Dynamic Plot Generation by Continual Multiagent Planning

(Extended Abstract)

Michael Brenner
Institute for Computer Science
Albert-Ludwigs-University
Freiburg, Germany
brenner@informatik.uni-freiburg.de

ABSTRACT

We describe how, by modelling plot generation as a Continual Multiagent Planning process, dynamic stories can be generated in which characters not only inteleave perception, action and interaction, but in which also beliefs and motivations may change repeatedly, thus driving the plot forward.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**] Artificial Intelligence—*Misc.*

General Terms Algorithms, Design, Languages

Keywords Narrative intelligence; Multiagent planning; Reactive vs deliberative behaviour

1. INTRODUCTION

Stories can be described as sequences of temporally and causally related events. Based on this view, many AI approaches to story generation have modelled stories as plans and story generation as planning. However, this view focusses only on the *plot* of a story. Looking at stories from a *character*-centric perspective, we can also say that stories are about autonomous agents acting in a multiagent system (MAS). Combining both perspectives, in our work we model story generation as a *multiagent planning* (MAP) task. Planning in MAS is difficult, for artificial agents as well as for humans: when several agents act concurrently in the same partially observable environment the individual agent's knowledge quickly becomes incomplete and uncertain. Thus, perfect deliberative planning is virtually impossible in MAS, and agents have to revise their plans continually. While this is unfortunate from a computational perspective, it highlights an important similarity to narrative: in a stories (as in real life), beliefs, plans, and motivations change continually, too. Indeed, it may be argued that an important aspect of what makes a plot "interesting" is such character development, i.e. a character trying to realise her plans, experiencing the actual results, and altering her beliefs, plans, and even motivations in consequence. Therefore, in this work we treat storytelling as *continual multiagent planning*, i.e. we describe a representational and algorithmic framework that models how characters continually switch between planning, acting, sensing, and interacting in a story.

Cite as: Dynamic Plot Generation by Continual Multiagent Planning (Extended Abstract), Michael Brenner, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1529-1530

2. REPRESENTATIONAL AND ALGORITHMIC FRAMEWORK

A story world is complex multiagent environment. To enable a planner to reason about it (and the beliefs and motivations of the characters acting in it), we need to model it formally. To this end we have integrated representations from several AI subfields into a formal representation language, MAPL (Multiagent Planning Language). Our approach to planning for this representation is built on our previous work on continual collaborative planning among distributed agents, which we have successfully applied in other areas, e.g., Human-Robot Interaction [2, 1]. However, characters in a story are not necessarily collaborative. (This would often lead to rather dull stories!) Therefore we have developed the Continual Multiagent Planning algorithm (CMP) which permits non-collaborative relationships between agents, too. Most importantly, the CMP process may not only lead to a revision of the beliefs of an agent, but also of its *goals*. Again, there is an obvious analogy to narrative: a character usually is better characterised by its *motivations* than its current knowledge; a change in its goals might be considered a substantial development in its personality that will drive a story into new directions. During CMP, the history of planned, executed or failed actions of the CMP agents, their change in desires and intentions is recorded in a data structure, the *plot graph*.

3. WORKED EXAMPLE

CMP and the underlying representation language MAPL are designed to be usable as the planning component in larger storytelling systems, e.g., for interactive drama. For our own evaluation of the approach across as many story domains as possible, we have developed a simulation tool, MAPSIM, in which CMP agents (and human users) can plan, act, and interact. MAPSIM automatically generates a simulation using the same formal MAPL description of the story world that the characters employ to determine their plans and actions in this world. In that manner, plots can be generated for different story domains just by switching to another MAPL description. An example story autonomously created by the system is shown in figure 1. It highlights several features of our approach.

Multimodal interaction Note first that characters' behaviour, generated by CMP, seamlessly interleaves physical action, sensing, and communication, e. g. in lines 6–8. Due to the explicit representation of epistemic states and information-gathering actions, the characters will plan which gaps in their knowledge they need to fill in order to further detail their plans. This may result in active observation (as in line 5, where Arthur checks whether the cave is empty) or in information-seeking *subdialogues* (as in lines 10–12).

Plan dynamics When Arthur arrives at the cave, he observes that the dragon, Smaug, is there. Arthur knows that he cannot take the treasure while the dragon is present. Thus, CMP detects, in its

Figure 1: A sample story involving three CMP agents, created non-interactively in MAPSIM.

1 This is a story about Smaug, King Arthur and Prince Valiant.

2 King Arthur was in the castle. 3 The treasure was in the cave. 4 King Arthur rode to the cave. 5 King Arthur saw that Smaug was in the cave.

6 King Arthur rode to the castle. 7 King Arthur saw that Prince Valiant was in the castle. 8 'Please bring me the treasure, Prince Valiant,' King Arthur said. 9 'As you wish, King Arthur,' Prince Valiant replied. 10 'Where is the treasure, King Arthur?' Prince Valiant asked. 11 'The treasure is in the cave, Prince Valiant,' King Arthur said. 12 'Thank you,' Prince Valiant said.

13 Prince Valiant rode to the cave. 14 Prince Valiant saw that Smaug was in the cave. 15 Smaug tried to kill Prince Valiant - but failed! 16 Prince Valiant saw that Smaug was not dead. 17 Prince Valiant killed Smaug.

