$\frac{\mathbf{P}, \mathbf{D} \vdash (\exists) (\alpha \otimes \psi')}{\mathbf{P}, \mathbf{D} \vdash (\exists) \psi}$$

where ψ' is obtained from ψ by deleting $\odot \alpha$.

Each inference rule consists of two sequents, and has the following interpretation: if the upper sequent (G_{i+1}) can be inferred, then the lower sequent (G_i) can also be inferred.

4.4 Interpreting Constraint Concurrent Horn Rules

Based on the constraint formalism presented in the above sections, here we initially present how the *TP-CTR* interpreter translates constraint concurrent Horn rules into concurrent Horn rules. Then we show how to formally execute prioritized constraint concurrent Horn rules using the SLD-style resolution procedure. Finally we illustrate how the *TP-CTR* compiler translates the concurrent Horn rules with timing constraints.

4.4.1 Compiling Priority Constraints

To illustrate how the translation takes place, we use the following example: let the program below define transactions s , p , and q .

$$\begin{aligned} s &\leftarrow p \mid q : \text{priority}(p, 2) \\ p &\leftarrow \text{ins}(r(a)) \\ q &\leftarrow \text{ins}(r(b)) \end{aligned} \tag{4.3}$$

Notice that the first rule specifies: to execute s , execute p and q concurrently. Its constraints specify: in the concurrent sequential goal $p \mid q$, p has higher priority than q . The second and third rules specify: to execute p , insert the atom $r(a)$ in the database; to execute q , insert the atom $r(b)$ in the database.

In essence, the compilation consists of translating the priority predicates occurring in the constraint formula into function terms. These function terms are added as an extra argument to the respective predicates in the head and body of the concurrent Horn rules.

Below we present the result of the translation of program (4.3).

$$\begin{aligned} s &\leftarrow p(\text{priority}(2)) \mid q(\text{priority}(1)) \\ p(\text{priority}(2)) &\leftarrow \text{ins}(r(a)) \\ q(\text{priority}(1)) &\leftarrow \text{ins}(r(b)) \end{aligned} \tag{4.4}$$

Notice that the compiler assigned the default priority level 1 to transaction q .

4.4.2 Using the Inference System to Execute Transactions

With the *TP-CTR* four inference rules, we can deduce the transaction s defined in program (4.3) as in the Table 4.1. Notice in program (4.3) and on Table 4.1 that if no priority constraints is used in the specification of s , then there are two possible sequence of database states when s executes: one in which $r(a)$ is inserted first and then $r(b)$ is inserted, and the other in which $r(b)$ is inserted first and then $r(b)$. However, because of the priority constraints bound to transaction p and q , they force the transaction s to run in a definite sequence, which causes $r(a)$ to be inserted before $r(b)$ is inserted in the database.

Table 4.1: A deduction for program (4.3)

Sequents	Inference rule	Hottest component
$P, \{\} \vdash s$	1	s
$P, \{\} \vdash p(\text{priority}(2)) \mid q(\text{priority}(1))$	1	$p(\text{priority}(2))$
$P, \{\} \vdash \text{ins}(r(a)) \mid q(\text{priority}(1))$	3	$\text{ins}(r(a))$
$P, \{r(a)\} \vdash q(\text{priority}(1))$	1	$q(\text{priority}(1))$
$P, \{r(a)\} \vdash \text{ins}(r(b))$	3	$\text{ins}(r(b))$
$P, \{r(a), r(b)\} \vdash ()$	<i>Axiom</i>	$\{\}$

4.4.3 Compiling Timing Constraints

The program below specifies a timing constraint for transaction b .

$$\begin{aligned}
 a &\leftarrow b \mid c \mid d : \text{timeElapsed}(100) \rightarrow \text{triggerEvent}(b) \\
 b &\leftarrow \text{ins}(r(e)) \\
 c &\leftarrow \text{ins}(r(f)) \\
 d &\leftarrow \text{ins}(r(g))
 \end{aligned}$$

The *TP-CTR* interpreter translates it into the following program

$$\begin{aligned}
 a &\leftarrow b \mid c \mid d \\
 b &\leftarrow \text{timeElapsed}(100) \otimes \text{ins}(r(e)) \\
 c &\leftarrow \text{ins}(r(f)) \\
 d &\leftarrow \text{ins}(r(g))
 \end{aligned} \tag{4.5}$$

where $timeElapsed(t)$ is a built-in predicate, which is satisfied when t seconds have elapsed since the execution of concurrent rule a has started. This predicate is interpreted by the $TP-CTR$ prototype interpreter in a specific way in the chapter 5. Here, we present the logical definition of the predicate $timeElapsed$ below:

$$\begin{aligned}
 timeElapsed(DelayTime) \leftarrow & \\
 & \quad startTime(InitialTime) \\
 & \quad \otimes now(PresentTime) \\
 & \quad \otimes (PresentTime \geq InitialTime + DelayTime) \\
 timeElapsed(DelayTime) \leftarrow & \\
 & \quad startTime(InitialTime) \\
 & \quad \otimes now(PresentTime) \\
 & \quad \otimes (PresentTime < InitialTime + DelayTime) \\
 & \quad \otimes timeElapsed(DelayTime)
 \end{aligned}$$

In the above definition, you can notice that if pre-condition $(PresentTime \geq InitialTime + DelayTime)$ fails, then $timeElapsed(DelayTime)$ stays in a sort of while-loop until the condition $(PresentTime \geq InitialTime + DelayTime)$ is satisfied, and the corresponding transaction is then triggered.

Chapter 5

The $TP\text{-}CTR$ Prototype


THE $TP\text{-}CTR$ prototype is an implementation of the Horn fragment of $TP\text{-}CTR$ and the inference system, both introduced in Chapter 4. In Chapter 3 we have presented the outline of the CTR prototype. Since the $TP\text{-}CTR$ prototype is an extension of the CTR prototype, here we focus on their differences.

As the CTR prototype shown in Figure 3.2, the $TP\text{-}CTR$ prototype also has two major components: the compiler and the interpreter. Both have been extended to handle the timing and priority constraints introduced in $TP\text{-}CTR$. In this chapter, we first present how the $TP\text{-}CTR$ prototype compiler translates a $TP\text{-}CTR$ program. Then we show how the $TP\text{-}CTR$ prototype interpreter works.

