Making Stories Player-Specific: Delayed Authoring in Interactive Storytelling

David Thue[1], Vadim Bulitko[1], and Marcia Spetch[2]

Departments of Computing Science[1] and Psychology[2]
University of Alberta
Edmonton, Alberta, Canada, T6G 2E8
{dthue | bulitko | mspetch}@ualberta.ca

Abstract. Of all forms of storytelling, interactive storytelling presents authors with a unique opportunity: while knowledge of the story's audience remains unavailable to traditional story authors, the nature of interactive stories to encourage player interaction allows information concerning the player's state to be automatically inferred. Given access to such information, an author's decision-making would become more informed, giving better chances of the story having the impact that the author intends. In this paper, we present an analysis of the decision-making process in interactive storytelling, and construct a method for characterizing storytelling systems based on features of their design. We demonstrate our method by characterizing four recently published systems, and review the related literature on inferring player information. Finally, we present Delayed Authoring, a new perspective on the design of interactive storytelling systems intended to take advantage of their opportunity to make stories player-specific.

1 Introduction

The primary goal of storytelling is to have some impact on an audience, be it emotional, educational, or both. The author's success in achieving their desired impact, however, can greatly depend on each viewer's personal state, formed from her prior life experience and current state of mind. In traditional forms of storytelling (e.g., books, movies), authors have no access to this information; they are forced to take a generalized approach, relying on an average understanding of their audience as a whole to guide their decision-making while they write. Interactive storytelling, on the other hand, allows a more specific, customized approach, as information about the current viewer can be inferred via the mechanism that supports her interaction [1–4]. If an author were given such information, his writing decisions would become more informed, allowing the creation of a story for each viewer that would be tailored to suit her particular state and have the impact that the author desired. While it would likely not be feasible to have a human author create a unique story for every potential viewer, the technology behind interactive digital storytelling presents a viable alternative approach: express the potential elements of a story in a generic form, and construct an automated system to determine their details dynamically, informed by both an author's knowledge of the creation of stories, and the viewer's current state as inferred from her interactions. Such a system would effectively serve as a decision-making proxy for the authors of a story, ideally choosing between potential elements of story content just as the authors would have, if they had been given a model of their viewer's state to guide their story decisions.

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-540-89454-4_30

Although having access to information about the viewer *while* a story is being both created and told is a benefit that is unique to interactive storytelling, the questions of both how to learn this information and how it might be well used (within the context of an interactive story) have only recently begun to be explored [1–7]. Based on this observation, the goals of this paper are twofold: to stimulate further research in learning the current state of an interactive story's viewer, and to promote a new perspective on the decision-making process of interactive story design. To meet these ends, the remainder of this paper is organized into the following three contributions: (i) we present an analysis of decision-making in interactive storytelling, and use it to construct a new method for characterizing interactive storytelling systems based on the decisions made during their design and operation; (ii) following a review of recent methods for inferring a viewer's personal state, we demonstrate our characterization method on four recently published storytelling systems, highlighting the ways in which they may be extended to take even further advantage of learned viewer states; (iii) we present a new design perspective, named *Delayed Authoring*, whose intent is to recognize the opportunities that arise when story decisions are delayed right up to the moment at which their results are needed. Finally, we conclude the paper with brief suggestions for a set of future research, showcasing the potential benefits that can be obtained by making interactive stories highly specific to a single, current viewer.

2 Decision-making in Interactive Storytelling

Before beginning our analysis of decision-making in interactive storytelling, we define the primary subject of the decisions in a story, namely, story events, as follows.

2.1 Story Events

Fundamentally, stories are a sequence of events, each of which involves some form of action. In journalism, it is common to report the occurrence of an event by answering each of six questions: "*Who?*", "*What?*", "*When?*", "*Where?*", "*Why?*", and "*How?*". Considering the types of facts that result from answering these questions, we propose that events can be well-described by six general properties, as given in Table 1.

Property	**Description**
Idea	A brief description of the action that occurs.
Actors	The people/creatures/forces that either perform some action or are acted upon.
Time	The time at which the action begins.
Place	The environment(s) in which the action occurs.
Actions	The changes that actors make to themselves, other actors, or their environment.
Reasons	The notions held by actors that prompt their actions.

