

# Coordination of Systems with Real-Time Properties in Manifold

George A. Papadopoulos<sup>†</sup>

*Department of Computer Science  
University of Cyprus  
75 Kallipoleos Str, P.O.B. 537  
CY-1678 Nicosia, Cyprus  
E-mail: george@turing.cs.ucy.ac.cy*

Farhad Arbab

*Department of Interactive Systems  
Centre for Mathematics and Computer Science (CWI)  
Kruislaan 413, 1098 SJ Amsterdam  
The Netherlands  
E-mail: farhad@cwi.nl*

## Abstract

*This paper combines work done in the areas of asynchronous timed computations and coordination models in order to derive a framework able to express real-time coordination without adhering to special architectures or real-time programming languages. In particular, it is shown how the coordination language MANIFOLD can be used to support coordinators with bounded response time. The only assumption made about the coordinated components is that they adhere to some constraints as imposed by timed asynchronous computational models derived for concurrent constraint languages. The model can be used as a basis for building complex software and hardware systems with the distinct advantage of enhancing reusability of existing components. Indeed, it is already being used in a project aiming to develop distributed multimedia applications.*

**Keywords:** Real-Time Coordination; Timed Asynchronous Languages; Reactive Systems; Software Development for Distributed and Parallel Processing Systems; Interoperability in Systems and Tools; Large-Scale Software System Integration.

## 1 Introduction

The concept of coordinating a number of activities, possibly created independently from each other, such that they can run concurrently in a parallel and/or distributed fashion has received wide attention and a number of coordination models have been developed ([3]). Most of the proposed coordination frameworks however are suited for environments where the sub-components comprising an application are conventional ones in the sense that they do not adhere to any real-time constraints. Those few that are addressing this issue of real-time coordination either rely on the ability of the underlying architecture apparatus to provide real-time support ([7]) and/or are confined to using a specific real-time language ([6]). However, all the issues pertaining to

conventional coordination frameworks are still valid in those exhibiting real-time properties. Even more, the use of a separate coordination formalism with real-time capabilities “gluing together” real-time (software and/or hardware) components helps in separating issues related to specifying the computational part of the components from those concerned specifically with their real-time behaviour ([7]).

In this paper we address the issue of real-time coordination but with a number of self imposed constraints which we feel, if satisfied, will render the proposed model suitable for a wide variety of applications. These constraints are:

- The coordination model should not rely on any specific architecture configuration supporting real-time response and should be implementable on a variety of systems including distributed ones.
- Language interoperability should not be sacrificed and the real-time framework should not be based on the use of specific language formalisms.

In addition, the proposed framework should be able to enjoy the properties of state-of-the-art coordination models proposed for conventional (non real-time) environments but also address issues peculiar to the presence of real-time behaviour such as the separation of the computational part of an application from that of real-time coordination.

We attempt to meet the above mentioned targets by combining work done in the area of coordination models with that done for asynchronous timed computations. Our starting points are the coordination event-driven language MANIFOLD ([1,2]) and the tcc computational model ([8]) developed for concurrent constraint languages. More to the point, we show how MANIFOLD coordinators can exhibit real-time behaviour in the sense that reaction to events can always be done in bounded time, even if not instantaneously. This combination yields a coordination environment exhibiting real-time properties without imposing any difficult to meet constraints on the implementation of the programming languages used or the underlying architecture.

## 2 The IWIM Model and the Language MANIFOLD

Most of the message passing models of communication can be classified under the generic title of TSR (Targeted-Send/Receive) in the sense that there is some asymmetry in

<sup>†</sup> This work was done while the author was visiting CWI as part of the ERCIM-HCM Fellowship Programme financed by the Commission of the European Community under contract no. ERBCHBGCT930350.

the way the sending and receiving of messages between processes is performed: it is usually the case that the sender is in general aware of the receiver(s) of its message(s) whereas a receiver does not care about the origin of a received message. Alternatively, the IWIM (Ideal Worker Ideal Manager) communication model aims at completely separating the computational part of a process from its communication part, thus encouraging a weak coupling between worker processes in the coordination environment.

