GENERATING DYNAMIC STORYLINES THROUGH CHARACTERS' INTERACTIONS

Fred Charles, Steven J. Mead and Marc Cavazza
School of Computing and Mathematics
University of Teesside
Middlesbrough, TS1 3BA,
United Kingdom
E-mail: {f.charles, steven.j.mead, m.o.cavazza}@tees.ac.uk

KEYWORDS
Interactive Storytelling, Artificial Actors, AI-based Animation, Computer Games.

ABSTRACT

In this paper we describe a fully implemented prototype for interactive storytelling using the Unreal™ engine. We describe the important mechanisms involved in the variability of plot instantiations, within a scenario of sitcom genre. We also provide an evaluation of the concepts of how the dynamic interactions between agents and/or the user influence the generation of story, with first results of examples

INTRODUCTION

In this paper, we present results from a fully implemented storytelling prototype, which illustrate the generation of variants of a generic storyline. These variants result from the interaction of autonomous characters with one another, with environment resources or from user intervention.
The development of artificial actors and AI-based animation naturally leads to envision future interactive storytelling systems. A typical interactive storytelling system would be based on autonomous virtual actors that generate the plot through their real-time interaction. Besides, the user should be allowed to interfere with the ongoing action, thereby altering the plot as it unfolds.
Many interactive storytelling models have been proposed: user-centred plot resolution (Sgouros et al. 1996), character-based approaches (Young 2000) (Mateas 2000), anytime interaction (Nakatsu and Tosa 1999) and the need for narrative formalisms (Szilas 1999). Previous work has identified relevant dimensions and key problems for the implementation of interactive storytelling, among which: the status of the user, the level of explicit narrative representation and narrative control, the modes of user intervention, the relations between characters and plot, etc (Cavazza 2001).
Some of these problems derive from the inherent tension between interaction and narrative (Young 2000) (Mateas 2000). Interactive systems demand user involvement but often at the expense of a real storyline; on the other hand, a strong narrative dimension is traditionally conceived with a user as spectator rather than being actively involved. Our solution to this problem consists in limiting the user involvement in the story, though interaction should be allowed at anytime. This is achieved by driving the plot with autonomous characters' behaviours, and allowing the user to interfere with the characters' plans. The user can interact either by physical intervention on the set or by passing information to the actors (e.g., through speech input).
In the next sections, we will introduce the important concepts of character-centred storytelling as well as a brief description of our implementation. Results of variants in story generation are illustrated with an example.

CHARACTER-BASED STORYTELLING

The storyline for our prototype is based on a simple sitcom-like scenario, where the main character "Ross" wants to invite the female character "Rachel" out on a date. This scenario tests a narrative element (i.e. "Will he succeed?") as well as situational elements (the actual episodes of this overall plan that can have dramatic significance, e.g., how he will manage to talk to her in private if she is busy, etc.). Our system is driven by characters' behaviours. These actually "compile" narrative content into characters'

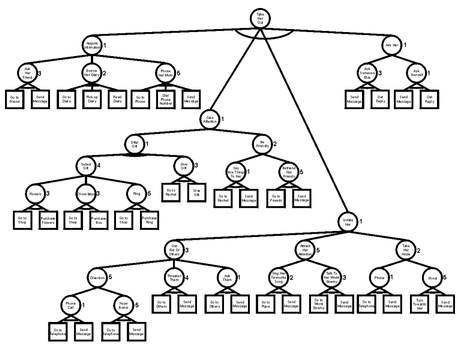

Figure 1: An HTN for the "Ross" Character Role

behaviours, by defining a superset of all possible behaviours, represented by a plan for each character. Dynamic choice of an actual course of action within this superset is the basis for plot instantiation (Young 2000). In that sense, this addresses the causality/choice duality described by Raskin (1998) in storytelling, though this choice takes place within the limits of the formalism used to represent possible behaviours, which is a plan-based formalism (Young 2000). This can be illustrated by considering the overall plan for the character Ross (see Figure 1).

In order to invite Rachel, he must for instance acquire information on her preferences, find a way to talk to her, and finally formulate his request (or having someone acting on his behalf, etc.). These goals can be broken into many different sub-goals, corresponding to various courses of action, each having a specific narrative significance.