Table 1. Six general properties of story events.

For example, consider this event: "*John rescues his friend Fred from the burning building by carrying him outside.*" In this case, the *Idea* might be "Rescue" or "Rescue from Fire". The *Actors* are John and Fred, the *Places* are the building and the area outside, and the *Action* is John carrying Fred outside. Although the *Time* of the event is not stated explicitly, one might assume that the *Time* is "now". One might also assume that

John's *Reason* for rescuing Fred is based on a desire to save his life. Although the given example fits well with the intuitive notion of story events being the significant elements of a story's plot, our definition is more general: a story event may be created from *any* action, regardless of its significance in terms of the plot. For example, one might say that the *Reason* for an apple falling from a tree is the failure of its stem to resist the pull of gravity. Due to this generality, the following analysis, which focuses primarily on decisions related to a story's plot, can be readily applied to character behaviour, dialogue generation, or any other story decision.

2.2 Story Decisions and Design Decisions

At its core, creating a story is about making decisions. For each of the six properties of an event, an author must decide what its value should be; that is, they must make six *story decisions* for every story event, as shown in Table 2. While the six questions shown all refer to a particular story event, consider also asking a similar set of questions with reference to each story decision itself; this set of questions is shown in the "Design Decision" column of Table 2. Note that the question asking "*where was the decision made?*" has been intentionally left out, as it is not relevant to this analysis.

Property	Story Decision	Property	Design Decision
Idea	**What** should happen?	*Result*	**What** was decided?
Actors	**Who** should be involved?	*Chooser*	**Who** made the decision?
Time	**When** should it happen?	*Time*	**When** was the decision made?
Place	**Where** should it happen?		
Actions	**How** should it happen?	*Method*	**How** was the decision made?
Reasons	**Why** should the actors act?	*Justification*	**Why** was the decision made in that way?

Table 2. At left: six decisions to make for every story event. At right: five design decisions to make for every story decision.

Just as the answers to the six story decisions in Table 2 yielded the properties of a given story event, answering the given five design decisions yields the properties of each story decision: *Result*, *Chooser*, *Time*, *Method*, and *Justification* (see Table 3). These properties describe details of the design of a given story experience, granting a better understanding of the storytelling system being considered. When considered in future tense (e.g., "*who will make the decision?*"), the questions from which they arose can be thought of as decisions that are made during the design of a storytelling system; we refer to them as *design decisions* (in contrast to story decisions) for this reason.

Property	Description
Result	The result of the decision being considered.
Chooser	The party who made the decision - either the player or the author.
Time	The time at which the decision was made - either offline (before the story) or online (during the story).
Method	The mechanism used to make the decision - this may be author imagination, a particular (computer) algorithm, or an in-game player action.
Justification	The author's or player's reason for using the method that they chose.

Table 3. Five general properties of story decisions.

3 Characterizing Storytelling Systems

While story decisions determine the content of a story, design decisions determine a process by which stories can be made. In this section, we combine these two types of decisions to construct a method for characterizing storytelling systems.

For each of the six properties of a given story event, consider the properties of the decision that determined its value: What was the *Result*? Who was the *Chooser*? At what *Time* was the decision made, by what *Method* was it made, and what was the *Chooser*'s *Justification* for making it in that way? Answering this set of questions give rise to Table 4, which contains sample values to aid in explaining how it should be read.

	Result	**Chooser**	**Time**	**Method**	**Justification**
Idea	Rescue	Author	Offline	Imagination	Unrestricted
Actors	John & Fred	Author	Online	Friend Finder	Plot Consistency
Time	Act 3	Player	Online	Persuaded John	Event Trigger
Place	Burning Building	Author	Offline	Imagination	Unrestricted
Actions	J Carry F Outside	Author	Online	AI Planner	Goal Satisfaction
Reasons	Save Life	Author	Offline	Imagination	Unrestricted

Table 4. Sample values for the five properties of each story decision.