In IWIM there are two different types of processes: *managers* (or *coordinators*) and *workers*. A manager is responsible for setting up and taking care of the communication needs of the group of worker processes it controls (non-exclusively). A worker, on the other hand, is completely unaware of who (if anyone) needs the results it computes or from where it itself receives the data to process. This suggests that a suitable (albeit by no means unique) combination of entities a coordination language based on IWIM should possess is the following:

- *Processes*. A process is a *black box* with well defined *ports* of connection through which it exchanges *units* of information with the rest of the world. A process can be either a manager (coordinator) process or a worker. A manager process is responsible for setting up and managing the computation performed by a group of workers. Note that worker processes can themselves be managers of subgroups of other processes and that more than one manager can coordinate a worker's activities as a member of different subgroups. The bottom line in this hierarchy is *atomic* processes which may in fact be written in any programming language.
- *Ports*. These are named openings in the boundary walls of a process through which units of information are exchanged using standard I/O type primitives analogous to read and write. Without loss of generality, we assume that each port is used for the exchange of information in only one direction: either into (*input* port) or out of (*output* port) a process. We use the notation  $p.i$  to refer to the port  $i$  of a process instance  $p$ .
- *Channels*. These are the means by which interconnections between the ports of processes are realised. A channel connects a (port of a) producer (process) to a (port of a) consumer process. We write  $p.o \rightarrow q.i$  to denote a channel connecting the port  $o$  of a producer process  $p$  to the port  $i$  of a consumer process  $q$ .
- *Events*. Independent of channels, there is also an event mechanism for information exchange. Events are *raised* by their sources in the environment (by means of executing the command `raise(e)`), yielding *event occurrences*. In principle, any process in the environment can pick up a broadcast event; in practice though, usually only a subset of the potential receivers is interested in an event occurrence. We say that these processes are *tuned in* to the sources of the events they receive. We write  $e.p$  to refer to the event  $e$  raised by a source  $p$ . Alternatively, a source  $p$  can *post* an event  $e$  locally, within its own environment (by means of executing the command `post(e)`). The occurrence of  $e$  will not be observed by any process other than  $p$  itself.

MANIFOLD is a coordination language which can be seen as a concrete version of the IWIM model just described where:

- Communication is asynchronous and raising and reacting to events (or signals) enforce no synchronisation between the processes involved.
- The separation between communication and computation, i.e. the distinction between workers and managers, is more strongly enforced.

Activity in a MANIFOLD configuration is *event driven*.

A coordinator process waits for the raising of some specific event (usually generated by the worker processes it coordinates) which trigger it to enter a certain *state* and perform some actions. These actions typically consist of setting up or breaking off connections of ports via channels. It then remains in that state until it observes the occurrence of some other event which causes the *preemption* of the current state in favour of a new one specified by the event. Once an event has been raised, its source generally continues with its activities while the event occurrence propagates through the environment independently and is observed (if at all) by the other processes according to each observer's own sense of priorities. The following trivial program shows the basic structure of a MANIFOLD process and should suffice for the time being. Space limitations preclude a fuller description of the language; nevertheless, in the following sections we take the opportunity to introduce more details about those features of the language relevant to this work.

```
manifold PrintUnits() import.
auto process print is PrintUnits.
```

```
manifold Main()
{
  begin: "Hello World!" -> print.
}
```

The program redirects (connects) the (default output port of the) character string (which in MANIFOLD is itself a process) to the default input port of the process `print` which is defined to be an instance of a predefined library process. `Main` has just one state and observes just one event, the predefined `begin` which is raised immediately upon activation of the process.

More information on MANIFOLD can be found in [1,2]. Note that the language has already been implemented on top of PVM and has been successfully ported on a number of platforms including Sun, Silicon Graphics and IBM SP1/2.

### 3 Real-Time Coordinators in MANIFOLD

The event driven mechanism of MANIFOLD is well suited to our purposes and provides a natural programming metaphor for expressing real-time behaviour. We recall that reacting to raised events is done asynchronously by the processes that observe them. We would like to retain this asynchronous behaviour, so that we do not have to impose extra constraints on the implementation of MANIFOLD or the underlying architecture, while providing bounded-time response. The timed asynchronous model proposed for concurrent constraint languages ([8]), as opposed to the perfect synchronous one ([4]) advocated by state-of-the-art real-time languages like ESTEREL, LUSTRE, or SIGNAL ([5]), provides the basic principles for achieving our goal.

The fundamental difference between timed asynchronous languages and ordinary asynchronous ones is that recursion (or iteration for that matter) is “guarded” in the sense that it can appear only in the next clock tick. Also, variables are treated as signals with a life span equal to a clock tick. These characteristics effectively guarantee that computation within a time instance is bounded and reaction to events happens in a very small amount of time (although not instantaneously).

