

# **Interval Scripts: a Programming Paradigm for Interactive Environments and Agents**

**Claudio Pinhanez**

IBM Research, T.J. Watson  
Emerging Interactive Spaces Group  
P.O. Box 218  
Yorktown Heights, NY 10598 – USA  
pinhanez@us.ibm.com

**Aaron Bobick**

Georgia Institute of Technology  
GVU Center  
801 Atlantic Dr.  
Atlanta, GA 30332 – USA  
bobick@cc.gatech.edu

# Interval Scripts: a Programming Paradigm for Interactive Environments and Agents

## Abstract

In this paper we present *interval scripts*, a new paradigm for the high-level programming of interactive environments and computer characters. In this paradigm, actions and states of the users and the system computational agents are associated with temporal intervals. Programming is accomplished by establishing temporal relationships as constraints between the intervals. Unlike previous temporal constraint-based programming languages, we employ a strong temporal algebra based in Allen's interval algebra with the ability to express mutually exclusive intervals and to define complex temporal structures. To avoid the typical computational complexity of strong temporal algebras we propose a method — *PNF propagation* — that projects the network implicit in the program into a simpler, 3-valued (*past, now, future*) network where constraint propagation can be conservatively approximated in linear time. The interval scripts paradigm is the basis of *ISL*, or *Interval Scripts Language*, that was used to build three large-scale, computer-vision-based interactive installations with complex interactive dramatic structures. The success in implementing these projects provides evidence that the interval scripts paradigm is a powerful and expressive programming method for interactive environments.

**Keywords:** interactive spaces; programming paradigms; ubiquitous computing; system architecture; temporal reasoning.

## 1 INTRODUCTION

Recent years have seen an increasing interest in the development of interactive environments where users interact with devices, data, images, computational agents, and computer characters. Such environments can be virtual spaces as it is the case of virtual reality systems (VR), or real physical spaces with embedded computing devices and control. Examples of the later include location-based entertainment spaces [1-4], interactive offices [5-7], classrooms [8], and homes [9, 10].

However, little attention has been paid to the special particularities imposed by interactive environments on programming methods. In particular, in highly immersive situations both the actions and states of the environment and its computational agents, as well as the actions of the users are not events, but take




Figure 1. Examples of difficulties when programming interactive environments.

time to be performed. This makes difficult the use of the traditional event-based programming paradigm commonly used in desktop interaction. Although the problem of system actions with non-zero duration has been the object of a great deal of research in the multimedia community, interactive environments extend significantly the problem since they include situations where the user's actions have also duration. Unlike desktop systems where the interaction occurs mostly through more clicks, in interactive environments the users talk, walk, make complex gestures, and interact with objects, computational agents, computer characters, and other users. However, the programming systems for interactive environments proposed so far [11, 12] use simple paradigms that do not provide adequate ways to handle this kind of long-duration interactivity.

Let us examine an example to characterize better the issues involved in programming an interactive environment. Consider a simple interactive environment in a museum context, composed of three different areas A, B, and C, where pre-recorded tapes have to be played whenever a user enters the corresponding area. Let's also assume that the system is able to differentiate among users and that the tapes cannot be played simultaneously because the overlapping makes them unintelligible. Figure 1 shows some typical difficulties to the programming of such system. In instant  $t_0$ , user 1 enters space A, triggering the playing of tape A. Notice that there is a *delay* in the start of the tape. But most important, if we do not want the tape to play again while user 1 is in the area A (for instance, after instant  $t_3$ ), then the control system cannot be purely reactive, i.e., based on a simple event structure such as "if user is in area A, play tape A", because that event would be repeatedly activated by the presence of user 1 in area A. Here, the control program needs some way to represent and access notions of *continuous states* and *interaction history* in order to activate an action.

When user 2 enters space B in instant  $t_1$ , the system normally would start playing tape B. However, since tape A is playing, the start of tape has to be delayed, because the two actions, as described before, are *mutually exclusive*. However, since tape A lasts longer than the presence of user 2 in area B, the

playing of tape B has to be, ultimately, ignored. The moving of user 2 to area C, in instant  $t_2$ , however, does not initiate the playing of the tape C either. Only after tape A ends, at  $t_3$ , it is possible to start tape C. Instant  $t_5$  illustrates another common problem: although tape C is playing, there is no user in area C. This is a typical situation where it would be desirable to have ways to *stop an action abruptly*, in this case by triggering a fadeout or similar effect. Finally, when tape C ends on instant  $t_6$ , the system is faced with a choice between *equivalent but mutually exclusively actions*: it can either play tape A or tape B, but not both.

As we see in the above example, a programming paradigm for interactive has to address complex temporal structures which arise from the fact that users' and computer actions and states have duration. Although it is theoretically possible to represent such temporal structures using simple models such as state machines or event-loop systems, this approach requires the enumeration of all possible combinations of states or events, creating programs that are hard to write, debug, and maintain. Moreover, since such systems require that each combination of states to be represented as a single machine state or event, the addition of new actions or users often causes an exponential increase in the number of machine states or events. This is a direct consequence of having mutually exclusive parallel actions or states (see [13]).

This paper proposes a method to program interactive environments based on temporal constraints. Called *interval scripts*, this paradigm is based on the encapsulation of actions and states in temporal intervals constrained by the temporal primitives of Allen's algebra [14]. Our objective is to provide the programmer of an interactive environment with a tool that combines expressiveness and simplicity and, at the same time, produces control code that can be run in real time. Among the desired programming features, we include in interval scripts not only facilities for simple handling of time durations but also features such as the definition of mutually exclusive actions, the ability to impose restrictions on what can happen (negative scripting), easy encapsulation of actions, basic handling of interaction history, and automatic recovery from errors and unexpected situations.

To provide these programming facilities, an interval script contains descriptions of how to initiate and stop each action of the interactive environment or its agents and how to determine the occurrence of actions and states of the users, environment states, and the states of computational agents. The control of activation of an action is determined by verifying the current state of the action, comparing it with the desired state according to the constraints with other actions, and issuing commands to start and stop

action as needed. Also, as we describe in detail later, interval scripts de-couple the actual state of an action from its desired state, allowing for the introduction of methods for recovery from sensor errors and increasing the robustness of the system.

Traditional constraint propagation methods cannot be used to run the interval algebra networks associated to the interval scripts due to the typical size of the networks and the requirement of real-time computation [15]. To solve this problem, our system employs a novel method for constraint propagation in such networks that overcomes such limitations by projecting the network implicit in the script into a simpler, 3-valued (*past, now, future*) *PNF-network* where constraint propagation can be approximately computed in linear time (initially proposed in [16])

Interval scripts proved to be essential to manage the complexity of three large-scale projects of interactive, story-based spaces we have developed: the computer theater plays “*SingSong*” and “*It / I*”, and the interactive art installation “*It*”. For instance, in “*It*” we created interval scripts that generated networks with more than 200 nodes and 350 constraints. Nevertheless, the network could be completely evaluated 20 times per second using a run-time engine based on the *PNF propagation* algorithm described here.

This paper starts by reviewing the current programming paradigms and identifying their shortcomings. In section 3 we introduce the interval scripts paradigm through some simple examples. The core of the paper is the description of the PNF theory needed to allow the efficient computation of temporal constraint propagation algorithms, and of the run-time engine architecture that uses it, as described in sections 4 and 5. An implementation of interval scripts paradigm into a programming language called *Interval Script Language*, or *ISL*, is described in section 6, together with examples of its use in the context of the art installation “*It*”. Section 7 explores the three concrete cases of interactive environments that have been built using the paradigm and in section 8 we discuss future directions in our research.

## **2 PROBLEMS WITH CURRENT PROGRAMMING TECHNIQUES**

The idea of interval scripts was spawned by our dissatisfaction with the tools for the integration of I/O devices and software modules in interactive environments and for the control of computer agents. As described by Cohen, systems to control interaction tend easily to become “*big messy C program(s)*” ([17], fig. 2). Also, from our experience with “*The KidsRoom*” [2], it became clear that one of the major

hurdles to the development of interesting and engaging interactive environments is that the complexity of the control program grows faster than the complexity of the programmed interaction.

“*The KidsRoom*” was programmed using an event-loop paradigm [2]. However, during the development of “*The KidsRoom*” we encountered many situations where unanticipated and undesirable co-occurrences of events disrupted considerably the interaction. A typical case, similar to the one described in the introduction, was when multiple characters were triggered to talk or act by different events. The problem is that the characters, in many trials, ended up being triggered almost simultaneously, creating unintelligible overlapped speech or action. Since the environment used a control structure based in event-loops, it became necessary to examine each possible co-occurrence and to explicitly program how to handle every individual case. However, for each new action or character added to the environment, it was necessary to explore in the code the interaction between its actions with every action of other characters. In other words, the result is a *de facto* exponential increase in the number of different situations to be handled by the programmer.

## 2.1 Finite-State Machines and Event Loops

The most common technique used to program and control interactive applications is to describe the interaction through finite-state machines. This is the case of the most popular language for the developing of multimedia software, *Macromedia Director's Lingo* [18]. In *Lingo* the interaction is described through the handling of events whose context is associated with specific parts of the animation. There are no provisions to remember and reason about the history of the interaction and the management of story lines. The same situation occurs with *Max* [19], a popular language for control of music devices (see [20]) and adopted in some interactive art installations [21]

Videogames are traditionally implemented through similar event-loop techniques [22]. To represent the interaction history, the only resort is to use state descriptors whose maintenance tends to become a burden as the complexity increases. But most of all, a fundamental problem with the finite state model is that it lacks appropriate ways to represent the duration and complexity of human action, computational agents, and interactive environments: hidden in the structure is an assumption that actions and occurrences are pinpoint-like events in time (coming from the typical point-and-click interfaces for which those languages are designed).

## 2.2 Constraint-Based Programming Languages

The difficulties created by the finite-state based model have sparked a debate in the multimedia research community concerning the applicability of constraint-based programming (starting with the works of Buchanan and Zellweger [23] and Hamakawa and Rekimoto [24]) versus procedural descriptions such as state machines (for example, see [25]). In general, it is believed that constraint-based languages are harder to learn but more robust and expressive (as discussed in [26]).

For example, Bailey et al. [27] defined a constraint-based toolkit, *Nsync*, for constraint-based programming of multimedia interfaces that uses a run-time scheduler based on a very simple temporal algebra. The simplicity of the temporal model and, in particular, its inability to represent non-acyclic structures, is a major shortcoming of *Nsync* and similar systems such as *Madeus* [28], *CHIMP* [29], *ISIS* [30], and *TIEMPO* [31].

André and Rist [32] have built a system, called *PPP*, that designs multimedia presentations based on descriptions of the pre-conditions of the media actions, their consequences and on the definition of temporal constraints. The *PPP* system uses a strong, hybrid temporal algebra that combines intervals and points as its primitives (based on Kautz and Ladkin's work [33]). The descriptions of media actions are used to plan a multimedia presentation that achieves a given goal. The multimedia presentation is represented as a constraint network. During run-time, the scheduler has to build and choose among all the possible instances of the plans. However, since Kautz and Ladkin's propagation techniques are exponential in time, and the number of possible plans can be exponentially large, the applicability of André and Rist's *PPP* [32] is restricted to simple cases with limited interaction (otherwise the number of possible plans increases too quickly).

Finally, Borning [34] has been a strong proponent of constraint-based languages for animation in computer graphics; similar approaches can be found in *TBAG* [35] and in Kakizaki's [36] work on deriving animation from text.

## 2.3 Programming Interactive Environments

Starting with the pioneer work of Myron Krueger [37] the interest on building interactive environments, especially for entertainment, has grown in the nineties (see [38] for a review). In the case of virtual reality (VR), it seems that the user-centered and exploratory nature of most interfaces facilitates the programming of the interface by using state machines and event loops, using languages such as MR [11]

and *Alice* [12]. The latter was used in Disney’s “*Alladin*” ride (described in [1]). It is a language that allows rapid prototyping but can hardly describe any non-trivial temporal structure.

In [17] Cohen proposes a distributed control structure, the *Scatterbrain*. In this paradigm, used in the control of a sensor-loaded room, knowledge and behavior is dispersed among multiple agents that negotiate among themselves. Although an appealing model, the use of multiple agents without centralized control makes authoring extremely difficult in practice. In the opposite side of the spectrum, Dey [39] proposes a framework for the prototyping of context-aware environments. Both proposals lack models of temporal control capable of efficiently handling even modestly complex temporal structures such as the ones described in the introduction section.