5.1 The $TP\text{-}CTR$ Prototype Compiler

Before executing a $TP\text{-}CTR$ program, a user first has to compile the transaction base and database files. In a $TP\text{-}CTR$ program, priority constraints and timing constraints are, in fact, a syntax sugar, which we call *constraint concurrent Horn rules*. Unlike the CTR prototype, the $TP\text{-}CTR$ prototype compiler first translates these constraint concurrent Horn formulas into CTR recognizable formulas. Figure 5.1 shows how this process takes place.

This translation consists of two stages, carried on by compiler I and compiler II,

Figure 5.1: The *TP-CTR* compiler.

respectively, as illustrated in Figure 5.1. Compiler I translates the constraint concurrent Horn rules in a *TP-CTR* program into an intermediate *CTR* program. Then compiler II translates this intermediate program into an internal notation format.

The reason why the *TP-CTR* prototype compiler translates a *TP-CTR* program into a *CTR* recognizable program instead of using *CTR* recognizable program at the beginning is because the translated *CTR* recognizable syntax looks clumsy, unlike a typical neat Prolog-like syntax. For example, assigning every transaction a priority level and distributing timing constraints to the triggered transactions. To make formulas neat and clear to users, the *TP-CTR* prototype puts all these constraints in the syntax sugar format and leaves the compiler to do the translation. Moreover, compiler I takes care of this part of the translation, only. Hence, introducing more complicated timing constraints and priority constraints into the system would require only extending compiler I. This improves the flexibility of the *TP-CTR* prototype compiler.

The example below illustrates the idea of compiler I. Suppose we have the following rules in the transaction base:

```

task :- initTime * taskA : timeElapsed(10) → triggerEvent(taskA).
taskA :- monitor('taskA complete').

```

Compiler I creates a temporary transaction-base file, which is added an extension *.temp* to the end of the old transaction-base file. Below are the corresponding translated rules:

```

task :- initTime * taskA.
taskA :- timeElapsed(10) * monitor('taskA complete').

```

If a rule in the original transaction base does not have a constraint, the translated rule is the same as the original one.

The compiling process carried on by compiler II is more elaborate, and thus deserves a more detailed analysis of the processes and notations used within. The next two sections provide this analysis.

5.1.1 Transaction Base Internal Syntax

Like a *CTR* program, a *TP-CTR* program consists of a transaction-base and a database too. We saw in Section 5.1 that ultimately a *TP-CTR* program is translated into an internal notation format. This subsection introduces the internal notation format.

The transaction base of a *TP-CTR* program has two kinds of components: atomic formulas and rules. Atomic formulas do not change after translation.

On the other hand, constraint concurrent Horn rules in transaction-base are in the form of $\alpha :- \beta : \gamma$, where α is an *atomic formula*, β is a *transaction formula*, and γ is a *constraint formula*, which can include both priority constraints and timing constraints. The head α is translated into the internal notation $trans(\alpha)$. Transaction formula notations involving sequential conjunction and concurrent conjunction, are translated into internal notations as follows:

- Sequential conjunction: $seq([a_1, a_2, \dots, a_n])$ is the internal representation of transaction formula $a_1 * a_2 * \dots * a_n$.
- Concurrent conjunction: $conc([a_1, a_2, \dots, a_n])$ is the internal representation of the transaction formula $a_1 \# a_2 \# \dots \# a_n$.
- Modality operator: $isolate(\theta)$ is the internal representation of $o(\theta)$.

Constraint formula notations involving priority and timing constraints are translated as follows:

- Priority constraint: $a(\text{priority}(\text{Level}))$ is the internal representation of the transaction a with the priority constraint $\text{priority}(\text{Level}, a)$.

- Timing constraint:

$$\text{seq}([\text{timeElapsed}(a, \text{Time}), \text{or}([\text{seq}([\text{isolate}(\text{seq}([\text{event}(a, \text{Time}), \text{now}(X), \\ \text{db}(\text{start_time}(Y)), X \geq \text{Time} + Y, \text{del}(\text{event}(a, \text{Time}))))], \\ , b)], \text{seq}([\text{event}(a, \text{Time}), a])])])$$

is the internal representation of the rule body of $a :- b$, where a is triggered by the timing constraint $\text{timeElapsed}(\text{Time}) \rightarrow \text{triggerEvent}(a)$. Note: the predicate $\text{timeElapsed}(\text{Time})$ in the translated intermediate file is replaced by this internal notations with the *TP-CTR* prototype interpreter recognizable built-in predicate $\text{timeElapsed}(a, \text{Time})$.

The next examples illustrate these translation methods of internal syntax.

Example 4 (Translating sequential and concurrent goals) Suppose a goal connected with sequential and concurrent conjunctions is as the form of:

$$p :- a * b \# c.$$

Like the *CTR* compiler, *TP-CTR* compiler translates it into the following internal syntax:

$$\text{trans}(p) :- \text{conc}([\text{seq}([a, b]), c]).$$

□

If a rule has priority constraints, it is translated as shown in the example below:

Example 5 (Translating priority constraints) Suppose in a transaction-base, there is the following simple priority constraint concurrent Horn rule:

$$p :- a \# b : \text{priority}(a, 8).$$

which has the following intermediate and internal representations during and after compilation:

$p :- a(priority(8)) \# b.$ - Intermediate syntax, i.e., equivalent *CTR* rule
 $trans(p) :- conc([a(priority(8)), b]).$ - *TP-CTR* internal syntax

□

The following example shows the original, intermediate and internal syntax in the case of a timing constraint concurrent Horn rule.

Example 6 (Translation of timing constraint rules) The rules below specify the concurrent execution transactions a and b . For the sake of easy understanding, the two transactions are just assigned simple tasks. The transaction formula a is scheduled to execute ten seconds after the execution begins.

$p :- initTime * (a \# b) : timeElapsed(10) \rightarrow triggerEvent(a).$
 $a :- monitor('transaction a completed').$
 $b :- monitor('transaction b completed').$

As the first step, the *TP-CTR* compiler I translates the *TP-CTR* program above into the following *CTR* program:

$p :- initTime * (a \# b).$
 $a :- timeElapsed(10) * monitor('transaction a completed').$
 $b :- monitor('transaction b completed').$

Then the *TP-CTR* compiler II translates the *CTR* program into an internal notation.