Although MANIFOLD’s features were designed with other purposes in mind, we have found them to be suitable in expressing the above mentioned behaviour. In particular, a MANIFOLD configuration exhibiting real-time behaviour in the above mentioned sense consists of the following components:

- A MANIFOLD coordinator process (the clock) responsible for monitoring the status of the coordinated processes, detecting the end of the current time instance, and triggering the next one. The coordinator process is also responsible for detecting the end of the computation.
- A set of MANIFOLD coordinated processes, each one monitoring the execution of some group of atomic processes. Each such coordinated process performs a bounded amount of work between the ticks as dictated by the coordinator process (thus any loops in such a process “spread over” the next one or more ticks).
- A set of groups of atomic processes (i.e., processes written in some language other than MANIFOLD), each group being monitored by a coordinated process. In order for the whole configuration to exhibit real-time behaviour, these atomic processes must also produce results in bounded time. There are two approaches possible here: (i) enforce the constraint that there are no loops within these processes and instead, put these loops in their respective coordinated processes, or (ii) treat them as asynchronous parallel components that take an unbounded amount of time.

The overall configuration is a hierarchical one with the MANIFOLD coordinator process on the top, monitoring a number of MANIFOLD coordinated processes, themselves possibly monitoring groups of atomic (non MANIFOLD) processes. One can regard MANIFOLD as being the “host language” for writing the control structures of reactive systems, while most of the actual computation (data handling, interfaces with any embedded systems) are done in other more conventional languages, typically C. This fits nicely into the spirit of real-time coordination models as we perceive them and separates the real-time coordination requirements from the rest of the performed activities.

An application featuring timed asynchronous behaviour takes the general form:

```
< application> ::= <coordinator>
                  <coordinated>+
                  <atomic>+
```

The general behaviour of a coordinator (clock) process is shown, as a first approximation, below (note that the construct  $(A_1, \dots, A_n)$  denotes a block where all  $n$  activities will be executed concurrently; there is also a ‘;’ separator imposing sequentiality).

```
manifold Clock()
port in term_in, next_in.
```

```
port out term_out, next_out.
{
  event tick, next_phase, end_comp.

  begin: (<set up network of initial activ procs>;
         terminated(void)).

  next_phase: (raise(tick), terminated(void)).

  end_comp: (<perform clean up>; post(end)).
}
```

clock first sets up the initial network of coordinated processes. It then suspends waiting for either of the following two cases to become true (one way to achieve suspension in MANIFOLD is by waiting for the termination of the special process void which actually never terminates):

- The coordinated processes have completed execution within the current time instance and are waiting for the next clock tick. clock raises the appropriate event (or signal) and suspends again.
- The computation has terminated in which case clock terminates, possibly after performing some clean up.

Detecting the completion of both the current phase and the end of the computation is done in a distributed fashion, provided some constraints regarding the organisation and communication protocols between the participating coordinated processes are imposed. We elaborate further on the exact nature of the work done by clock once we describe the activities performed by a coordinated process.

The general behaviour of a coordinated process is as follows:

```
manifold Process()
port in term_in, next_in.
port out term_out, next_out.
{
  begin: (<raise event>;
         <wait until input event received>).

  input_event: (<perform data transfer p1->p2>,
               <generate new processes>;
               terminated(void)).

  tick.clock: (<perform further actions>;
              post(end)).
}
```

where clock is a global process instance of the manifold Clock defined as follows (for the sake of simplicity we will be using from now on the names clock and Clock interchangeably when describing their functionality):

auto process clock is Clock.

A typical behaviour of a timed asynchronous coordinated process, as understood in the MANIFOLD world, is to post some events, possibly wait until the presence of some event in the current time instance is detected and then react by producing some data transfer between a group of atomic processes that it itself coordinates (say from  $p_1$  to  $p_2$ ), post more events and/or generate further processes. Upon termination of its activities within the current time instance, the process suspends waiting for the next tick event from the coordinator (Clock) process, in which case it performs

more activities of similar nature, or alternatively, simply terminates within the current time instance.

Interestingly, the concept of activities within some time instance provides a basis for enhancing MANIFOLD with event patterns comprising *negative* operators detecting the *absence* of some event, something that otherwise would have no meaning.