## 2.4 Programming of Interactive Agents

In many interactive environments populated with computer characters, the control of the interaction is left to the characters themselves (seen as creatures living in a world). This is the case of Blumberg’s work in *ALIVE* [40], where the mood of the CG-dog *Silus* commands the interaction; and also in *Swamped* [41] where a user-controlled chicken plays cat-and-mouse with a raccoon. In both cases, however, there is no easy way to incorporate complex control structures into the control system.

Programming computer graphics characters and, in particular, humanoids, raises issues of duration and parallelism similar to those present in interactive environments. Many researchers have worked with languages for programming movements and reactions of characters, like Blumberg and Galyean [42], Perlin [43], Cassel et al. [44], and Thalman [45]. In particular, Perlin and Goldberg [46] have developed a language, *Improv*, that nicely integrates low-level control of geometry with middle-level primitives such as “*move*”, “*eat*”, etc. A major shortcoming of *Improv* is the lack of simple methods to synchronize and to temporally relate parallel actions.

Bates et al. [47] created an environment composed of several modules that encompass emotions, goals (using the language *Hap*), sensing, language analysis, and generation. Character actions are represented as AND/OR trees where AND nodes contain conjunctions of goals and OR nodes represent disjunctions of plans [48]. During run-time, the system grows an *active plan tree* considering the goals that are achieved and excluded by the current state of the world, and selects appropriate sub-goals and plans to be executed. Although *Hap* provides basic forecast of the future effect of actions, it is hard to express complex temporal structures in the formalism. By the own nature of AND/OR trees, *Hap* imposes many restrictions in the ways parallel actions can be synchronized.

Another line of research stems directly from Brooks' works with autonomous robots [49]. The objective is to create characters with their own behaviors and desires. Significant examples are Blumberg and Galyean's behavior control system [42], and Tosa et al.'s model for emotional control [50]. The problem with purely behavior-based characters is that they are unable to follow the high-level, complex structure of most stories and many interactive environments. Since the integration of high-level reasoning into behavior-based control is still an open problem, it is unlikely that behavior-based agents can be coordinated successfully to perform interesting stories and, even less, to manage complex interactive environments.

Further, Funge et al.[51] proposed the use of situation calculus to integrate reasoning into the behavior of characters. However, the applicability of the approach is arguably limited by the *frame problem* [52], that is, the need to declare every possible change in the environment resulting from a character's action.

### **3 THE INTERVAL SCRIPTS PARADIGM**

In traditional desktop-based interfaces and programs system and user actions are assumed to have no significant duration, except in some cases of multimedia programming. In the same manner, traditional programming is based on the concept that the states of all variables remain the same as long as no event occurs. When an event occurs, a piece of program is executed, normally assuming "instantaneous" response, and then a new state is reached. In other words, states are a-temporal, i.e., they last as long as no new detected event happens.

As described in the example of the introduction section, interactive environments and agents perform actions that often take a non-trivial amount of time, such as playing a sound file, showing a character animation, or closing a door. Similarly, user actions many times also have non-zero duration, such as when performing a gesture, walking around the space, or saying a sentence to a speech interface. These observations drove us to design a programming paradigm where all elements of a program, including its states and actions, are assumed to start and end, i.e., to have a non-zero duration.

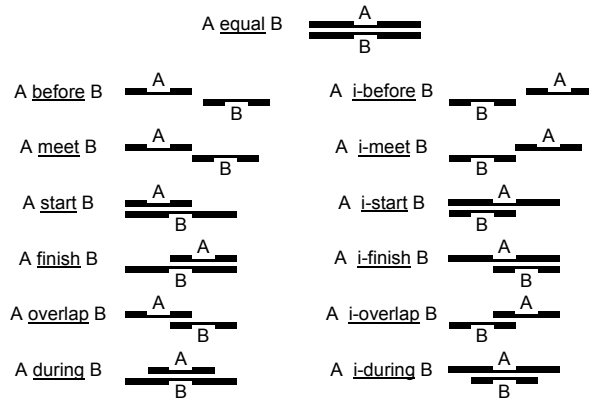
The fundamental concept of the interval script paradigm is that all actions and states are associated with the temporal interval in which they occur. However, the actual beginning and end time of the intervals are not part of the script. Instead, the script contains a description of the temporal constraints that the intervals should satisfy during run-time. These constraints are enforced by a run-time engine that examines the current states of all intervals and then coordinates the system's reaction to ensure that the interval's states respect the constraints.

For example, suppose we want to describe the situation where a computer graphics (CG) character appears on the screen whenever the user makes a particular gesture. To do this with interval scripts, we associate the user gesture to an interval  $A$  and the appearance of a CG character to another interval  $B$ . We then put a temporal constraint stating that the end of the gesture  $A$  should be immediately followed by the entrance of the character  $B$ . During run-time, the interval script engine monitors the occurrence of the gesture. When it finishes, the engine starts the entrance of the CG character in an attempt to satisfy the temporal constraint. Developing an interface or interactive system with interval scripts is a classical case of programming by constraints as discussed in section 2.2.

The idea of using temporal constraints to describe interaction in interactive environments initially appeared in Pinhanez et al. [16], in a different formulation from the one described in this paper. That work drew from a technique for representing the temporal structure of human actions by Pinhanez and Bobick [53]. The formalism and theory presented here considerably improve the syntax and semantics of the paradigm, as well as introduce a much more consistent, efficient, and robust run-time engine than in that work, as we will see in the next sections.

### **3.1 Allen's Interval Algebra**

To model the time relationships between two actions, a program in interval scripts employs the interval algebra proposed by Allen [14]. The representation of temporal constraints in Allen's algebra is based on the 13 possible primitive relationships between two intervals as shown in fig. 2. Given two actions or states in the real world, all their possible time relationships can always be described by a disjunction of those primitive time relationships. For instance, in the preceding example, the relation  $A$  meet  $B$  describes the temporal constraint between the gesture interval  $A$  and the interval associated with the entrance of the CG character  $B$ . If the entrance of the character could start before the end of the gesture, the relation between the two intervals would be described by the disjunction overlap OR meet. Of course, in any single actual occurrence of the intervals, only one of the disjunctive elements actually happens.



**Figure 2. Allen's 13 primitive relationships between two time intervals.**


To simplify the manipulation of the intervals and their constraints, a set of actions and states and their respective temporal constraints can be represented by a network in which each action or state is associated with a node and the temporal constraints between them are represented by arcs. The networks composed of such nodes and disjunctions of Allen's primitive relationships as arcs are called *interval algebra networks*, or simply, *IA-networks* [14]. It is important to notice that the nodes are, in fact, variables that can assume any temporal interval representing the actual interval of time of when an action or state occurs.

### 3.2 Strong Interval Algebras

Allen's algebra allows the expression of mutually exclusive intervals. For instance, to state that a computer agent does not perform actions  $C$  and  $D$  at the same time we simply constrain the relation between the intervals  $C$  and  $D$  to be meet OR i-meet. That is, in every occurrence of the intervals, either  $C$  comes before  $D$  or after  $D$ , with no overlap.

The ability of expressing mutually exclusive intervals defines different classes of temporal algebras [54]. In general, algebras without mutual exclusiveness allow fast constraint satisfaction but are not expressive [15]. However, in our approach we are able to employ Allen's *strong* algebra in a real time system because we have developed a fast method to compute approximations of the values that satisfy the constraints — PNF propagation — described later in this paper.

André and Rits [32] have also used a strong temporal algebra to script multimedia presentations. Their work employs a more expressive temporal algebra that combines Allen's [14] and Villain's [15] proposals as suggested by Kautz and Ladkin [33]. However, constraint propagation in this temporal



**Figure 3. Diagram of temporal constraints showing a sequence of three intervals.**

algebra is basically  $O(n^2(e+n^3))$ , where  $n$  is the number of nodes and  $e$  is the time required to compute the minimal network of the IA-network; typically,  $e$  is exponential in the number of nodes,  $O(e) = O(2^n)$ , making the whole system basically  $O(n^2 2^n + n^5) = O(2^n)$ . Moreover, their system requires the enumeration in computer memory of all possible temporal orders of actions implied by the constrained network, which can grow exponentially, especially in the case of multiple and frequent user interaction.

### 3.3 An Example of Programming with Temporal Constraints

To understand how to use interval algebras in interaction programming, let us examine an example from the script of one of our experiments, the art installation called “*It*” described in subsection 7.3. In that installation a CG object that look like a photographic camera, referred here as “*camobject*”, appears on the screen and interacts with the user.

Suppose we want to describe a situation where “*camobject*” moves front, that is, moves forward towards the user, produces a clicking sound (as it was taking a picture), and then moves back to its previous position. The first step, not described here, is to associate each of these three actions to nodes in a IA-network, and name them, respectively, “*camobject moves front*”, “*camobject clicks*”, and “*camobject moves back*”. To program the desired sequencing of the three actions, we define meet constraints between each consecutive pairs:

“*camobject moves front*” meet “*camobject clicks*”

“*camobject clicks*” meet “*camobject moves back*”

Figure 3 shows the desired relation between the three actions. The first constraint states that the action “*camobject moves front*” should meet with the action “*camobject clicks*”, that is, the latter action should start immediately after the end of the former. The second constraint establishes a similar constraint between the actions “*camobject clicks*” and “*camobject moves back*”.






Figure 5. Possible occurrences of “user is moving”.

state of the user as detected, for instance, by a sensing device. To program that the “clicking” action should not happen if the user is moving, we can simply enforce the following constraint:

“*camobject takes a picture*” before OR i-before “*user is moving*”

Figure 5 shows three possible occurrences of the sensor associated with “*user is moving*”. Cases (1) and (2), show situations where it happens respectively before or after “*camobject takes a picture*”, and therefore are compatible with the occurrence of this sensed activity. However, cases (3) and (4) show hypothetical situations where “*user is moving*” overlaps with “*camobject takes a picture*” and thus it is not compatible.

Our goal is to design a run-time engine that prevents the occurrence of case (3) by not allowing the “*camobject takes a picture*” action to start while “*user is moving*” is not finished; and by immediately (maybe prematurely) finishing “*camobject takes a picture*” if the user starts moving as in case (4). To understand better how this is accomplished, it is necessary to examine how the run-time engine actually works, as it is described in the next two sections.

### 3.5 The Case for Allen’s Algebra

As can be seen from the above examples, there are several reasons to use Allen’s algebra to describe relationships between actions as the programming paradigm for interactive environments. First, no explicit mention of the interval duration or specification of relations between the intervals’ starting and endings are required. Second, the notion of disjunction of interval relationships can be used to declare multiple paths and interactions in a story. Third, it is possible to declare mutually exclusive intervals that describe actions and states whose co-occurrence is not desirable.

Fourth, as discussed in the next section, it is only necessary for the programmer to declare the relevant constraints. By applying *Allen’s path consistency algorithm* [14] to the IA-network associated with the script, as described in this paper, it is possible to generate a more constrained version of the network that

contains all and only the constraints that are consistent with time. Fifth, it is possible to determine in linear time which computational agents' and characters' and environment's reactions to trigger in response to users' actions by properly propagating occurrence information from one node of the IA-network to the others, using the PNF propagation method described next.

#### **4 A METHOD FOR FAST CONSTRAINT PROPAGATION IN STRONG ALGEBRAS**

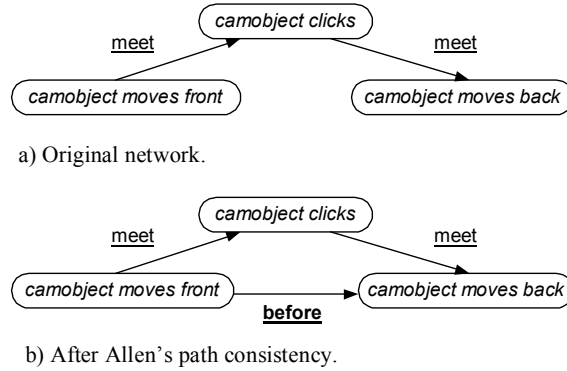
Programming with interval scripts involves the association of actions and world states of an interactive environment to a collection of nodes in an IA-network with temporal constraints between them. However, in order to control an interactive environment or agent based on such descriptions, it is necessary to employ a run-time engine that examines which intervals are happening, considers the constraints between them, and controls the start and stop of intervals as needed to enforce the constraints. To perform such actions, it is necessary to use a process generally known as *constraint propagation*.

In general, performing constraint propagation in an IA-network is exponential (as proved in [15]). We overcome this problem by using a simplification of constraint propagation in IA-networks called *PNF propagation*. A complete account of the theoretical foundations of PNF propagation as well as a more detailed analysis of the algorithms to be presented here can be found in [13]. In this paper we provide the basics of the method that are necessary to appreciate its possibilities and strengths as the backbone of a constraint propagation run-time engine.

