

Interactive Storytelling: From Computer Games to Interactive Stories

Fred CHARLES, Steven J. MEAD and Marc CAVAZZA

University of Teesside, School of Computing and Mathematics,
Borough Road, Middlesbrough, Cleveland, TS1 3BA, United Kingdom

f.charles@tees.ac.uk, steven.j.mead@tees.ac.uk, m.o.cavazza@tees.ac.uk

Abstract Interactive storytelling can either be based on explicit plot representations or on the autonomous behaviour of artificial characters. In such a character-based approach, the dynamic interaction between characters generates the actual plot from a generic storyline. Characters' behaviours are implemented through real-time search-based planning techniques. However, the top-down planning systems that control artificial actors need to be complemented with appropriate mechanisms dealing with emerging ("bottom-up") situations of narrative relevance. After discussing the determinants that account for the emergence of narrative situations, we introduce additional mechanisms for coping with these situations. These comprise situated reasoning and action repair: we also illustrate the concepts through detailed examples.

Keywords Interactive Storytelling, Virtual Humans, Autonomous Agents, Computer Games, Artificial Intelligence.

1 INTRODUCTION

One of the main challenges for future computer games is to provide an extensive narrative experience. This is the key to extending the status of computer games towards that of a proper medium, and to attract new audiences across age and gender barriers. Interactive storytelling is a long-term objective that is attracting researchers from various backgrounds (Sgouros et al., 1996), (Young, 2000), (Mateas, 2000), (Nakatsu and Tosa, 1999), (Szilas, 1999), considering its potential for digital media convergence.

There is a strong inter-dependency between the sophistication of the Artificial Intelligence (AI) in a game and the possibility of implementing interactive storytelling. Interactive storytelling can be seen as a natural extension of the implementation of autonomous actors. As virtual characters become more intelligent, the action can increasingly rely on their automatic behaviour, generating a larger diversity of story than with current authoring methods. This dynamic computation of the action also makes possible various forms of user intervention, whose consequences on the story can then be propagated, as the plot is re-computed.

The implementation of more intelligent actors will also support new interaction modalities such as speech and natural language (Cavazza and Palmer, 1999), which will have a significant impact on interaction paradigms. Indeed, one major limitation of current computer games is that they only allow physical interaction, which explains in part the dominance of simulation and action games. The development of language technologies will make possible realistic interaction with autonomous actors and will be one of the enabling components of future interactive storytelling.

In this paper, we describe the implementation of an interactive storytelling prototype, developing specific AI techniques on top of a game engine. After introducing the overall design and the AI techniques used, we present some example results from our system.

2 SYSTEM ARCHITECTURE AND OVERVIEW

We are developing an experimental platform for interactive storytelling and have already completed a first proof-of-concept prototype (Figure 1).

To support interactive storytelling, systems must provide realistic graphical visualisation of on-stage situations and support interaction with the virtual entities. A 3D game engine provides essential technical features to develop engaging virtual environments, where virtual actors can perform on a virtual stage. In addition, the dramatisation of situations can be enhanced by the realism of character animations. For instance, skeletal animations and body postures emphasise visual representation of

character behaviours. The engine also provides mechanisms for physical intervention, e.g. with objects that bear narrative significance. And the *mise en scene* of sequences can be enhanced using cinematographic-like camera movements.

The system is implemented using the Unreal Tournament™ (UT) game engine, which provides a high-quality graphic environment for animation and physical interaction. The AI components implementing interactive storytelling have been developed in C++ and UnrealScript™, the game engine’s scripting programming language. The search-based planning system supporting autonomous actors’ behaviour (see next section) has been implemented as C++ code included in the Unreal engine through dynamic link libraries. The code involved with low-level actions in the environment has been developed in UnrealScript™ directly within the game environment. The system can also process spoken commands using the EAR™ SDK speech recognition system from Babel Technologies; the recognised string is passed to a command processing module using the UDP socket interface to UT.