We now present in more detail the way detection of the end of the current phase, as well as the whole computation, is achieved. Due to space limitations only the most essential parts of the MANIFOLD code are shown below. The techniques we are using are reminiscent of the ones usually encountered within the concurrent constraint programming community based on short circuits. We recall that the coordinator and each one of the coordinated processes have (among others) two pairs of ports: the `term_in/term_out` pair is used to detect termination of the whole computation whereas, the `next_in/next_out` pair is used to detect termination of the current clock phase. Upon commencing the computation, the `Clock` process sets up a configuration where all the processes involved are connected by means of short-circuits. This is achieved by means of the following MANIFOLD constructs set up between a `Clock` process and `n P` coordinated processes:

```
(C.next_out->P1.next_in,...,Pn.next_out->C.next_in)
(C.term_out->P1.term_in,...,Pn.term_out->C.term_in)
```

Any process wishing to further generate other processes is also responsible for setting up the appropriate port connections between these newly created processes. Detecting termination of the whole computation is done as follows: a process `P` wishing to terminate, first redirects the stream connections of its input and output `term` ports so that its left process actually bypasses `P`. It also sends a message down the `term_in` port of its right process. If `P`'s right process is another coordinated process the message is ignored; however, if it happens to be the `Clock` controller, the latter sends another message down its `term_out` port to its left process. It then suspends waiting for either the message to reappear on its `term_in` port (in which case no other coordinated process is active and computation has terminated) or a notification (by means of raising the `got_token` signal) from its left coordinated process (which signifies that there are still active coordinated processes in the network). The basic MANIFOLD code realising the above scenario for the benefit of the `Clock` controller is shown below.

```
// Clock code

begin: (guard(term_in, transport, check_term1),
        terminated(void)).

check_term1: ("token" -> term_out,
              post(check_term2)).

check_term2: (guard(term_in, transport, end),
              terminated(void)).

got_token: post(begin).
```

A `guard` process is set up to monitor the activity in the `term_in` port. Upon receiving some input in this port, `guard` posts the event `check_term1`, thus activating `Clock` which then sends `token` down its `term_out` port

waiting to get either a `got_token` message from some coordinated process or have `token` reappear again. The related code for a coordinated process is as follows:

```
// Coordinated process code

begin: guard(term_in, transport, check_term).

check_term: (<term_in->void,
             if data in port is "token"
             raise(got_token)>).
```

Detecting the end of the current time instance is a bit more complicated. Essentially, quiescence, as opposed to termination, is a state where there are still some processes suspended waiting for events that cannot be generated within the current time instance. We have developed two methods that can detect quiescent points in the computation. In the first scheme, all coordinated processes are connected to a `Clock` process by means of reconnectable streams between designated ports. A process that has terminated its activities within the current time instance breaks the stream connection with `Clock` whereas a process wishing to suspend waiting for an event `e` first raises the complementary event `i_want_e`. Provided that processes wishing to suspend but are also able to raise events for the benefit of other processes, do so before suspending, quiescence is the point where the set of processes still connected to `Clock` is the same as the set of processes that have raised `i_want_e` events. The advantage of this scheme is that processes can raise events arbitrarily without any concern about them being received by some other process. The disadvantage, however, is that it is essentially a centralised scheme, also needing a good deal of run-time work in order to keep track of the posted events.

An alternative approach requiring less work that is also distributable is a modification of the protocol used to detect termination of the computation. More to the point, a process wishing to suspend waiting for an event performs the same activities as if it were about to terminate (i.e. have itself bypassed in the port connections chain) but this time using the `next` input/output ports. A process wishing to raise an event before suspending (or terminating for that matter) does so, but waits for a confirmation that the event has been received before proceeding to suspend (or terminate). A process being activated because of the arrival of an event, adds itself back into the `next` ports chain. Quiescence now is the point where the `Clock` detects, as before, that its `next_out` port is effectively connected to its own `next_in` port, signifying that no event producer processes are active within the current time instance. Note that unlike the case for detecting termination, here the short circuit chain can shrink and expand arbitrarily. Nevertheless, it will eventually shrink completely provided that the following constraints on raising events are imposed:

- Every raised event must be received within the current time instance so that no events remain in transit. An event multicast to more than one process must be acknowledged by all receiver processes whose number must be known to the process raising the event; this latter process will then wait for a confirmation from all the receiver processes before proceeding any further.
- A process must perform its activities (where applicable) in the following order: 1) raise any events, 2) spawn any new processes and set up the `next` and `term` port

connections appropriately, 3) suspend waiting for confirmation of raised events, 4) repeat the procedure.