The basic idea of our approach is to remove all information about actual interval duration (i.e., time intervals) from an IA-network by making its nodes assume only three values: *past*, *now*, and *future*, representing respectively the situation where a node has occurred, is occurring, or is still to occur. The value associated with a node at a certain moment of time is called the *PNF-state* of the node, or simply, the *state* of the node.

##### **4.1 First Step: Allen's Closure**

As mentioned above, the collection of temporal constraints between nodes is represented by an *interval algebra network* [14], which is a network where the nodes correspond to variables on temporal intervals and the temporal constraints are links between them. Figure 6.a displays the IA-network associated with the three nodes connected through meet constraints of the "camobject" example described before.



**Figure 6. Interval algebra network associated with two meet constraints, before and after path consistency.**

To prepare an interval algebra network for run-time execution, we start by tightening the temporal relations by running Allen's *path consistency algorithm* (see [14]) on the network. The result is a network where implied temporal relations are detected and explicitly incorporated to the network. For example, by looking at fig. 6.a, we can see that since the node “*camobject moves front*” meet “*camobject clicks*”, and that “*camobject clicks*” meet “*camobject moves back*”, there is only one possible relation between “*camobject moves front*” and “*camobject moves back*”: the former must happen before the latter. As shown in fig. 6.b, this is automatically detected by Allen's algorithm.

In fact, detecting every possible implied temporal relations is *NP-hard* (as shown in [15]). However, Allen's path consistency algorithm computes a very good, conservative, approximation in most practical cases [14], and it is polynomial in the number of constraints. Furthermore, when used in the programming process of an interactive environment or agent, the closure needs to be computed only once by an offline initialization procedure, making its computational time less important.

#### 4.2 Minimal Domains of IA-Networks

Traditional constraint satisfaction in interval algebra networks tries to determine for each network node the sets of time intervals (pairs of real numbers describing segments of the time line) that are compatible with the constraints given the known occurrence of some intervals. An assignment of time intervals that satisfies the constraints is called a *solution* of the network. Given a network node, the set of all time intervals for that node that belongs to at least one solution is called the *minimal domain* of the node. Notice that if there is at least **one** interval in the minimal domain of a node that contains a given instant

of time then the node **can** be happening at that time. If **every** time interval of the minimal domain of a node contain that given instant of time, then the node **must** be happening at that given moment.

This observation constitutes the basic principle of our run-time engine. However, direct constraint propagation is *NP-hard* due to the combinatorial explosion involved in the manipulation of sets of intervals (see [15]). To overcome this difficulty we devise an optimization scheme called *PNF propagation* comprising two mechanisms, *PNF restriction* and *time expansion*.

### 4.3 PNF Restriction

The idea of simplifying interval algebra networks was first proposed by Pinhanez and Bobick [55], in a work motivated by the problem of action recognition in a computer vision system. The authors' key observation is that for control and recognition purposes there is little information coming from the duration of the nodes. Instead, the focus in such situations is to simply determine if an interval has happened (*past*), is happening (*now*), or may still to happen (*future*).

Based on this criterion, it is possible to project the interval network into a similar network where the nodes corresponding to the intervals can assume only one of the three symbols, *past*, *now*, or *future* — a *PNF-network*. In practice it is very difficult to determine exactly the state of every node in a PNF-network. For instance, suppose a node that is associated to a binary sensor. If the sensor is *on*, we represent this by associating the symbol *now* to the node. On the other hand, if the sensor is *off* we assign to the node a disjunction of values, stating that either the sensor was *on* in the *past*, or will be *on* in the *future*, represented by the set  $\{past, future\}$ .

To simplify notation, let us start by defining the set  $M$  of all subsets of  $\{past, now, future\}$

$$M = \{\emptyset, \{past\}, \{now\}, \{future\}, \{past, now\}, \{now, future\}, \{past, future\}, \{past, now, future\}\}$$

whose elements, called *PNF-values*, are abbreviated as

$$M = \{\emptyset, P, N, F, PN, NF, PF, PNF\}$$

Also, we call any tuple  $W = (W_1, W_2, \dots, W_n)$  of PNF-values associated to the  $n$  nodes of a PNF-network a *component domain* of the PNF-network, where each  $W_i \in M$ , and  $W \subseteq M \times M \times \dots \times M = M^n = U$ . Since each element of a component domain represents possible states of its associated node, we abuse our notation and also call each  $W_i$  the *PNF-state* of the node  $i$ .

**Table 1. Admissible values for primitive temporal relations in a PNF-network.**

admissible values of B when A r B			
$r$	A={P}	A={N}	A={F}
equal	P	N	F
before	PNF	F	F
i-before	P	P	PNF
meet	PN	F	F
i-meet	P	P	NF
overlap	PN	NF	F
i-overlap	P	PN	NF
start	PN	N	F
i-start	P	PN	F
during	PN	N	NF
i-during	P	PNF	F
finish	P	N	NF
i-finish	P	NF	F

It is easy to see that, given a primitive temporal relation  $r$  between two nodes  $A$  and  $B$  of a PNF-network, the admissible states of the two nodes must satisfy the constraints depicted in table 1. For instance, if  $A$  is happening *now* ( $N$  state), and  $A$  meet  $B$ , then  $B$  can only happen in the *future* ( $F$  state, as highlighted in table 1). The construction of the table is formally justified in [13] but it should be noticed that the table correspond to our intuitive notions of *past*, *now*, and *future*.

If a temporal constraint is a disjunction of primitive temporal primitives, we simply take the union of the PNF-values corresponding to the primitives. For instance, suppose  $A$  and  $B$  are mutually exclusive,  $A$  meet OR i-meet  $B$ , and  $A$  is happening *now*. Table 1 yields that for the meet constraint,  $B$  may happen in the *future* ( $F$ ), while for the i-meet constraint  $B$  must have already happened ( $P$ ). By taking the union of the two constraints in the disjunction,  $PF = P \cup F$ , we determine that the actual state of  $B$  must belong to  $PF$ , and therefore, that  $B$  is not happening now.

On the other hand, if the node  $A$  is associated to a non-singular PNF-value (i.e.,  $PF$ ,  $PN$ ,  $NF$ , or  $PNF$ ), it is straightforward to see that the possible values for  $B$  are simply the union of the admissible values for each element of the set associate with  $A$ . For instance, if  $A$  finish  $B$  and the PNF-state of  $A$  is  $PN$ , then the PNF state of  $B$  is the PNF-value  $PN = P \cup N$ .

Given an initial set of PNF-values for each node represented by a component domain  $W = (W_1, W_2, \dots, W_n)$ , we define a *solution* of a PNF-network under  $W$  to be any assignment of one of the

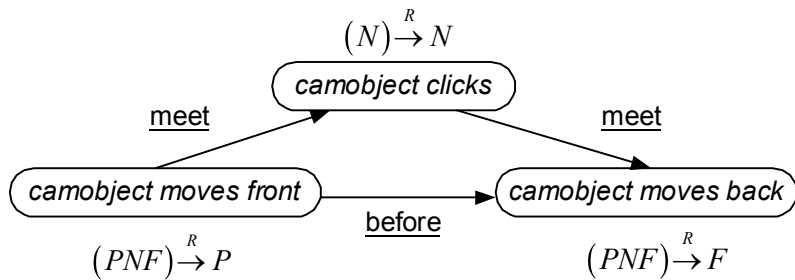


Figure 7. An example of PNF restriction.

basic  $\{past, now, future\}$  values to each node  $i$  of the network in which each assigned value belongs to  $W_i$  and that each pair of nodes satisfies the constraints of table 1. We call the *minimal domain* under  $W$  of a node in a PNF-network as the set of all values that belong to at least one solution under  $W$ . Notice that the minimal domain of a node is also a PNF-value. Similarly, the *minimal domain* of a PNF-network under  $W$  is the component domain composed of the minimal domain of each node of the network.

We call the *restriction of a PNF-network* under  $W$ ,  $R(W)$ , the process of computing the minimal domain of the network under  $W$ . That is, the restriction process eliminates all the values from each component of  $W$  that are incompatible with at least one solution under  $W$ . Notice that the restriction of a component domain  $W$  is also a component domain  $R(W) \subseteq U$ , and that restriction is maximal, i.e.,  $R(R(W)) = W$ . Notice that if the minimal domain of any node is empty, i.e.  $W_i = \emptyset$  for any  $1 \leq i \leq n$ , then there are no solutions for the network.

Figure 7 shows a simple example of PNF restriction. The initial set of values for each node is shown between parenthesis and the values after restriction appears on the right side of the arrows. In this case, we assume that it is known that the “*camobject clicks*” action is happening ( $N$ ). Since this action must happen after “*camobject moves front*” because of the meet constraint, the only possible state for the later node is *past* ( $P$ ), as shown in fig 7. Similarly, “*camobject moves back*” cannot have happened because it must follow “*camobject clicks*” and therefore its possible state is restricted to *future* ( $F$ ).

Computing the exact minimal domain of a PNF-network is, however, still a *NP-hard* problem (see [56]). Instead, we employ in our applications a conservative approximation to the minimal domain based on the computation of the *arc-consistency* (as proposed by Mackworth [57]) of the PNF-network. The big

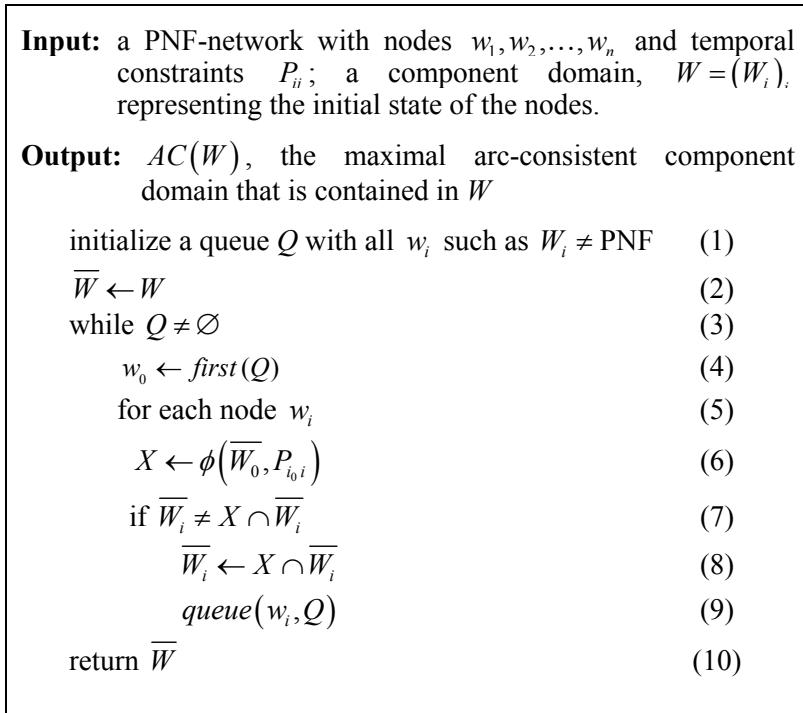


Figure 8. Algorithm to compute arc-consistency.

advantage is that arc-consistency is  $O(c)$  in the number of constraints  $c$ , that is, at most  $O(n^2)$  in the number of nodes  $n$ .

The procedure in fig. 8 shows an algorithm that computes the maximal arc-consistent domain of an  $n$ -node PNF-network under a component domain  $W = (W_1, W_2, \dots, W_n)$ . This is a version of the arc-consistency algorithm AC-2 proposed in [57] and adapted to the component domain notation. The algorithm uses the function  $\phi$  which given a PNF-value and a set of primitive relations, returns a PNF-value that satisfies table 1. In [13], Pinhanez proves that the result of this algorithm always contains all possible solutions of the network (that is, it is conservative) and it is linear in the number of constraints. Pinhanez also shows that arc-consistency produces a reasonable approximation of the minimal domain. In fact, in our experiments we have encountered few situations where they were actually different (see [13] for details).

#### 4.4 Time Expansion

PNF-restriction deals exclusively with determining feasible options of an action **at a given moment of time**. However, information from the previous time step can be used to constrain the occurrence of intervals in the next instant. For example, after an interval is determined to be in the *past*, it should be

impossible for it to assume another PNF-value, since, in our semantics, the corresponding action is over. Similarly, if the current value of the interval is *now*, in the next instant of time it can still be *now*, or the corresponding action might have ended, when the interval should be *past*. To capture these ideas, we define a function that time-expands a component domain into another that contains all the possible PNF-values in the next instant of time for each node.

