

Temporal Trajectories in Shared Interactive Narratives

ABSTRACT

Temporal trajectories can represent the complex mappings between story time and clock time that are to be found in shared interactive narratives such as computer games and interactive performances. There are three kinds. Canonical trajectories express an author's intended mapping of story time onto clock time as part of the plot and schedule of an experience. Participant trajectories reflect a participant's actual journey through story time and clock time as they interact with the experience. Historic trajectories represent the subsequent selection and reuse of segments of recorded participant trajectories to create histories of past events. We show how temporal trajectories help us analyse the nature of time in existing experiences and can also generate new approaches to dealing with temporal issues such as: disengagement and reengagement, adapting to different paces of interaction, synchronising different participants, and enabling encounters and travel across time.

Author Keywords

Time, narrative, CSCW, synchronization, games, mobile applications, performance, schedule, plot, story

ACM Classification Keywords

H.5.3 [Information Systems] Group and Organization Interfaces – *Collaborative Computing*.

INTRODUCTION

The relationship between time and interaction has been a longstanding concern within HCI, spanning wide ranging discussions of responsiveness, pace and interaction [5], revealing and communicating delays [18], synchronizing multiple users' interactions [14,19], and visualizing, browsing and synchronizing convergent and divergent histories of interaction [8,9,17]. This paper extends HCI's concern with time to address the challenges raised by a new generation of *narrative-driven* experiences such as

computer games, artistic performances and hypermedia stories. A distinctive feature of such experiences is that an author first creates an underlying narrative, a guiding temporal structure for events that drives the experience. Participants subsequently interact with this structure at different times according to their desire and availability.

Single player computer games typically provide players with very flexible control over time and pacing. They can pause and resume the game's narrative at their convenience and can move backwards and forwards in game time, replaying previous levels. Game time may also be flexibly sped up or slowed down as in simulation games where players pause the game to make changes to the simulation and then rapidly advance time to see them played out. In contrast, Massively Multiplayer Online Role Playing Games (MMPORGS) tend to adopt a more seemingly objective approach to time. Here, groups of players share a common narrative timeline that progresses linearly in relation to clock time, often advancing even when they are not playing. This gives each individual far less control over the progression of the narrative but enables players to share a social experience. An emerging generation of mobile experiences further complicates this picture as participants rapidly dip in and out of an experience while on the move, engaging in the 'downtime' between other activities, being subject to frequent disconnections and reconnections, and fitting the experience to the patterns of their daily lives [2].

A further interesting characteristic of time in narrative-driven experiences is the participants' desire to experience or tell stories of what happened to them in the past, replaying highlights, reviewing histories, drawing on recordings of their interactions to make animations and videos, and even mixing live experiences with recordings of previous ones (e.g., the Gran Turismo racing game in which players can race against previous versions of themselves).

In short, designing the temporal structure of a shared interactive narrative is a complex business that raises new and difficult challenges for HCI. How can we negotiate the pace and timing of interaction between authors who wish to impose a driving narrative and participants who wish to experience this at times that suit them? How can participants fit an ongoing long-term narrative into the patterns of their daily lives, flexibly engaging and disengaging as required? How can we extend this negotiation to reflect the needs of multiple participants in a

COPYRIGHT NOTICE

shared experience? Finally, how can participants access recordings of past interactions to create new stories or to weave them in with ongoing ones?

This paper proposes a conceptual framework for answering these questions. This framework introduces the new concept of *temporal trajectories*, which is further specialized into *canonical*, *participant*, and *historic* trajectories, as a mechanism for reasoning about time and interactivity in shared narratives. This framework is intended to help researchers and designers analyse existing experiences, relate them to previous work in the HCI and narrative literature, and also to envisage some radical new approaches to the treatment of time.

The framework has emerged from our attempts to analyse the nature and experience of time in our own mobile game called Day of the Figurines (DoF) that has its own distinctive temporal structure. However, it soon became apparent that the framework could also help us envisage ways in which DoF might be extended and redeveloped. This paper introduces the framework through these two stages, using DoF to illustrate its key concepts while also relating them to other examples and to the wider literature. We therefore begin with a short introduction to DoF.

DAY OF THE FIGURINES

Day of the Figurines is a narrative-driven text messaging game for mobile phones [10]. Players send and receive SMS to control a ‘character’ – their figurine – as it lives through a day in the life of a fictional town, visiting destinations, observing events, using objects, responding to dilemmas, undertaking missions and chatting with others.

The design of time in DoF

DoF was consciously designed to unfold in the background of players’ daily lives, reflecting the way in which mobile phones enable us to finely interweave many different threads of activity. It was also designed to accommodate the distinctive nature of SMS (slow, costly and infrequent messages) by requiring players to send and receive only a few messages each day as part of a slow game. As a result, DoF adopted a distinctive treatment of narrative and time.

DoF balances pre-scripted narrative with interactivity. The game is fundamentally narrative driven, following a pre-scripted storyline. Players are refugees who are dropped off in the town in the early morning. As the day unfolds they experience a sequence of scheduled events including a fete at the recreation ground, two dead lovers being found at the cemetery, a riotous gig at the Locarno nightclub, and an army sweeping into town. These scripted events are interspersed with interactive elements such as multiple-choice dilemmas and missions that require players to visit destinations, find and use objects, and maintain their health. Each element has a temporal scope that constrains when in game time it becomes available to players. It may also have a timeout after which the game server assumes a default answer or action so as to keep the narrative moving on.