The code for the Clock controller is very similar to the one managing the term ports, with the major difference that upon detecting the end of the current phase Clock raises the event tick, thus reactivating those coordinated processes waiting to start the activities of the next time instance.

```
// Clock code

begin: (guard(next_in, transport, check_term1),
        terminated(void)).

check_term1: ("token" -> next_out,
              post(check_term2)).

check_term2: (guard(next_in, transport, raise_t),
              terminated(void)).

raise_t: raise(tick).

got_token: post(begin).
```

The code for a coordinated process is as follows:

```
// Coordinated process code

some_state: {begin: (raise(e),
                    <possibly spawn other processes>;
                    terminated(void)).
             i_got_e:
             }
<continue>
```

We have elaborated extensively on how the detection of termination or quiescence is achieved, because it is an important issue in realising timed asynchronous models.

## 4 Examples

In this section, we present some examples of real-time MANIFOLD coordinators where, for the sake of both clarity and brevity, we only show the real-time part of the code; however, the management part in each process, as described in the previous section, should be considered as being present in the code.

### 4.1 Real-Time Constructs

The framework presented in the previous section can be used to model the functionality of standard constructs found in the state-of-the-art real-time languages such as ESTEREL. A small collection of them is shown below.

```
manifold Whenever_Do(event e, manifold P)
{
  begin | wait: terminated(void).

  e: (P, terminated(void)).

  tick.clock: {ignore *.
               begin: post(wait).}.
}

manifold Always(manifold P)
{
  begin: (P, terminated(void)).

  tick.clock: post(begin).
}
```

```
manifold Do_Watching(manifold P, event e)
{
  begin: (P, terminated(void)).

  e: {begin: terminated(void).
      tick.clock: raise(abort).}.

  tick.clock: terminated(void).
}
```

Note that ignore \* clears the event memory of the manifold executing this command. By using ignore a “recursive” manifold can go to the next time instance without carrying with it events raised in the previous time instance. Note also that ‘|’ is the “or” operator (e1 | e2 is satisfied if either of the events e1 or e2 have been raised or posted).

### 4.2 Coordinating Real-Time Components

The next example illustrates the use of real-time manifolds in coordinating the activities of concurrently executing multimedia objects. This is a promising area for real-time coordination models because they provide a platform allowing different temporal coordination patterns to be specified (and tested) between sets of active objects encapsulating media processing activities ([6]).

```
manifold Decompress(process frame,
                    port in tolerance)
{
  process alarm is Alarm(tolerance).

  begin: (activate alarm, frame->decompressor,
          terminated(void)).

  tick.clock: {ignore *.
               begin: terminated(void).
               decomp.decompressor:decompressor->display.
               timeup.alarm: (disp_prev_frame->display,
                              terminated(void).)}.
}
```

where decompressor, disp\_prev\_frame and display are process instances of respective manifolds. The above is a simplified version of some manifold responsible for decompressing and displaying a video frame. Initially, the manifold redirects the (compressed) frame to a decompressor process, activates an alarm process and waits for the next time instance. It then checks whether either the decompressor has finished executing, in which case the frame is sent to a display process, or the alarm has raised an event signifying that the time during which the decompression should have taken place is over, in which case the manifold displays again the previous frame. Note that the detection of either of the two events is made at a pace dictated by the application’s Clock manifold and not by, say, the alarm process (which may well be a true real-time device). Thus, response by Decompress is not done instantaneously from alarm’s point of view but, nevertheless, it is within a small period of time.

