Exploring Scalability of Character-based Storytelling

Fred Charles and Marc Cavazza
School of Computing, University of Teesside
Middlesbrough TS1 3BA, United Kingdom
{f.charles, m.o.cavazza}@tees.ac.uk

Abstract

Interactive Storytelling is establishing itself as a major application of virtual embodied characters. To achieve further progress in the field, some authors have suggested that it was necessary to break the 10-minute barrier for story duration, while preserving story pace.

In this context, understanding scalability issues is an essential aspect of the development of future Interactive Storytelling technologies. Scalability can be defined as the production of a richer narrative which follows the scaling-up of the Artificial Intelligence representations for plot structure or characters' roles.

We have formalised narrative events in terms of "film idioms" which are dynamically recognised as the story is generated. This enabled us to stage a number of experiments in which we modified several determinants of scalability, such as the number of feature characters or the complexity of their roles and recorded subsequent narrative extension, through the number of film idioms generated.

1. Introduction

Interactive Storytelling (IS) is developing rapidly as a major application for autonomous characters. Early research in the field [8] [14] [16] has identified key problems, such as narrative control, the duality between character and plot, the tension between interactivity and storytelling, etc. In a similar fashion, there appears to be a convergence in underlying techniques, such as planning systems [3] [14]. Now that the field is gaining in maturity, an increasing number of IS prototypes is being developed, such as [3] [6] [9] [11] [12] [15].

Yet, one essential condition to develop IS as a technology is the ability of current approaches to scale up to implement more complex interactive narratives.

Mateas and Stern [7] were the first to introduce some simple but widely accepted metrics to characterise scalability. Their first proposal consists in being able to generate a 10-minute story with approximately one beat per minute. The pace constraint of one beat per minute introduces some kind of normalisation over the duration of actions, which otherwise could be arbitrarily extended, i.e. by increasing the size of the virtual stage on which the actors are playing the story. The 10-minute duration is a good empirical assessment of the current state-of-the-art. Of course, one could argue that generating a 4-minute high-quality story is already a formidable challenge: however, the real meaning of the 10-minute metric consists in an assessment of the robustness and scalability of the Artificial Intelligence techniques involved, which include the knowledge acquisition and authoring aspects.

The first version of our IS prototype [3] only produced stories that were approximately 3 minutes in length. Although this was also a practical consequence of the effort put into authoring the narrative representations supporting the specific AI approach used, only a scaled up version could evaluate the appropriateness of some of its underlying hypotheses.

By contrast, Mateas and Stern's *Façade* system [9] is able to generate much longer stories, though there is significant amount of authoring and scripting involved in their system, which is justified by design considerations.

In this paper, we explore the determinants of IS scalability within our own approach, described as character-based storytelling [3]. Very little work to date has been dedicated to either scalability or evaluation (the former being in our view a pre-condition for the latter). Some previous work has laid the foundations for evaluating story progression in IS, though it has not been used explicitly to that purpose. These include: the neo-Aristotelian model of story progression of Mateas [7], the suspense model of Young [13] and, although

AAMAS'04, July 19-23, 2004, New York, New York, USA.

not specific to storytelling, emotional models based on plan progression [5]. The latter have potential applications precisely because of the popularity of planning techniques in IS.

However, work on scalability remains largely unpublished to date. It seems essential to progress on these aspects and, as there are no accepted formal frameworks for studying IS scalability, we will be promoting a rather empirical approach in the first instance.

2. Character-based interactive storytelling

We have developed in previous work an approach to IS known as character-based storytelling [3]. This reflects the fact that the interactive narrative is not driven by a single plot model controlling the artificial actors: instead, the story is described through the individual roles of each actor. These roles, rather than the overall plot, are formalised using planning techniques, in our case Hierarchical Task Network (HTN) planning [10]. The various sub-tasks of the HTN correspond to various components of the narrative and terminal actions of the HTN give rise to the animations which are actually staged and constitute the story presentation (Figure 1). This approach has however led us to adopt a more nuanced view on the duality between character and plot as traditionally described in the IS literature. Within character-based storytelling, the central character's role is actually determinant for the overall plot although other feature characters also contribute to it (or even add elements of parallel plots).