$trans(p) :- seq([initTime, conc([a, b])]).$
 $trans(a) :- seq([timeElapsed(a, 10), or([seq([isolate(seq([event(a, 10), now(X),$
 $\quad db(startTime(Y)), X \geq 10 + Y, del(event(a, 10))])),$
 $\quad monitor('transaction a completed')]), seq([event(a, 10), a])])]).$
 $trans(b) :- monitor('transaction b completed').$

□

When a *TP-CTR* program combines timing constraints and priority constraints, it is translated in the way as shown in the example below.

Example 7 (Translating priority and timing constraint) Suppose we add a constraint to the rule *p* shown in Example 6, resulting in the following rule:

$$p :- \text{initTime} * (a \# b) : \text{timeElapsed}(10) \rightarrow \text{triggerEvent}(a) \\ \wedge \text{priority}(a, 8).$$

First, the *TP-CTR* compiler I translates them into the following *CTR* program:

$$p :- \text{initTime} * (a(\text{priority}(8)) \# b). \\ a(\text{priority}(Z)) : - \text{timeElapsed}(10) * \text{monitor}('transaction a completed'). \\ b :- \text{monitor}('transaction b completed').$$

Then the *TP-CTR* compiler II translates the *CTR* program above into the following internal notation:

$$\text{trans}(p) :- \text{seq}([\text{initTime}, \text{conc}([a(\text{priority}(8)), b])]). \\ \text{trans}(a(\text{priority}(Z))) :- \text{seq}([\text{timeElapsed}(a(\text{priority}(Z)), 10), \text{or}([\text{seq}([\text{isolate}(\text{seq}([\\ \text{event}(a(\text{priority}(Z)), 10), \text{now}(X), \text{db}(\text{start_time}(Y)), X \geq 10 + Y, \\ \text{del}(\text{event}(a(\text{priority}(Z)), 10)])), \text{monitor}('transaction a completed')]) \\ , \text{seq}([\text{event}(a(\text{priority}(Z)), 10), a(\text{priority}(Z))])])]). \\ \text{trans}(b) :- \text{monitor}('transaction b completed').$$

□

By recursively applying the translation rules above, the *TP-CTR* compiler translates a transaction-base file into a *TP-CTR* transaction-base object file.

5.1.2 The *TP-CTR* Prototype Database

The syntax and semantics of the *TP-CTR* prototype database is the same as in the *CTR* prototype database. The only difference is that internally, the system uses a couple of

other system predicates. In the *TP-CTR* prototype databases, there are three system-required database tuple declarations in *TP-CTR* prototype databases, namely:

- *updatable event/2.* - the declaration of timing-event-to-be-triggered list
- *updatable priorityList/1.* - the declaration of concurrent transaction priority list
- *updatable start_time/1.* - the declaration of execution start time

5.2 The *TP-CTR* Prototype Interpreter

The *TP-CTR* prototype interpreter consists of inference engine and relational oracle, as the *CTR* prototype interpreter shown in Figure 3.3. Although both the *CTR* prototype and the *TP-CTR* prototype use the same relational oracle, the *TP-CTR* prototype interpreter use a different inference engine. Compared to the *CTR* inference engine, it implements the notion of *hottest component* instead of *hot component*, as shown in Chapter 4. During execution, the *TP-CTR* inference engine picks up the *hottest* transaction component, i.e., the first highest-priority-level *element* in the *hot component* queue.

The interpreter is the key module of the *TP-CTR* prototype. The user interacts with the prototype by presenting transaction goals to be executed, *execute(Goal)*. The interpreter then picks up the *hottest component* of the goal for execution. If the program in the transaction base does not include any constraints, the priority level of all transactions is considered as the default lowest priority level. In this case, the notion of *hottest component* introduced in *TP-CTR* system is analogous to the notion of *hot component* introduced in *CTR* system because the underlying simplified Rate-Monotonic algorithm handles the non-priority-constraint goals in the same way as that of the Round-Robin algorithm.

The inference engine of *TP-CTR* uses the underlying scheduling algorithm to pick up the *hottest component* in a concurrent transaction execution. It also interprets the priority and timing constraints and drives the execution of the transactions. The rest of this

action is organized as follows: Next we introduce the adopted underlying scheduling algorithm - *simplified Rate Monotonic Algorithm*(*SRMA*)[10]. Then we present how priority constraints are interpreted, and how the *SRMA* schedules the concurrent transactions. Finally, we present how the *TP-CTR* timing constraints are interpreted.

5.2.1 Underlying Scheduling Algorithm in the *TP-CTR* Interpreter

Before introducing the underlying scheduling algorithm in the *TP-CTR* interpreter, we first give a brief introduction of the Round-Robin scheduling algorithm used in the *CTR* interpreter.

Round-Robin Scheduling Algorithm

The Round-Robin algorithm is one of the oldest, simplest and most widely used scheduling algorithms, designed especially for time-sharing systems. It assigns each concurrent process an unit of time in a one-by-one sequence and feeds the process that has used up its share of time back to the end of the sequence. In *CTR*, Bonner made a small adjustment: using one step transaction execution instead of one unit of cpu time, to simplify transaction execution in a concurrent sequence. Thus, the Round Robin scheduling algorithm in the *CTR* system assigns one execution step time to the selected transaction, which is in the head of the concurrent transaction sequence. In *CTR*, one execution step means one inference step. In this way, the *CTR* scheduler picks up the first hot component, executes one inference step, feeds the unfinished transaction component back to the end of the sequence, and releases CPU resource to the scheduler in the mean time.

The Round-Robin scheduling algorithm is a simple and effective algorithm for a concurrent transaction system without priority demands. However, in reality, many concurrent transaction systems need to assign priority level to some special transactions. For example, in a classified-customer financial system, a transaction request from a higher level customer always has preference over that from a lower level customer. When faced with such requirements, the Round-Robin scheduling algorithm shows its application limita-

tion. In order to handle such types of applications, the *TP-CTR* system uses a simplified version of *Rate monotonic algorithm*.

Simplified Rate-Monotonic Scheduling Algorithm

In the SRMA, each concurrent transaction is assigned a fixed priority by user or by default. Assuming task *a* and task *b* are concurrent execution transactions, during execution, the SRMA schedules tasks *a* and *b* in the following way:

- Task *a* can not be executed before task *b* finished, if task *b* has a higher priority level than task *a*.
- Task *a* has an equal chance to execute as task *b*, if task *a* and *b* have the same priority level.