Consider the first row of values in Table 4 (*Idea*). Reading from left to right, we learn that the event in question concerns some kind of rescue, and that this decision was made offline by the author, drawing freely from his imagination. The second row (*Actors*) shows that John and Fred are the actors in this event, and that this author decision was made at run-time (i.e., online) using an algorithm to automatically select two actors based on the consistency requirements that the actors be in the same area as the event and friends with one another. The third row (*Time*) shows that the rescue event occurred in Act 3 of its story, and that it happened at this time because the player (perhaps being unable to save Fred herself) found John and persuaded him to help. The fourth row (*Place*) shows that the rescue event happened inside a burning building, and that the author chose this location from his imagination, offline. The fifth row (*Actions*) shows that the rescue event consisted of John carrying Fred out of the burning building, and that the decision for this action to occur was made online by an algorithm designed to automatically plan John's actions toward achieving a given goal. Finally the sixth row (*Reasons*) shows that John's motive for rescuing Fred was to save his life, and that this motive was chosen offline by the author by drawing from his imagination.

By filling out this table with respect to a representative event from a given storytelling system (and providing additional details in the "Method" column), one can concisely summarize that system's operation, characterizing it in the process. When an example event (such as the rescue event of Table 4) is made common to multiple systems, one can also highlight the inter-system differences and similarities that exist; we take this approach later on in this paper. Even though there are systems for which the rescue event could not have occurred exactly as we have described it above (e.g., the "carry" action may not be defined), we will nevertheless describe how this event might occur within the design of every system that we consider.

4 Inferring Player States

As we described briefly in the introduction, the past experience and current state of mind of any particular viewer causes them to approach a given story from a unique personal perspective. While traditional forms of storytelling offer little to no access to this information until after every aspect of a story has been fixed, the interactive nature of interactive storytelling gives viewers the opportunity to convey their state (whether consciously or not) through the actions that they take as part of their interaction. Henceforth, we will refer to the viewers of an interactive story as *players*, to better capture the facts that: 1) they are expected to perform actions *while* a story unfolds, and 2) their actions are meant to have some effect on the system that facilitates their story experience.

By observing the actions of their players, several interactive storytelling systems have begun to model aspects of their players' personal states. What might it be useful for such models to represent? That is, what information about individual players would it be helpful for authors to have? Within the domain of Interactive Storytelling, the majority of relevant research to-date has focused on inferring one or more of three aspects of a current player's state, namely, her knowledge, her preferences, or her goals. We review each of these aspects in the following sections, citing examples from recent research.

4.1 Inferring Player Knowledge

Having a representation of what the player knows concerning the current story is especially important when the story is set in a virtual environment that affords a high degree of player exploration, for players might easily miss seeing an element of content that is crucial toward understanding the circumstances of a subsequent story event. If this lack of player knowledge were correctly inferred, steps could be taken to either encourage the player to notice the content that she missed (perhaps via auditory or visual cues), modify the upcoming dependant event to rely on an element of player-known content instead, or switch to a different element entirely.

For example, Magerko's Interactive Drama Architecture (IDA) aims to infer its player's knowledge of the stories it presents by automatically tracking her movements between rooms in a virtual world [5]. More specifically, at all times during the story, a hypothesized knowledge base is maintained as a set of factual statements about the story's virtual environment, such as "*the cleaver is in the kitchen*". Whenever the player enters a new location or the state of the player's current location changes, the hypothesized knowledge base is updated under the assumption that the player will always be aware of the entire state of any location that she inhabits.

4.2 Inferring Player Preferences

When an interactive storytelling system is aimed at providing entertainment, having some notion of the current player's preferences can aid the creation of a satisfying playing experience; the variation of styles of play or even basic interests from one player to another are typically too great for a one-size-fits-all approach. The common solution to this problem in commercial video games is to include a wide variety of content that is designed to appeal to different groups of players. The way in which this content is presented, however, often requires every player to either (i) experience every element of content in a largely sequential fashion (e.g., *Half-Life 2* [8]), or (ii) manually search through a massive environment full of content to find the elements that they prefer (e.g., *Oblivion* [9]). If information about the player's preferences were correctly inferred, well-suited elements of the story's content could be automatically selected or brought to the attention of the player.

Sharma et al. have focused on learning the interests of a current player [3]. Using a database of interest-annotated logs of the experiences that previous players had with its stories, Sharma et al.'s system attempts to infer the interests of its current player by matching her trajectory through the space of possible story events with a trajectory in the annotated database, hoping that players who progress through a story in similar ways will have similar interests as well.