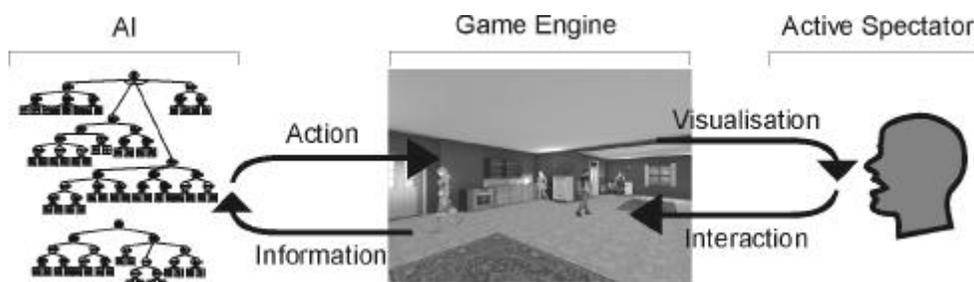


Figure 1: System Architecture

The scenario we are using for this first demonstrator is inspired from a popular sitcom, where the main character “Ross” wants to invite the main female character “Rachel” out on a date. The rationale for using a sitcom is that the plot structure can be kept relatively simple, while at the same time both the ending and intermediate situations are of narrative interest. Hence, this constitutes a good test case for a system whose main objective is to generate plot diversity, because it can be assessed on both the situations created and the final ending of the story produced.

The story appears as a real-time 3D animation in the UT engine: it can be displayed from any characters’ perspective, though the default mode is played through Ross’ perspective, as he is the main character. The user can also follow the action in first-person mode, and freely navigate on the set while action is in progress. This is also the basis for his physical intervention, for instance by hiding or stealing objects of narrative significance, which play a role in the story (e.g. a diary, letter, keys, handgun, etc.).

3 CHARACTERS’ BEHAVIOURS AS NARRATIVE DRIVE

The characters’ plans are based on narrative content rather than generic desires or intentions. As a consequence, the characters are actually playing a role rather than improvising. In that sense, the role-playing aspect serves as a narrative drive for the whole action: characters will pursue their long-term goals in a consistent fashion even if they are distracted from them. Their intermediate sub-goals, on the other hand, can differ depending on various circumstances and this constitutes the basis for the generation of plot diversity. This narrative drive distinguishes interactive storytelling from state-of-the-art simulation games such as *The Sims*™, in which the character behaviour is essentially reactive and the user intervention looks more like character design.

To understand the examples of this paper, a brief outline of Ross’ plan is necessary. In order to take Rachel out, Ross must first acquire information on her, such as her availability and likes and dislikes. He should then find a way to talk to her in private, which requires that she is accessible (i.e., not busy, alone, etc.). In the meantime, he should have gained her friendship or at least not alienated her (e.g. by upsetting her friends, etc.). He can offer her various gifts that will have to match her preferences, and finally will have to ask her out, to which she will give her answer. We will not detail here the various ways in which these sub-goals can be satisfied, nor the executability conditions for the actions that form part of the sub-plans, but this will become self-explanatory in the examples described in the forthcoming sections.

The definition of characters’ roles corresponds to the authoring part of interactive storytelling. The story genre prescribes the relations between the actors: for instance, in the sitcom genre, Ross’ plan is entirely dedicated to inviting Rachel, while her plan is based in her daily activities, unrelated to Ross’ quest. This asymmetry is part of the genre itself and helps defining the various roles corresponding to each character’s stereotype.

3.1 Underlying techniques

There is a wide consensus both in behavioural animation (Webber et al., 1994) and interactive storytelling (Young, 2000) on the use of AI planning techniques for the implementation of autonomous actors' behaviour. However, there exists many different techniques and implementing them efficiently is always a challenge. In this section, we describe our approach, which is inspired from search-based planning (Korf, 1990), (Bonet and Geffer, 1999).

We started with a generic description of a characters' role in terms of a Hierarchical Task Network (HTN). A HTN comprises both goals, sub-goals and actions organised hierarchically. While the top-level goals of the HTN correspond to the constant part of a character's role, the lower levels contain alternative sub-goals, which are exclusive within a single plan instantiation (e.g., the various ways in which Ross can acquire information about Rachel). As such, the HTN subsumes many possible plans within a single representation. HTN can be represented using AND/OR graphs: we initially relied on this formalism mainly for descriptive purposes. In the general instance, it is possible to generate plans from HTN through a formal operation called serialisation. However, when the plan is decomposable (i.e. there are no dependencies between sub-goals), the plan can be generated by directly searching the HTN. We are making a decomposability assumption on our characters' roles, as part of the generic properties of simple stories' development. As a consequence, our system generates characters' plans by directly searching the HTN representing the agent's role and producing a sub-graph corresponding to the agent instantiated behaviour (Figure 2).