Turning to its treatment of time, DoF deliberately slows down fictional time so that the twenty four hours of time in the narrative are played out over twenty four days of the players’ real lives. The authors intend for all players to share the same objective game time, which unfolds linearly in relation to actual time no matter how much or little each individual plays. Indeed, a player’s figurine remains active in the game even when their phone is switched off or loses its network connection, with the player receiving SMS notifications of any missed events the next time they reconnect to the game. Finally, DoF is delivered as an event-based touring artistic performance, being booked to run at hosting venues for fixed periods of time. Each performance runs for twenty four days, between set start and end dates for ten hours a day (while the venue is open), with the game being suspended outside of these times.

The experience of time in DoF

To date, DoF has been performed in Berlin, Singapore and three times in the UK, being experienced by over 750 players. Overall, the game has been well received; over 70% of the 100 players who responded to our post-game questionnaire said that they would pay to play again. However, the experience of playing the game did reveal some interesting issues with regard to its temporal structure.

First, the majority of participants played episodically, dipping in and out of the game, sometimes not actively engaging for several days before becoming active again. Over half of respondents to the questionnaire reported that they played ‘occasionally’ as opposed to ‘regularly’ or ‘seldomly’. Patterns of engagement also varied greatly among players due to variations in their preferences for where and when to play, different patterns of phone use, and shifting personal circumstances over nearly a month of play. Several observed that they would like to be able to explicitly suspend and resume the game.

Messages could be delayed for hours before being delivered to players due to network congestion, lack of coverage, or phones being switched off. Some players were irritated by the sudden flood of messages that could arrive as they switched their phone back on after a long break. In extreme cases, these messages might even inform them they had died while their phone had been unavailable, for example if another player had repeatedly hit them with a weapon. Some players reported frustration with playing across time zones as this led to a mismatch between the game’s scheduled opening hours and their own waking hours. For example, UK players in the Singapore game would tend miss the first half of each day’s play and consequently suddenly become very active towards the end. This combination of delays, episodic engagement and multiple time-zones made it difficult for players to engage in conversation and maintain social relationships. A common complaint was that of being ignored by other players.

The game’s orchestration interface enabled its operators to control the rate of progression of time in the fictional

narrative relative to time in the real world. This was used once in Singapore when the hosting venue was closed for a day. In response, the game’s operators speeded up game time over the following few days until the missing fictional hour had been recovered.

Finally, players enjoyed reviewing their histories of play in order to plan what to do next or just for the pleasure of reflection. This could be achieved by scrolling through saved messages on their phones. However, many players lacked sufficient storage for their entire history and so ended up reviewing incomplete histories, typically consisting of a combination of recently received messages and more ‘important’ older messages that they had saved. Alternatively, players could review their whole history of interaction on the web, but only after the performance had finished. Even then, they reported wanting to be able to compare their own history with those of others, which was not possible.

USING TEMPORAL TRAJECTORIES TO UNDERSTAND THE DESIGN AND EXPERIENCE OF TIME IN DOF

We now introduce the concept of **temporal trajectories** as a way of analyzing the design and experience of time in DoF and other narrative-driven experiences. In general, a temporal trajectory expresses a mapping or path between fictional time in an underlying story universe which we refer to as **story time (ST)**, and the actual time at which this could be or actually is experienced by participants, which we call **clock time (CT)**. We need to clarify these key terms a little more before moving on.

We adopt the term story time from existing accounts of time and narrative in drama and literature [16]. Story time is determined by the author of the story and describes the time span defining the narrative outline of the story, in chronological order. In other words, story time refers to the span, timing and order of the underlying events in the implied story universe. Important characteristics of story time discussed in literary and dramatic theory include the historic epoch in which the story is set and its subdivision into primary, secondary and tertiary time (see [16]). Games introduce their own distinct temporal structures, including the distinction between ‘result time’ (games played until someone wins) versus ‘set time’ (games played for a fixed period of time) [20]. For example, in the case of DoF, story time encapsulates the life of a contemporary town over a period of twenty four fictional hours, and the experience operates according to an overall set time.

We chose the term clock time to express the idea of time as measured by a clock in the real world. We might have chosen the term ‘real time’ instead, but in some quarters this term is taken to refer to time as actually *experienced* by a participant in the real world, which is a subtly different idea. For example, participants may experience the passage of time differently when they enter a ‘flow state’ during the playing of a game or similar intense pleasurable experience [3]. In a similar vein, inspired by Husserl’s

phenomenology, Francisco Varela talks about experienced time as having a three-part structure based on now, retention and protention [22]. Retention is described as belonging to the past even though it is happening now, whereas protention is ‘the expectation or the construction of the future’. Whereas ‘flow’ indicates a perception of time experienced as duration, the ‘now, retention and protention’ structure point to the possibility of an experience based on tenses or trajectories. Previous HCI research has also highlighted key factors that may affect our experience of time including psycho-cognitive factors (reaction and co-ordination times), the duration of short-term memory, and the influence of circadian and other natural rhythms [6].