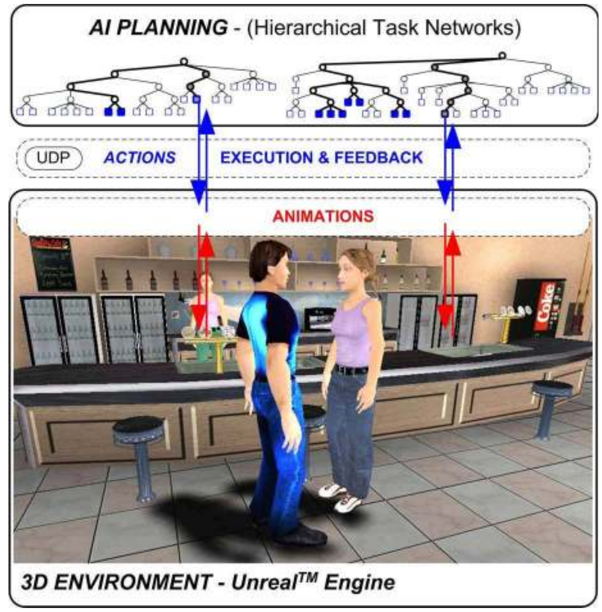

Figure 1. Architecture of the character-based storytelling system.

One key aspect of IS is to go from event sequences, i.e. the actions generated by the actors' plans to the staging of an actual 3D animated story eventually indistinguishable from an authored 3D animation, which responds dynamically to the unscripted interaction between characters (or to user interventions) in a way which is not directly predictable by the spectator. The latter corresponds to the actual staging of the story, which, beyond the sequence of narrative actions, controls the dramatisation of narrative events through dynamic camera positioning and "real-time editing".

The importance of these staging elements varies

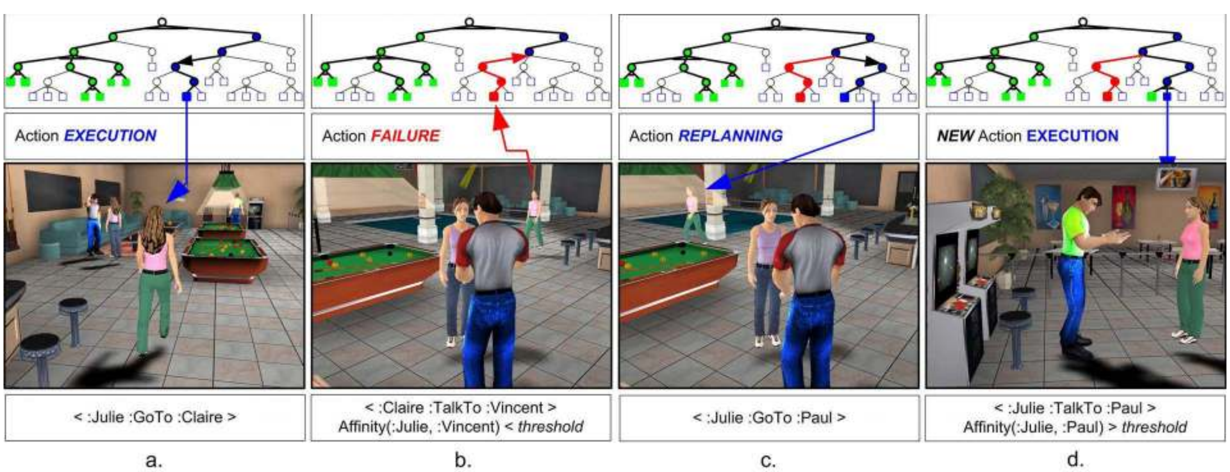

Figure 2. Situations arising from the dynamic interactions between characters.

according to story genres and we have dedicated our research to sitcoms, with an emphasis on short-term humorous situations. This can be refined into two essential aspects: i) the creation of situations, described as the emergence of non-scripted actions involving specific actors (for instance in Figure 2, Julie wants to invite Claire to her party, but refrains from talking to her as she is with Vincent, whom she dislikes) and ii) action failure, a negative outcome for a feature character's action, which has to be properly displayed and dramatised to contribute to the spectator's experience (e.g., a character has left the bottle he was bringing to a party unattended, and this bottle has been drunk by another actor).

