Efficient Intent-Based Narrative Generation Using Multiple Planning Agents

Jonathan Teutenberg
Independent Researcher
jono.teutenberg@gmail.com

Julie Porteous
Teesside University
Middlesbrough, UK
j.porteous@tees.ac.uk

ABSTRACT

In Interactive Storytelling (IS) the prevailing approach for the automatic generation of plausible narratives that meet global author goals is intentional planning. However, existing approaches suffer from limited expressiveness and poor scalability. We address this by replacing single intentional planners with multiple agents representing the characters of a narrative, which can reason about the relevance of narrative actions given their individual intents. These are then combined using a state-based forward search procedure that results in a significantly smaller search space.

Unlike other multiagent approaches, these agents calculate all reasonable plans in a state.This allows a search of a wide range of narrative possibilities prior to execution as in planner-based approaches, rather than agents making early plan commitments in a simulation.

We demonstrate that this not only produces the same forms of narrative as single intentional planners but can be extended to generate narratives that are beyond their scope. We also present a search heuristic that exploits the agents' relevant actions to further reduce the size of the explored search space. Experimental results demonstrate system performance that makes it suitable for use in real-time applications such as IS.

Categories and Subject Descriptors

H5.1 [**Multimedia Information Systems**]: Artificial, augmented and virtual realities

General Terms

Algorithms

Keywords

Interactive Storytelling, Agents in games and virtual environments, Narrative Modelling, Planning

1. INTRODUCTION

In Interactive Storytelling (IS) systems an important factor in audience narrative understanding is that the virtual characters must be *believable* – the audience must suspend belief and the actions of virtual characters must not threaten this [2]. In other words the audience must perceive them as intentional agents [7]. As well as maintaining a believable cast, narratives generated for IS systems must strive for quality in terms of such things as author goals or preferences [15].

Narrative generation that features multiagent simulation neatly solves the problem of endowing agents with intentionality. By treating each character as an autonomous agent with its own beliefs and intents, believable character interactions can emerge from a simulation [1]. The drawback for an IS system based on such an approach is in the fulfilling of global narrative goals and author preferences. A character-only simulation relies on emergence and there is no guarantee of even an approximation of criteria other than believability. Systems have been proposed that include director agents [11] that can interact with the simulation to try and satisfy author goals. In many ways this turns narrative generation into a hill-climb through the space of possible narratives with the director guiding its trajectory as best as possible at each point in time.

In contrast, the planning approach to narrative generation searches large expanses of the narrative space but struggles with producing an appearance of intentionality. Narrative generation by a single classical planner that has available to it all character actions and plot events is able to compose sequences of character actions to fulfil the global goals that make no sense from any individual character's perspective. To make character-based planning suitable for IS, *intent-based planning* has been proposed with the IPOCL planner[17]. This generates narratives that ensure that characters' actions have intent, as though planned and performed by autonomous agents, by explicitly representing intent in the narrative and requiring that each action is assigned to some intent. Effectively this explores the space of narratives by considering not only every action available, but every action-intent pairing available. Unfortunately the applicability of an IPOCL style planner to narrative generation for IS is limited due to its poor performance. For example, a figure of 12.3 hours is given for the generation of a sample narrative in [17], making it unsuitable for real-time IS use.

This is compounded by the fact that IPOCL has a non-standard representation language which prevents experimentation with more state-of-the art planning approaches whose performance might be faster. Interestingly, Haslum has recently introduced a compilation for IPOCL narrative domains to classical planning [8]. This meant that the intent-based planning problems could now be tackled by a wide range of classical planners, and for IS it was hoped that this

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

would enable real-time narrative generation. However, we conducted a series of experiments with versions of IPOCL domains created using Haslum's compilation and the results, which we report in this paper, suggest that narrative generation was still too slow to realistically be used in an IS system. Further, the coupling of planning for characters' goals with planning for narrative goals prohibits the generation of certain narratives, most notably those narratives in which one or more characters become unable to complete their intent. Thus the motivation for our work was to develop an approach that would be capable of performing an extensive search of the space of narratives, in order to be able to generate these types of interesting narratives but which would also be capable of performing within the time constraints of a real-time IS system.