The initial storyline should actually determine not only the main character's plan, but those of other characters as well. The problem of dependencies between characters' roles has actually been described within modern narratology, though not to a formal level. Narrative functions can be refined into bipolar relations between a couple of actors, emphasising the asymmetry in their roles (Barthes 1966). We have adopted this framework to define the respective behaviours of our two leading characters. We started with the overall narrative properties imposed by the story genre (sitcoms). In terms of behaviour definition, this amounts to defining an "active" plan for the Ross character (oriented towards inviting Rachel) and a generic pattern of behaviour for Rachel (her day-to-day activities).

OVERVIEW OF THE PLANNING SYSTEM

Previously, we introduced the need for describing characters' plan in a planning formalism that can incorporate narrative content. We have opted for Hierarchical Task Networks (HTN) as the most appropriate planning formalism, and have implemented a form of HTN planning (Nau et al. 1998). In this section, we briefly introduce HTN and discuss their application to characters' behaviours in interactive storytelling.

HTN can be seen as networks representing (generally ordered) tasks decomposition. The top-level task is also the main goal; each task can be decomposed into sub-tasks, which can be further refined until description of primitive actions. HTN are commonly represented as AND/OR graphs.

HTN planning is based on forward search, while being goal-directed at the same time, as the top-level task is the main goal. An important consequence is that, since the system is planning forward from the initial state and expands the sub-tasks left-to-right, the current state of the world is always known, which makes it possible to perform reasoning about it. In interactive storytelling, each primitive action, specifying the sub-task elements, is performed on stage by the character, and perceived from the spectator's point of view as a sequence of meaningful events.

HTN planning is well suited to story representation regarding that sub-tasks are largely independent as they represent various stages of the story. Decomposability of the problem space derives from the inherent decomposition of the story into various scenes (i.e. sub-tasks). This satisfies a classical representation for stories (Schank and Abelson 1977), as being a succession of on-stage situations. Moreover, a single HTN corresponds to several possible decompositions of the main task, because OR nodes correspond to alternative sub-tasks and the sub-tasks considered by an AND node can be subject to several orderings.

In the present context, each ordered decomposition will constitute the basis for a character's plan. Figure 1 illustrates the behaviour representation for the main character Ross. His main goal is to "ask Rachel out". This task can be decomposed into various sub-tasks, from "finding information about her" to "offering her gifts", etc. Each of these sub-tasks is recursively decomposed until primitive actions are obtained. For instance, in the graph of Figure 1, the "acquire information" node can be expanded into different sub-goals, such as "read Rachel's diary" or "ask one of her friends". In turn, "Read Rachel's diary" is decomposed into primitive actions (locating the diary, going to the diary, picking it up, reading it). These are the animation sequences carried out on-stage by the virtual actors, relying on Unreal™'s built-in mechanisms and animation libraries.

In interactive storytelling, several characters may interfere with each other's plans as their sub-tasks may generate