In the design of *TP-CTR*, our main goal is to verify the logic of time-event-based prioritized transaction systems, while introducing features of real-time systems, i.e., priority constraints and timing constraints. However, we do not include a great variety of real-time features at current stage. More precisely, we ignore an important feature of real-time systems, *deadline*, which improves the efficiency of concurrent systems by giving up some tasks running over *deadlines*. Moreover, the fixed priorities of the system are assigned by users when writing *TP-CTR* programs. These adjustments simplify the Rate Monotonic algorithm in our system by ignoring the transaction execution period assessment and corresponding automatic priority assignment, which is not important in a timing-event-based transaction verification system.

5.2.2 Interpreting and Scheduling Priority Constraints

Section 5.1 shows how the *TP-CTR* compiler translates a concurrent Horn rule with priority constraints. This subsection presents how the *TP-CTR* prototype interpreter interprets priority constraints represented in internal notations.

After the translation by the *TP-CTR* prototype compiler, a priority constraint is added as the first parameter of the corresponding transaction, as shown in the example 5. Then the *TP-CTR* prototype interpreter uses a SRMA to schedule constraint concurrent Horn rules. Each concurrent transaction in a *TP-CTR* program has a priority level, which is used to determine their execution priority. To find the first highest priority level transaction from a concurrent sequence, the SRMA needs to know both the priority level of a transaction and the highest priority level of the concurrent transaction sequence. In the *TP-CTR* prototype, lists are used to store these variants.

- The list *priorityList* stores all priority levels of the concurrent transactions in a priority descending order. For example, *priorityList* [3, 2, 2] means there are three transactions in the concurrent sequence, and their priority levels are 3, 2 and 2, respectively.
- During interpretation, a transaction's priority level is indicated by the last element of the transaction list, e.g., [seq([a, b]), 2] means the seq([a, b]) has priority level 2. The priority level will not disappear before this sequential transaction execution finishes.

The example below illustrates how priority constraints are interpreted by the *TP-CTR* interpreter according to the simplified RMA.

Example 8 (Interpreting a *TP-CTR* program) Consider the following transaction base:

```

task      :- taskA#taskB#taskC
           : priority(taskA, 1) ∧ priority(taskB, 2) ∧ priority(taskC, 2).
taskA     :- ins(in(a)).
taskB     :- ins(in(b)).
taskC     :- ins(in(c)).

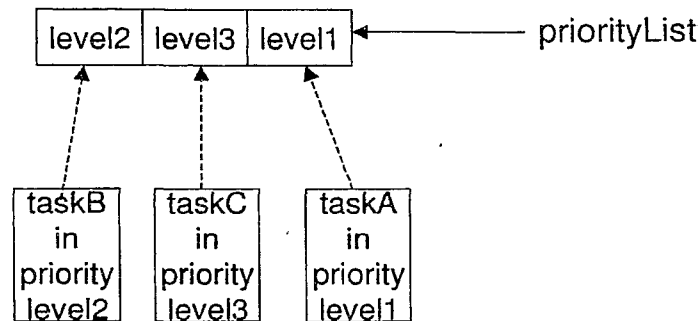
```

When interpreting the concurrent transaction *task*, the sub-transaction *taskA*, *taskB* and *taskC* are converted to the following *conc* list in the *TP-CTR* prototype interpreter:

$conc([ins(in(a)),1],[ins(in(b)),2],[ins(in(c)),2]])$

where the respective priority level of each execution transaction is the last element of their list. The list *priorityList* is used to store the three transaction priority levels, as shown in Figure 5.2:

Figure 5.2: The list *priorityList* in *TP-CTR* interpreter.



where *level1* is the taskA's priority level 1, *level2* is the taskB's priority level 2, and *level3* is the taskC's priority level 2. The elements of the *priorityList* sort in an descending order, from higher priority level to lower priority level. \square

During interpretation, the SRMA uses the following steps to determine the *hottest component*:

1. Picks up the first element of the execution transaction list
2. Gets the priority level of the first element from the execution transaction list
3. Gets the highest priority level from the *priorityList*(the first element)
4. If the current element's priority level is not the highest priority level, puts it in the end of the execution transaction list, then returns to the first step, and repeats steps 1 to 4.

5. Otherwise, the element is the *hottest component* and the interpreter should execute one step of this element.

- If this step is the last step of the belonged transaction execution, deletes the first element of *priorityList*. Then returns to the step 1 if there is still transactions in the execution queue; or ends execution if no other transaction is left in the execution sequence.
- Otherwise, returns the remaining of this element to the end of the execution transaction list, and then returns to the step 1.

In such a way, the *TP-CTR* interpreter implements the inference system using a SRMA to formally execute *TP-CTR* programs.

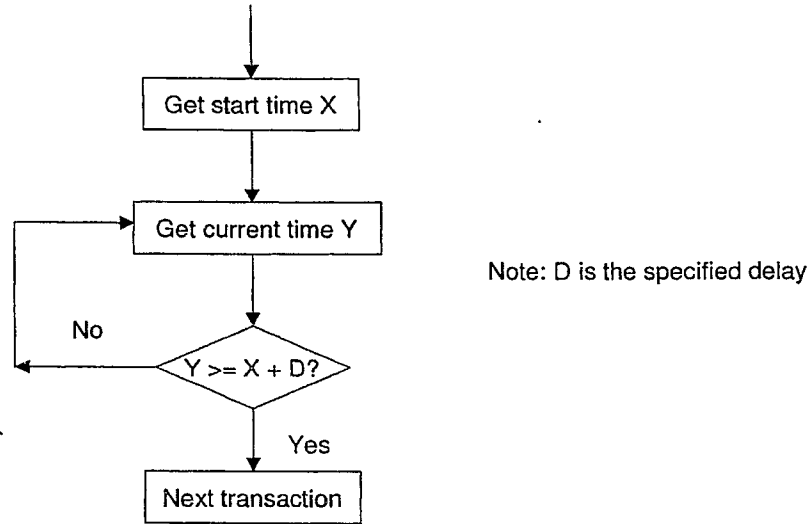
5.2.3 Interpreting Timing Constraints

At current stage, the *TP-CTR* prototype has two timing constraints: *timeElapsed/1* and *delay/1*. *timeElapsed* is used to set an absolute timing relation, where a transaction is set to be triggered at an absolute time instant. *delay* is used to set a relative timing relation, where a transaction is set to start a period of time after the end of another transaction.