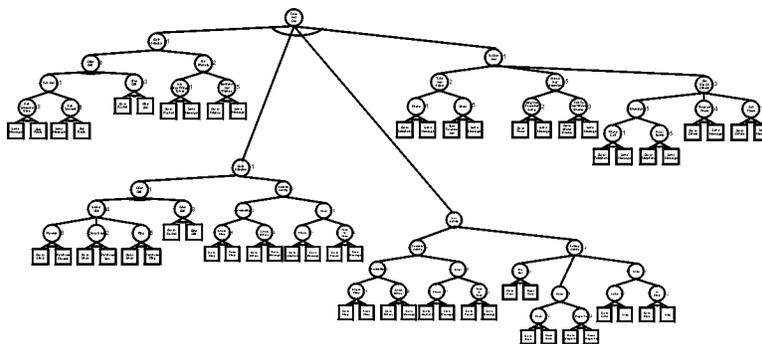


Figure 2: A Character's Plan

Important aspects of planning for interactive storytelling are that planning should be interleaved with execution, and that re-planning abilities should be included. The rationale is of course that the world in which action takes place is constantly changing under the influence of interaction between actors, and between the user and the actors, and plans established off-line might not stay valid when they are to be executed. As the HTN is formalised as an AND/OR graph, we thus use a "real-time" variant of the AO* algorithm (Pearl, 1984), (Knight and Rich, 1991), (Nilsson, 1998). While AO* is a standard, generic algorithm for searching AND/OR graphs, it is not appropriate to the current context, which requires fast computation as well as interleaving planning and execution. Our real-time variant essentially uses the forward component of AO*, which computes a solution basis (selection of the top-node of a sub-graph) using the heuristic function. It searches the graph in a depth-first, left-to-right fashion, not unlike the "MinMin" algorithm of Pemberton and Korf (Pemberton and Korf, 1994). Each time a terminal node is reached, its associated action (e.g., moving to a given location, talking to another character) is dramatised on the virtual stage by using the relevant character animation from the game engine. Should the action fail, for instance due to another character or the user having modified the on-stage resources availability, evaluation functions will be revised through back-propagation and search will be resumed, thus implementing re-planning on action failure. Re-planning is an important element of interactive storytelling, illustrating the dependency between the AI planning techniques and the on-going situations defined by both characters' behaviour and user intervention in the virtual environment. Furthermore, it supports the generation of story variants based on narrative incidents that are themselves dramatised.

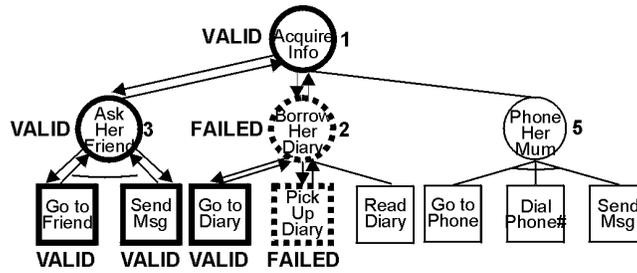


Figure 3. Re-planning on action failure.

For instance, Figure 3 shows a fragment of Ross’ sub-plan for acquiring information about Rachel to know more about her likes and availability to invite her out. His initial plan consists in reading Rachel’s diary, but the user has stolen it. On reaching the diary’s default location Ross realises that it is missing and needs to re-plan a solution to find information about Rachel, which in this case consists in asking Phoebe. Using the mechanism of back-propagating the action failure to the corresponding sub-goal, the search process will resume and produce an alternative solution.

From a narrative perspective, the user has contrasted Ross’ visible goal. But, apart from the immediate amusement of doing so, because failure of Ross’ action is dramatised and part of the plot (see Figure 4d), the real impact lies in the long-term consequences of the resulting situations. For instance, in the above example, when asking Phoebe about Rachel, Ross might be seen by Rachel, who would misunderstand the situation and become jealous.

3.2 Representational aspects

Because we are using heuristic search, it is important to associate an appropriate meaning to the evaluation function attached to sub-goal nodes in the task network. In order to do so, the various sub-goals or actions are indexed according to some narrative criteria. One possible criterion is a social characterisation of the action, like its “harshness”: for instance, if the sub-goal consists in isolating Rachel from the group, in order to be able to talk to her in private, this can be achieved in various ways: calling Rachel aside, interrupting the conversation and asking her friends to leave, etc. There can thus be a direct mapping between these criteria and personality profiles for the character operating the plan.