The situations so generated are metaphorically speaking the "cross-product" of the individual plans representing the characters roles. These plans accommodate for possible interaction between characters: through social actions, at the initiative of one of the actors, or through the sharing of resources (narrative objects and places). But the occurrence of these encounters is not predictable by the spectator as it depends on several factors: initial spawning of the virtual actors, choice of actions, characters' personality traits and affinity with other characters, actions' starting time and duration, etc. Random initial conditions associated with a deterministic baseline plot have the potential to generate many unpredictable situations. To maintain the consistency of our experiments, story variability arises from interactions between actors rather than user intervention (however, the same user interaction mechanisms than those described in [2] are available). An in-depth technical description of the system and the formalism used can be found in [2] [3].

We have devised a new IS system to support our investigations of scalability in storytelling. This extends our previous work on virtual sitcoms [2] by providing an entirely new cast, storyboard and 3D environment, with a greater number of actors and more complex storylines.

This sitcom takes place in a single environment (a bar/disco) and features up to three main characters and four supporting characters. By contrast, our previous work featured only one main character with two supporting characters. Most importantly, the average complexity of the characters plan has been increased. Yet, the objective for this new setting is to reach a sufficient complexity to explore the phenomena accompanying the scaling-up of an interactive story. A gross estimation of the overall system complexity can be gained by the total count of HTN's terminal actions which appear as characters' behaviours (182 possible actions), as compared with our earlier system (48 possible actions). The principal character's role consists in organising a party, for which it has to find a suitable venue, decide whom to invite and whom to prevent from coming, get food and drinks, etc. Other feature characters include a barmaid running the day-to-day business in the bar, as well as a male character looking for a date.

3. An experimental approach for scalability

Our experiments explored the development of the generated narrative as a function of the baseline plot description, represented through characters' roles as outlined above.

Figure 3. <Subject-Verb-Object> syntax is used to recognise the situations taking place within the virtual environment.

The first step is to find a method to formalise the recognition of situations from the presentation perspective, i.e. as the story appears in the eye of the spectator. The unfolding of the story could be tracked by following terminal actions activated by the various HTN, but this would not capture the actual story appearance. This requires a representation specific to the actual staging of events, which should also be related to the nature of generated situations. The notion of idiom, as used in film studies [1], captures a situation taking place on stage, whatever its origin. It consists in essence in a grammar-like representation for story events using a *S-V-O* syntax, where *S* is the main character featured, *V* the action it is performing and *O*

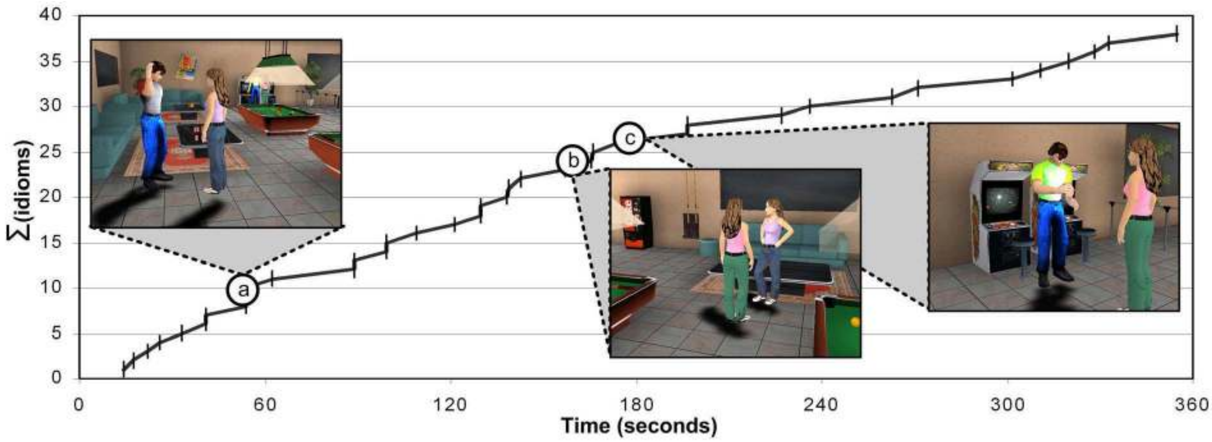

Figure 4. Our experiments are based on the interpretation of the evolution of the cumulative sum of idioms, related to situations where the feature characters are involved.