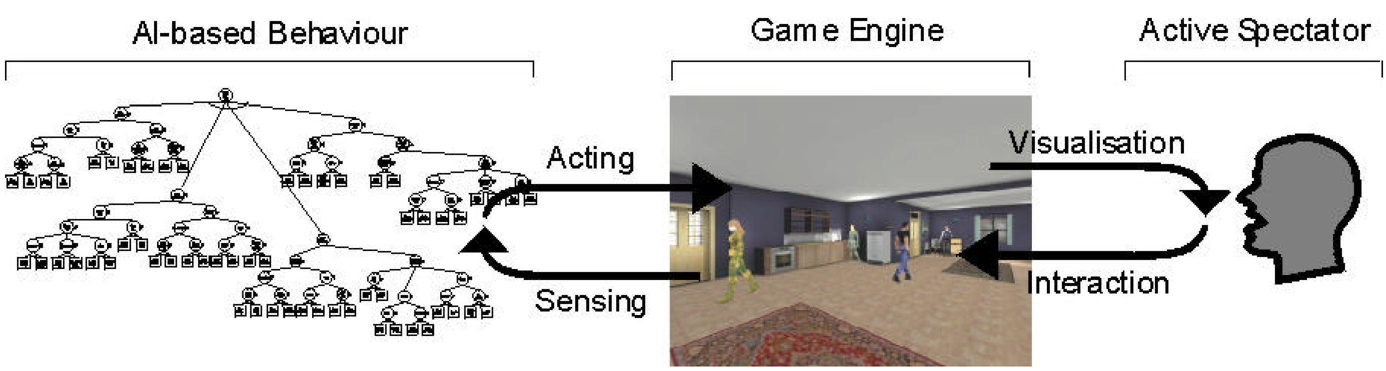

Figure 2: System Overview

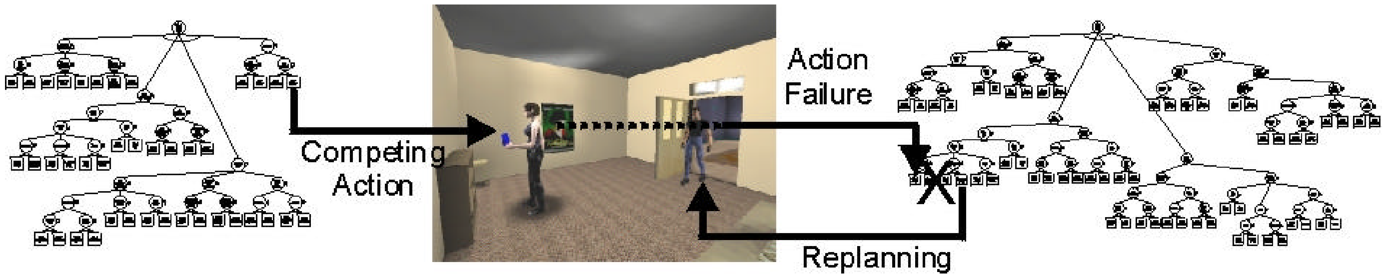

Figure 3: Characters' Interaction On Stage

encounters with other characters. The main catalyst of such interaction is competition for action resources, these resources being pre-conditions for actions. In this case, the story can only carry forward if the character has re-planning capabilities. When planning and execution are interleaved, re-planning can take place through direct backtracking in the HTN. For instance, Ross is executing the sub-task of reading Rachel's diary, though she might be using it. Failure of a primitive action is back-propagated to its parent task and the search process will then backtrack to generate an alternative decomposition until a new primitive action is generated and executed, in this case, "going to ask her friend" (see Figure 3 above).

In any case, failed primitive actions (looking for an object that is missing, giving an unanswered phone call, etc.) cannot be undone, as they have been played on stage. Action failure is indeed part of the story itself, as well as an important aspect of interactivity. This is why the dramatisation of actions must take their possible failure into account and store corresponding animations.

The "virtual sitcom" prototype described in this paper has been developed using the Unreal™ game engine. The Unreal™ environment provides most of the user interaction features required to support user intervention in the plot, such as navigating about and interacting with objects within the virtual set and its use has been previously reported in prototyping interactive storytelling (Young 2000). The system has been fully implemented as a set of template C++ classes, which can be used as native functions from within UnrealScript™, Unreal™'s scripting language.

DETERMINANTS OF NON-PREDICTABILITY IN STORY GENERATION

Though the mechanisms behind story generation are formally deterministic, there are many factors that contribute to variations throughout a story unfolding.

THE INITIAL SPATIAL ALLOCATION OF VIRTUAL ACTORS

To illustrate factors in story variation, we must first consider the impact of initial conditions. The virtual actors are allocated initial positions at various locations on the virtual set. Their first sub-task's primitive action is then triggered from these initial positions and they might direct themselves towards the on-stage resources required for their plan.

For instance, if Phoebe's first activity is to do some shopping, she would have to traverse the whole flat before reaching the exit. Ross will then have many opportunities to go and talk to her, if such is his plan (e.g., to get information about Phoebe). On the other hand, if her initial position is closer to the exit door, she will leave the set in no time (Figure 4) and Ross would have to replan another alternative. Though, the above example describes a negative outcome to the success of Ross' plan, initial positions could in contrary be more beneficial if his intention was to avoid meeting with another character. A similar example, depending on Rachel's initial location and activity, she might spot Ross talking to Phoebe and mistakenly believe that they are having an affair, leaving the flat in anger.

Whilst selecting sub-tasks during search process, as production of solution sub-plans, each secondary character, governed by generic activity plans (i.e., in our scenario, all characters but Ross), has his first task selected randomly.

The evident inter-dependence between initial spatial location and characters' displacements (determined by the initial task, such as shopping, meeting friends, watching TV, etc.) has a great potential for creating many kinds of encounters between characters, or having characters missing one another.