The personality profile of a character can be defined a priori to the story unfolding (Raskin, 1998). For each sub-goal node of the HTN graph, the character’s personality will influence the forward heuristic search according to the sub-goal’s relevance towards the personality trait. On the other hand, emotional states, or “moods”, arise from the succession of situations taking place as the story unfolds. A possible extension to the system will make an allowance for different emotional states, with regard to consider the complexity of emotional concepts for the characters. For instance, a character may be happy, but not sociable, so he would favour any solitary activity to group or even one-to-one activity. In this way, it is possible to take into account personality and “mood” variables in the description of the characters, with a direct mapping to their dynamic behaviour and to story generation.

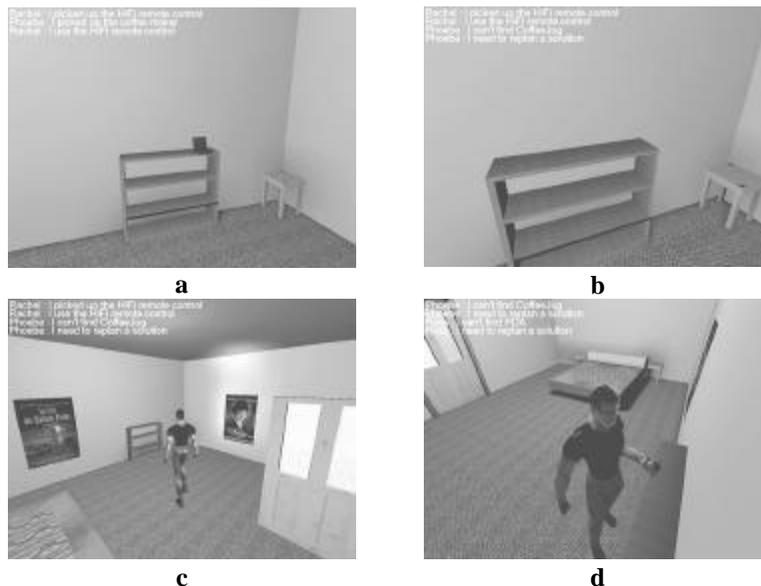


Figure 4: Physical User Intervention (a, b) and Situational Dramatisation (c, d)

4 THE INTERACTION CYCLE

We have seen that the terminal actions of a character's plan are dramatised in the virtual environment provided by the game engine, through real-time animation sequences. Action failure itself is dramatised and is indeed part of the story. The user thus has a consistent and continuous view of the unfolding plot. Because every action taken by the characters is narratively meaningful, the user can decide, from the understanding he has of these actions, to contrast them. For instance, if he sees Ross trying to reach Rachel's diary he can decide to steal it (Figure 4a, 4b) When Ross reaches the location of the sought-after resource, he stops, fruitless, then re-plans a new solution basis for the current sub-goal (Figure 4c, 4d). The overall cycle for interactive storytelling can be summarised as follows i) role playing by character and action dramatisation, ii) user understanding of the plot: considering that every action taken by a single character corresponds to some narrative goal, the user can anticipate the overall objective within a given story genre iii) anytime intervention by the user and iv) dramatisation of the specific consequences of user intervention.

The user can interfere with the action in two ways: either by physically interacting with on-stage objects or by addressing the characters using speech input (Charles et al., 2001). Many objects constitute resources for action, though at different levels. Some narrative objects play a great role: for instance, possible sources of information (such as letters, diaries, telephones) or potential gifts are important resources that directly impact on the unfolding of the main character's plan. Acting on these objects has generally a significant impact on the unfolding plot, whether in the short-term (e.g. forcing interaction with specific characters) or in the long-term (e.g. influencing Rachel's gift). Other objects, such as a coffee machine or TV set, play a role in the localisation of actors. Interacting with these objects might cause some secondary characters to move to other locations, potentially triggering new situations that can impact on the plot.