the narrative object on which it takes place (which can be another character). This formalisation facilitates situation recognition from the real-time tracking of the various characters by the system (Figure 3).

The formalism of idioms, especially as it distinguishes between characters and actions, however allows them to be linked to the elements that have a narrative value (either by identifying the key actions, taking part in narrative functions, or by identifying the feature characters). The proof of the validity of idioms to detect the elements of the story is in fact in their use to control the camera, i.e. the focus of narrative presentation, as adopted by several researchers in the field of IS [4] [17].

This representativity of idioms is however gauged to the quality of measures they are associated with, i.e. the level of refinement at which we can classify the idioms according to their relevance. A rough measure of idioms overevaluates story beats or narrative actions. In our experiments, we have tried to limit this overevaluation by concentrating on the main characters. With this adaptation, idioms constitute a satisfying basic measure for the type of measurements of scalability we are interested in (it would be insufficient to measure the intrinsic quality of the story, for instance while including dynamic elements such as suspense).

To represent story progression on a graph, we have chosen to trace the cumulated number of idioms as a function of time. This kind of representation displays the overall plot progression as determined by the baseline plot included in the individual actors' roles but, most importantly, also relates the emergence of situations to the plot progression.

4. Results and discussion

The results of these scalability experiments have to be interpreted from both a quantitative and qualitative perspective: the various scalability curves essentially provide quantitative information, on the number of situations generated or the overall story duration. Possible hypotheses have to be confirmed by observing the corresponding animations and other systems logs, such as HTN planning logs indicating action failure and re-planning (Figure 4). We have explored three main determinants of scalability: i) the number of feature characters ii) the complexity of their roles and iii) the number of secondary characters.

4.1. Scalability from feature characters

Throughout the experiments we have retained a classical approach in which the narrative features one prominent or central character. However, the addition of more feature characters improves the overall richness of the story, and contributes not only additional elements to the main plot but also elements of secondary plots that are common in sitcoms. What remains to be explored is the form and extent of the narrative development brought by the addition of feature characters (these behaving as supporting roles for the main characters).

The experiment we have designed consists in measuring the variation in narrative complexity (again measured through the idioms produced across the story duration) when more feature characters are added to the system (Figure 5). The baseline version consists, very much like in our former "Friends" system [2], in the main character and in four secondary characters

AAMAS'04, July 19-23, 2004, New York, New York, USA.

populating the environment. The main character (named Julie) has a complex role (108 possible actions) which consists in organising a party, while secondary characters have simpler HTN but can still interact with feature characters. The addition of a second feature character (named Vincent whose role is to find a date) results in an increased pace of the story and in a 33% increase in story duration. Because the story ending is determined by the main character's role, this increase in story duration can be attributed to additional interaction between the feature characters' roles. When the cast is further enriched by a third actor (a barmaid named Charlotte), one observes a substantial increase in the situations generated, while the overall story duration remains unchanged. This can be interpreted as the constitution of parallel plots, which is confirmed by an observation of some of the generated stories (the extent of action failures and re-planning remaining stable). The overall action is still largely dictated by the role of the main character, which also has the most sophisticated role description.