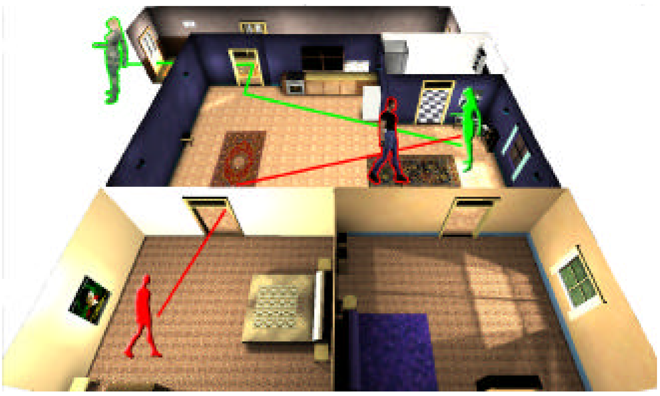

Figure 4: Influence of spatial location.

THE DURATION OF ACTORS' ACTIONS

From the previous description of spatial distribution, we can notice that the duration of actions also becomes an important factor in the eventuality of characters' encounters. To extend the previous example when Ross wants to talk to Phoebe, he might or not be able to catch up with her before she leaves the flat for some shopping. This will in turn

affect him obtaining information about Rachel, from which a different course of action will follow.
The Unreal™ game engine provides a resonably accurate mean of defining the duration of actions (based on fractions of a second). Durative actions include activities that characters perform on environment resources, such as using objects (e.g. read book, make coffee, watch TV, etc.) or interacting with other characters (e.g. conversation face-to-face, over the phone, etc.)
Although, specification of duration of various actions is included to the system, speed of actions (e.g. walking pace) might become relevant in future developments. For instance, we could imagine Ross running after Phoebe to catch up before she leaves the flat.

THE RANDOM OUTCOME OF PRIMITIVE ACTIONS

Actions performed on-stage by the characters have direct impact on the state of action resources. A random factor can be introduced on actions that do not require specific on-stage interactions with characters, despite making use of resources as a communication medium (i.e. virtual interaction). These mechanisms are possible on primitive actions' outcomes that are envisaged a priori, in order for the system to select one out of the pre-defined set.
For instance, when Ross organises his evening out, he may phone a restaurant to book table. Different cases, modifying the executability conditions of the primitive action, may occur such as the phone is engaged, nobody answers, the correspondent puts the phone down on him, etc. From each outcome, a different new plan could result: he might decide to try again after a while, to go and book the table in person instead, to change venue (see Figure 5).

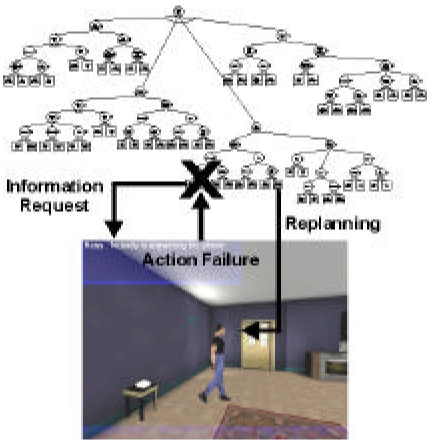

Figure 5: Random Outcome of Action

THE INTERACTION BETWEEN ACTORS' PLANS

We have seen previously that characters' behaviours are represented by HTN plans, which are independently defined for each character. In our approach to interactive storytelling, the basis for characters' interaction is inherent to their competing for action resources, as part of the execution of their plans.
For instance, Ross might want to learn more about Rachel by reading her diary, but she might be using it herself, preventing him to do so (see Figure 6). He could try to talk to Phoebe, but she could be busy talking to Monica. This makes characters themselves become resources for actions from others. This competition for resources has the potential to trigger a causal "chain reaction" (though causality is not explicitly represented in a character-based approach).
A typical example consists in competition for action resources used in entertainment activities, such as multimedia systems or vending machines, which are part of all character's sets of behaviours. Resources have specific locations on stage emphasising their role in the localisation of actors on stage. If a character is prevented from having access to an item, because it is used by another character, or has been (re)moved by the user, it will have to re-plan a new activity. Consequently, the character must move across the stage, increasing the probability of dynamic encounters with other characters, resulting in a whole range of situations.