The other mode of interaction consists in influencing actors using speech recognition. This form of intervention includes:

- providing information needed by the actors to complete their plans (e.g. Rachel's preferred gifts). In this way, the user is actually influencing the progression of a character's plan by directly satisfying information-seeking sub-goals. It should be noted that, very much like artificial characters, the user can lie and provide the wrong information, whose impact will only become apparent at a later stage of the story.
- giving "doctrine advice" (Webber et al., 1994) that influences the personality of an actor (i.e. recommending a friendly behaviour towards certain characters, such as "be nice to Phoebe").
- getting actors to perform certain actions that have narrative consequences, such as moving to a certain location that increases the probability of meeting other characters and triggering a cascade of consequences.

5 DYNAMICS AND STORY EMERGENCE

5.1 *Situated Reasoning and Action Repair*

The basis for story generation rests with the interaction between the various characters' plans (this can be metaphorically described as the "cross product" of the various characters' plans). As the actors are evolving in the same environment, they are essentially sharing the same resources, whether these are objects of narrative significance or other actors. In other words, the basis for their interaction at plan level consists in them competing for resources for action. For instance, Ross might want to access Rachel's diary to acquire information about her, but Rachel could be using the diary herself. Or he might need to talk to Phoebe but she would be busy talking to Monica. However, this sole mechanism would not suffice to produce a diversity of plots. Besides the "top-down" component of story generation represented by the actors' plans, there exists a "bottom-up" aspects that corresponds to situations emerging from the initial conditions. For instance, actors are allocated random initial positions on the set. Considering the duration of their initial actions, various situations can arise depending on their initial positions. For instance, depending on their respective initial positions, Ross might be able to talk to Phoebe before she leaves the flat for some shopping. Or, in a similar fashion, Phoebe will meet Monica at an early stage and will appear busy talking while Ross needs information from her. One essential aspect is that the consequences of actions are causally propagated, even though causality might not be explicitly represented, unlike with plot-based representations (Sgouros et al., 1996), (Grabson and Braun, 2001).

In order to cope with these emerging situations, we have developed two mechanisms, following Geib and Webber (Geib and Webber, 1993): situated reasoning and action repair. Situated reasoning consists in generating a specific goal-oriented behaviour to respond to emergent situations that cannot be explicitly represented in the baseline plans, were it only for practical reasons of plan complexity. One example consists in Ross and Rachel bumping into each other at an early stage of the plot, before Ross has acquired information about her. Just following the baseline plan would result in Ross passing by Rachel without noticing her, which is not realistic from a narrative perspective. Situated reasoning should thus be used to respond to this kind of situations: in the following example, Ross could for instance hide from Rachel and then resume his initial plan. From an implementation perspective, situated reasoning is implemented through a separate, "local" plan, which passes its post-conditions to the main plan when it resumes. Situated reasoning may actually influence the unfolding of the main plan in several ways. Firstly, the returned

post-conditions may trigger re-planning in the original plan. Secondly, the actions corresponding to situated reasoning could in principle generate further bottom-up situations, even though we have not yet encountered this case in our simulations.

The other mechanism developed to cope with the bottom-up aspects of interactive storytelling is action repair (Geib and Webber, 1993). Action repair designates remedial action that is carried out to restore the executability condition of a failed action. For instance, if Ross wants to read Rachel's diary but she is using the diary herself, rather than re-planning another way of finding information about her, Ross can just wait for Rachel to leave the room (Figure 5).

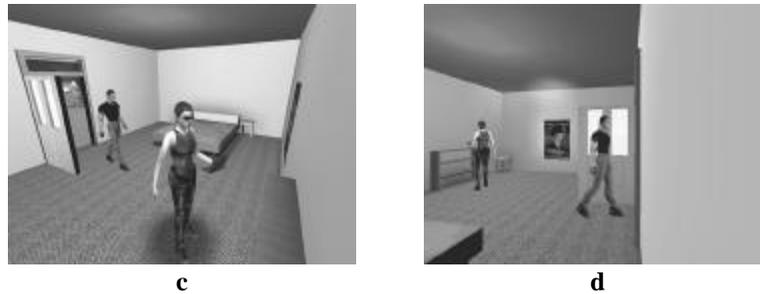


Figure 5: Example of Action Repair (a, b)

However, one of the fundamental questions is when to choose action repair over re-planning. We can consider another example where Ross wants to read Rachel's diary, but this time he is prevented from doing so because Phoebe is in the room. He thus cannot steal the diary, which essentially leads to two options: one is to look for another source of information about Rachel (re-planning), the other one is to repair the original action by waiting for Phoebe to leave or other actions on Phoebe. Repair should be based on generic and principled knowledge about action pre-conditions and action properties (e.g. duration), such as the fact that presence has a limited duration in time (unlike absence). This problem also arises because re-planning in our storytelling context is essentially a short-range change of action focus, rather than a radically new strategy. Action repair can be used to restore executability conditions under certain circumstances, especially in the case of competition for action resources, which is typical of interaction between actors.