### 4.3 Programming Synchronous Algorithms - Timed Fibonacci Sequence

The final example is a timed version of the Fibonacci sequence, adapted from the one presented in [8]. It is by no means the most efficient timed version one can write in our

framework but it has some interesting features like spawning dynamically new processes over a number of clock ticks. Recursion is unfolded and “spread over” two consecutive clock ticks and each call to a `Fib` process lasts 3 clock ticks.

```

auto process v0 is variable(0).
auto process v1 is variable(1).
auto process print is PrintUnits.
auto process sum(port in i1,i2) is Add.

event v.

manifold Fib0()
{
  begin: (raise(v), v0->(->print, ->output)).
}

manifold Fib1()
{
  begin: terminated(void).

  tick.clock: (raise(v), v1->(->print, ->output)).
}

manifold FibN()
{
  begin: terminated(void).

  v.*p: {process x is variable. // work at time T
  begin: (p->sum.i1, x->sum.i2, sum->y);
  terminated(void)}.

  tick.clock: // work at time T+1
  {begin: (FibN, terminated(void))}.

  v.*p: {process y is variable.
  begin: ((p->sum.i1, x->sum.i2, sum->y);
  terminated(void)).

  tick.clock: // work at time T+2
  (FibN, raise(v),
  y->(->print, ->output))}}.
}

```

At every time instance a new `FibN` process is created. It then waits until it receives the event `v` signifying that some other process has created the first of the two numbers needed to compute the next Fibonacci number (note here that the construct `event . *p` binds `p` to the id of the process raising event). After storing locally the number, `FibN` waits for the next time instance and then gets, in the same fashion, the second number, spawns a copy of itself and computes the next Fibonacci number. In the following and final (as far as it is concerned) time instance, `FibN` passes the result to its output port as well as to the printing process, raises the event `v` for the benefit of the other `FibN` processes waiting for the result, spawns another copy of itself and terminates.

Note that in the spirit of the MANIFOLD model, processes have a minimal awareness about the activities performed in their environment; an incarnation of `FibN` for instance, passes its result to its output port without any concern as to whom will get it (aside from the printing process).

## 5 Conclusions and Further Work

We combined work done in developing timed asynchronous models for concurrent constraint languages ([8]) with that done for developing coordination models. We

showed that if the latter models are enhanced with the properties imposed by the former ones we can derive coordination models with real-time capabilities which do not have to resort to special real-time languages ([5]) or impose constraints on the underlying architecture or operating system to meet the demands of the more traditional synchrony hypothesis based real-time models ([4]). We believe that the model is not only easier to implement (especially when distributed environments are being considered) but can also be used for a variety of application areas where *soft* real-time deadlines suffice.

Compared to, say, [6] or [7] our approach is essentially bottom up: we paid emphasis more on *how* such a framework can be developed within the functionality offered by a particular coordination language (MANIFOLD) and less on what should be the exact nature of the top-level end-user programming constructs or how these can be used to separate the real-time part of the computation from the rest of the activities performed (although we got a glimpse of this in section 4).

We are currently developing suitable high-level abstractions able to express real-time synchronisation constraints as defined in this work. We are also working on a more efficient implementation of the apparatus required to support the real-time constraints. Finally, we are using this model in a project aiming to derive a framework for distributed multimedia applications.

## References

- [1] F. Arbab, “The IWIM Model for Coordination of Concurrent Activities”, *Coordination '96*, Cesena, Italy, 15-17 April, 1996, LNCS 1061, Springer Verlag, pp. 34-56.
- [2] F. Arbab, I. Herman and P. Spilling, “An Overview of Manifold and its Implementation”, *Concurrency: Practice and Experience* 5(1), Feb. 1993, pp. 23-70.
- [3] N. Carriero and D. Gelernter, “Coordination Languages and their Significance”, *Communications of the ACM* 35(2), Feb. 1992, pp. 97-107.
- [4] N. Hallbwachs, *Synchronous Programming of Reactive Systems*, Kluwer Academic Publis., 1993.
- [5] IEEE Inc, “Another Look at Real-Time Programming”, special section of the *Proceedings of the IEEE* 79(9), Sept. 1991.
- [6] M. Papathomas, G. S. Blair and G. Coulson, “A Model for Active Object Coordination and its Use for Distributed Multimedia Applications”, *Object-Based Models and Languages for Concurrent Systems*, Bologna, Italy, July 5, 1994, LNCS 924, Springer Verlag, pp. 162-175.
- [7] S. Ren and G. A. Agha, “RTsynchronizer: Language Support for Real-Time Specifications in Distributed Systems”, *ACM SIGPLAN Workshop on Languages, Compilers and Tools for Real-Time Systems*, La Jolla, California, June 21-22 1995.
- [8] V. Saraswat, R. Jagadeesan and V. Gupta, “Programming in Timed Concurrent Constraint Programming”, in *Constraint Programming*, NATO Advanced Science Institution Series, Series F: Computer and System Sciences, LNCS, Springer Verlag, 1994.