Actually, the timing constraints *timeElapsed* and *delay* are similar from a logic viewpoint. The only difference is the start time instant. Figure 5.3 illustrates the underlying logic for these two timing relations.

Both timing constraints are handled using such logic within the *TP-CTR* interpreter. The timing constraint *timeElapsed* uses the *start_time(X)* to retrieve the start time, which stores the start time instant of the whole execution. In *TP-CTR*, we use timing-event-based mechanism for this timing constraint: satisfaction of the constraint precondition is used to trigger another transaction event. The *TP-CTR* interpreter uses an atom *event* to specify the triggered transaction. After the timing condition is satisfied and the

Figure 5.3: Timing logic handled in the *TP-CTR* interpreter.



specified transaction is triggered, the atom *event* corresponding to this transaction is deleted from the database system.

The other timing constraint of the current *TP-CTR* prototype, *delay*, is used to define a delay relationship between two transactions. Thus, its start time is not the start time of the execution, but the start time of itself. When the specified delay time goes, the delay loop will end and it will go to the second transaction execution.

Chapter 6

Time-based Prioritized $TP\text{-}CTR$ Program Examples

TO illustrate how $TP\text{-}CTR$ can be used to specify priority and timing constraints in timing-event-based prioritized transaction system, this chapter presents two examples in different application areas. More specifically, the first example regards financial transaction. The second example regards real-time logic control system.

6.1 A Time-based Financial Transaction Application

Example 9 is an extension of the *financial transaction application* presented in [1], where Bonner and Kifer illustrate how CTR can be used to specify the atomicity of financial transactions. Here we show how timing features and priority features in $TP\text{-}CTR$ is used to schedule tasks in a long-run financial transaction application.

Example 9 (Specifying time-based prioritized financial transactions) For the sake of understanding, first we briefly introduce the rules used in Bonner and Kifer's financial transaction example in [1].

$$\begin{aligned} transfer(Amt, Acct1, Acct2) \leftarrow & \odot (withdraw(Amt, Acct1) \\ & \otimes deposit(Amt, Acct2)) \end{aligned}$$

$$\begin{aligned}
& \text{withdraw}(Amt, Acct) \leftarrow \text{balance}(Acct, Bal) \\
& \quad \otimes Bal \geq Amt \\
& \quad \otimes \text{change_balance}(Acct, Bal, Bal - Amt) \\
& \text{deposit}(Amt, Acct) \leftarrow \text{balance}(Acct, Bal) \\
& \quad \otimes \text{change_balance}(Acct, Bal, Bal + Amt) \\
& \text{change_blance}(Acct, Bal1, Bal2) \leftarrow \text{del}(\text{balance}(Acct, Bal1)) \\
& \quad \otimes \text{ins}(\text{balance}(Acct, Bal2))
\end{aligned}$$

The first predicate, *transfer*, specifies how a money transfer transaction is accomplished, i.e., by withdrawing an amount *Amt* from one bank account *Acct1* and then *depositing* the *Amt* into another bank account *Acct2*. The next two predicates, *withdraw* and *deposit*, define the *withdraw* transaction and the *deposit* transaction, respectively. Both of them use the predicate *change_balance* to update an account's balance via the built-in elementary updates *del(balance(Acct, Bal))* and *ins(balance(Acct, Bal))*.

Based on the above basic transaction rule definitions, the *TP-CTR* program below simulates a long-run time-based financial transaction scenario: to a source account, higher priority level transfer requests can always be answered even when normal transfer transactions are executing. In the program, a timed prioritized transfer transaction is assigned with a higher execution priority. Another transfer transaction process, assigned with the default priority level, runs recursively every two seconds until the remaining amount in the account is not enough for another transfer.

$$\begin{aligned}
& \text{transfer_process} \leftarrow \text{transfer_queue}(Fee1, Client, Broker) \\
& \quad | \text{transfer_urgent}(Fee2, Client, Urgent_Acct) \\
& \quad : \text{timeElapsed}(100) \rightarrow \text{tirggerEvent}(\text{transfer_urgent} \\
& \quad \quad (Fee2, Client, Urgent_Acct)) \\
& \quad \wedge \text{priority}(\text{transfer_urgent} \\
& \quad \quad (Fee2, Client, Urgent_Acct), 2) \\
& \text{transfer_queue}(Amt, Acct1, Acct2) \leftarrow \text{balance}(Acct1, Bal)
\end{aligned}$$

$$\begin{aligned}
& \otimes (Bal \geq Amt) \\
& \otimes transfer(Amt, Acct1, Acct2) \\
& \otimes delay(2) \\
& \otimes transfer_queue(Amt, Acct1, Acct2)
\end{aligned}$$

$$transfer_queue(Amt, Acct1, Acct2) \leftarrow monitor('transfer\ ends')$$

$$transfer_urgent(Amt, Acct1, Acct2) \leftarrow transfer(Amt, Acct1, Acct2)$$

The first constraint concurrent Horn rule specifies the simulation process. It consists of two predicates, $transfer_queue(Fee1, Client, Broker)$ and $transfer_urgent(Fee2, Client, Urgent_Acct)$, which runs concurrently. The constraint formula specifies a timing event for $transfer_urgent$, which has priority in this concurrent execution and is scheduled to take place 100 seconds after the simulation starts. The transaction $transfer_queue$ implements a long-run series of bank account transfer transactions, in which an amount $Fee1$ is transferred from one account $Client$ to another account $Broker$, and the transfer process is called recursively every two seconds until the balance of $Client$ is smaller than the amount $Fee1$. In this example, we adjust the initial amount of the account $Client$ and the transferred amount $Fee1$ to ensure the transaction $transfer_urgent$ is triggered during the execution of the transaction $transfer_queue$.

The result shows that the specified prioritized event takes place right at 100 seconds, and the transaction $transfer_urgent$ is triggered and then interrupts the normal transfer process by becoming the *hottest component* of the concurrent conjunction. As expected, the amount $Fee2$ is thus transferred from the $Client$'s account to the urgent account $Urgent_Acct$ with priority. Only after this prioritized transfer transaction is completed, the interrupted normal transfer process can then resume its normal execution.