Action repair is implemented through a separate mechanism, which however shares the same formalism and techniques as standard behaviours. HTN graphs are used to describe the set of sub-plans, which will be used for repair by the character. The outcome of the sub-plan, developed as a consequence to the action repair, does not need to be explicitly transferred back to the characters' main plan. In action repair, unlike situation reasoning, there is no need to return post-conditions to the original plan when it resumes. The reason is that in most cases studied so far, repair was targeting executability conditions, which depended on narrative objects. Communication between the action repair modules and the generic plans takes place through the side effects of action resources, whether these are narrative objects or characters.

5.2 Examples

In its current status, the system is able to generate short complete stories up to three minutes in duration. The dramatic action appears from Ross perspective (though the user can switch viewpoints to either of the characters' or even freely explore the stage while the plot is unfolding) and progresses until he asks Rachel out in what is the final scene. The story concludes with Rachel's (positive or negative) answer.

The following story instantiation presents the four characters, two as main role (Ross and Rachel) and two as secondary role (Phoebe and Monica). As presented previously, the characters are engaged in activities defined by plans represented by their own HTNs. Below is a sample story produced by the system (Figure 6).

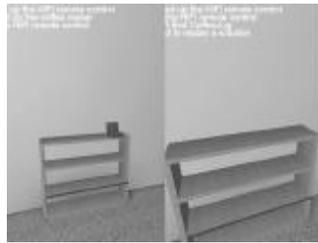
Ross wants to use Rachel's PDA to acquire relevant information regarding her preferences (a). In the meantime, Phoebe is going to prepare coffee (a). The user, watching the story unfolding, decides to interfere with both Phoebe's and Ross' plans by removing the coffee jug (b) and Rachel's PDA (c) from the virtual set. Unaware of user intervention, Phoebe goes to the coffee maker, though she cannot find the coffee jug. She thus has to re-plan a new activity (d). Phoebe decides to go to the shop (e, f, g) to look for a magazine (h). While Phoebe carries her own activities, Ross is going to read Rachel's PDA (i). After re-planning, Ross makes the decision (j) to go and talk to Phoebe (k), as she may provide him with the relevant information. When he arrives in the shop (l), Phoebe is already engaged in reading a magazine, and talking with Monica (m). As Ross does not want to jeopardise his plan by upsetting Phoebe, he decides to go and get himself a drink (n) until Phoebe finishes talking to Monica (o). As Ross was rather kind to Phoebe, she responds positively to his request by telling him Rachel's preferences, i.e. to offer Rachel a bouquet of roses (p). After succeeding in gathering important information, Ross goes to purchase his gift for Rachel (q). After buying the bouquet of flowers, he goes back to the flat (r) to offer them to Rachel (s). As she is alone, he goes and asks her out, which she *inevitably* accepts (t).



a



b



c



d



e



f



g



h



i



j



k



l



m



n



o



p



q



r



s



t

Figure 6: Example of a Full Story Instantiation

Another story can unfold before the user by modifying other factors. The second example we describe here is based on the same main storyline as the previous sample (Figure 7).

Ross wants to use Rachel's PDA to retrieve relevant information regarding her preferences. He goes to Rachel's bedroom, unseen by Phoebe, who is preparing some coffee. The user removes Rachel's PDA from the virtual set to alter the on-going storyline. After re-planning, Ross makes the decision to talk to Phoebe. Though, since the Ross' personality was pre-defined as being "tactless", Ross interrupts Phoebe regardless of what she is doing (a). As Ross was rather unkind to Phoebe, she decides to lie to him concerning Rachel's preferences, telling him to offer Rachel a box of chocolates (b). After succeeding in gathering important information, Ross goes to purchase his gift for Rachel from the shop (c). After buying the box of chocolates, he goes back to the flat to offer them to Rachel. As she is alone, he goes and asks her out, which she inevitably refuses (d).