Formally, we define a function  $T$ , called the *time expansion function*, that considers a component domain  $W^t$  at time  $t$  and computes the component domain  $W^{t+1} = T(W^t)$  at time  $t+1$  by considering what the possible states are in  $t+1$ . To define that function, we start by defining a time expansion function  $\tau_m$  for each element of  $\{past, now, future\}$ , such as:

$$\tau_m(past) = P \quad \tau_m(now) = PN \quad \tau_m(future) = NF$$

Notice that  $\tau_m$  assumes that between two consecutive instants of time there is not enough time for an interval to start and finish. That is, if a node has a value *future* at time  $t$ , it can only move to *now*, but never straight to the *past* state at time  $t+1$ . Based on  $\tau_m$ , we define the function that expands the elements of  $M$ ,  $T_m : M \rightarrow M$  as being the union of the results of  $\tau_m$ ,

$$T_m(\Omega) = \bigcup_{\omega \in \Omega} \tau_m(\omega)$$

and the time expansion of a component domain  $W$ ,  $T_m : U \rightarrow U$  (abusing the notation), as the component-wise application of the original  $T_m$  on a component domain  $W = (W_i)_i$ ,

$$T_m(W) = (T_m(W_1), T_m(W_2), \dots, T_m(W_n)).$$

In [13], Pinhanez shows that it is possible to define a similar time expansion function that allows an interval to start and finish between two consecutive instants of time. For simplicity, we limit the discussion in this paper to the function defined above.

#### 4.5 PNF Propagation

Now that we have methods both for constraint propagation and for incorporating time information, we are ready to define the temporal reasoning foundation of our run-time engine, so that for each instant of time it is possible to determine which PNF-states are compatible with the current information coming from sensors, the previous state of the system, and the constraints between the intervals.

Let us consider an initial component domain  $W^0$  as being composed only of *PNF* *PNF*-values,  $W^0 = (PNF)_i$ . Then, for each instant of time  $t$ , we can determine through sensor information or external sources the *PNF*-state of some of the nodes. With this information, we create the component domain  $V^t$  containing all these known values and assigning *PNF* for the other nodes. Then, given the previous component domain  $W^{t-1}$ , we can compute an upper bound of the current minimal domain of the *PNF*-network (see the proof in [13]) by making

$$W^t = R\left(T_m(W^{t-1}) \cap V^t\right)$$

We call this process *PNF propagation*. Notice that the more information contained in  $V^t$ , the smaller is  $W^t$ . In the extreme case, if  $V^t$  is the minimal domain, then  $W^t$  is also the minimal domain. In most cases, however, we will have  $V^t$  providing new information that is filtered through the intersection with the past information (provided by  $T_m(W^{t-1})$ ). Then, *PNF*-states incompatible with the structure of the problem are removed by restriction, through the computation of the minimal domain. In practice, we can employ the arc-consistency algorithm as defined in fig. 8 instead of full restriction to assure that computation occurs in linear time.

Let's examine a simple example of *PNF* propagation where the IA-network is composed of two intervals,  $A$  and  $B$ , and the constraint  $A$  meet  $B$ . Suppose also that  $A$  corresponds to a binary sensor. The standard way to represent the *on/off* states of  $A$  in the *PNF* formalism is to assign the *PNF*-value  $N$  to the *PNF*-state of  $A$  when  $A$  is *on*, and  $PF$  when  $A$  is *off*. Similarly,  $B$  corresponds to an action of the system that must happen immediately after the occurrence of  $A$ ; however, the system has no direct way to determine the state of  $B$ . This example shows that, using *PNF* propagation, it is possible to compute when  $B$  starts to occur (as a response to  $A$ ).

Suppose that the initial state of the sensor  $A$  is *off* ( $PF$ ) while  $B$ 's is unknown ( $PNF$ ), represented by  $W^0 = (PF, PNF)$ . Assume that in the next instant of time,  $t=1$ ,  $A$  turns *on*, represented by  $V^1 = (N, PNF)$ . Using *PNF* propagation,

$$\begin{aligned} W^1 &= R\left(T_m(W^0) \cap V^1\right) = R\left(T_m((PF, PNF)) \cap (N, PNF)\right) = \\ &= R\left((PNF, PNF) \cap (N, PNF)\right) = R\left((N, PNF)\right) = (N, F) \end{aligned}$$

The result,  $W^1 = (N, F)$ , yields that since  $A$  is *on*, then  $B$  has not occurred yet.

Now, let's assume that in the next instant of time,  $t = 2$ ,  $A$  turns *off*,  $V^2 = (PF, PNF)$ . The computation of the current state using PNF propagation determines, as expected, that  $B$  must be occurring in instant 2,

$$\begin{aligned} W^2 &= R(T_m(W^1) \cap V^2) = R(T_m((N, F)) \cap (PF, PNF)) = \\ &R((PN, NF) \cap (PF, PNF)) = R((P, NF)) = (P, N) \end{aligned}$$

This result is simply a consequence of the fact that, since  $A$  meet  $B$ ,  $B$  must start immediately after the end of an occurrence of  $A$ . However, the time when  $B$  ends cannot be determined, as shown here in the computation of the state of the next instant of time:

$$W^3 = R(T_m((P, N)) \cap (PF, PNF)) = R((P, PN) \cap (PF, PNF)) = R((P, PN)) = (P, PN)$$

In fact, since  $W^3 = (P, PN)$ ,  $B$  is either happening ( $N$ ), or has already happened ( $P$ ). To determine the end of  $B$ , it is necessary to have another sensor or a constraint from another node.

## 5 THE INTERVAL SCRIPT RUN-TIME ENGINE

Programming in interval scripts is achieved by defining temporal constraints between nodes associated to sensors and actuators of the interactive environment, agent, or computer characters. The constraints describe the desired temporal structure of the occurrences of the actions and states associated with the sensors and actuators. However, an interval script only declares how the actions and sensed states should and should not happen when the interactive environment or agent is running. In order to actually control the environment or agent, it is necessary to use a run-time system that collects the information from the sensors and, based on the interval script, actually triggers and stops the actions of actuators — the *interval script run-time engine*.

Although there may be different ways to implement a run-time engine for intervals scripts (for instance, using Andre and Rist's constraint propagation method [32]), we describe here an implementation based on the concept of PNF propagation introduced in the previous section. As discussed before, the main advantage of PNF propagation over similar methods is its fast computation time and linear memory requirements.

The basic cycle of our run-time engine is composed of four stages. First the engine computes the states of all sensors. Second, it applies PNF propagation to determine how the system can act so in the next

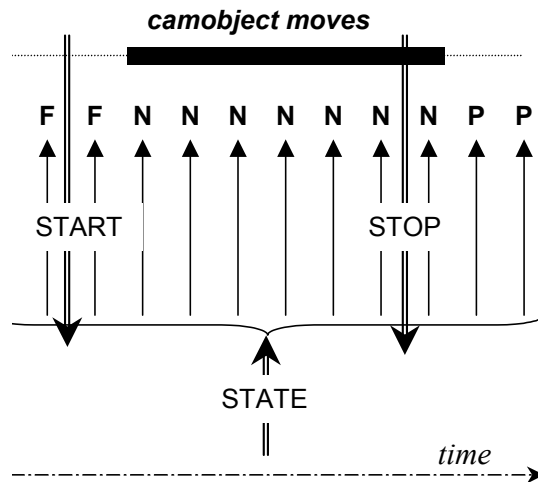


Figure 9. The execution of an action node.

cycle the states of all sensors and intervals are compatible with the constraints. Third, the engine selects a particular combination of goal states for all actuator nodes. And fourth, the engine tries to start or stop actions calling appropriate functions, aiming to achieve a state in which all constraints are satisfied.

One way to understand our run-time engine is to regard the engine as a player in a two-player game, and assuming the other player to be the environment and its users. Every play is composed of two “moves”, the first performed by the environment and/or its users (sensed/computed by the first stage of the engine), and the second performed by the fourth stage of the run-time engine. The second and third stages can then be seen as a planning process in which the system explores the effects of different combinations of actions, trying to find a set of actions that results in states of the nodes that satisfy the constraints. This planning process, whose computation is based on PNF propagation, assumes that the environment remains constant except for the system’s actions. We will later discuss how to overcome this limitation within this formalism.

### 5.1 Start, Stop, and State Functions

To connect an interval script to a real environment or agent, it is necessary to associate each node of the interval script to actual sensors and actuators. In an interval script, this is accomplished by associating three types of functions to a node:

- **START:** a function that determines what must be done to start the action associated with a node; however, after its execution, it is not guaranteed that the action has actually started.
- **STOP:** a function defining what must be done to stop the action associated with a node; similarly to START, the interval may continue happening after the STOP function is executed.

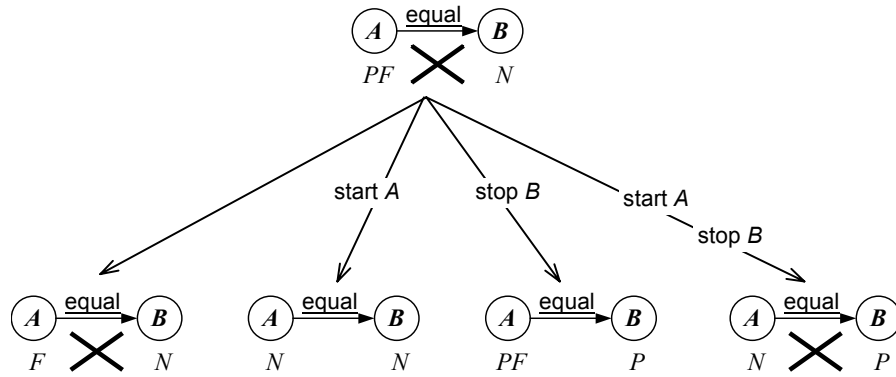


Figure 10. Example of an inconsistent state and the effect on possible actions.

- **STATE:** a function that returns the actual current state of the interval (in our implementation as a disjunction of the three values, *past*, *now*, or *future*) corresponding respectively to the situation where the node has already happened, is happening in that moment of the time, or it has not happened yet.

The STATE functions should be oblivious to START and STOP functions and should be defined so they provide a real assessment of the actual situation of the interactive environment and the on-going actions of the system and its users. STATE functions are used both to define nodes attached to sensors and to provide the actual state of execution of an action.

As we can see, in interval scripts we de-couple the goal of starting or stopping an action (represented by the START and STOP functions) from the actual state of its associated node. That is, an interval script describes how an interaction should happen and does not assume that response is perfect. That is simply because it is impossible to predict how devices, users, or computer agents actually react in run-time due to delays, device constraints, or unexpected interactions.

Figure 9 shows a typical execution of an action in interval scripts. This action corresponds to a movement of an object on a screen. Time in this diagram proceeds from left to right. Notice that the START function is executed some time before the STATE functions actually detects the occurrence (*N*) of the action. Similarly, it takes time for the effects of the STOP call to be sensed and/or take effect. In fact, fig. 9 depicts the normal case where the action actually happens some time after the START call.

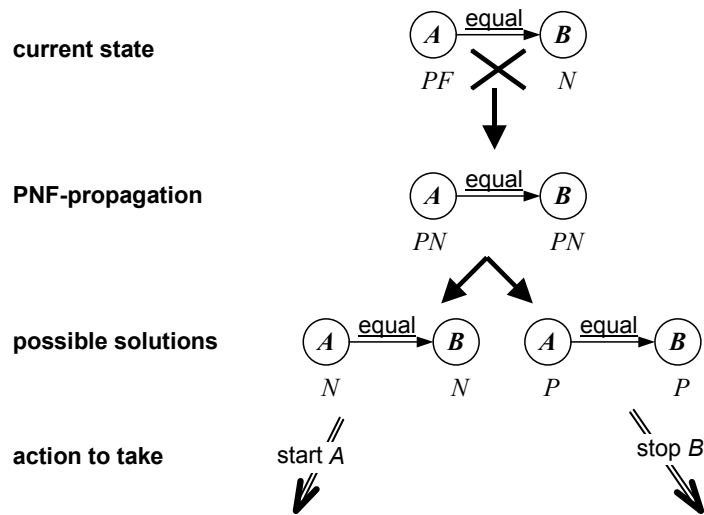


Figure 11. Finding a globally consistent state for the situation depicted in fig. 10.

## 5.2 Choosing Globally Consistent States for Actions

We now detail how the state of the intervals is determined during run-time and the process that triggers the call of START and STOP functions. At each time step the goal is to select which START/STOP functions to call in order to make the state of the world in the next cycle, represented by the PNF-state of the nodes, consistent with the temporal constraints of the PNF-network associated to the script.

Let us first examine a very simple case of a network of two intervals, *A* and *B*, temporally constrained to be equal. Figure 10 shows a situation where the PNF-state of interval *A* is *PF* and *B* is *N*, and therefore their values do not satisfy the constraint between them. In that situation the system should take one of the two actions: try to start interval *A* (so its PNF-state becomes *N*) or to stop interval *B* (so its PNF-state becomes *P*). However, as shown in fig. 10, there are four options of actions to be taken, considering that the no action can be done, one of the two, or both. In particular, notice that if the system decides both to start *A* and to stop *B*, a new inconsistent state is reached.