Having introduced these core terms, we now return to the definition of three kinds of temporal trajectory, beginning with canonical trajectories.

Canonical trajectories express the author’s intent

As noted previously, a distinctive feature of collaborative games, performances and stories when compared to other collaborative applications such as shared editors and drawing tools is the presence of an underlying narrative that has been created by an author in order to drive the experience along one or more pre-scripted paths. Such narratives invariably involve imposing an ordering, schedule and deadlines on events.

This intended timescale of interaction is expressed by a **canonical trajectory**, an authoritative mapping between ST and CT that is intended to steer participants through the temporal structure of the experience. A canonical trajectory represents a guideline which participants will generally follow, although they may occasionally diverge from it due to delay, disconnection or perhaps through choice. Figure 1 shows the general canonical trajectory underlying DoF that defines its linear mapping between twenty four fictional hours of ST onto twenty four actual days of CT.




Figure 1: the core canonical trajectory underlying DoF

In general, ST might advance at variable rates relative to CT, speeding up or slowing down as we see in many computer games. It might also jump up and down the vertical axis, representing (fictional) time travel in the story

universe, for example flashbacks or revisiting previous levels in a game. ST could even progress downwards vertically so that it unfolds in reverse as CT advances. Canonical trajectories might also include branching narrative structures [15] that reflect interactive decisions, for example expressing the idea of moving to the next level or repeating the current level according to some success criterion in a game. However, *it is not possible for a canonical trajectory to jump or move backwards in CT* – time travel is not possible in the real world and so CT always moves forwards.

In terms of narrative theory, canonical trajectories represent **plot time**, the temporal structure of the narration of the story, i.e., the timing and ordering of events in their representation. This need not be the same as their timing and order in the story universe and hence plot time is not the same as story time. For example, plays may adopt so-called ‘open’ structures compressing story time into a shorter plot time which then refers to events outside of plot time that influence the course of the plot’s events, while others assume ‘closed’ structures in which the story is self-contained, with no background events influencing the beginning, and the ending being final [16]. In DoF, unconventionally, plot time expands story time, mapping twenty four hours onto twenty four days. Plots may contain subplots defined through different temporal constructions. Film and television narrative structures also compress time (‘ellipsis’) and alter the ordering of events as they are narrated (e.g., ‘flashbacks’) [1]. In their turn, interactive media have introduced their own distinctive structures of plot time in the form of multi-threaded hypertext plots [11,15] and the looping structures of computer games [13] which may be combined with filmic elements.

Adopting film terminology, the structure of plot time is determined by the director, although they may be the same person as the author or even the performer in the case of performance practice. The more interactive the work, the more the participant will also be able to author their own plot, determining its temporal construction and influencing its perception by other players.

Canonical trajectories also have a technical meaning as they reflect one of the most common approaches to realizing collaborative and distributed experiences, the client-server model. Here, a central server maintains the agreed state of an experience (including its canonical trajectory) which is replicated at different clients that may maintain their own local versions that might diverge and re-converge due to network delays and other factors as discussed below.

Canonical trajectories express performance schedules

Although figure 1 captures the author’s overall intent for DoF, its actual canonical trajectory is more complex than this due to the need to schedule multiple performances of the experience to take place at particular times on particular days. Figure 2 therefore shows two more refined canonical trajectories that express the detailed scheduling of two

distinct performances of DoF. Each consists of twenty four segments (only five of these are actually shown in the figure), where each segment specifies how one hour of ST is intended to unfold over ten hours of a particular day of CT. Each distinct performance of DoF is represented by its own canonical trajectory, potentially defining its own local schedule of opening hours and days. These canonical trajectories will overlap in CT if multiple performances ever occur simultaneously.

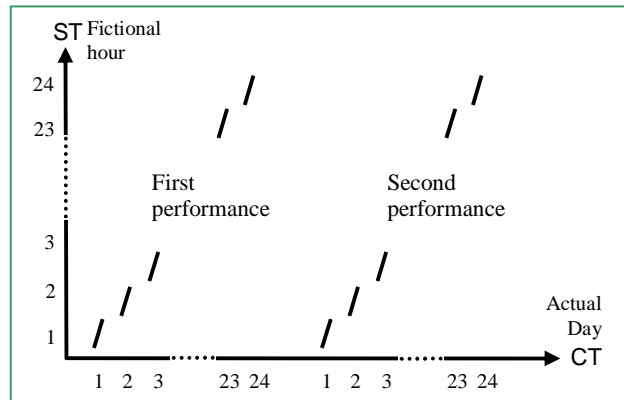


Figure 2: Canonical trajectories express scheduling

In general then, a canonical trajectory also expresses the **schedule time** at which a narration is made available to participants, be they readers, viewers or players. All media involve a schedule time (books are published and plays and films are shown at set times), but it is in television that schedule time, through its relationship to channels and advertising sponsorship, takes on a particular significance to the point where it strongly influences the form of the underlying story and plot, as seen by the rise of TV series and serials [1]. Schedule time is controlled by the scheduler or publisher, who may or may not be separate from the author and/or director.