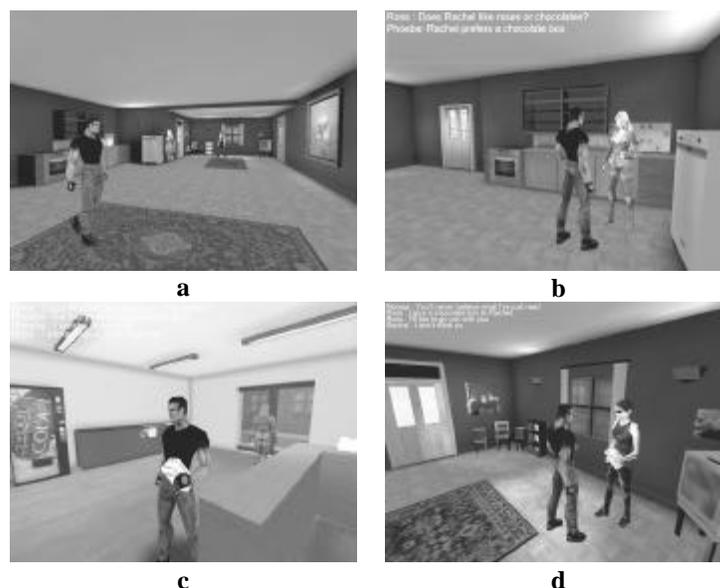


Figure 7: Example of a Story Variation

5.3 Discussion

With the inherent tension between interactivity and storytelling (Young, 2000), (Mateas, 1997), (Mateas, 2000), dramatisation of the unfolding story should convey meaningful details to the user through the generation of novel situations. For instance, the relevance of objects should be obvious from a narrative perspective, such as letters, flowers or a dagger. All these would be mainly targets for user intervention. However, the precise outcome of user intervention should not be directly accessible to the user, in order to preserve the richness of story generation. In other words, stealing the dagger from the set should lead either to the character abandoning its murder plans or to it finding other means, depending on the current circumstances. This, in our view, constitutes yet another principle of interactive storytelling: that interventions can have an impact at the plot development level (substantial alteration of the unfolding and the ending, such as murder vs. no murder) or at the situational level (the means by which goals are pursued and their dramatisation). Even though the individual actors' behaviours, as supported through their original plans, are a priori deterministic, there are several factors that contribute to the non-predictability of the plot from the users' perspective. These are: i) the initial spatial allocation of the virtual actors, ii) the duration of actors' actions and iii) the interaction between actors' plans.

At the beginning of the story, the characters are allocated random initial positions on the virtual set. Their plans are then triggered from these initial positions: hence, they might direct themselves towards the on-stage resources required at the early stage of their plan. For instance, if Phoebe's first activity were to go shopping, she would have to walk across the whole flat before leaving via the main door. This would leave many opportunities for Ross to intercept her if necessary. On the other hand, if her initial position happened to be very close to the main door, she would leave the set in no time and, if Ross needed to talk to her, (e.g., to get information about Rachel) he would have to find another alternative (i.e., through re-planning).

In addition, initial spatial location and characters' intrinsic speed in carrying elementary actions and displacements, are clearly inter-dependent. For instance, according to their initial locations and respective speeds, Ross might or not be able to catch up with Phoebe before she leaves the flat. As a consequence, obtaining information about Rachel would be affected, and thus Ross might have to follow a different course of action.

Because characters' behaviours are determined by plans, characters' interaction is emphasised through competition for resources of action, as part of the execution of their plans (Cavazza, 2001). This competition for resources has the potential to trigger a "chain reaction" of causal events. One classical example consists in competition for resources used in entertainment

activities, such as magazines or coffee machines, which have specific locations on stage and hence play an important role in the localisation of actors on stage. If a character is prevented from having access to resources (because it is used by another actor, or has been moved or removed by the user), it will have to re-plan another course of action involving a different activity. In doing so, it will often move across the virtual set, which increases the probability of spontaneous encounters with other characters, hence rising the potential for narratively meaningful situations. The latter point is supported by the fact that characters are aware of each other at various stages of the narrative action.

The existence of multiple actors naturally increases the probability of competition for resources and the generation of situations. In our current prototype, we have incorporated four autonomous actors, each with their own plan-based behaviour. Apart from competition for action resources, the interaction between characters' plans results in dynamic on-stage encounters between virtual actors that have the potential to create situations of narrative relevance.