Barber et al. and El-Nasr have explored the task of learning the personality (preferences in behaviour) of their current players [1, 2]. Both of the resulting systems (GADIN and Mirage, respectively) model the personality of their players as vectors of values representing attributes such as their honesty, selfishness, or cowardice. The models are updated via pre-specified annotations on potential player actions; for example, if a player chose to flee from a presented confrontation, the model's representation of the player's cowardice would increase.

Thue et al.'s PaSSAGE also attempts to learn aspects of its player's preferences in behaviour, focusing in particular on the player's preferred styles of play (e.g., fighting, solving puzzles, or amassing power and riches) [4]. For each style of play, the player's inclination toward playing in that style is maintained by pre-specified annotations on potential player actions, similarly to the operation of GADIN and Mirage.

4.3 Inferring Player Goals

In spite of the unique advantages that Interactive Storytelling provides, the very notion of allowing interactivity at all still causes great discomfort among traditional story authors; the introduction of a character who is beyond their control (namely, the player) is often viewed as an obstacle to overcome, rather than of a source of information to enjoy. Toward alleviating the reluctance of such authors, inferring a player's goals and intentions offers clues as to what the player is likely to do next, based on her history and current situation.

One of the earliest examples of research toward modelling the player in an interactive story is Albrecht et al.'s work on keyhole plan recognition in multi-user-dungeons [10]. In their work, they used various abstractions of both the state of the virtual world and the player's history of actions as the input to Bayesian networks, the output of which were predictions of the player's current goal and her likely subsequent action.

5 Using Player States

In this section, we employ the characterization presented in Section 3 to describe four recent interactive storytelling systems, all of which take advantage of inferred player state to guide their storytelling decisions. We refer back to our example of the "Rescue from Fire" event in Section 2 throughout, toward highlighting the similarities and differences between two pairs of systems: IDA and PaSSAGE, and Façade and Mirage.

5.1 The Interactive Drama Architecture

Table 5 shows the story decision properties for Magerko's Interactive Drama Architecture (IDA) [5] for the "Rescue from Fire" event. Reading across the rows, the following information can be learned. For the given event to be possible in an IDA story, it would need to be created by an author and included in IDA's library of plot points. The actors taking part in this event might vary; preconditions can be authored with constrained variables, allowing actors to be selected subject to those constraints while the story is being told. The timing of events in IDA is based on preconditions as well; as soon as an event's constraints are satisfied, its content begins to occur. Similarly to the case with actors, the location of events can also be authored generically via preconditions, and for all of *Actors*, *Time*, and *Place*, the player's role in satisfying the event's preconditions could affect the decisions that are made. The *Actions* of IDA's actors are primarily driven by a goal-based planner, but when the player is expected to violate a set of constraints imposed by the author, reactive and/or pre-emptive behaviours are triggered to help avoid the pending problem. The actors' *Reasons* for taking their actions are predetermined by the author to lend believability to their behaviours.

	Chooser	**Time**	**Method**	**Justification**
Idea	Author	Offline	Imagination	No Restrictions
Actors	Player & Author	Online & Offline	Variable Preconditions - The identity of the character in need of rescue could have been left undetermined by the author, but constrained to be a friend of John's.	Promote Variable Content
Time	Player & Author	Online & Offline	Satisfied Preconditions - The preconditions for this event might have been for Fred to be inside the building, with John and the player standing outside, and the building being on fire.	Use Available Plot Point
Place	Player & Author	Online & Offline	Variable Preconditions - The location of the fire may have been left as a variable by the author, allowing IDA to select the building that Fred is in to satisfy one of the event's preconditions.	Allow Plot Point Flexibility
Actions	Player & Author	Online & Offline	AI Goal Selection or Reactive/Pre-emptive Direction - pre-scripted actions occur unless the player causes (or is predicted to cause) a boundary problem, at which time reactive or pre-emptive direction occurs.	Preserve Story Coherence
Reasons	Author	Offline	Imagination	Character Believability

Table 5. Story decision properties for the Interactive Drama Architecture.