We saw in the Singapore performance of DoF how the game’s operators adjusted the rate of passage of ST relative to CT in order to make up for a lost day of performance. This shows that canonical trajectories need not only be defined in advance of a performance; they might also be manipulated ‘on the fly’ during the performance in order to adapt to the contingencies of changing schedules.

Multiple canonical trajectories can express complex nested temporal structures

Multiple canonical trajectories can also represent the more complex nested temporal structures of some Massively MMPORGS such as Ultima Online [21]. At the core of Ultima is a shared virtual world that provides a space for social interaction with time unfolding linearly in relation to clock time. However, Ultima also includes additional narrative-driven elements in the form of quests that must be completed to particular timescales and which are initiated and shared by groups of players. Quests are then associated with their own localized canonical trajectories that serve to

group together a sub-set of players for a limited duration of CT. Interestingly, these are also created – or scheduled – on the fly in response to players’ actions within the game. In terms of narrative theory, this is an example of multiple canonical trajectories representing plots and sub-plots.

Participant trajectories capture individual experience

So far we have considered the intention of the author (or director or scheduler) in defining the temporal structure of an interactive narrative experience. We now consider how participants actually interact with such an experience.

We propose that each participant in an interactive narrative follows their own **participant trajectory** that captures the temporal nature of their individual experience, relating where they actually are in ST as CT unfolds. Whereas a canonical trajectory describes where an author wants the participant to go (in time), a participant trajectory describes where they actually do go. While we might expect participant trajectories to generally follow canonical trajectories, they may occasionally diverge from them, either as a matter of choice or due to technical constraints such as network disconnections or delays. Participant trajectories therefore represent **interaction time**, the times at which participants engage with the story once it has been made available to them. Interaction time is primarily determined by the player and previous HCI research has produced various accounts of factors that shape interaction time in narrative experiences for example studies of how players coordinate their engagement in MMPORGS [7] and how their engagement with long-term mobile games adapts to the patterns of their daily lives, for example commuting [2]. As a concrete example, interaction time in DoF is characterized by episodic patterns of play, with players frequently disengaging and reengaging.

So far, we have introduced canonical and participant trajectories to express the mapping of story time onto clock time, capturing key features of the intermediate layers of plot time, schedule time and interaction time. We now begin to put these concepts to work, initially to help us analyse an existing interactive narrative experience – DoF.

Understanding disengagement and reengagement

We begin by analysing what happens when a participant disengages from and subsequently reengages with an ongoing experience, for example by losing and then regaining their network connection to a remote server that is maintaining the canonical trajectory for their experience.

In DoF, story time appears to freeze for the participant as they will no longer receive information from the server. However, the server actually continues to operate their character in their absence (albeit it in a passive mode) so that others can still see it, speak to it, and use objects on it, although it will not respond. This explains why participants often feel ignored by others (who may be disengaged) and why a few experience the more severe problem of ‘waking up dead’ where others have killed them in their absence.

Figure 3 uses temporal trajectories to represent DoF’s approach to disengagement. Initially the participant trajectory follows the canonical trajectory until the point of disengagement when we show it splitting into two threads. One of these represents the server actively maintaining the character along the canonical trajectory (shown by the solid blue line). The other represents the state of the participant’s phone which initially receives no updates and so its local story time is frozen (the dashed blue line). It then receives a series of delayed messages from the server in quick succession as it reconnects, resulting in its local story time rapidly catching up with the canonical story time (the dotted blue line)

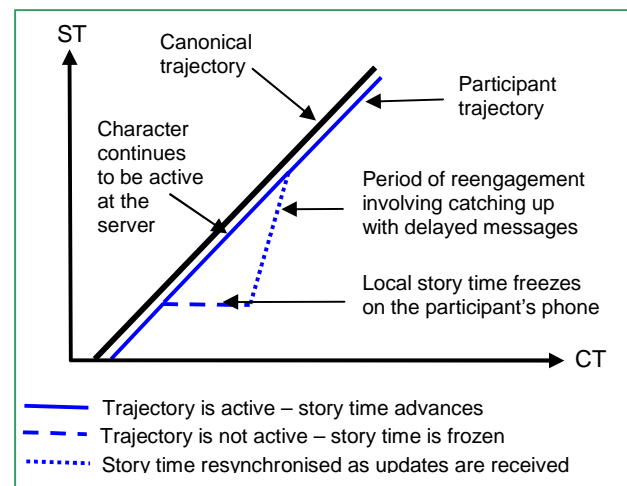


Figure 3: Disengagement and reengagement in DoF

Some participants experienced floods of messages on reengagement. The slope of the dotted line is important in this regard. If it is steep (i.e., delayed messages are delivered quickly on reengagement) then they experience a flood of messages. On the other hand, if it is shallow (they receive the delayed messages gradually) then there is a risk that they will try to respond to earlier messages before they see any later ones, trying to interact with the experience as if they were at an earlier point in canonical time than they actually are, which could cause confusion.