This example shows that choosing an action involves looking for sets of actions whose outcome lead to the global consistency of the network. If we look only locally in each node for a way to make the network consistent, it might happen the nodes change their state but keep the network inconsistent as, in the example of fig. 10, it is the case when both *A* is started and *B* is stopped. However, this example also shows that there are multiple, often contradictory, ways to satisfy the constraints of a network.

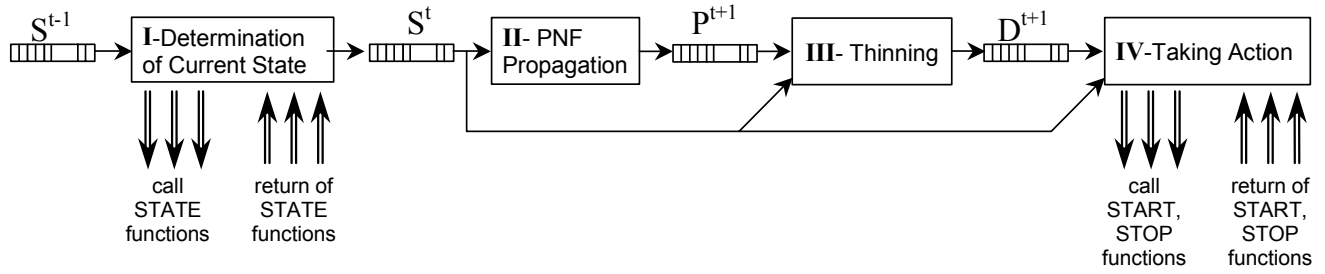


Figure 12. A cycle of the interval engine.

In our run-time engine, we choose which intervals to start or to stop by looking for global solutions for the IA-network corresponding to the achievable PNF-states. However, as shown in fig. 11, running PNF propagation on the current state generates a state  $(PN, PN)$  which represents not only the two different possible solutions,  $(N, N)$  and  $(P, P)$ , achievable by starting the  $A$  interval or stopping the  $B$  interval, respectively but also the PNF-states  $(P, N)$  and  $(N, P)$  which are not compatible with the constraint equal.

To decide which action to take, we could examine every possible combination of PNF-values contained in the result of the PNF propagation, and discard those who are not solutions. This is clearly an exponential search, so in practice we have used heuristics that “thin” the result of PNF propagation, as described in the next subsection, although still looking for a globally consistent solution. Based on the result of the thinning process, the system then chooses the set of appropriate actions that potentially lead to that state.

### 5.3 The Basic Cycle of the Run-Time Engine

Figure 12 shows a diagram of one cycle of our run-time engine. The main component is the current state  $S^t$  at time  $t$  that contains the current state of all intervals in the associated PNF-network.

Each cycle  $t$  is composed of the following four stages:

- I. Determination of the Current State: all the STATE functions are called and the engine waits until all the results are reported back and stored in the component domain  $S^t$ . For this computation, the STATE function is allowed to use the previous state of the intervals, available in the previous component domain  $S^{t-1}$ , and the current states of all the intervals that had already been computed. After this stage, the current state remains unchanged for the rest of the cycle.

II. PNF Propagation: in this stage the engine tries to determine what actions the system can perform to satisfy the constraints in the next cycle  $t+1$ , noted as  $P^{t+1}$ , assuming that the state of the environment does not change. For this, we use a variation of the idea of PNF propagation where we consider for temporal expansion only those intervals that can be started or stopped,

$$V_i^{t+1} = \begin{cases} T_m(S_i^t) & \text{if the interval } i \text{ has START or STOP functions} \\ S_i^t & \text{otherwise} \end{cases}$$

By computing PNF propagation only on the nodes that can be started or stopped, the system only examines changes that can happen as a result of its own actions, ignoring all the possible changes in the environment that could be detected by the sensors. In practice, the run-time engine behaves as it was in a two-player game against the environment and its users. While in state I it detects the “play” of the environment, in this stage it explores all its own “moves”, eventually leading to a decision in stage III, and to action in stage IV.

Using the definition above, this stage then PNF-propagates the current state,

$$P^{t+1} = R(T_m(S^t) \cap V^{t+1}) = R(V^{t+1})$$

If the computation of  $P^{t+1} = R(V^{t+1})$  detects a conflict,  $P^{t+1} = \emptyset$ , the engine tries to enlarge the space of possibilities by using simple PNF propagation, that is, expanding all states regardless of the existence of START or STOP functions,

$$P^{t+1} = R(T_m(S^t))$$

Although this normally handles most problems, there are situations where sensors report incorrectly and produce a state with no possible solutions. In those extreme cases, the engine simply time-expands the current state without computing the restriction,

$$P^{t+1} = T_m(S^t)$$

III. Thinning: the result of stage II is normally too big and undetermined, although it contains all the globally consistent solutions. To have a more specific forecast of the next stage without doing search, we apply a heuristic based on the principle that minimal system action is desirable: the current state of an interval should remain the same unless it contradicts the result of the PNF propagation. This is accomplished by taking a special intersection operation between the forecast

result of the actions in the next state  $P^{t+1}$  and the current state  $S^t$ . For each node the special intersection is computed by

$$S_i^t \bar{\cap} P_i^{t+1} = \begin{cases} S_i^t \cap P_i^{t+1} & \text{if } S_i^t \cap P_i^{t+1} \neq \emptyset \\ P_i^{t+1} & \text{otherwise} \end{cases}$$

As seen in the formula, this special intersection operation simply computes a component domain where the PNF-values of each interval is exactly the intersection of its current value and the possible globally consistent values it can assume if action(s) is(are) taken by the system. If the intersection is empty, it means that the current state is incompatible with the result of any possible action of the system, and therefore, that change is necessary, as represented in  $P_i^{t+1}$ . If the intersection is non-empty, we have a node whose current state (or a fraction of it) is possibly compatible with the predicted results of possible actions. Following the minimal change heuristics, the thinning process makes the goal of the engine to keep the PNF-state of this node  $i$  to be the intersection of the current state  $S_i^t$  and the forecast results  $P_i^{t+1}$ .

The component domain that results from the component-wise intersection is then passed through PNF restriction to assure that there are solutions and to remove impossible states, yielding the desired state  $D^{t+1}$

$$D^{t+1} = R(S^t \bar{\cap} P^{t+1})$$

If the computation of  $D^{t+1} = R(S^t \bar{\cap} P^{t+1})$  yields a state with no solutions we ignore the intersection,

$$D^{t+1} = P^{t+1}$$

In practice ignoring the intersection normally prevents any action to be taken since the states of  $P^{t+1}$  tend to be not as small as required to call start and stop functions as described next. Notice that it is not necessary, in this case, to apply the PNF restriction to  $P^{t+1}$ , since this was already been done in the previous stage.

IV. Taking Action: The result of thinning,  $D^{t+1}$ , is compared to the current state  $S^t$ , and START and STOP functions are called if the need to change the state of an interval is detected. To choose among equally possible alternatives such as in the case of fig. 11, we use a heuristics that

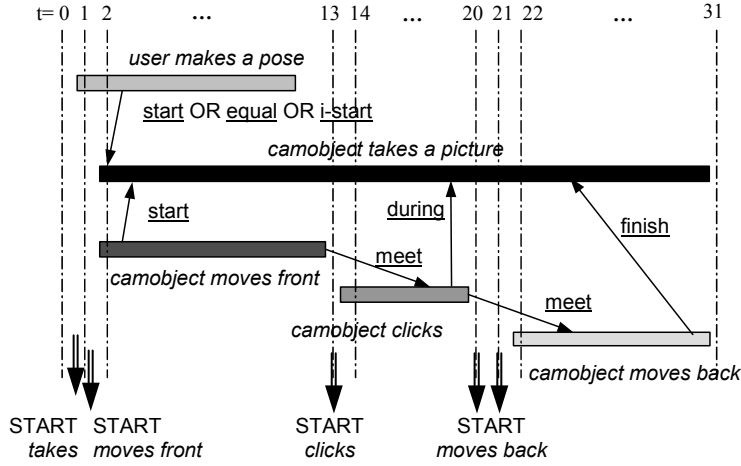
makes the system more willing to start an action than to stop one. Actions are performed by the system according to the following table:

$x \subseteq S_i^t$	$D_i^{t+1}$	action
$F$	$N, PN$	START
$N$	$P$	STOP
$F$	$P$	STOP

For example, if the current state of interval  $i$  can be future,  $F \subseteq S_i^t$ , and the desired state is *now*,  $N = D_i^{t+1}$ , then the START function of the interval is called. Notice that actions are started both if the desired state is  $N$  or  $PN$ , while they can be stopped only if the desired state is exactly  $P$ . In practice, this simple distinction is sufficient to make more likely the starting than the stopping of an action.

Notice also that the START function is never called if the current state of its corresponding interval does not contain the *future* state. That is, an interval is started only if there is a possibility that interval is in the *future* state. However, a STOP function is sometimes called even if the current state of the interval does not contain the *now* state. In this case we employ a mechanism by which if a STOP function is called and the state is exactly *future*, then the state of the interval is automatically set to *past*, without actually calling its STOP function.

To perform the search for globally consistent solutions instead of the thinning process of stage III, this stage would be substituted by the search of all solutions and by a mechanism to arbitrary choose one solution among them (for instance, by considering the first solution found) As discussed in [13], it is possible to use arc-consistency to speed up considerably such search. In practice, the run-time engine with the minimal change heuristics described above has been able to run all our applications without halting, deadlocking, or flip-flopping.



interval	t=0	t=1	t=2	t=13	t=14	t=20	t=21	t=22	t=31
	$S^0 P^1 D^1$	$S^1 P^2 D^2$	$S^2 P^3 D^3$	$S^{13} P^{14} D^{14}$	$S^{14} P^{15} D^{15}$	$S^{20} P^{21} D^{21}$	$S^{21} P^{22} D^{22}$	$S^{22} P^{23} D^{23}$	$S^{31} P^{32} D^{32}$
user makes a pose	F F F	N N N	N N N	P P P	P P P	P P P	P P P	P P P	P P P
camobject moves front	F F F	F N N	N N N	P P P	P P P	P P P	P P P	P P P	P P P
camobject clicks	F F F	F NF F	F NF F	F N N	N PN N	P P P	P P P	P P P	P P P
camobject moves back	F F F	F NF F	F NF F	F NF F	F NF F	F N N	F N N	N PN N	P P P
camobject takes a picture	F F F	F N N	N N N	N N N	N N N	N N N	N N N	N N N	P P P

Figure 13. Example of a run of the interval script of fig. 17.

Let us examine the behavior of the run-time engine in the example of fig. 10 where intervals  $A$  and  $B$  are constrained by an equal relation and their initial states are  $PF$  and  $N$ , respectively. To simplify notation, let us assume that  $t=0$  and the current state determined at stage I is  $S^0 = (PF, N)$ , where  $PF$  corresponds to the state of  $A$  and  $N$  to the state of  $B$ . As shown in fig. 11, the result of PNF propagation is  $P^1 = (PN, PN)$ . The first attempt of thinning the result of PNF propagation yields  $S^0 \bar{\cap} P^1 = (P, N)$ , which is not consistent,  $D^1 = R(S^0 \bar{\cap} P^1) = R((P, N)) = \emptyset$ . By the recovery heuristics, we define the desired state to be  $D^1 = (PN, PN)$ . According to the stage IV of the run-time engine, only the START function of  $A$  would be called. The choice of starting  $A$  over stopping  $B$  is, in this case, an arbitrary by-product of the taking action heuristics of stage IV that tends to favor starting actions over stopping them.

#### 5.4 An Example of a Basic Run Without Conflicts

Figure 13 shows an example of a run of the interval script of “ $It$ ” described in section 3. In the first instant of the run  $t=0$ , we assume that the state of all intervals is  $F$ . Because “*user makes a pose*” is a sensor, i.e., without START/STOP functions, its PNF-state  $F$  is not expanded to  $NF$  during the temporal expansion process of stage I. Although the result of time expansion would allow some of the action

intervals to be also in the  $N$  state, the  $F$  state of “*user makes a pose*” constrains all other intervals to be in the  $F$  state. This is detected by the result of the PNF propagation process,  $P^1$ , which keeps the state of the system as it is.