Although IDA does take advantage of inferred player knowledge to make part of its experience player-specific (namely, the time and place of its events along with some of the actions and actors involved), the "Time" column in Table 5 clearly shows the ways in which it could be even further customized (marked by "Offline"). For example, in the IDA testbed *Haunt 2*, while the murderer's *Reason* for having killed player's character is completely pre-determined, some players might find an alternative motive to be substantially more interesting and compelling. It may be possible to learn these interests using work similar to Sharma et al.'s [3], and use them to further guide the selection of the events in the plot themselves (Nakasone and Ishizuka's have recent work on the topic of interest-based decision-making in interactive storytelling [6]).

5.2 Façade

Table 6 shows the story decision properties for Mateas and Stern's Façade [11] for the "Rescue from Fire" event. Even though Façade has enjoyed much praise from both the academic and consumer communities alike, Table 6 shows that it takes advantage of inferred player state in only two respects: determining the times at which events occur and the actions that actors take. A technique similar to the variable preconditions in IDA could be used to add variety to the actors and locations of Façade's events (given additional content), and certain players' interest in Façade could be improved by selectively introducing a more varied set of topics to discuss with the story's actors.

	Chooser	**Time**	**Method**	**Justification**
Idea	Author	Offline	Imagination	No Restrictions
Actors	Author	Offline	Imagination	No Restrictions
Time	Player & Author	Online & Offline	Tension Arc / Player Interest - a recent, sharp increase to the tension arc may have caused a high tension event to be sequenced. Alternatively, the player may have examined a fire alarm on a nearby wall, prompting John to begin his rescue.	Follow Dramatic Principle / Respond to Player
Place	Author	Offline	Imagination	No Restrictions
Actions	Player & Author	Online & Offline	Interruptible Scripts - High-level actor behaviours are pre-determined, but are authored in a way that they can handle player interruptions and then resume the original behaviour.	Allow Player Interaction
Reasons	Author	Offline	Imagination	No Restrictions

Table 6. Story decision properties for Façade.

5.3 PaSSAGE

Table 7 shows the story decision properties for Thue et al.'s PaSSAGE [4], with respect to the "Rescue from Fire" event. Similarly to IDA, PaSSAGE makes significant use of inferred player information, basing the *Time*, *Place*, and *Actors* of its events on its model of the player along with her position in the story's virtual world. PaSSAGE differs from IDA in that it additionally chooses what should happen (i.e., the *Idea*s) in its stories, similarly to Barber's GADIN [1] and other generative systems. The clear way in which PaSSAGE can further inform its decision-making lies in the *Reasons* for the actions of its actors; current work on PaSSAGE is focused on modelling personality attributes for its actors, toward motivating the actions that they take.

	Chooser	Time	Method	Justification
Idea	Player & Author	Online & Offline	Encounter Selection via Player Model - This event would have been selected over others due to the current values in the player model.	Cause Enjoyable Events
Actors	Player & Author	Online	Role Passing - The identities of the characters in this event would be determined dynamically at run-time based on proximity and relationship constraints.	Satisfy Role Passing Constraints
Time	Player & Author	Online & Offline	Triggers for Encounter Specification - The fire may have been authored to start only when the player approached a suitable location.	Ensure Event Visibility
Place	Player & Author	Online & Offline	Triggers for Encounter Specification - The location of the fire may have been specified as the next building approached by the player which had one actor outside and another inside.	Bring Interesting Events to the Player
Actions	Player & Author	Online & Offline	Encounter Refinement via Role Passing with Hinting - Actors satisfying the encounter's trigger conditions would assume the behaviours that were authored for this event, tailored to encourage the player's preferred styles of play.	Preserve System Flexibility
Reasons	Author	Offline	Imagination	No Restrictions

Table 7. Story decision properties for PaSSAGE.

5.4 Mirage

Table 8 shows the story decision properties for El-Nasr's Mirage [2], with respect to the "Rescue from Fire" event. Although Mirage infers substantially more player information than Façade, it uses its information to make primarily the same decisions: the *Time* of its events and the *Actions* of its actors. The main difference between the two, however, is that Mirage's actors choose between different tactics based on predicted player behaviour, while Façade's characters only cope with interruptions before returning to the same tactic as before. Mirage might take better advantage of its inferred player information by choosing actors for events based on an estimate of their tactics' success against the current player's character.