Alternative approaches to managing disengagement

Further inspection of Figure 3 reveals other approaches to managing temporary disengagement. In general, a disengagement and subsequent reengagement creates a triangular area of uncertainty, a ‘Bermuda Triangle’ if you like, in our trajectory diagram as the participant trajectory first becomes separated from the canonical trajectory and then subsequently rejoins it. During this period, there is inconsistency between the participant’s notion of story time and that of the central server (and hence other participants), potentially leading to confusion and the need to resynchronise. Figure 4 shows how DoF’s approach is just one of four general ‘routes’ by which this triangle of uncertainty can be negotiated.

In route (a) the participant is totally removed from the game for the duration of their disengagement. They no longer interact with it on their phone and their character disappears from the server. This avoids floods of update messages on reengagement as well as the potentially ‘nasty surprises’ that can be experienced in DoF, although at the cost of the participant missing a chunk of the narrative. This requires a technical mechanism to detect disengagement which is not reliably possible with SMS, although can be achieved with many other communication technologies. This approach is typical of MMPOGRS such as World of Warcraft.




Figure 4: four ‘routes’ around the ‘Bermuda triangle’

Route (b) is the current version of DoF that we have already discussed above. Route (c) represents the case where the participant continues to interact with a local version of their character that runs on their phone but that is no longer visible to other participants via the server. The participant controls an independent participant trajectory with its own local story time which only gets resynchronised with the canonical trajectory as they reengage, at which point they become visible to other players again. This requires that a version of the game can be run independently on the participant’s local device, which was not the case in DoF.

Finally, route (d) represents that case where both the canonical and participant trajectories progress independently. The participant continues to interact with a local version of their character while the server also continues to maintain a separate version that is visible to others. Both branches of the trajectory have to be resynchronised with each other on reengagement. This has the advantage of allowing everyone to progress, but at considerable risk of inconsistency.

These four routes can be mapped on to existing approaches to concurrency control in distributed simulation [12]. Pessimistic concurrency control requires that one component of the system exclusively acquires and locks a resource (the character) before interacting with it. This is the case in route (b) where the server takes control of the character and route (c) where the client does. Route (d) on the other hand represents optimistic concurrency control where both client and server update the character independently and resolve any inconsistencies later on.

Historic trajectories capture views of past experience

The final issue raised by DoF concerned how participants reviewed their history of interactions, either partially through often incomplete records on their phones, or in its

entirety through the web, but only after each performance had finished. The issues involved in reviewing the past are expressed through our third and final kind of temporal trajectory, **historic trajectories**. If we assume that all interactions are recorded by the system as they are in DoF, then it becomes possible to synthesise different historic views of what took place. A historic trajectory encapsulates such a synthesis, i.e., defines one particular view of history.

A historic trajectory involves a selection of segments from participant trajectories. This might be based upon various criteria such as the participants who were involved, the importance of particular events, or even how recently they occurred. These selected events are then made available to participants, perhaps being scheduled at a particular time and maybe even presented (narrated) in a particular order (which could be different to the one in which they originally occurred). They may also be presented in a different medium. Thus, historic trajectories act rather like canonical trajectories, but involve a new plot being synthesized from out of previous participant trajectories.

As examples, DoF involved two kinds of historic trajectory. The first was where participants reviewed text messages on their phones. The selection here was that each participant could only see their own messages and these were often further selected (due to limited storage) to be only important or recent messages. These were made available as soon as they had been received, so that the historic trajectory was interwoven with their ongoing participant trajectory, and they were presented through the same medium, their mobile phone. The second was the history of events published on the web. Again, the selection was by participant, but this time all messages were shown. These many different historic trajectories (one for each participant) were scheduled to appear after the event and were translated into a new medium (web versus phone).

In fact, DoF’s notion of historic trajectories is relatively simple and there are much richer ways in which they might be used. We therefore suspend further discussion of them for the time being and instead turn our attention in the second part of this paper to how temporal trajectories can open up entirely new possibilities for interactive narrative.

USING TEMPORAL TRAJECTORIES TO ENVISAGE NEW POSSIBILITIES FOR TIME AND INTERACTION

So far we have seen how temporal trajectories in their different forms can express the temporal structure of an existing experience. In this section, we put them to work more proactively, showing how they also have the power to generate new insights and design approaches.

A rolling narrative experience

Our first step is to envisage a rolling experience; one that is still driven by a strong narrative, but where participants can join at many different times. For example, rather than having just a few widely separated performances of DoF as we currently do, we could launch a new performance on

each day of clock time, enabling participants to join on any day and experience the game over the subsequent twenty four days. The result would be a rolling structure of twenty four concurrent games each with its own separate group of participants. This can be represented by twenty four parallel canonical trajectories (Figure 5), with a new one starting and an old one finishing on each day. This becomes more interesting when we also allow for different paces of play.

Accommodating different paces of interaction

In his extensive writing about time and HCI, Alan Dix has emphasized the importance of pacing within interactive and collaborative systems. Dix has argued that communication breakdowns often occur when the pace of participants’ interactions or communication is considerably slower than that of the task that they are trying to achieve [5]. Pacing was clearly an issue in DoF where we saw participants adopting different styles and paces of play over twenty four days. A few played intensively, sending messages nearly every day. However, the majority, played episodically, increasing and decreasing the intensity of their interactions according to shifting personal circumstances. How might we support this kind of episodic play?