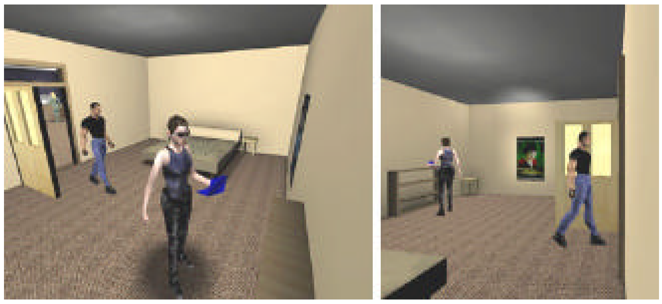

Figure 6: Competition for action resources.

The existence of several characters naturally increases the prospect of competition for action resources and the generation of situations. In the 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 may also result from dynamic on-stage encounters between characters that have the potential to create situations of narrative relevance. These constitute the "bottom-up" aspect of interactive storytelling: as it is not taken into account by plan-based behaviours.
Thus, there is a need for specific mechanisms as an extension to HTN planning, such as situated reasoning and action repair, in order to deal with dynamic situations arising from characters' interactions.
For instance, if Ross encounters Rachel at an early stage of the plot, before he has acquired information on her, his plan has to be interrupted: he can choose to hide from Rachel until he knows enough about her. Another example consists in simple action repair as an alternative to re-planning. For instance, if another actor is competing for action resources, Ross can simply wait for this actor to complete its current action. The combined use of situated reasoning and action repair with search-based planning offers additional flexibility.

USER INTERVENTION

The user watches the story as a spectator. At this stage he can follow the story from any character's perspective or navigate on the virtual set while the action is in progress. From his understanding of the current action, he can choose whether to interfere or not with the characters' goals. Characters' actions are dramatised through the timing of appropriate animations. Because the actors are playing a role rather than improvising, their actions are always narratively meaningful. Hence, if a character moves towards a given object, it is likely to bear significance for the story and can be the target for user intervention. For instance, if the user sees Ross moving towards Rachel's diary, he can choose to steal or hide that diary (see Figures 7 and 8).

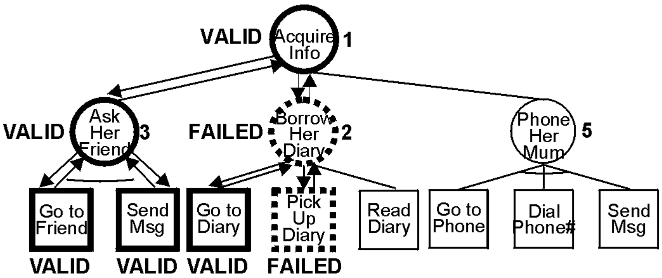

Figure 7: Re-planning on Action Failure

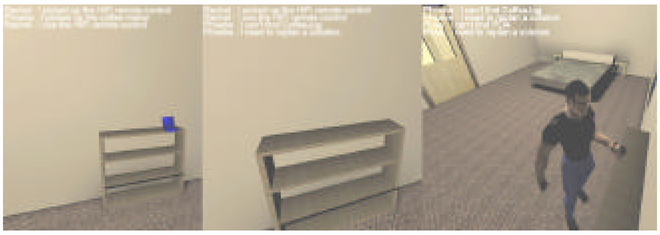

Figure 8: Dramatisation of Action Failure

The user can intervene by either acting on physical objects on-stage that bear narrative relevance (and are often obvious, such as keys, letters, gifts, weapons, etc.). These objects being resources for actions will force the character into re-planning or action repair, which, being dramatised as well, will create a new course for the plot. The other mode of interaction consists in influencing actors using speech recognition. This form of influence will become the main one in further developments of the system and will include:

1. providing information needed by the actors to complete their plans (e.g. Rachel's preferred gifts) (see Figure 9)
2. giving doctrine advice that influences the personality of an actor (i.e. recommending a friendly behaviour towards certain characters)
3. trying to alter the mood of a character
4. 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