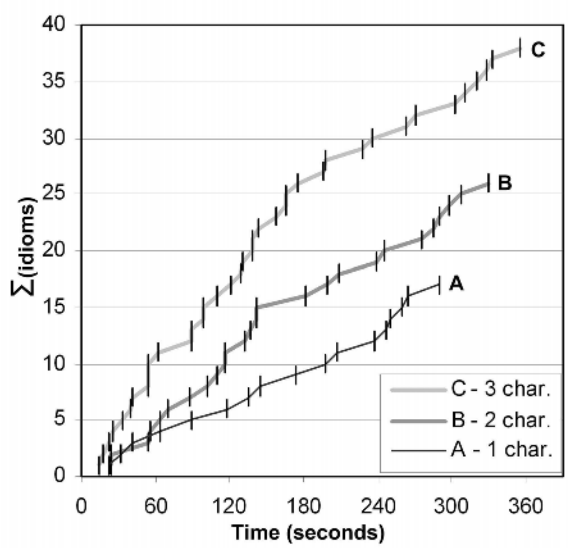

Figure 5. Scalability experiment: influence of the number of *feature* characters on story complexity.

4.2. Scalability from secondary characters

Secondary characters essentially interfere with the feature characters' actions, thereby also contributing to the emergence of humorous situations whenever feature characters plans are contrasted. A traditional example is when secondary characters conflict with the main character for the use of a story resource (for instance, the secondary character Paul may take the drinks left unattended by the main character Julie to fulfill his own task) or for interacting with another feature character.

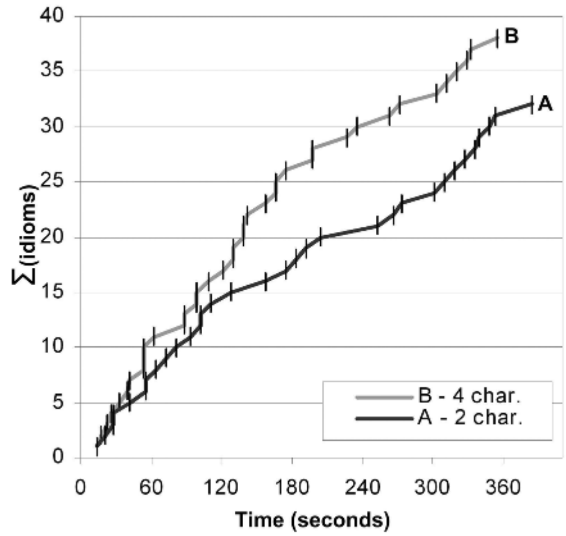

Figure 6. Scalability experiment: influence of the number of *secondary* characters on story complexity.

The complexity of the generated narrative in terms of cumulated film idioms was measured in two different configurations (Figure 6). In the first one, three feature characters constituted the story cast, together with two secondary characters. In the second configuration, this number was extended to four secondary characters. The effect is an increase in the number of situations generated, without any significant impact on the overall story duration. When observing the generated narrative, one notices the increased interactions between feature characters and secondary characters through random encounters, or deliberate interaction, for instance when the secondary character Jane's role is to socialise with the main character Julie.

There are also intrinsic limits to the number of secondary characters that can be included in a story. On one hand, adding too many secondary characters will soon impair the story's "readability"; on the other hand, this will also lead secondary characters to compete with each other for action resources, thus reverting their natural function in the narrative, which is to interact with the feature characters.

Finally, in contrast with additional feature characters, additional secondary characters do not contribute any significant new plot elements of their own, as their role-plans are restricted to less complex activities (e.g. occupying their free time, socialising, etc.) with no proper narrative goals.

4.3. Scalability from roles' complexity

It is possible to attribute some form of narrative interpretation to the various levels of an HTN describing a character's role: for instance, a top level *AND* node will subsume sub-tasks which can be seen as various scenes of the narrative (this is especially true for the principal character). Deeper *OR* nodes often correspond to options in the pursuit of a narrative goal and support the interactive components of the story.