As a starting point imagine extending our rolling game with the following simple pacing rule: *any day in CT during which a participant sends a message to the game is deemed to be an ‘engaged day’, causing ST to advance by one hour; whereas any day when they do not send a message is deemed to be a ‘disengaged day’ during which ST does not advance.* A participant who sent a message every day would complete the game in twenty four days as before. Slower paced participants however would take longer due to missed days of play during which their ST would be paused. They would still require twenty four days of play to finish, but these would no longer have to be contiguous.

They would also encounter different groups of players and possibly different states of the game as they passed through different instances. Figure 5 represents this using twenty four parallel canonical trajectories as discussed above. The red and blue participants join the experience on the same day. Each time one of them disengages for a day they fall back into the next less temporally advanced instance of the game (the next canonical trajectory coming along) where they will encounter a new set of players.




Figure 5: A rolling experience with variable pacing

Each time such a participant joined a new canonical trajectory they would encounter a new group of participants

who were currently following it and who therefore shared the same current point in story time. However, while it might be interesting to meet many new participants, it could also be frustratingly difficult to maintain friendships over the course of a game. Imagine two friends who join on the same day and play regularly, but then one of them skips a day and so falls into the next canonical trajectory. If they both continue to play regularly after that they will be locked into parallel trajectories, separated by one hour in story time and will never meet each other again.

Figure 6 therefore shows how we might use multiple canonical trajectories to more subtly manage pacing. Here there are two canonical trajectories with different slopes, one for fast-paced participants and another for slow ones. The blue participant is initially assigned to the faster paced trajectory. They disengage briefly at one point and are resynchronised when they re-engage (as in Figure 3). They then disengage for a much longer period. When they re-engage for the second time, they are synchronised with the slower canonical trajectory and are henceforth grouped with participants who adopt a slower pace. This dynamically divides participants into differently paced groups.




Figure 6: Using multiple canonical trajectories to group participants according to their pace of interaction

A further alternative involves the system managing pacing without an author or scheduler intervening. Figure 7 shows a case in which the system automatically synchronises two different participants’ trajectories.




Figure 7: synchronizing two participant trajectories

Red and Blue initially describe similar but slightly divergent participant trajectories. At CT (1), the system begins to gently steer them back together by subtly

speeding up the rate at which Blue consumes ST while slowing down Red, until they meet at 2, from which point on, their trajectories are held together, enabling them to share the narrative. This approach of dynamically synchronising unfolding participant trajectories relates to previous discussions of the WYSIWIS (What You See is What I See) and relaxed-WYSIWIS techniques that afford different degrees of relaxation of synchronisation in groupware tools such as shared editors [14, 19].

Managing encounters across time

Another more radical possibility is to enable participants to encounter each other even when they are at different times. Figure 8 shows two participant trajectories that describe quite different temporal journeys through an experience. An inspection of this figure reveals some interesting potential encounters between these participants, labeled (1) to (4).




Figure 8: different kinds of encounters in time

At point (1) our two participants share a common temporal trajectory for a while, being coincident in both CT and ST. In a shared interactive narrative such as DoF it would make sense for these participants to encounter one another (subject to other factors such as spatial proximity) as they share a common temporal context; that is they are seeing the same fictional events unfold around them.

At point (2) their paths briefly cross again in both ST and CT, however this time the blue participant is disengaged while the red participant remains engaged. If we allow them to encounter one another at this point then we need to be aware that one may appear to ignore the other or may affect the other in their absence as we saw in DoF which allows encounters involving disengaged participants.

A more unusual possibility is at point (3) where two participants encounter one another even though they are at different story times. This would enable participants from a relative fictional future and past to meet and exchange information. For example, in our proposed rolling version of DoF, you and I might simultaneously (in CT) be in Kath's Café, but for you it would be 11:00 AM in ST, whereas for me it would be 15:00 PM. What would it then mean for us to encounter one another? This unconventional scenario opens up new narrative possibilities. For example, a player might use a 'time portal' object to allow them to

discover information about what is going to happen in the future or to leave a warning or hint for those in the past.

Finally, point (4) shows an encounter between two participants who are at different clock times. This represents asynchronous communication between participants who actually share a common fictional time. For example, the blue participant might see a recording of what the red participant did at this moment in ST earlier on in this performance (or even in a previous performance).

To generalise, we propose that encounters between participants in shared interactive narratives can be mediated according to their proximity in a combination of ST and CT, and also according to their ongoing levels of engagement. Combining these factors in various ways can establish new narrative possibilities, including the extent to which communication is synchronous or asynchronous (proximity in CT), whether it reflects a notion of fictional time travel (proximity in ST) and whether or not participants can be affected in their absence (engagement). If we allow such encounters then it may be helpful to convey the nature of any time differences to the participants involved, perhaps by extending previous proposals for awareness widgets that reveal the presence and nature of network delays between participants in groupware [18].

Multiple histories and time travel

For our final discussion we return to the subjects of historic trajectories and time travel. In general, ST, like the fictional space in which a story is set, is a construct of the author's imagination and so there is no reason why participants cannot move backwards and forwards within it, provided that the author can satisfactorily explain this within the story. Traveling in CT however is a different matter. Outside of relativistic effects (which we can ignore in this paper), participants cannot change or diverge their positions in CT (though they might try to spoof the system by resetting their local system clocks). What they can do, however, is replay events from previous clock times that have been recorded by the system.