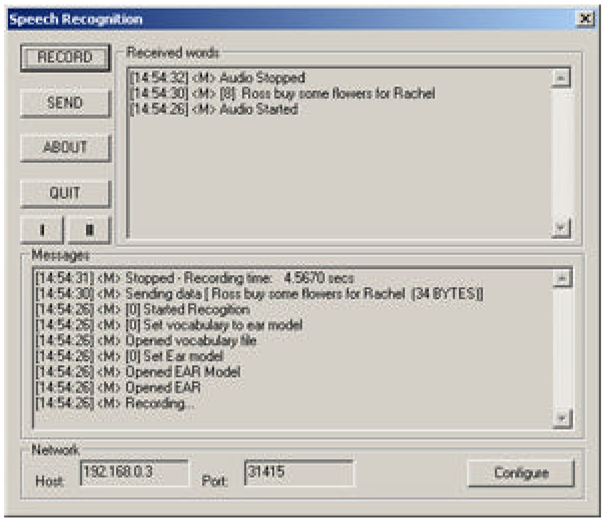

Figure 9: Speech-driven user intervention

RESULTS

While the conditions for character interaction lie in the on-stage spatio-temporal instantiation of the storyline, additional mechanisms are required to recognise these interactions and propagate their consequences.

Figure 10 (see next page) illustrates the chain of events that successively constitute the overall story. To acquire information about Rachel, Ross goes to her bedroom to read her diary, though discovers (a) that Rachel is already using it, so he replans (b) and goes to the phone (c). Rachel's mum lies to him, telling Ross that Rachel prefers chocolates. The user oversees the scene, and removes the resource (the box of chocolates) from the set. Ross arrives in the shop, where Monica and Phoebe are chatting. He engages in a temporary activity avoiding their attention (d), going to the vending machine. After the girls have left the shop, he goes to collect the gift (e), but has to replan as the resource was removed from the stage. He goes to get a bouquet of roses (f). He then needs to organise the night out by phoning the restaurant (g). The phone might be engaged, so Ross tries another venue, successfully this time. On his way to see Rachel, he discovers she is engaged in a conversation with Monica, so he waits for them to have finished (h). As Rachel becomes available, Ross goes to her, offers her the bouquet, and asks her out (i).

This example illustrates the interaction of the two main characters' plans. These plans are designed from narrative principles. It appears that exploring actors' behaviour in storytelling is more feasible within narrative genres that display the simplest storylines; as such developing "virtual sitcoms" seems a relevant first step in the pursuit of interactive storytelling. As its own name suggests, sitcom standing for "situation comedy", a significant fraction of the story interest arises from the situations into which the actors find themselves. For instance, the fact that Rachel could misunderstand the situation where Ross was talking to Phoebe, then triggering the emotional reaction of jealousy. Her state of mind being modified (i.e., she gets upset), Rachel will then leave the room. The succession of "small" interesting situations is a mechanism for cause-and-effect