18 Prince Valiant took the treasure. 19 Prince Valiant rode to the castle. 20 Prince Valiant gave King Arthur the treasure. 21 'Thank you for bringing me the treasure, Prince Valiant,' said King Arthur.

22 King Arthur and Prince Valiant lived happily ever after. Smaug did not.

monitoring phase, that Arthur's plan has become invalid. Arthur generates a new plan, this time a multiagent plan in which Valiant is supposed to help him get the treasure. Switching to the new plan, Arthur leaves the cave and returns to the castle. We claim that it would be quite difficult to describe the plot so far with a single plan, let alone generate it with a single planner run. Continual planning, on the other hand, seems like the natural way to model how a character reacts to the obstactles she encounters.

A form of *proactive* continual planning is exemplified in lines 8-13. Prince Valiant initially does not know the location of the treasure. Thus he could normally not find a plan to get it and therefore would have to decline Arthur's request. However, using the CMP concept of *assertions*, Valiant deliberately postpones part of the planning process and first engages in the short subdialogue of lines 10–12 in order to gather the missing information [2].

Goal dynamics The standard use of continual planning is to adapt plans to changing conditions in the outside world or an agent's belief about it. However, in real life (and thus in stories) *motivations* change, too. CMP agents can adopt temporary subgoals, e. g., when accepting a request by another agent, as in lines 9 and 12 of figure 1. Such changing goals usually lead to more substantial changes in the plot than mere plan adaptation for the same goal. Only after Arthur's request (line 8), Valiant gets involved in the story at all. In particular, this prompts a nice example of *mixed-initiative* behaviour (line 10), where Valiant immediately asks back to get more information necessary to achieve his new goal.

Characterisation by affective goal activation As noted above, changes in an agent's goals may lead to substantial changes in her behaviour. Indeed, it can be argued that a character is better characterised by her motivations than her (fairly volatile) current state of knowledge. However, a complex character has many motivations, which depending on its internal (or some external) context may be active, i. e. they drive her current behaviour, or inactive. Such context-dependent *goal activation* allows for a more fine-grained characterisation, e. g., in our story, the dragon will only want to kill humans when they have entered its lair. MAPL permits modelling the conditions for goal activation and deactivation for a character, thereby describing a multi-faceted personality whose concrete intentions may change, but who will show consistent, believable behaviour. It is important for storytelling that the conditional goals characterising an agent can refer to *emotions* or mental attitudes directed towards other agents and objects, e. g., `angry(a)`, `loves(a,b)`, etc. For example, if the dragon only attacked intruders when angry, but was willing to share the treasure when happy, another story might tell how Prince Valiant charmed the dragon and thus could acquire the treasure without violence. This also opens CMP for use in affective storytelling.

Beliefs, desires, intentions CMP forces characters to commit to a goal/desire before they can actively pursue it, i.e. make it an intention first. In the multiagent case this means that if character A is not directly controllable by another agent B, B must first somehow persuade A to commit to a new goal before B can assume A's working towards it. In our story, Arthur knows that he cannot "use" Valiant in his plan to get the treasure unless Valiant commits to that goal himself, i. e. makes it an intention of his own. Here, CMP finds a plan for Arthur to achieve this using a simple request (lines 8–9), since in the MAPL description Valiant has been modelled as being cooperative towards Arthur. On the other hand, before CMP could include actions of the dragon into Arthur's plans, it would first have to indirectly activate one of the dragon's own desires.

False beliefs are important for plots, as they result in misunderstandings, in misguided or unsuccessful behaviour. Again, continual planning can be used to reason about the consequences of believing something wrongly. To show this, the example domain is set up such that the "stronger" character always wins a fight. Here, Smaug falsely believes to be stronger than Prince Valiant and attacks him (line 15), which eventually leads to his own death.

Chekhov's gun The plot graph describes which facts and objects are used in the plot. Those facts should be mentioned so that the reader/player can follow the reasoning of the characters. Crucially, the plot graph does not only point to preconditions of actions that characters actually perform, but also to those beliefs never used in an executed plan (because of the plan or the belief becoming obsolete first), but that are necessary to *explain* the changing motivations and plans of the characters.

4. SUMMARY

We have developed a representational and algorithmic framework for modelling and reasoning about complex story worlds in a continual multiagent planner. It emphasises the dynamic aspects of plots, the changing beliefs, motivations and plans of characters, which are hard to reason about in classical "one-shot" planning settings. Our prototype implementation shows how the representation and algorithm can be used in storytelling applications to generate complex plot structures, autonomously or interactively.

5. ACKNOWLEDGMENTS

Funded under grant agreement No. 215181, CogX, in the European Community's 7th Framework Programme [FP7/2007-2013].

6. REFERENCES

[1] M. Brenner, N. Hawes, J. Kelleher, and J. Wyatt. Mediating between qualitative and quantitative representations for task-orientated human-robot interaction. In *Proc. IJCAI*, 2007.

[2] M. Brenner and B. Nebel. Continual planning and acting in dynamic multiagent environments. *Journal of Autonomous Agents and Multiagent Systems*, 19(3), 2009.