□

6.2 A Simplified Elevator Logic Control Application

Besides the area of database transaction applications, *TP-CTR* can also be used to simulate real-time control systems. Working with priority constraints, a *timeElapsed* timing-event-based constraint can create an *interrupt* functionality, which is scheduled to take place at specified time. The *interrupt* is one key feature of a real-time control system. Example 10 presents how to use *TP-CTR* to simulate the *interrupt* in an elevator logic control system.

Example 10 (Simulating a simplified elevator controller) For the sake of understanding, below we present the assumptions for a simplified elevator controller model and basic control logic rules:

- The elevator is for a ten-floor building with one stop button on each floor. When the button is pushed, a stop request is sent to the controller.
- The elevator cage takes about 3 seconds to go up or down one floor.
- Every floor has a location sensor to indicate the elevator's position.
- In any exceptional case, the elevator should take *stop* action immediately.

Based on these assumptions, the program below specifies the simplified control logic for the elevator controller:

```

simulate ← moving_control(stop)
  | user_request | accident
  : timeElapsed(31) → triggerEvent(accident)
  ∧ priority(accident, 2)
user_request ← req(5) | req(8) | req(2)
  : timeElapsed(10) → triggerEvent(req(5))
  ∧ timeElapsed(30) → triggerEvent(req(8))
  ∧ timeElapsed(50) → triggerEvent(req(2))
req(Level) ← ins(stop_req(Level))

```



```

moving_control(Status) ← Status = stop
    ⊗ stop_req(Level)
    ⊗ moveto(Level)
    ⊗ del(stop_req(Level))
    ⊗ moving_control(stop)
moving_control(Status) ← moving_control(stop)
moveto(Level) ← accSignal
    ⊗ stop_handle
moveto(Level) ← not accSignal
    ⊗ getLocation(X)
    ⊗ X <> Level
    ⊗ moveto(Level)
moveto(Level) ← not accSignal
    ⊗ getLocation(X)
    ⊗ X = Level
    ⊗ stop_handle
accident ← ins(accSignal)

```

where basic atomic formulas and built-in predicates in the specification have the following definition:

- The atomic formula *stop_req(Level)* is used to store the stop requests to be answered.
- *stop_handle* is a built-in predicate to stop the elevator.
- *getLocation(X)* is a built-in predicate to retrieve the current elevator location.
- The atom *accSignal* is used to denote an urgent exceptional event.

In the program, the first constraint concurrent Horn rule specifies the entire simulation process. It consists of three concurrent transactions, *moving_control*, *user_request* and *accident*. The constraint formula specifies an exceptional event: an *accident* which is scheduled to take place at 31 seconds after the simulation starts. The predicates *moving_control* and *user_request* implement the normal elevator control logic. The predicate *user_request* simulates use case scenario: three stop requests are scheduled to take

place at specified times. The predicate *moving_control* is used to control the elevator: moving to a floor according to the stop requests. The predicate *moveto(Level)* performs the elevator moving action, and monitors any exceptional accident as well.

When the specified accident takes place at the 31 seconds time instant, because of its priority, the predicate *accident* works as an interrupt and notifies the system at once by inserting the atom *accSignal* into the database. In our case, the elevator is still moving from 5th floor to 8th floor at that time. The predicate *moveto(Level)* receives this signal, then correspondingly executes the predicate *stop_handle* to stop the elevator at once. \square

This example demonstrates that *TP-CTR* can be used to simulate some intrinsically real-time phenomena, e.g., the interrupt effect and timing relations. By elaborately combining priority constraints and timing constraints in different ways, we can verify different elevator logic control cases with the *TP-CTR* program.

Chapter 7

Conclusions and Future Works

IN the previous chapters, we presented an extension of Concurrent Transaction Logic, a formalism originally designed to handle state changes in deductive databases. We extended *CTR* by defining a logic framework in which the user can not only specify concurrent transaction processes as logic programs, but also define priority and timing constraints on these processes. This increases the flexibility and power of the language. Users are no longer forced to schedule the transaction order only by the sequential conjunction and concurrent conjunction. The execution order of concurrent transactions can also be specified or changed by extra ways: assigning timing constraints and priority constraints to these transactions. The interrupt effect created by combining the event-driven feature of the timing constraint *timeElapsed* with priority constraints schedules transactions in a more flexible and powerful way, and thus opens the logic to a new range of advanced application, e.g., real-time domain application, long-run time-related application, digital circuit design with clock-driven and continuous-simulation of the designed digital circuit, etc.

To allow the formal execution and thus the simulation of such programs, we introduced an inference system that is able to handle priority constraints with the underlying simplified Rate-Monotonic scheduling algorithm. The formal execution of programs works as SLD-style refutation mechanism. By this way, our approach provides a high-level framework for specifying and executing transaction logic programs involving priority and

timing properties. Also, we have shown how logic programming techniques can be used to implement real-time database application domains.

To make theory meet practice, we have implemented our *TP-CTR* prototype in XSB Prolog and have tested the examples presented in this thesis.

There are still some issues left open, especially real-time issues. Below we elaborate on them.

- *Sophisticated Timing Constraints.* The timing properties in our timing constraint are quite straight-forward at present stage. It only handles the timing property of a transaction *when to be triggered*, and a delay relation between two transactions. However, in reality, timing properties of a real-time system are much diverse, such as described in [4, 5, 6]. Some of them relate to the *deadline* of a transaction as well as the initial time instant. Adopting these ideas into the *TP-CTR* timing constraint will be valuable to expand its real-time application domain.
- *Underlying scheduling algorithm.* The current *TP-CTR* possible prototype uses a simplified Rate Monotonic Scheduling algorithm to schedule concurrent transactions. In real-time domain, this algorithm is straight-forward and may be inefficient in some cases. It should be possible to adopt a more sophisticated scheduling algorithms, to improve the performance of the *TP-CTR* prototype in real-time domains.