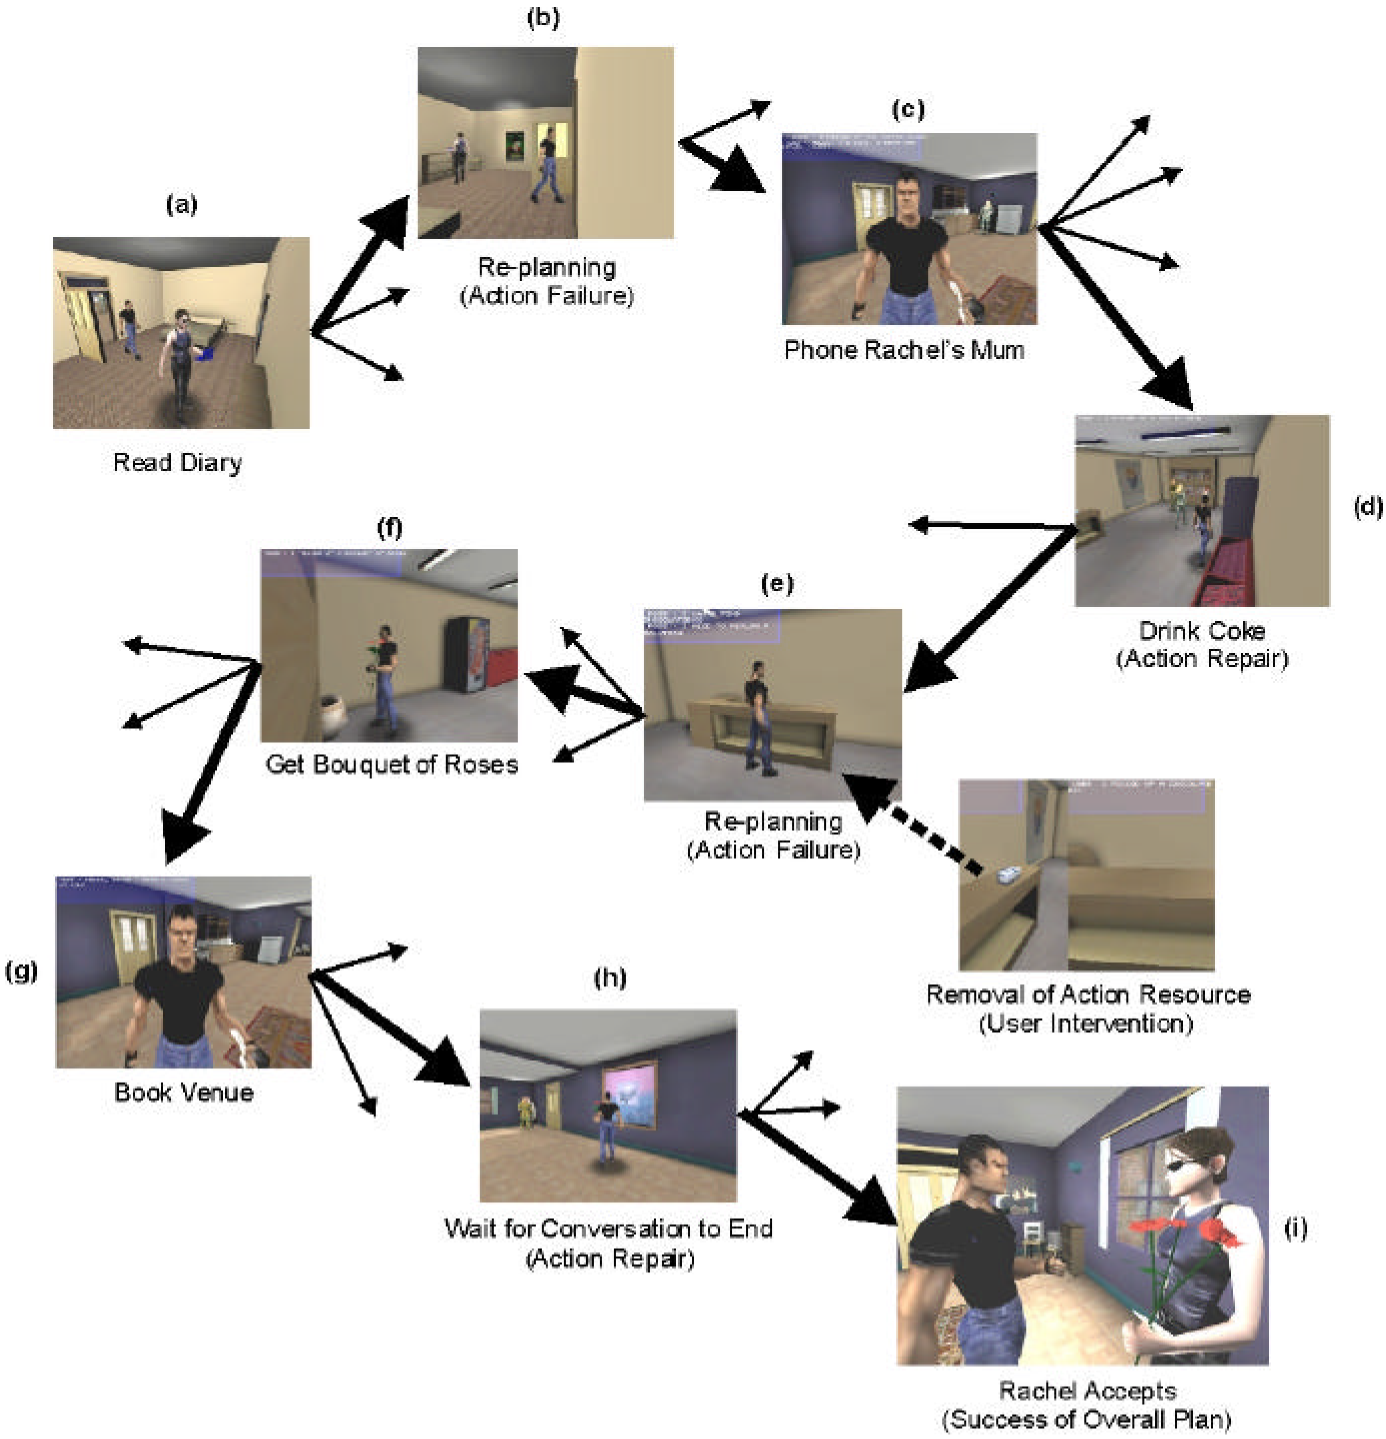

Figure 10: Causal Chain of Events Emerging as the Story Unfolds

relationships (Raskin 1998), providing the basis for dramatic story generation.