	Chooser	Time	Method	Justification
Idea	Author	Offline	Imagination	Use Dilemmas
Actors	Author	Offline	Imagination	No Restrictions
Time	Player & Author	Online & Offline	Actor Improvisation / Player Model - John's tactic for rescuing Fred may have been to convince the player save to him, but after predicting that this goal would fail, John's tactic would change to rescuing Fred himself.	Character Believability
Place	Author	Offline	Imagination	No Restrictions
Actions	Player & Author	Online	Goal-directed Behaviours - Actors continuously monitor the potential success their goals, and choose new behaviours if failure is predicted.	Character Believability
Reasons	Author	Offline	Imagination	No Restrictions

Table 8. Story decision properties for Mirage.

6 Delayed Authoring

Having presented our analysis of decision-making in interactive storytelling and reviewed the ways in which players' personal states can be learned while they play, we now present *Delayed Authoring*, a new perspective from which to approach the design of interactive storytelling systems.

6.1 Player Desire vs. Author Intent

In the field of Interactive Storytelling at present, many researchers believe that a degree of tension exists between two stereotypically opposing forces: player desire, and author intent. The argument behind this point of view is straightforward; players wish to act freely in a virtual world, while authors wish to tell the best story that they can without contending with the whims of an uncontrollable main character. In truth, however, the unique ability of interactive storytelling to take player interaction into account is what allows this tension to be reduced, via the mechanism of player modelling. Making effective use of this mechanism is the primary goal of Delayed Authoring, as guided by three key principles of interactive story design; these principles are described in the following three sections.

6.2 A Decision-making Proxy

The heart of the player versus author divide lies in the desire for control over the story experience, and this control is manifested in the types of story decisions that each party wishes to make. In the extreme case, players wish to choose when and where they travel in the virtual world and which characters they interact with, and they generally wish for their actions to have a significant influence over the story that unfolds. Authors, on the other hand, wish to choose all of the details that the players would choose, along with a myriad of others as well, including the behaviour of secondary characters, the phrasing of lines of dialogue, and more. Recall, however, the information on which such author-driven story decisions are traditionally based: the author's knowledge of the creation of stories, and his general understanding of his audience as a whole. If the author knew more about his audience, his story decisions would be more informed, allowing his story to more effectively have the impact that he intended. This fact motivates the acquisition of the audience's personal state, and we have already discussed methods by which a player's knowledge, preferences, and goals may be inferred from her actions in a virtual world. It also motivates the *use* of the audience's state when making story decisions, but information about this state is only available after an interactive storytelling system has already been designed, built, and distributed to its players. The solution is to build an Artificial Intelligence (AI) component into the design of an interactive storytelling system which automatically both infers aspects of the player's state as the story is presented, and decides on subsequent elements of story content just as the story's authors would have, had they been asked to write a story for the particular current player themselves. While many existing interactive storytelling systems include AI to make story decisions, viewing the AI component as a decision-making proxy for the author is the first principle of Delayed Authoring.

6.3 A Just-in-time Approach

We stated in the previous section that the greater an author's knowledge of his audience, the more informed his decisions would be. At what point in time would his decisions be *most* informed? Given that potentially useful information about the player's state could become available at any time, it should be desirable to delay the making of each decision for the maximum amount of time possible, i.e., right until the moment at which its result is needed. This just-in-time approach to decision-making is the second principle of Delayed Authoring.

6.4 Recognizing Opportunities

Referring back to the "Chooser" and "Time" columns of Table 4, three combinations of values are possible: Author/Offline, Author/Online, and Player/Online (assuming that players have no effect on the system before the story begins). For every Author/Offline pair, one might ask the question as to why it is not Author/Online instead; that is, for every decision that the author makes offline, one can ask if it could be made online, potentially benefiting from inferred player information in the process. The third principle of Delayed Authoring is to ask this question for *every* story decision encountered during the authoring process. Given the unique advantage of interactive storytelling to inform story decisions with learned player information, this principle encourages authors to recognize such opportunities whenever they arise. Table 9 lists the three principles of Delayed Authoring in summarized form.