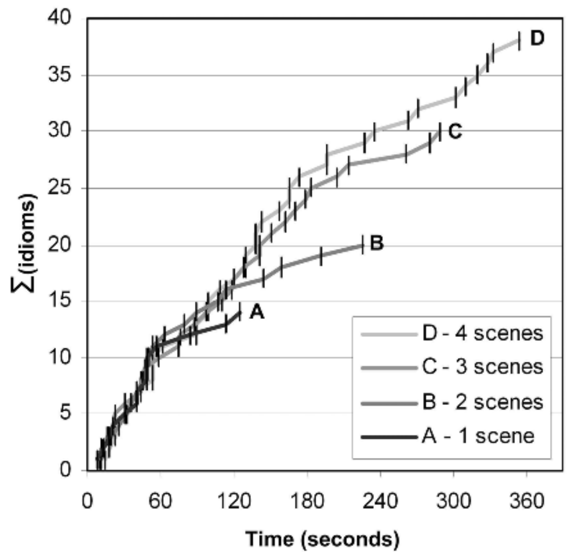

Figure 7. Modification of the number of scenes defined for the story instantiations illustrating one aspect of HTN scalability.

The complexity of a character's role in the context of IS, which can be approximated by the total number of possible terminal actions can thus be increased in several fashions: i) by adding extra scenes to the role (top-level *AND* node) ii) by adding more options to the tasks representing the various stages of the story (internal OR nodes) and iii) by making the description of sub-tasks (internal *AND* nodes) more complex, which will expose that sub-task to a greater risk of failure through dynamic interaction with other characters.

Because of the task decomposability assumption used in our approach, adding more scenes is a straightforward way of increasing the story complexity, as reflected on Figure 7. This shows a steady increase in both number of idioms and story duration when additional scenes are defined for the baseline plot, with the number of idioms produced increasing faster than the story duration itself, as additional scenes bring more potential for interaction between characters.

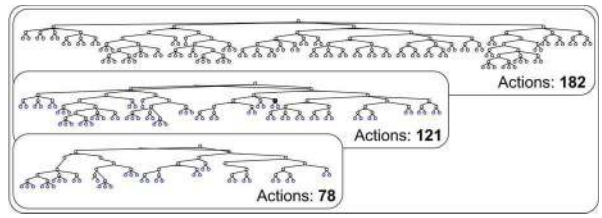

Figure 8. Increasing the complexity of a character's role by growing their HTN.

However, a better measure of scalability is obtained by growing the HTN in terms of the various strategies an actor can use to achieve its narrative goals. In these experiments, the HTN for all feature characters are progressively extended, essentially by adding more options at the level of *OR* nodes and in some cases by providing more complex descriptions for certain sub-tasks (internal *AND* nodes). Figure 8 shows different HTN of increasing complexity for a given character.

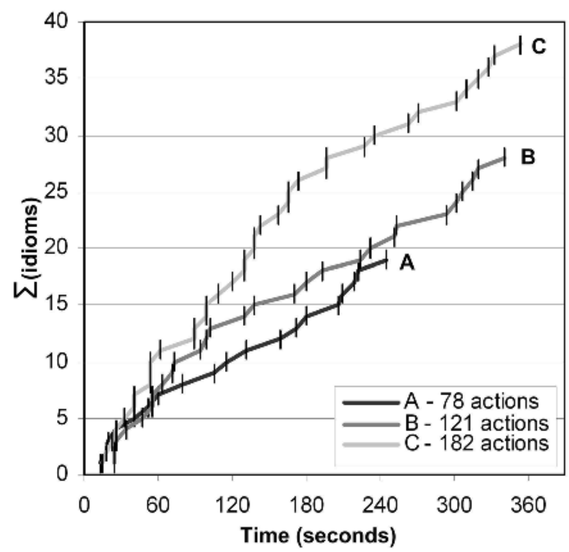

Figure 9. Variation on the number of actions as part of the main characters' roles illustrating one aspect of HTN scalability.

Figure 9 displays the evolution of the narrative in term of film idioms generated in three different configurations, in which the HTN describing the feature characters' roles are progressively made more complex. To obtain a global measure of that complexity, we have counted the total number of terminal actions available to the cast (feature characters only). There is a clear correlation between the complexity of the HTN and the number of situations generated, with no singularity or saturation being observed. The modification of story duration observed

AAMAS'04, July 19-23, 2004, New York, New York, USA.

between curve A (78 actions) and B (121 actions) is due to the fact that some task descriptions have been made richer (internal *AND* nodes), hence increasing the duration of some sub-tasks. On the other hand the difference between curve B (121 actions) and C (182 actions) can be interpreted mostly in terms of additional options being provided as well as additional interactions between characters. Throughout these experiments, care has been taken to always provide the characters with one "safe" option for those tasks that are compulsory, to avoid premature story termination.