Though plans are designed from global narrative principles, considering the current story genre, they are run independently. The bipolarity between the characters' plans was defined to emphasise the asymmetry in their roles (Barthes 1966). The overall narrative properties imposed by the story genre defined interaction between the main character's "Ross" and its supporting role's "Rachel" behaviours. The generic pattern of Rachel's day-to-day activities may interfere with Ross' "active" plan, as illustrated when Ross want to read Rachel's diary while she is already using it. This interactivity between characters' behaviours must be emphasised visually when it demonstrates narrative relevance.

As part of the story believability, mechanisms in action recognition will help to make the characters' emotional status visible to the user, so he can understand their interactions. The variations in characters' emotions and moods must emerge from situations relevant to the story genre without changing their overall personality profile. For instance, Rachel's mood towards other characters can vary according to the meaning of their actions for Rachel (e.g., jealousy). The next generation of real-time animation engine (i.e., Unreal2™) will help representing facial expressions, or detailed non-verbal behaviour (e.g., body postures) to improve the dramatisation of events through physical characterisation.

Above the planning and interleaving of actions, explicit situated mechanisms for reactive behaviours (Geib 1994) are needed in order to deal with specific situations (e.g., Ross suddenly meets Rachel on his way). This implies high-level action recognition of interactions between characters' behaviours. If a narratively meaningful (considering the story genre) situation arises, the mechanism would act on the character's current plan by ordering a re-planning of its action.

CONCLUSION

We have shown that, although actor's behaviours are deterministic, the interaction between actors could considerably contribute towards story variability. This degree of unpredictability conditions the generation of dramatic situations. The character-centred approach has the advantage of being modular and extendable to many actors.

Further work is to be dedicated to developing more complex storylines within differing genres, scaling up using multiple plans for each actor to increase characters' interactions and narrative function recognition.

Acknowledgements

We are indebted to Eric Jacopin for his advice on AI planning formalisms: any remaining misconceptions are the authors' sole responsibility.

The following credits are given for the models used for the actors:

"Ross" created by Brian Collins, converted from Half-Life to Unreal by Kempel.

"Rachel" created by Austin, converted from Half-Life to Unreal by Usaar.

"Phoebe" and "Monica" created by Roger Bacon.

REFERENCES

Barthes, R. 1966. "Introduction à l'Analyse Structurale des Récits" (in French), Communications, 8, pp. 1-27.

Cavazza, M., F.Charles and S.J. Mead. 2001. "Characters in Search of an Author: AI-based Virtual Storytelling". First International Conference on Virtual Storytelling, Avignon, France.

Geib C., "The Intentional Planning System: ItPlans". Proceedings of the 2nd Artificial Intelligence Planning Conference, AIPS-94, 1994, pp. 55-64.

Kambhampati S. and Hendler J.A. A Validation Structure Based Theory of Plan Modification and Reuse, Artificial Intelligence, 55 (2-3), pp. 193-258, June 1992

Mateas, M.. 2000. "A Neo-Aristotelian Theory of Interactive Drama", AA,AI Spring Symposium in Artificial Intelligence and Interactive Entertainment, AAAI Press.

Nakatsu R. and N. Tosa. 1999. "Interactive Movies". In: B. Furht (Ed), Handbook of Internet and Multimedia – Systems and applications, CRC Press and IEEE Press.

Nau D.S., Smith S.J.J., and Erol K. 1998. Control Strategies in HTN Planning: Theory versus Practice. In AAAI-98/IAAI-98 Proceedings, pp. 1127-1133.

Pearl, J., 1984. "Heuristics: Intelligent Search Strategies for Computer Problem Solving". Reading (Massachusetts), Addison-Wesley.

Raskin, R. 1998. "Five Parameters for Story Design in the Short Fiction Film", P.O.V., n. 5.

Schank, R.C. and Abelson, R.P., 1977. Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures. Hillsdale (NJ): Lawrence Erlbaum.

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

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

Tsuneto, R., D. Nau and J. Hendler. 1997. "Plan-Refinement Strategies and Search-Space Size", Proceedings of the European Conference on Planning, pp. 414-426.

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