6 CONCLUSION

We have described the various elements of an interactive storytelling system as well as the first results obtained from our prototype. Our system is essentially character-based, which largely determines the kind of AI mechanisms that support the implementation (i.e. planning systems). The dynamic behaviour of the prototype has shown the importance of both top-down behaviours and bottom-up control mechanisms that cope with emerging situations. This will lead us to revisit the notion of narrative control as introduced for instance by Young (Young, 2001), when scaling up the system. Further work will also be dedicated to the authoring aspects of interactive storytelling, with the prospect of getting feedback from authors and scriptwriters. Finally, we will be investigating more complex stories with multiple storylines, which, beyond the authoring and narrative control problems, will also bring specific requirements on the real-time presentation of the story, such as automatic camera control.

REFERENCES

- Bonet, B., and Geffner, H. (1999) "Planning as Heuristic Search: New Results", *Proceedings of the European Conference on Planning (ECP'99)*, pp. 360-372
- Cavazza, M., and Palmer, I. (1999) "Natural Language Control of Interactive 3D Animation and Computer Games", *Virtual Reality*, 4, pp. 85-102
- Cavazza, M., Charles, F. and Mead, S.J. (2001) "Agents' Interaction in Virtual Storytelling", *Intelligent Virtual Agents 2001 Conference*, Madrid, Spain
- Charles, F., Mead, S.J., and Cavazza, M. (2001) "User Intervention in Virtual Interactive Storytelling", *Proceedings of VRIC 2001*, Laval, France
- Geib, C., and Webber, B. (1993) "A consequence of incorporating intentions in means-end planning", Working Notes – AAI Spring Symposium Series: *Foundations of Automatic Planning: The Classical Approach and Beyond*. AAAI Press
- Grabson and Braun (2001) "A Morphological Approach to Interactive Storytelling", *Proceedings of Cast01, Living in Mixed Realities*, Sankt Augustin, Germany, pp. 337-340
- Knight, K., and Rich, E. (1991) *Artificial Intelligence*, 2nd Edition, McGraw Hill
- Korf, R.E. (1990) "Real-time Heuristic Search", *Artificial Intelligence*, 42:2-3, pp. 189-211
- Mateas, M. (1997) "An Oz-Centric Review of Interactive Drama and Believable Agents", *Technical Report CMU-CS-97-156*, Department of Computer Science, Carnegie Mellon University, Pittsburgh, USA
- Mateas, M. (2000) "A Neo-Aristotelian Theory of Interactive Drama", *AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2000*, AAAI Press
- Nakatsu, R., and Tosa, N. (1999) "Interactive Movies", In: B. Furht (Ed), *Handbook of Internet and Multimedia – Systems and applications*, CRC Press and IEEE Press
- Nilsson, N.J. (1998) *Artificial Intelligence: A New Synthesis*, Morgan Kaufmann, San Francisco
- Pearl, J. (1984) *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, Addison-Wesley, Reading, Massachusetts, USA
- Pemberton, J.C., and Korf, R.E. (1994) "Incremental Search Algorithms for Real-Time Decision Making", *Proceedings of the 2nd Artificial Intelligence Planning Systems Conference (AIPS-94)*
- Raskin, R. (1998) "Five Parameters for Story Design in the Short Fiction Film", *P.O.V.*, n. 5

Sgouros, N.M., Papakonstantinou, G., and Tsanakas, P. (1996) "A Framework for Plot Control in Interactive Story Systems", *Proceedings AAAI'96*, Portland, AAAI Press

Szilas, N. (1999) "Interactive Drama on Computer: Beyond Linear Narrative", *Papers from the AAAI Fall Symposium on Narrative Intelligence, Technical Report FS-99-01*, AAAI Press

Young, R.M. (2000) "Creating Interactive Narrative Structures: The Potential for AI Approaches", *AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2000*, AAAI Press

Webber, B.N., Badler, N.I., Di Eugenio, B., Geib, C., Levison, L., and Moore, M. (1994) "*Instructions, Intentions and Expectations*", IRCS Technical Report 94-01, University of Pennsylvania

Young, R.M. (2001) "An Overview of the Mimesis Architecture: Integrating Narrative Control into a Gaming Environment", *AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2001*, AAAI Press