Principle #1:	The AI component of an interactive storytelling system should be viewed as a decision-making proxy for the interactive story's authors.
Principle #2:	Any story decision that is made online should be delayed for as long as possible, to maximize its chance of being informed by new player information.
Principle #3:	For every story decision that arises during the authoring process, one should ask if it could be better informed by inferred player information.

Table 9. The three principles of Delayed Authoring

7 Caveats and Future Work

Although the content of this paper has been focused on highlighting the potential advantages of a player-specifc approach to interactive storytelling, there remain several important obstacles to overcome. Perhaps the most daunting of these is the problem of obtaining a sufficient amount of content to effectively *use* inferred player states at all; it is certain that traditional means of authoring will be unable to meet the varied demands of a large and diverse group of players. It may be the case, however, that the mechanisms for player-specific adaptation that have arisen in recent years can help to solve this problem, for the constructs underlying their ability to adapt are effectively representations of story events in an abstract form. Applying these adaptation mechanisms can vary the details of such events combinatorially, allowing the procedural generation of a large amount of content with comparatively little authoring effort. Furthermore, in addition to a story's content itself, the expression of that content can be made player-specific as well (e.g., Mirage, Façade [2, 11]). The principles of Delayed Authoring are meant to apply to *every* decision in storytelling, whatever those decisions may be.

8 Conclusion

In this paper, we made the following contributions. First, we presented an analysis of the decision-making process in interactive storytelling, conceptually distinguishing between story decisions as the determiners of story events, and design decisions as the determiners of how story decisions are made. We used this analysis to present a method for characterizing storytelling systems with respect to their decision-making processes, and demonstrated its use on four recent interactive storytelling systems. We provided a comprehensive review of the literature on inferring player state from a player's actions during a story, distinguishing between the player's knowledge, preferences, and goals as different types of information to infer. We then presented Delayed Authoring, a new perspective on the design of interactive storytelling systems devised to alleviate the tension between player desires and author intent, encouraging authors to recognize the many opportunities to make their stories player-specific.

Acknowledgements

This research was funded by grants from NSERC and iCORE. Parts of this text appear in revised form from David Thue's M.Sc. dissertation at the University of Alberta.

References

1. Barber, H., Kudenko, D.: Dynamic generation of dilemma-based interactive narratives. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007), Palo Alto, California, AAAI Press (June 6-8 2007)
2. El-Nasr, M.S.: Interaction, narrative, and drama creating an adaptive interactive narrative using performance arts theories. Interaction Studies **8**(2) (August 2007)
3. Sharma, M., Ontañón, S., Mehta, M., Ram, A.: Drama management evaluation for interactive fiction games. In: AAAI Fall Symposium 2007 on Intelligent Narrative Technologies, AAAI Press (November 2007)
4. Thue, D., Bulitko, V., Spetch, M., Wasylishen, E.: Interactive storytelling: A player modelling approach. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007), Palo Alto, California, AAAI Press (June 6-8 2007)
5. Magerko, B.: Player Modeling in the Interactive Drama Architecture. PhD thesis, Computer Science and Engineering, University of Michigan (2006)
6. Nakasone, A., Ishizuka, M.: ISRST: An interest based storytelling model using rhetorical relations. In: 2nd International Conference of E-Learning and Games (Edutainment 2007). Volume 4469/2007 of Lecture Notes in Computer Science., Hong Kong, Springer Berlin / Heidelberg (June 2007) 324–335
7. Peinado, F., Gervás, P.: Transferring game mastering laws to interactive digital storytelling. In: 2nd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2004), Darmstad, Germany, Springer (December 4-6 2004) 1–12
8. Valve Corporation: Half-life 2. http://www.half-life2.com/ (2004)
9. Bethesda Softworks: The Elder Scrolls IV: Oblivion. http://www.elderscrolls.com/games/oblivion_overview.htm (2006)
10. Albrecht, D.W., Zukerman, I., Nicholson, A.: Bayesian models for keyhole plan recognition in an adventure game. In: User Modeling and User-Adapted Interaction 8(1-2), Special Issue on Machine Learning in User Modeling, Kluwer Academic Publishers (1998) 5–47
11. Mateas, M., Stern, A.: Procedural authorship: A case-study of the interactive drama Façade. In: Digital Arts and Culture (DAC), Copenhagen (November 2005)