Role complexity can thus be a significant determinant of story scalability. Most importantly, it also supports the ability to generate a larger set of different stories, something not measured by our current experiments.

Figure 10 illustrates a sequence of situations taken from an actual story instantiation. The baseline plot for this experiment is based on one feature character, Julie, and four secondary characters, which are the barmaid and three other characters with roles of similar complexity. The role of the main character Julie consists in organising a party, for which she has to find a suitable venue, decide whom to invite and whom to prevent from coming, get food and drinks, etc. The roles of the secondary characters include activities, such as buying items from the vending machine or playing arcade games, as well as talking to other characters. The first task for the main character is to make a booking for a venue for the party. Julie will ask the barmaid regarding the availability of the café (Figure 10a). After having successfully organised a venue, she will then invite people for her party. As seen in Figure 2, she may encounter a situation whereby she goes to talk to a friend, but discovers that s/he is already in conversation with a person she dislikes (her level of affinity with the other characters is dynamically modified throughout the story). In this instance, she goes to talk to Jane who is on her own, and as Julie's affinity with this character is quantified as positive (Figure 10b). After having invited friends, Julie will order food and drinks for the party by ordering them from the counter (Figure 10c). She then goes to talk to another friend, passing by in the café, leaving the goods unattended on the counter. A secondary character, Paul, is about to buy a drink, but notices the drinks left at the bar by Julie and take them (Figure 10d).

5. Conclusions

While more IS prototypes are being developed, there is a need to develop systems reaching a critical mass, beyond proof-of-concept demonstrators. Understanding the determinants of scalability is an essential aspect for the development of IS systems, in particular from the perspective of authoring interactive narratives, which has been recently recognised as an essential problem.

Our experiments have explored different aspects of scalability, suggesting that increasing the number of feature characters and the complexity of their role descriptions were equally promising in generating richer narratives. The former approach seems to take advantage of the "modular" nature of character-based storytelling, especially for its ability to include elements of parallel plots.

On the other hand, if we consider story duration, our results also suggest that the effort in system development grows significantly with the target duration. Over our experiments, doubling the average duration of the story has required a 4-fold increase in HTN complexity; this was however assisted by the

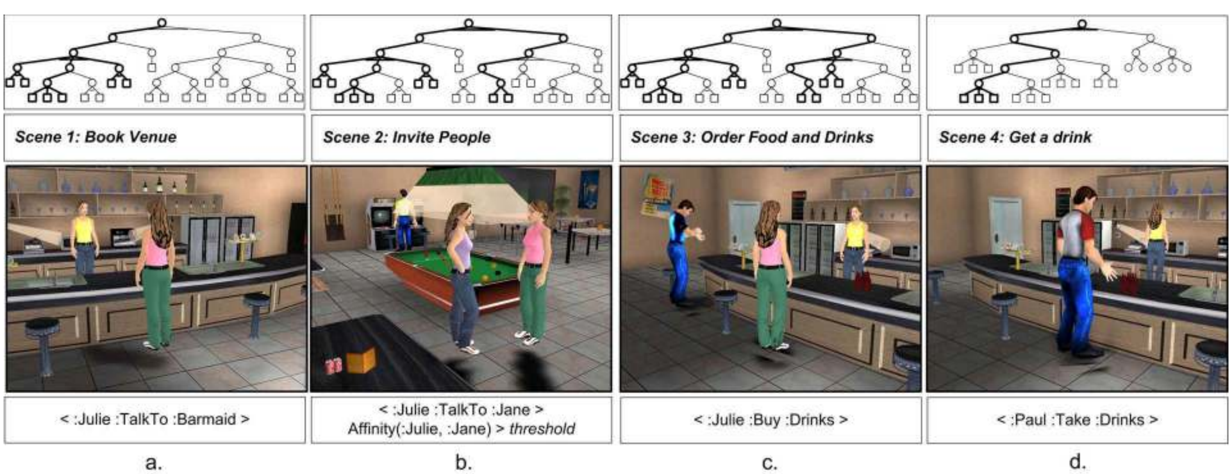

Figure 10. A sequence of consecutive story events taken from our experiments.

development of authoring tools for HTN description.