Figure 9 shows various possibilities for the red participant to view recorded events. At 1, 2 and 3 they view events that were recorded by the blue participant at past, present and future story times respectively. The act of viewing a recorded event can itself be recorded, so that at 4 they see a recording of themselves viewing the recorded event from 1.

Further narrative possibilities arise when we combine recorded and live events. At any moment, a participant may experience a combination of live events that have been scripted to appear at this point in CT and ST and recorded events that have already happened at some previous CT. Depending on how these are mixed, the participant will have quite different experiences of apparent 'time travel'. If they only see live events, then the experience is like replaying a level in a conventional computer game – they travel back to a suitable moment in ST (but not of course in

CT) to retry scripted interactions. If they see only recorded events, then they can only view the history of their own or others' activities from previous clock times. If they see a mixture of the two, then they will be able to interact with some elements while only viewing others, for example racing against recorded 'ghost' versions of themselves.




Figure 9: accessing recorded events from the past

It is important to note that recorded events are immutable, that is then they cannot be changed as this would require time travel in CT, which is not possible. How then can we support interesting narratives that exploit classic time travel paradoxes that arise from altering the past? The answer lies in our concept of historic trajectories that synthesize fragments from different recorded participant trajectories.

Figure 10 includes a participant trajectory where the participant repeatedly revisits different story times, for example replaying two levels in a computer game.

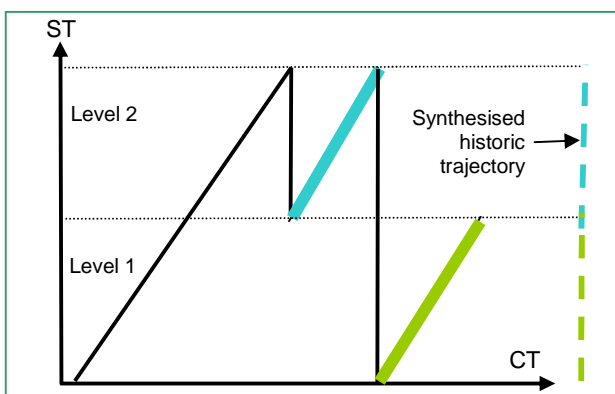


Figure 10: synthesizing a historic trajectory

In this example, the system generates a historic trajectory by choosing the most recently (in CT) recorded version of events at each ST (the highlighted segments). Under this selection rule, the experience works like a multi-track recorder, with participants being able to overwrite the apparent history of events by replaying previous levels. Some care must be taken to avoid jarring anomalies that might become apparent at the transitions between segments, for example the cast of participants (e.g., other players) suddenly changing. This might be achieved by ensuring that

any such transitions occur at natural breaks in the structure of story time (e.g., between chapters, acts or levels).

There are other interesting selection rules for synthesizing a historic trajectory from many recorded participant trajectories. For example, the 'best' version of events (e.g., with the highest score) might be selected, participants might see their own subjective versions of history composed from only their recorded trajectories, or in the most extreme case they might be able to explore all recorded histories, following different participants' stories and seeing alternative versions and outcomes.

Once again, this use of temporal trajectories has its parallels in the HCI and CSCW literature, in this case in proposals for managing divergent and reconvergent histories in asynchronous shared document editors. For example, Timewarp enabled participants to create and navigate divergent edit histories, including moving backwards and forwards in a document's history (story time in our terms) before creating a final definitive version (an example of synthesizing a historic trajectory) [9]. A later extension to this work involved multiple local edit histories that could be edited independently and/or related to a common overall document history [8], a more complex example of synchronizing multiple trajectories. The MASSIVE-3 collaborative virtual environment platform included a nested recording facility where scenes in one virtual world could be enacted, recorded and then played back within another live world. This enabled '3D flashbacks' within an interactive narrative [4], providing an example of creating a historic trajectory (recording) that subsequently becomes part of a canonical trajectory later on in an experience.

CONCLUSION

Narrative-driven experiences such as computer games, interactive performances and hypermedia stories raise new challenges for the treatment of time in HCI. How can we balance an author's desire to impose an ordering and timing on events as part of plot and schedule with a participant's need to fit interaction with the ongoing patterns of their daily life? How can we manage different paces of interaction and establish appropriate alignments between different participants' experiences of time? And how can we access recordings of past interactions to create new stories or to weave them in with ongoing ones?

Inspired by our experience of designing and staging Day of the Figurines, we have developed the concept of temporal trajectories to express the complex mappings between fictional story time and actual clock time. We have introduced three kinds of temporal trajectory, canonical, participant and historic, and have shown how a systematic comparison of these enables the analysis of existing experiences and also generate new approaches to the treatment of time in future experiences. The key points to emerge from this discussion are as follows.

Canonical trajectories capture an author's intended mapping between story and clock time as part of a plot. They can also express the scheduling of an experience, which can be defined in advance or managed on the fly. Multiple canonical trajectories can reflect more complex nested temporal structures such as those of some MMPORGS.