Appendix A

TP-CTR Tutorial

A.1 *TP-CTR* Prototype File System

The *TP-CTR* prototype in this thesis is extended on the base of Concurrent Transaction Logic with Recovery prototype available in [17]. It consists of the following models:

- *ctr.P* - the basic *TP-CTR* interpreter
- *parser.P* - a parser for *TP-CTR* rules
- *updates.P* - the code for back-trackable updates
- *load.P* - startup routine that loads the prototype modules
- *upload.P* - a module including rules to load a *TP-CTR* transaction-base and a database into *XSB* system.
- *timer.P* - including *TP-CTR* system rules regarding timing constraints

A.2 Getting Started

The *TP-CTR* prototype was implemented in *XSB* Prolog, and consists of the modules introduced above. To run the implementation, at query prompt, the *load* module must first be consulted, then invoke the predicate *ctr_init*. The predicate *ctr_init* is defined

in the module to load all necessary files into XSB Prolog and initialize the necessary parameters. Below we illustrate how the user interacts with the system.

```
$xsb -i
XSB Version 2.5 (Okocim) of March 11, 2002
[i686-pc-linux-gnu; mode: optimal; engine: slg-wam; gc: indirection; scheduling: local]
| ? - [load].
[load loaded]
yes
| ? - ctr_init.
[ctr loaded] - basic TP-CTR interpreter
[updates loaded] - code for back-trackable updates
[parser loaded] - parser for prototype rules
[upload loaded] - compilers for loading and translating rules
[timer loaded] - embedded timing constraints
[scrtutl loaded] - necessary XSB system module for system timing predicates
yes
```

A.3 *TP-CTR* Prototype Commands and Program Files

Like *CTR* programs, *TP-CTR* programs are also stored in the transaction-base and database with filename extensions *.ctr* and *.db*, respectively. Although in some cases the database file maybe empty, it is still needed in our prototype.

A.3.1 Compiling Commands in *TP-CTR*

The *TP-CTR* programs are compiled with the following three commands:

- *ctr_comp(program_name)*. - to compile both the transaction-base and database;
- *comp_trans(program_name)*. - to compile the transaction-base only;

comp_db(program_name). - to compile the database only.

where *program_name* is the filename of the *TP-CTR* program without extension. After compiled, a *TP-CTR* transaction-base file generates two files:

program_name.ctr.temp

program_name.ctr.o

where *program_name.ctr.temp* is an intermediate compiled file, in which the timing constraints and priority constraints have been translated into a recognizable *CTR* program. The *program_name.ctr.o* file is the final transaction object file via the underlying *TP-CTR* compiler.

Like the *CTR* database object file, a *TP-CTR* compiler creates a *TP-CTR* database object file with the name:

program_name.db.o

where the extension *.db.o* denotes the file to be a database object file.

A.3.2 Execution Command

To execute a transaction in the *TP-CTR* program after compilation, the user uses the following command:

execute(transaction_name).

where *transaction_name* is the head of one of the rules in the transaction-base file.

A.4 The *TP-CTR* Prototype Syntax

A.4.1 Transaction Rules

The constraint concurrent Horn rules consists of three parts: a rule head, a rule body, and a constraint body.

Head :- Body : Constraint Body.

or

Head :- *Body*.

or

Head.

where

- *Head* is an atomic formula;
- *Body* is a sequence of transaction actions (or queries) connected by any of the following conjunction operators:
 - sequential conjunction(*),
 - concurrent conjunction(#),
 - isolation(*o*).
- *Constraint Body* is a sequence of the following constraints connected by conjunction(\wedge):
 - timing constraint *timeElapsed*(*X*) \leftarrow *triggerEvent*(*Transaction_name*),
 - priority constraint *priority*(*Transaction_name*, *Priority_Level*)

A.4.2 Database Rules

A database is any Prolog rule or atom in the forms of format:

Head :- *Body*.

or

Head.

.5 Built-in $\mathcal{TP}\text{-CTR}$ Predicates

The prototype has some built-in predicates, which are used in the transaction-base or database. These predicates should not be re-defined by the user.

A.5.1 Database Declaration

Database atoms represent tuples in a relation. Like table in relational database, a relation should be declared before it can be accessed. The statement below declares a database relation *Name* with arity *NArgs*:

updatable Name/Nargs.

By default the following basic system database declarations and tuples are always declared for any $\mathcal{TP}\text{-CTR}$ program:

updatable event/2 - timing driven event list
updatable start_time/1 - start time instant of an execution
updatable priorityList/1 - a current executed transaction priority list in descending order
priorityList([]) - the initial empty list

A.5.2 Built-in Predicates in Transaction-base

For the declared database atoms, there are three built-in predicates for accessing them:

- *db(p(x))* and *empty(p(x))* - query atom *p(x)* in the database
- *ins(p(x))* - inserts atom *p(x)* into the database
- *del(p(x))* - deletes atom *p(x)* from the database

At present stage, there are the following built-in constraint predicates:

- *timeElapsed(X)* - a timing constraint based on an absolute timing instant
- *delay(X)* - a timing constraint based on a relative timing instant
- *priority(Transaction_name, priority_level)* - priority constraint

Besides, in the *TP-CTR* prototype we have kept the useful *CTR monitor* command below, which can be used to trace a transaction execution.

monitor(Task).

where *Task* is a name or a term. Note that *monitor(Task)* does not execute *Task*, it only displays messages. When the monitor executes, it displays one of the messages:

- *Completed Task* - when *monitor(Task)* is executed.
- *undoing Task* - when *monitor(Task)* is rolled back.

A.6 Programming Examples

The following examples illustrate some simple *TP-CTR* programs and their execution. They are executed in XSB Prolog, with the loaded *TP-CTR* prototype.

Example A.1 Concurrent Processes with Priority Constraints

The following program executes three concurrent non-interacting transactions. The three transactions have different priority levels and has two tasks respective. The program name is *interleave*. Below are the contents of the database file and transaction file.

Database file *interleave.db*:

```

updatable event/2.
updatable priorityList/1.
priorityList([]).

```

Transaction Base file *interleave.ctr*:

```

parent :- taskA#taskB#taskC : priority(taskA,2) ∧
                                priority(taskB,2) ∧
                                priority(taskC,1).

taskA :- monitor(taskA1) * monitor(taskA2).
taskB :- monitor(taskB1) * monitor(taskB2).
taskC :- monitor(taskC1) * monitor(taskC2).