An additional hypothesis we are making is that the experimental approach described, based on film idioms, applies independently of the specific techniques used for narrative generation. As a consequence, it would not be restricted to character-based storytelling, provided specific determinants of scalability can be identified for other approaches.

Acknowledgements

We would like to thank Guillaume Révy, Steven J. Mead and Andrei Shires for their contributions to this project. We also would like to give credits to the designers of the character models, Demiurge Studios (available from the Unreal Developer Network website).

References

[1] Arijon, D., *Grammar of the film language*, Communication Arts Books, Hastings House: New York, 1976

[2] M. Cavazza, F. Charles and S.J. Mead, "Interacting with Virtual Characters in Interactive Storytelling", *First ACM Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, pp. 318-325, 2002.

[3] M. Cavazza, F. Charles and S.J. Mead, "Character-based Interactive Storytelling", *IEEE Intelligent Systems*, 17(4), 2002, pp.17-24.

[4] N. Courty, F. Lamarche, S. Donikian, E. Marchand, "A Cinematography System for Virtual Storytelling", *in Int. Conf. on Virtual Storytelling, ICVS'03*,, O. Balet, G. Subsol, P. Torguet (eds.), Toulouse, France, Novembre 2003.

[5] J. Gratch and S. Marsella, "Tears and Fears: Modeling emotions and emotional behaviors in synthetic agents", *in Proceedings of the 5th International Conference on Autonomous Agents*, Montreal, Canada, June 2001.

[6] B. Magerko, "A Proposal for an Interactive Drama Architecture", *AAAI 2002 Spring Symposium Series: Artificial Intelligence and Interactive Entertainment*, March 2002.

[7] M. Mateas, "A Neo-Aristotelian Theory of Interactive Drama", *Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, AAAI Press 2000.

[8] M. Mateas. and A. Stern., "Towards Integrating Plot and Character for Interactive Drama", In Socially Intelligent Agents: The Human in the Loop. *Papers from the 2000 AAAI Fall Symposium*, Technical Report FS-00-04, AAAI Press, 113-118.

[9] M. Mateas and A. Stern, "A Behavior Language for Story-based Believable Agents", IEEE Intelligent Systems, 17(4), 2002.

[10] D.S. Nau, S.J. Smith, and K. Erol, "Control Strategies in HTN Planning: Theory versus Practice," *Proc. AAAI/IAAI-98*, AAAI Press, Menlo Park, Calif., 1998, pp. 1127–1133.

[11] M. O. Riedl, C. J. Saretto, R. M. Young, "Managing interaction between users and agents in a multi-agent storytelling environment", *Second ACM Joint Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS 2003: pp. 741-748

[12] W. Swartout, R. Hill, J. Gratch, W.L. Johnson, C. Kyriakakis, C. LaBore,, R. Lindheim, S. Marsella., D. Miraglia, B. Moore, J. Morie, J. Rickel, M. Thiebaux, L. Tuch., R. Whitney and J. Douglas, "Toward the Holodeck: Integrating Graphics, Sound, Character and Story", *in Proceedings of the Autonomous Agents 2001 Conference*, 2001.

[13] R.M. Young, "Cognitive and computational models of suspense: Towards the automatic creation of suspense in interactive narratives", talk presented at *Interactive Frictions: The Conference Produced at the Pressure Point Between Theory and Practice: An International Conference on Interactive Narrative*, Friday, June 4, 1999, University of Southern California, Los Angeles, CA, 1999.

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

[15] R.M. Young, "An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment", *Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, AAAI Press 2001.

[16] D. Christian and R. M. Young , "Comparing Cognitive and Computational Models of Narrative Structure", *Liquid Narrative Technical Report TR03-001*, North Carolina State University, 2003.