In the next instant of time,  $t=1$ , the interval “*user makes a pose*” is detected, as reported by the  $N$  value coming from its STATE function. When PNF propagation is run, the fact that “*user makes a pose*” and “*camobject takes a picture*” should start together constrains the desired PNF-state ( $D^2$ ) of the latter to be  $N$ , generating a call for the START function of “*camobject moves front*”. Notice also that in this situation, the propagated PNF-state ( $P^2$ ) of the “*camobject clicks*” interval (as well as the state of “*camobject moves back*”) is  $NF$ , meaning that they may either be occurring or may occur in the future. However, they are not triggered because the thinning heuristics prevents change in intervals whose propagated PNF-state (in this case,  $NF$ ) contains the current state ( $F$ ).

In the next instant,  $t=2$ , we assume that “*camobject moves front*” successfully started and therefore no further action is required. The system remains in this state up to  $t=13$ .

When  $t=13$ , “*camobject moves front*” finishes and assumes the  $P$  state. Because of the meet constraint, “*camobject clicks*” should start immediately. Notice that the result of PNF propagation,  $P^{14}$ , shows exactly that configuration and the appropriate action is taken. In  $t=14$  the desired state is reached for the interval.

When “*camobject clicks*” finishes at  $t=20$  we have a similar situation where the START function of “*camobject moves back*” is called. However, this time we assume that, for some reason, the camera does not start moving back in the next instant of time,  $t=21$ . As shown in fig. 13, the discrepancy between the current state  $F$  ( $S^{21}$ ) and the desired state  $N$  ( $D^{22}$ ) persists, generating a second call for the START function of “*camobject moves back*”. At  $t=23$  the interval starts running and the situation is kept until the interval ends. At  $t=31$  the  $P$  state of “*camobject moves back*” is immediately detected by the STATE function of “*camobject takes a picture*”, and all the intervals go to the  $P$  state.

## 5.5 A Run With Conflicts

An example of a run with conflicts is show in fig. 14 for the same interval script. Here, we consider the case where the intervals “*camobject takes a picture*” and “*user is moving*” are defined as mutually exclusive and a conflict situation arises when both are expected to happen.

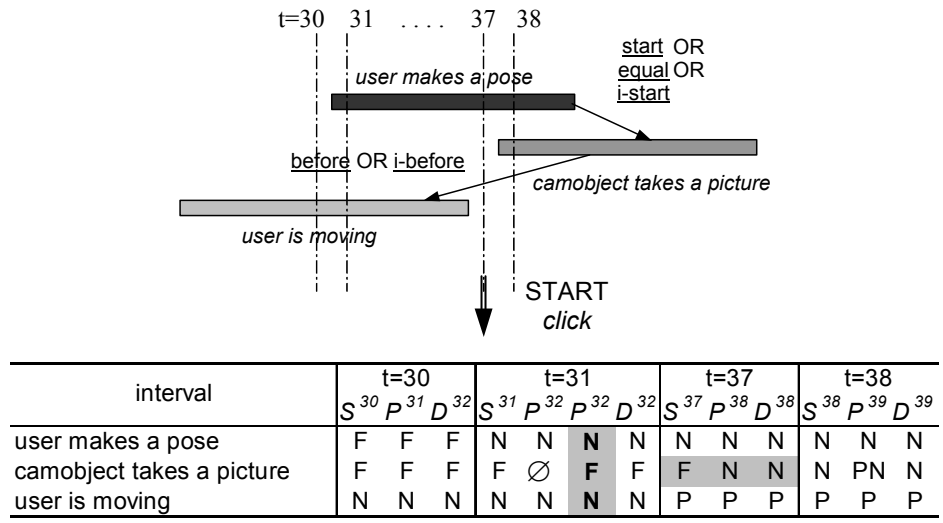


Figure 14. Example of a run with conflict.

At  $t=31$ , a conflict is detected by the PNF propagation stage since there is no way to satisfy the requirements that “camobject takes a picture” should start together with “user makes a pose” and that “camobject takes a picture” can not happen while “user is moving” is occurring. In this situation the first level of recovery succeeds and finds a set of states that is compatible with the constraints as shown in fig. 14. Basically, the system keeps the current situation as it is without taking any action. Later, when  $t=37$ , “user is moving” finishes, and since “user makes a pose” is still happening, “camobject takes a picture” is started.

As we see, the run-time engine of interval scripts was designed to avoid halts by searching for feasible states with the least amount of change. This strategy is not guaranteed to succeed but has proved to be robust enough to run quite complex structures. We are currently working on better methods to overcome contradictions such as keeping a history of previous states for backtracking purposes (such as those described in [13]) and devising mechanisms to detect and isolate the source of conflicts.

## 6 THE INTERVAL SCRIPT LANGUAGE (ISL)

The run-time engine described in the previous section has been implemented and used in the experiments described in section 7. For programming purposes, we have defined a text-based language called the *interval script language*, or *ISL*, that provides a simple way to describe an interval script. In the this section we show a programming example of ISL that also provides further insight on how to program an interactive environment with interval scripts.

```

"camobject moves front" =
{
START: execute [> camobject.MoveFront(1.0); <];
STOP:  execute [> camobject.StopMoving(); <];
STATE: set-state prf-expression
      [> (camobject.isMoving() ? _N_ : P_F) <];
}.

```

**Figure 15. Interval describing the movement of the “camobject” object.**

Although there are multiple ways to implement the concepts described in this paper, we chose to develop a text-based language containing the descriptions of the intervals and the temporal constraints between them. In our implementation, the file containing the language statements is run through an ISL compiler that produces a set of C++ files. The C++ files can be compiled using any standard compiler and executed at run-time through specially defined functions defined in another C++ file. It should be noticed that ISL is a particular implementation of the paradigm of interval scripts. As discussed later in section 8, it is possible to program with intervals and temporal constraints using, for instance, a graphical interface.

The complete grammar of ISL is described in the appendix. This section covers just the basic syntax and semantics of ISL, using as example the implementation of the interaction described in section 3. More details about the semantics of the language and other features of ISL can be found in [13].

### 6.1 Encapsulating Code in Intervals

We start our description of ISL by examining how the language communicates with low-level programming languages that provide basic functions and device access. In ISL, we simply allow the inclusion of C++ code directly into the script.

To begin our example, we show how to connect an interval in ISL to low-level C++ functions in the context of the example described in section 3. Figure 15 shows the definition of the interval “*camobject moves front*” in ISL where a camera-like CG object referred to as “*camobject*” moves to the front of the screen. The definition of the interval is comprised between curly brackets containing the definition of the interval’s basic functions. To include C++ code we use the command `execute` followed by the symbols “[>” and “<]”. For instance, when the `START` function is called, it executes the C++ code

```
"camobject moves front" = { ... }.
"camobject clicks" = { ... }.
"camobject moves back" = { ... }.
better-if "camobject moves front" meet "camobject clicks".
better-if "camobject clicks" meet "camobject moves back".
```

**Figure 16.** Interval describing the movement of the “camobject” object.

between those symbols, that is, a C++ method called “MoveFront” for the object “camobject” with parameter value 1.0. The STOP functions is defined in a similar way.

The definition of the STATE function is slightly different. In this case, the function is defined to set the state of the interval to be equal to the state returned by the C++ code inside the special symbols. In the case depicted in fig. 15, a method of the object “camobject” determines if the computer graphics camera is moving or not. If true, the state of the interval is set to *now*, referred in C++ code by the special constant “\_N\_”; otherwise, the state is set to be *past OR future*, symbolized by “P\_F”.

**6.2 Describing Constraints**

Using the example of section 3, we examine how constraints between intervals are defined in ISL. Figure 16 shows the script corresponding to the two constraints imposed over the three intervals “camobject moves front”, “camobject clicks”, and “camobject moves back”. The script first defines the intervals (here omitted for clarity), followed by two statements establishing the temporal constraints between the intervals. Notice the syntax better-if which was chosen to imply that this is a constraint that the system will try to enforce but that is not guaranteed to occur.

**6.3 Defining on Previous Intervals**

A key feature of ISL is the possibility of defining a new action or world state solely based on other intervals. With this we can create hierarchies, abstract concepts, and develop complex, high-level scripts.

Continuing with the example of section 3, we want to define the interaction between the user and the camera-like object as follows: when the user makes a pose (as detected by a computer vision system), the camera takes a picture, an action that is composed of moving the camera forward, closer to the user,

```

"user makes a pose" =
{
STATE:  set-state pnf-expression
        [> (camobject.isMoving())?_N_:P_F] <];
}.

"camobject takes a picture" =
{
STOP:   if current-state is NOW
        tryto start  "camobject moves back".
STATE:  if "camobject moves front" is NOW
        OR "camobject clicks" is NOW
        OR "camobject moves back" is NOW
        set-state NOW
        else if "camobject moves back" is PAST
            set-state PAST
        endif
        endif.
}.

better-if "user makes a pose" start OR equal OR i-start "camobject takes a picture".
better-if "camobject moves front" start "camobject takes a picture".
better-if "camobject clicks" during "camobject takes a picture".
better-if "camobject moves back" finish "camobject takes a picture".

"user is moving" = { ... }.

better-if "user is moving" before OR i-before "camobject takes a picture".

```

**Figure 17. Scripting based on previously defined intervals.**

clicking, and returning to its original position. Figure 17 shows the program corresponding to this interaction.

Initially the interval *"user makes a pose"* is defined as before, by a reference to C++ code that communicates with the vision system. Notice that this interval corresponds to a sensor and therefore has not START or STOP functions. To accomplish the clicking of the camera and the taking of the picture, the interval *"camobject takes a picture"* is defined, based on its three component intervals. First, a STOP function is defined so if the interval is happening, and it is made to stop, it first makes the camera-like object move back to its original position before actually stopping the main interval. Notice that, as showed in this example, stopping an interval may require actions to be executed and, therefore, some time is required before an interval actually stops. If the STOP function was not explicitly defined,

the interval still could be stopped, but in that case the interval “*camobject moves back*” would not be necessarily executed. In other words, the ISL language provides a great deal of extra detail to help in the programming task.

This is also the case of the STATE function that defines the computation of the current state of the interval “*camobject takes a picture*”. As shown in fig. 17, if any of the component actions is happening, the state of the interval is set to *now*. If the action of moving back the camera has occurred, that is, its current state is *past*, then the interval “*camobject takes a picture*” is determined to have already happened and its own state is set to *past*. Notice that, according to this definition, when the STOP function of the interval is called, the interval is not considered finished until the object completely moves back to its original position.

Finally, following that example, we want that whenever “*user makes a pose*” starts happening, the camera takes the picture. This is established by the constraint start OR equal OR i-start that forces the two intervals to start together. Also, we impose constraints that relate the interval “*camobject takes a picture*” to its three constituents. Notice that, it is not necessary to explicitly define a START function, since the constraint start between the intervals “*camobject moves front*” and “*camobject takes a picture*” establishes that the triggering of each interval immediately starts the other. And, completing the example of section 3, the interval “*user is moving*” is defined, and constraints are imposed that make it mutually exclusive to the interval “*camobject takes a picture*”.

The Interval Script Language includes many other features that have shown to be extremely helpful in the programming of the interactive environments described in section 7. See [13] for a detailed description of those language features.

## **7 WORKING WITH INTERVAL SCRIPTS**

Evaluating a programming method is always difficult. Although we have not performed formal studies on how easy is for programmers to understand and use the paradigm and the ISL language, we introduce here three systems built using interval scripts as evidence towards our belief that the language is an improvement over current paradigms and programming methods for interactive environments and agents. We do not see how these systems (especially the last two) could have been programmed using event-loops or state-machines in a reasonable amount of time and without major problems in debugging and maintenance. In fact, all the control structures of the systems described here were developed in very short periods of time. Based on these experiments, we believe that our main objectives when designing

intervals scripts were achieved, that is, that the paradigm provides expressiveness and simplicity beyond current programming methods.

The first two experiments with interval scripts, “*SingSong*” and “*It/I*”, are interactive theatrical performances where human and autonomous computer-graphics actors interact following a story. They constitute what is called *computer theater*, a term referring to live theatrical performances involving the active use of computers in the artistic process (see [58]). In these particular works, our focus was concentrated on building automatic, story-aware computer actors able to interact with human actors on camera-monitored stages. The third system, “*It*”, is an interactive art installation where interval scripts were used not only to program the environment but also its inhabitant computer characters and agents.

### 7.1 “*SingSong*”: a First Experiment

“*SingSong*” was our first experiment with interval scripts [16]. “*SingSong*” is composed of a large video screen that displays four computer graphics-animated characters that can “sing” musical notes (as produced by a MIDI synthesizer). A camera watches the user or performer determining the position of his/her head, hands, and feet. The body position is recovered by the software *pfinder* developed at the MIT Media Laboratory [59].