```

Compiling the program and executing a transaction:

```

| ?- ctr_comp(interleave).      - compile the program
yes

| ?- execute(parent).          - execute the parent goal
completed taskA1               - complete the first task of taskA
completed taskB1               - complete the first task of taskB
completed taskA2               - complete the second task of taskA
completed taskB2               - complete the second task of taskB
completed taskC1               - complete the first task of taskC
completed taskC2               - complete the second task of taskC
yes

```

Example A.2 Interrupt in Concurrent Processes

This example shows how the interrupt is created by combining a timing constraint and a priority constraint in a long-run program. The *initial* process is timed to start at 1 sec after the execution start, and the three concurrent sub-processes of the process *parent* have the same priority level. The subprocess of the *initial* process, *interrupt*, has higher priority level in the concurrent execution.

Database file *interrupt.db*:

```

updatable event/2.
updatable start_time/1.
updatable a/1.

```

```

updatable priorityList/1.
priorityList([]).
finished :- not(unfinished).
unfinished :- a(X).
a(100). a(99). ... a(1). a(0).

```

Transaction Base file *interrupt.ctr*

```

parent :- initTime * (long_run_transaction#initial#short_transaction)
        : timeElapsed(1) -> triggerEvent(initial)
        ^priority(long_run_transaction, 2)
        ^priority(initial, 2)
        ^priority(short_transaction, 2).

long_run_transaction :- o(del(a(X)) * monitor(X)) * long_run_transaction.
long_run_transaction :- finished.

initial :- monitor(interrupt_init_task) * interrupt : priority(d, 5).
interrupt :- monitor(interrupt_task_1) * monitor(interrupt_task_2).
short_transaction :- monitor(short_transaction_task_1)
                    *monitor(short_transaction_task_2)
                    *monitor(short_transaction_task_3)
                    *monitor(short_transaction_task_4).

```

Compiling the program and executing a transaction:

```

| ?- ctr_comp(interrupt).           - compile the program
yes
| ?- execute(parent).              - execute the parent goal
completed 100                       - shows one round of long_run_transaction is
                                   completed
completed short_transaction_task_1  - shows one round of short_transaction is
completed short_transaction_task_2  completed

```

completed 99
 completed short_transaction_task_3
 completed short_transaction_task_4
 completed 98
 .
 .
 .
 completed 75
 completed interrupt_init_task - *shows the initial task of the interrupt is*
 completed 74 *completed*
 completed interrupt_task_1 - *shows one task of the interrupt is completed*
 completed interrupt_task_2
 completed 73
 .
 .
 .
 completed 0
 yes

Bibliography

- [1] Anthony J. Bonner and Michael Kifer, "Concurrency and Communication in Transaction Logic", *Proceedings of the Joint International Conference and Symposium on Logic Programming (JICSLP)*, Sep. 2-6 1996, Pages 142-156. MIT Press.
- [2] Anthony J. Bonner, "Concurrency Transaction Logic", Presented at the *Dagstuhl Seminar on Transaction Workflows*, July, 15-19 1996, International Conference and Research Center for Computer Science, Schloss Dagstuhl, Wadern, Germany.
- [3] Anthony J. Bonner and Michael 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>.
- [4] P. Bellini, R. Mattolini, and P. Nesi, "Temporal Logics for real-Time System Specification," *ACM Computing Surveys*, Vol.32, No.1, March 2000.
- [5] B. Dasarathy, "Timing Constraints of Real-Time Systems: Constructs for Expressing Them, Methods of Validating Them", *IEEE Trans. Software Engineering*, Vol. SE-11, No. 1, January 1985, PP.80-86.
- [6] Horng-Yuan Chen and Jeffrey J.P. Tsai and Yaodong Bi, "An Event-Based Real-Time Logic to Specify the Behavior and Timing Properties of Real-Time Systems," *IEEE*, 0-8186-2300-4/91 (c)1991.

- [7] Kristian Torp, Christian S. Jensen, Richard T. Snodgrass, "Effective timestamping in databases," *The VLDB Journal*, (c)Spring-Verlag 2000.
- [8] Ögur Ulusoy, Geneva G. Belford, "Concurrency Control in Real-Time Database Systems," *ACM*, 089791-472-4/92/0002/0181, 1992.
- [9] Rajendran M. Sivasankaran, John A. Stankovic, Don Towsley, Bhaskar Purimetla, Krithi Ramamritham, "Priority assignment in real-time active databases," *The VLDB Journal* , (c) Spring-Verlag 1996.
- [10] Wayne Wolf, "Computer as Components: Principles of Embedded Computing System Design", *Morgan Kaufman Publishers*, 2001, ISBN 1-55860-541-X.
- [11] Neugass H., Espin G., Nunoe H., Thomas R., Wilner D., "VxWorks: an interactive development environment and real-time kernel for Gmicro," *TRON Project Symposium*, 1991. Proceedings., Eighth , 21-27 Nov. 1991 Page(s): 196 -207
- [12] Terrasa, A.; Garcia-Fornes, A , "Real-time synchronization between hard and soft tasks in RT-Linux" *Real-Time Computing Systems and Applications*, 1999. RTCSA '99. Sixth International Conference on , 13-15 Dec. 1999 Page(s): 434 -441
- [13] Amalia F. Sleghel, "An Optimizing Interpreter for Concurrent Transaction Logic" *Thesis for degree of Master of Science in University of Toronto*, 2000.
- [14] Jiwen Ge, Marcus Vinicius Santos, "Using logic programming techniques to handle timing and priority constraints in real-time systems" *Seventh World Conference On Integrated Design & Process Technology IDPT Vol.2* 2003, ISSN No. 1090-9389, Page(s): 800-806
- [15] Konstantinos Sagonas, Terrance Swift, David S. Warren, Juliana Freire, Prasad Rao, Baoqiu Cui, Ernie Johnson, "The XSB System Version 2.5 Volume I: Programmer's Manual" *SUNY at Stony Brook*, 2002.

- [16] Konstantinos Sagonas, Terrance Swift, David S. Warren, Juliana Freire, Prasad Rao, Baoqiu Cui, Ernie Johnson, “The XSB System Version 2.5 Volume II: Programmer’s Manual” *SUNY at Stony Brook*, 2002.
- [17] Anthony Bonner, “*CTR* prototype with recovery”, <ftp://ftp.cs.toronto.edu/cs/ftp/pub/bonner/papers/transaction.logic/prototype/recoverable/>.