Participant trajectories reflect a participant's actual journey through a combination of story and clock time. A discussion of how participants disengage from and subsequently reengage with an ongoing experience uncovers a metaphorical 'Bermuda Triangle' of uncertainty that is formed as their participant trajectory first diverges from and then reconverges with a canonical trajectory. Various routes around this triangle embody established approaches to synchronization between clients and servers.

A further comparison of canonical and participant trajectories reveals new approaches to dealing with variable pacing in long-term interactive experiences. Authors can use multiple canonical trajectories with different slopes to group participants according to their pace or the system might gently steer similarly paced participants together to form stable social groups.

A comparison of different participant trajectories suggests new ways in which we might facilitate encounters across a combination of story time and clock time, reflecting differences between synchronous and asynchronous communication and opening up new possibilities for participants meeting others from the fictional past or future.

Historic trajectories embody an approach to synthesizing different historical views of events by selecting information from potentially many participant trajectories and then making it available in a particular order at a particular time. By mixing canonical and historic trajectories we can create different combinations of live and recorded events, enabling participants to access flashbacks within a live experience and experience various forms of fictional time travel.

The expressive power of temporal trajectories may be due to several factors. They capture in an integrated way key aspects of time across the different time scales of story time, plot, schedule and interaction time, mapping these onto clock time in the real world. They also balance the needs of authors and participants and live and recorded events. Finally, we find that their diagrammatic expression supports easy comparison and identification of new possibilities. Whatever the cause, we propose that this combination of canonical, participant and historic trajectories provides a powerful tool for reasoning about time and interactivity in shared narrative experiences.

REFERENCES

1. Allen, R. C., *Channels of discourse, reassembled: Television and Contemporary Criticism*, Second Edition, Routledge, 1992
2. Bell, M., Chalmers, M., et al., Interweaving mobile games with everyday life, *CHI 2006*, Montréal, Canada, 417-426, ACM, 2006
3. Chen, J., Flow in Games (and everything else), *Communications of the ACM*, 50 (4), 31-34, April 2007
4. Craven, C., Taylor, I., et al, Exploiting Interactivity, Influence, Space, Time to Explore Non-Linear Drama in Virtual Worlds, *CHI 2001*, pp. 30-37, 2001, ACM
5. Dix, A., Pace and Interaction, *Proc HCI '92 People and Computers VII*, (edited by Monk, A., Diaper, D., Harrison, M.), 193-207, 1992, CUP.
6. Dix, A., Que Sera Sera: the problem of the future perfect in open and cooperative systems, *Proc HCI '94: People and computers IX*, (edited by Cockton, G., Draper, S and Weir, G), 397-408, Glasgow, 1994, CUP
7. Ducheneaut, N., Yee, N., Nickell, E., Moore, R, The life and death of online gaming communities: a look at guilds in world of warcraft, *CHI 2007*, 839-848, San Jose, CA, 2007, ACM
8. Edwards, W. K., Igarashi, T., LaMarca, A., Mynatt, E., A Temporal Model for Multi-Level Undo and Redo, *UIST 00*, 31-40, San Diego, CA, ACM
9. Edwards, W. K., Mynatt, E, D., Timewarp: Techniques for Autonomous Collaboration, *CHI 97*, 218-225, Atlanta, GA, ACM
10. Anonymised reference – overview of DoF
11. Landow, G. (ed), *HyperText/Theory*, Baltimore and London: The Johns Hopkins University Press, 1994.
12. Fujimoto, R. M., *Parallel and Distributed Simulation Systems*, John Wiley & Sons, 2000
13. Manovich, L., *The Language of New Media*, MIT, 2001
14. Munson, J., Dewan, P., A concurrency control framework for collaborative systems, *CSCW 96*, 278-287, Boston, MA, 1996, ACM
15. Murray, J., *Hamlet on the Holodeck The Future of Narrative in Cyberspace*, MIT Press, 1997
16. Pfister, M., *The Theory and Analysis of Drama*, Cambridge University Press, 1998 (English translation)
17. Rekimoto, J., Time-machine computing: a time-centric approach for the information environment, *UIST 99*, 45-54, Asheville, North Carolina, 1999, ACM
18. Stach, T., Gutwin, C., Pinelle, D., Irani, P., Improving recognition and characterization in groupware with rich embodiments, *CHI 2007*, 11-20, 2007, ACM.
19. Stefik, M., Bobrow, D., Foster, G., Lanning, S., Tatar, D., WYSIWIS revised: early experiences with multiuser interfaces, *ACM TOIS*, 5 (2), 147-167, 1987, ACM
20. Thomson, P., 'Games and Plays – An Approach to Ionesco', *Educational Theatre Journal*, 22:1, 1970, 60-70
21. Ultima website, www.uoherald.com, verified Sept 2007
22. Varela, F. J. 'The Specious Present: A Neurophenomenology of Time Consciousness' in Petiot, J. et al, (eds) *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*, Stanford University Press, 1999, pp. 266-314.

CONTRIBUTION AND BENEFIT STATEMENT

Temporal trajectories express complex mappings between story time and clock time in shared interactive narratives, supporting analysis of existing experiences and generating new techniques for managing time and interaction.