All the interaction is physical, non-verbal: the user or performer gestures and the CG-characters sing notes and move. “*SingSong*” was produced and performed in the summer of 1996. Although it is a short play of about 4 minutes, the interval script involves around 200 intervals, including a great deal of low-level control of I/O devices. It took about two days to write and debug the controlling program. Figure 18 shows a sequence of scenes from



Figure 18. Scenes from “*SingSong*”.

“*SingSong*”.

The paradigm of interval scripts was initially conceived to satisfy the scripting needs of “*SingSong*”. Although “*SingSong*” portrays a very simple, quasi-linear story, there is a large number of parallel actions and conditional interaction during the performance.

“*SingSong*” used the first version of the concept of interval scripts, described in [16]. The “*SingSong*” version of interval scripts already included the important idea of separating the desired action from the actual action. However, in that version, the idea was implemented using two implicitly defined intervals instead the simpler use of START, STOP, and STATE functions (see [16]). The use of two intervals required a more complicated model of the situation, making scripting more difficult than in the subsequent versions.

The main problem with the “*SingSong*” version of the interval scripts idea was that the script had to be written implicitly into C++ code. Although the interval script paradigm helped to avoid extremely complex and undebuggable control structures, after a while it became quite difficult to read the C++ file. This realization prompted for the development of the Interval Scripts Language introduced in the previous section.

## 7.2 “*It / I*”: A Computer Theater Play

Following “*SingSong*” we decided that it was necessary to develop a more comprehensive test for interval scripts. Particularly, we were interested in testing the robustness of the interval script paradigm in conditions where difficulties in scripting or run-time failures would be critical.

With these ideas in mind, one of the authors of this paper, Claudio Pinhanez, wrote the computer theater play “*It / I*”. The play is a pantomime where one of the characters, *It*, has a non-human body composed



Figure 19. Scenes from “*It / I*”. The computer graphics objects on the screen are autonomously controlled by the computer character *It*.

**Table 2. The composition of the interval scripts for the different scenes of “*It/I*”.**

	total number of nodes	total number of constraints*
scene I	120	88
scene II	80	70
scene III	92	66
scene IV	115	97

(\*) including implicit constraints (see [13]).

of CG-objects projected on screens (see fig. 19). The objects are used to play with the human character, *I*. *It* “speaks” through images and videos projected on the screens, through sound played on stage speakers, and through the stage lights.

The play is composed of four scenes, each being a variation of a basic cycle where *I* is lured by *It*, is played with, gets frustrated, quits, and is punished for quitting (see [13] for a detailed script of the play). It is important to note that any individual scene of “*It/I*” has a complexity level beyond of most previous full-body interactive systems such as [37, 40]. The play was produced in the summer/fall of 1997 and performed six times for a total audience of about 500 people. Each performance, lasting 40 minutes, was followed by an explanation of the workings of the computer-actor. After that members of the audience were invited to go up on stage and play the second scene from the play (see [13]).

There are major differences between the interval scripts of “*SingSong*” and “*It/I*”. First, the initial version employed C++ classes to interface the interval structure. Since we found that the resulting code ended up being difficult to read, we developed a language to describe intervals and their constraints, resulting in the first version of the Interval Script Language. By this time, we had also found that a better way to represent the distinction between wish and occurrence is the START, STOP, and STATE functions (as opposed to the double interval concept that founded the “*SingSong*” version). Using this structure, the triggering of the functions is accomplished simply by comparing the current and the desired state of all intervals, as explained in section 5.

For performance safety, we implemented each scene of “*It/I*” in a separate script, generating four different executable files that were loaded in the beginning of the corresponding scene. Table 2 summarizes the composition of the interval scripts of each scene of “*It/I*”, displaying the total number of intervals (i.e., the number of nodes of the PNF-network) and the number of constraints. As can be

seen in the table, we have more than 100 nodes per scene with an average of 80 constraints between them.

Although simple, the design of the run-time engine of the “*It / I*” version of interval scripts was not able to take full advantage of the prediction powers of constraint propagation as discussed in section 5. In particular, there were no recovery mechanisms and therefore constraints had to be carefully placed to avoid unexpected interactions that could, sometimes, stop all actions until a state change. This problem was solved in the following version of the run-time engine, where the idea of computing the desired state by PNF propagation and thinning (as described in section 5) considerably improved the robustness of the control, allowing, among others, the graceful survival of conflicts.

### 7.3 “*It*”: an Interactive Installation

After “*It / I*”, we developed an interactive art installation called “*It*” where users can experience the story of the play without having had previously watched it. In “*It*”, the user enters the area, is approached by the character *It*, starts playing with the CG-objects as *I* did in the original play, and like *I*, is not allowed to leave the space or stop playing. A major point in “*It*” is to make a pro-active CG-character that can drive an interactive experience with dramatic structure. The installation was finished in the spring of 1999 and demonstrated successfully during two open houses held at the MIT Media Laboratory (see fig. 20).

In “*It*” we employed the version of ISL introduced in section 6 and described in the appendix, and the run-time engine discussed in section 5. With such structures it was easy to describe the basic structure of an interactive story and to refine it by explicitly incorporating constraints to describe specific situations that must or must not occur, independently of the story sequence. Also, in “*It*” we used the ISL both to program the environment and its story and the computer agents that control the characters and agents that inhabit it. In fact, the control structure of “*It*” is distributed among four agents: the *story* agent, that controls the basic development of the environment; the *It* character and the *I* character agents, corresponding to the two main characters of the interactive story; and the *light designer* agent, that controls all the lighting effects of the environment.




Figure 20. Scenes from a run of “*It*”.

In table 3 we summarize the composition of the interval scripts used for the different agents of *It*. Table 3 shows that interval scripts can be used to construct and control large and complex collections of actions. For example, the *It* character agent employs 224 nodes for a total of more than 350 constraints. During run-time, we apply PNF propagation on this network at the rate of 20 Hz without any visible impact on the performance of the *SGI R10000* machine used in the installation. In fact, all four agents run simultaneously in one single machine, consuming at most 15% of the CPU time, including the communication overhead.

The major shortcoming of the current version of ISL, based on our experience with “*It*”, is the lack of a better visual interface for the programmer of the space. We envision as the next step of the development of interval scripts the implementation of a graphical user interface or a management system for the script files that could provide, for instance, visual manipulation of temporal constraints.

**Table 3. The composition of the interval scripts for the different agents in “It”.**

	number of nodes	total number of constraints
story	59	101
<i>It</i> character	224	355
<i>I</i> character	8	8
light designer	94	133

## 8 FUTURE DIRECTIONS

Although the interval script language has been fully implemented and intensively tested, there is still work to be done in many different directions. First, it would be interesting to get the system to a state where it could be released to other researchers and programmers interested in using interval scripts to build interactive systems. In this context it becomes possible to evaluate the difficulty in learning the language and which features different types of programmers would like to have added to the language.

An area we want to investigate is the best interface for the interval script programming paradigm: text-based or graphical. In the case of a graphical interface, a possible approach is to allow the user to construct temporal diagrams similar to fig. 4. However, since we allow disjunction of temporal constraints, there are always multiple possibilities for the graphical arrangement of the intervals and therefore the visualization may create more confusion than to provide help. Nevertheless, we have found very useful to draw temporal diagrams when programming in ISL.

We would also like to integrate into the interval script paradigm mechanisms that represent actions in forms that the machine can reason about. Currently, the run-time engine can only control the execution of scripts of but has no mechanisms to create action plans beyond the simple 1-step look-ahead performed in stage II. Comparing, for instance, to the *Hap* system developed by Loyall and Bates [48], the interval script run-time engine cannot plan a sequence of actions to achieve a goal in the future. In other words, the current interval script control system is basically reactive, and although it can run and control interactive plans with very complicated temporal structure, it has no way to construct plans from goals beyond the immediate one-step future. This limitation is partially due to the ideas employed to reduce the complexity of temporal constraint propagation as described in section 4.

We have considered different mechanisms to overcome this limitation. A basic idea is to perform multi-step look-ahead, by considering all possible different combinations of sensor input. In [13] we show that

it is possible to book keep multiple threads of sensor values, under some simplifying assumptions, because a majority of the generated threads become inconsistent and can be discarded after some cycles, due to contradiction with real sensor values. Since that work was focused on action recognition, this was done in the context of keeping multiple possibilities for the events in the past. Similarly, we could use multiple threads of sensor values to search the future for an activation pattern that achieve a determined goal expressed as a combination of PNF-states for the different intervals of a script.

To accomplish this, it is necessary to incorporate mechanisms that forecast the likely duration of actions and the likely delays in starting and stopping them, so the step-by-step PNF-style exploration of the future becomes feasible. Traditional plan mechanisms [48, 60] avoid dealing with duration issues by assuming strictly sequential sub-actions and states, and therefore they can search the future by ignoring the actual duration of the actions. In fact, those systems examine the state of the world exactly before or after the performance of an action, ignoring the possibility of other state changes while they are occurring.

A simple mechanism to incorporate the estimation of duration of actions in interval scripts would be to change the special temporal expansion in stage II of the run-time engine in the case of sensors to

$$V_i^{t+1} = \begin{cases} T_m(S_i^t) & \text{if the interval } i \text{ has START or STOP functions} \\ N(S_i^t, d_i, c) & \text{otherwise} \end{cases}$$

where we substitute  $S_i^t$  by the probabilistic function  $N(S_i^t, d_i, c)$  that estimates the PNF-state of node  $i$  given its current state  $S_i^t$  the length of time  $d_i$  that this node has been in this state, and the current context  $c$ . Using a function like this, it is possible to perform a multi-step look-ahead planning system that estimates the duration of the actions. Although it is likely that many of the threads representing sequences of future events result in inconsistent states as we observed in [13], there is no a-priori reason to expect the same degree of pruning in the prediction of the future. Therefore, to allow a reasonable depth in the search, it would be necessary to assume simplifying conditions for the sensors and actuators. The nature and actual pruning power of such conditions is a research topic that we plan to explore in the future.

## 9 CONCLUSION

In this paper we propose the interval script paradigm for programming interactive environments that is based on declarative programming of temporal constraints. Unlike previous constraint-based systems, we allowed the use a strong temporal algebra with mutually exclusive intervals. Among other things, we have created a paradigm that adequately handles the representation and control of some of the main issues in the programming of actual interactive environments: delays, continuity, and history of actions, and states; mutually exclusively actions; stopping of long actions; and recovery from conflicts and contradictions. Through examples we have shown that our paradigm allows good expressiveness and facilitates the management of context in an interactive environment.

This paper significantly extends the initial proposal of interval scripts outlined in [16]. Among other innovations, the run-time engine described in this paper is able to actively forecast future states that satisfy the temporal constraints, allowing the search for a consistent set of action triggerings. Also, the definition of the Interval Script Language greatly simplifies the use of temporal constraints in the programming process.

From the experience acquired in the use of the paradigm for the implementation of three different projects, we believe that programming systems incorporating the interval script paradigm can significantly ease the design and building of interactive systems in the future.

## ACKNOWLEDGEMENTS

*“SingSong”* was written, directed, and performed by Claudio Pinhanez. *“It / I”* was written and directed by Claudio Pinhanez, and produced by Aaron Bobick; the crew was composed by John Liu, Chris Bentzel, Raquel Coelho, Leslie Bondaryk, Freedom Baird, Richard Marcus, Monica Pinhanez, Nathalie van Bockstaele, Maria Redin, Alicia Volpicelli, Nick Feamster, and the actor Joshua Pritchard. *“It”* was designed and implemented by Claudio Pinhanez based on the elements produced for *“It / I”*.

Claudio Pinhanez was supported in different stages of this research by the scholarship from CNPq, process number 20.3117/89.1; by the DARPA contract DAAL01-97-K-0103; by the MIT Media Laboratory; and by the MIT Japan Program through the Starr Foundation. *“SingSong”* was sponsored by the Media Integration and Communication (MIC) laboratory of the Advanced Technology Research (ATR) laboratories in Kyoto, Japan. *“It / I”* and *“It”* were sponsored by the Digital Life Consortium of the MIT Media Laboratory.

## APPENDIX

### GRAMMAR OF THE INTERVAL SCRIPT LANGUAGE

The following is the list of production rules that define the Interval Scripts Language. This grammar expresses all the features of the language, including some that are not described in this paper; see [13] for the full semantics of ISL.

```
script:= list_of_declarations ;
list_of_declarations:= declaration '.' | list_of_declarations declaration
    '.' ;
declaration:= | cpp_dec | interval_dec | start_dec | stop_dec | state_dec |
    forget_dec | data_dec | condition | constraint | duration ;
cpp_dec:= DECLARE '['>' <cpluspluscode> '<';
interval_dec:= <name> '=' visible interval_specs;
visible:= | DISPLAY ;
interval_specs:= '{' list_of_declarations '}' | TIMER '(' number ',' number
    ')' | CYCLE interval | SEQUENCE list_of_interval ;
list_of_intervals:= interval | list_of_intervals ',' interval ;
interval:= <name> ;
condition:= WHEN expression TRYTO list_of_simple_execs ;
list_of_simple_execs:= simple_exec | list_of_simple_execs ';' simple_exec ;
simple_exec:= START list_of_intervals | STOP list_of_intervals | FORGET
    list_of_intervals ;
list_of_execs:= exec | list_of_execs ';' exec ;
exec:= START list_of_intervals | STOP list_of_intervals | FORGET
    list_of_intervals | SETSTATE pnfexpression | SETDURATION '(' number ','
    number ')' | EXECUTE CCODE | IF expression list_of_execs ifelse ENDIF ;
ifelse:= | ELSE list_of_execs ;
constraint:= BETTER-IF list_of_intervals time_relation list_of_intervals ;
time_relation:= time_relation_tag
    | time_relation OR time_relation_tag ;
```

```

time_relation_tag:= EQUAL | MEET | I-MEET | BEFORE | I-BEFORE | AFTER |
    DURING | I-DURING | OVERLAP | I-OVERLAP | SSTART | I-START | FINISH | I-
    FINISH | ONLY-DURING | NOT-DURING | DIFFERENT | PRECEDE | I-PRECEDE ;
start_dec:= START ':' list_of_execs ;
stop_dec:= STOP ':' list_of_execs ;
forget_dec:= FORGET ':' list_of_execs ;
data_dec:= DATA ':' '[' <code> '<' ;
state_dec:= STATE ':' list_of_execs ;
pnfexpression:= timetag | interval | CURRENTSTATE | EXPANDEDSTATE;
timetag:= PAST | NOW | FUTURE | UNDETERMINED | PAST-OR-NOW | PAST-OR-FUTURE
    | NOW-OR-FUTURE | P__ | _N_ | __F | PNF | PN_ | _NF | P_F ;
expression:= expression AND or_expression | or_expression ;
or_expression:= or_expression OR primary | primary;
primary:= pnfexpression | pnfexpression IS timetag | ISCONTAINED
    pnfexpression timetag | TRUE | '(' expression ')' | NOT primary |
    BOOLEXPRESSION '[' <code> '<' ;
duration:= DURATION '(' number ',' number ')' ;
number:= <number> | '[' <code> '<' ;

```

## REFERENCES

1. Pausch, R., et al. *Disney's Alladin: First Steps Toward Storytelling in Virtual Reality*. in *Proc. of SIGGRAPH'96*. 1996.
2. Bobick, A., et al., *The KidsRoom: A Perceptually-Based Interactive Immersive Story Environment*. PRESENCE: Teleoperators and Virtual Environments, 1999. **8**(4): p. 367-391.
3. Tosa, N., *Expression of Emotion, Unconsciousness with Art and Technology*, in *Affective Minds*, G. Hatano, N. Okada, and H. Tanabe, Editors. 2000, Elsevier. p. 183-205.
4. Pinhanez, C., et al., *Physically Interactive Story Environments*. IBM Systems Journal, 2000. **39**(3&4): p. 438-455.
5. Lai, J., et al., *BlueSpace: Personalizing Workspace through Awareness and Adaptability*. To appear in *Communications of the ACM*, 2002.
6. Crowley, J.L., J. Coutaz, and F. Berard, *Things that See*. *Communications of the ACM*, 2000. **43**(3): p. 54-64.

7. Raskar, R., et al. *The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays*. in *Proc. of SIGGRAPH'98*. 1998. Orlando, Florida.
8. Abowd, G.D., *Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment*. IBM Systems Journal, 1999. **38**(4): p. 508-530.
9. Mynatt, E.D., I. Essa, and W. Rogers. *Increasing the Opportunities for Aging in Place*. in *Proc. of ACM Conference on Universal Usability, CUU'00*. 2000. Arlington, Virginia.
10. Pentland, A., *Smart Rooms*. Scientific American, 1996. **274**(4): p. 68-76.
11. Shaw, C., et al., *Decoupled Simulation in Virtual Reality with the MR Toolkit*. ACM Transactions on Information Systems, 1993. **11**(3): p. 287-317.
12. Pausch, R., et al., *A Brief Architectural Overview of Alice, a Rapid Prototyping System for Virtual Reality*. IEEE Computer Graphics and Applications, 1995.
13. Pinhanez, C.S., *Representation and Recognition of Action in Interactive Spaces*, in *Media Arts and Sciences Program*. 1999, Massachusetts Institute of Technology.
14. Allen, J.F., *Maintaining Knowledge about Temporal Intervals*. Communications of the ACM, 1983. **26**(11): p. 832-843.
15. Vilain, M. and H. Kautz. *Constraint Propagation Algorithms for Temporal Reasoning*. in *Proc. of AAAI'86*. 1986. Philadelphia, Pennsylvania.
16. Pinhanez, C.S., K. Mase, and A.F. Bobick. *Interval Scripts: A Design Paradigm for Story-Based Interactive Systems*. in *Proc. of CHI'97*. 1997. Atlanta, Georgia.
17. Coen, M.H. *Building Brains for Rooms: Designing Distributed Software Agents*. in *Proc. of IAAI'97*. 1997. Providence, Connecticut.
18. *Director's User Manual*. 1990: MacroMind Inc.
19. *MAX Getting Started Manual*. 2002: Opcode.
20. Roads, C., *The Computer Music Tutorial*. 1996, The MIT Press: Cambridge, Massachusetts. p. 969-1016.
21. Rowe, R., *Interactive Music Systems*. 1993, Cambridge, Massachusetts: The MIT Press.
22. Rollings, A. and D. Morris, *Game Architecture and Design*. 2000, Scottsdale, Arizona: Coriolis Group.
23. Buchanan, M.C. and P.T. Zellweger. *Automatic Temporal Layout Mechanisms*. in *Proc. of ACM Multimedia'93*. 1993. Ahaheim, California.
24. Hamakawa, R. and J. Rekimoto. *Object Composition and Playback Models for Handling Multimedia Data*. in *Proc. of ACM Multimedia'93*. 1993. Ahaheim, California.

25. van Rossum, G., et al. *CMIFed: a Presentation Environment for Portable Hypermedia Documents*. in *Proc. of ACM Multimedia'93*. 1993. California.
26. Vazirgiannis, M., I. Kostalas, and T. Sellis, *Specifying and Authoring Multimedia Scenarios*. IEEE Multimedia, 1999. **6**(3): p. 24-37.
27. Bailey, B., et al. *Nsync - A Toolkit for Building Interactive Multimedia Presentations*. in *Proc. of ACM Multimedia'98*. 1998. Bristol, England.
28. Jourdan, M., et al. *Madeus, an Authoring Environment for Interactive Multimedia Documents*. in *Proc. of ACM Multimedia'98*. 1998. Bristol, England.
29. Selcuk, K.C., B. Prabhakaran, and V.S. Subrahmanian. *CHIMP: A Framework for Supporting Distributed Multimedia Document Authoring and Presentation*. in *Proc. of ACM Multimedia'96*. 1996. Boston, Massachusetts.
30. Agamapolis, S. and V.M. Bove Jr., *Multilevel Scripting for Responsive Multimedia*. IEEE Multimedia, 1997. **4**(4): p. 40-50.
31. Wirag, S. *Modeling of Adaptable Multimedia Documents*. in *Proc. of the European Workshop on Interactive Distributed Multimedia Systems and Telecommunications Services*. 1997. Darmstadt.
32. Andre, E. and T. Rist. *Coping with Temporal Constraints in Multimedia Presentation Planning*. in *Proc. of AAAI'96*. 1996. Portland, Oregon.
33. Kautz, H.A. and P.B. Ladkin. *Integrating Metric and Qualitative Temporal Reasoning*. in *Proc. of AAAI'91*. 1991.
34. Borning, A., *The Programming Language Aspects of ThingLab, A Constraint-Oriented Simulation Laboratory*. ACM Transactions on Programming Languages and Systems, 1981. **3**(4).
35. Elliott, C., et al. *TBAG: A High Level Framework for Interactive, Animated 3D Graphics Applications*. in *Proc. of SIGGRAPH'94*. 1994. Orlando, Florida.
36. Kakizaki, K.I. *Generating the Animation of a 3D Agent from Explanation Text*. in *Proc. of ACM Multimedia'98*. 1998. Bristol, England.
37. Krueger, M.W., *Artificial Reality II*. 1990: Addison-Wesley.
38. Bederson, B.B. and A. Druin, *Computer Augmented Environments: New Places to Learn, Work and Play*. 1995, Ablex: Norwood, New Jersey.
39. Dey, A.K., D. Salber, and G.D. Abowd, *A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications*. Human-Computer Interaction Journal, 2001. **16**(1).
40. Maes, P., et al. *The ALIVE System: Full-Body Interaction with Autonomous Agents*. in *Proc. of the Computer Animation'95 Conference*. 1995. Geneva, Switzerland.

41. Johnson, M., et al. *Sympathetic Interfaces: Using a Plush Toy to Direct Synthetic Characters*. in *Proc. of CHI'99*. 1999. Pittsburgh, Pennsylvania.
42. Blumberg, B.M. and T.A. Galyean. *Multi-Level Direction of Autonomous Agents for Real-Time Virtual Environments*. in *Proc. of SIGGRAPH'95*. 1995.
43. Perlin, K., *Real Time Responsive Animation with Personality*. IEEE Transactions on Visualization and Computer Graphics, 1995. **1**(1): p. 5-15.
44. Cassell, J., et al. *Animated Conversation: Rule-Based Generation of Facial Expression, Gesture & Spoken Intonation for Multiple Conversational Agents*. in *Proc. of SIGGRAPH'94*. 1994. Orlando, Florida.
45. Thalmann, N.M. and D. Thalmann, *Synthetic Actors in Computer Generated 3D Films*. 1990, Berlin, Germany: Springer-Verlag.
46. Perlin, K. and A. Goldberg. *Improv: A System for Scripting Interactive Actors in Virtual Worlds*. in *Proc. of SIGGRAPH'96*. 1996.
47. Bates, J., A.B. Loyall, and W.S. Reilly. *An Architecture for Action, Emotion, and Social Behavior*. in *Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World*. 1992. S. Martino al Cimino, Italy.
48. Loyall, A.B. and J. Bates, *Hap: A Reactive, Adaptive Architecture for Agents*. 1991, Carnegie Mellon University.
49. Brooks, R., *Intelligence without Reason*. Artificial Intelligence, 1991. **47**: p. 139-159.
50. Tosa, N., et al. *Network-Based Neuro-Baby with Robotic Hand*. in *Proc. of IJCAI'95 Workshop on Entertainment and AI/Alife*. 1995. Montreal, Canada.
51. Funge, J., X. Tu, and D. Terzopoulos. *Cognitive Modeling: Knowledge, Reasoning and Planning for Intelligent Characters*. in *Proc. of SIGGRAPH'99*. 1999. Los Angeles, California.
52. McCarthy, J. and P. Hayes, *Some Philosophical Problems from the Standpoint of Artificial Intelligence*, in *Machine Intelligence 4*, B. Meltzer, D. Michie, and M. Swann, Editors. 1969, Edinburgh University Press: Edinburgh, Scotland. p. 463-502.
53. Pinhanez, C.S. and A.F. Bobick. *Human Action Detection Using PNF Propagation of Temporal Constraints*. in *Proc. of CVPR'98*. 1998. Santa Barbara, California.
54. Meiri, I. *Combining Qualitative and Quantitative Constraints in Temporal Reasoning*. in *Proc. of AAAI'91*. 1991.
55. Pinhanez, C.S. and A.F. Bobick. *PNF Propagation and the Detection of Actions Described by Temporal Intervals*. in *Proc. of the DARPA Image Understanding Workshop*. 1997. New Orleans, Louisiana.
56. Dechter, R., *From Local to Global Consistency*. Artificial Intelligence, 1992. **55**(1): p. 87-107.

57. Mackworth, A.K., *Consistency in Networks of Relations*. Artificial Intelligence, 1977. **8**(1): p. 99-118.
58. Pinhanez, C.S. *Computer Theater*. in *Proc. of the Eighth International Symposium on Electronic Arts (ISEA'97)*. 1997. Chicago, Illinois.
59. Wren, C., et al., *Pfinder: Real-Time Tracking of the Human Body*. IEEE Trans. Pattern Analysis and Machine Intelligence, 1997. **19**(7): p. 780-785.
60. Nilsson, N.J., *Problem-Solving Methods in Artificial Intelligence*. 1971, New York, New York: McGraw-Hill.