The solution, which we present in this paper, is one where the task of reasoning about intentions is delegated to narrative agents who are responsible for the checking of "narrative action relevance" with regard to their individual intentions. Then a single narrative planner is used to generate narratives which satisfy the constraints obtained from the relevant actions. These ideas are fully implemented in a system, called IMPRACTical (**I**tentional **M**ulti-agent **P**lanning with **R**elevant **ACT**ions). The key contributions of this work are the design, implementation and evaluation of this approach.

The paper is organised as follows. We start in the next section with a detailed consideration of related work. Then in Section 3 we discuss a number of narrative examples which illustrate the scope of intent-based narrative generation. In Section 4 we overview our novel approach for the integration of multiple agent reasoning with state-based planning. This is followed in Section 5 with discussion of the implementation of these ideas. We present the results of an experimental evaluation in section 6 and finally present our conclusions.

2. RELATED WORK

The question as to what constitutes a criteria for 'good' narratives that are to be used in IS is still open. Current state-of-the-art requires input from human authors such as assigning interesting authorial goals [15] the design of appropriate initial and goal states [16], or describing desired tension arcs [13]. Minimum length criteria also plays a part, as this implies a tight narrative without superfluous actions. In addition to a quality measure, narrative plans must be produced with hard constraints on their plausibility: every action seen to be performed by a character must appear to have some motivation or intent [17].

Simulation as narrative generation, also referred to as "character-based" narrative generation relies on a system of independent planning agents (characters) with a view to generating interesting emergent interactions [5]. This has evolved to model characters as agents with ever-increasing sophistication, such as social influence [6], recursive beliefs [19], and social planning [4]. These enable a greater scope of possible interactions amongst agents and therefore a greater space in which to find interesting narratives. However the existence of a vast space alone is insufficient – if solutions of high quality are to be guaranteed this space must be searched intelligently to find a suitable narrative. Pure simulations such as the Continual Multiagent Planning of [4] effectively produce a single walk through the space of world states, indeed it was noted that a total of only 20 plans were evaluated by characters and the simulation to create their example narrative. It should be noted however that multiagent narrative generation need not extract uninformed paths from the narrative space. Adding external guidance in the form of modifications to agent behaviour [18] or director agents that can affect the world [12, 11] can hill-climb by estimating future narrative quality at each point in the simulation's execution.

Intentional planning with the IPOCL algorithm [17] was originally devised as an extension to a causal link planner that searches the space of partial plans, working backward from the narrative goals. The appearance of intent was ensured by including, in parallel, the planning of *frames of commitment* that span contiguous sub-sequences of actions within a plan. The intent for a frame is defined by a character and a target fact that they intend to achieve. The backward chaining nature of the search ensures that characters' actions are causally linked to, and therefore appear relevant to, the final action that closes their frame of commitment. While the extensive search of the narrative space can find solutions that closely match global goals, a notable drawback of the IPOCL planner is the high complexity of the search and long run-time. A method of compiling intent-based problems so that the frames of commitment are represented as a classical planning problem was proposed in [8]. By transforming the problem to a standard format, advances in the broader area of generic planning can be applied to intent-based narrative generation. The concept behind the compilations is that the possible choices of intent for each action can be determined prior to planning, and a new action made for each action-intent pair.

For the remainder of this paper we shall refer to approaches as being *intentional* when both intent and action causality are modelled in the narrative; as *monolithic* when this is performed by a single planner (thus [17] is both intentional and monolithic); and as *multi-agent* when planning is distributed amongst characters that manage their own intent. Our IMPRACTical approach, described in the following sections, lies somewhere in the middle: it is intentional in that it performs global heuristic search, but is not monolithic as it delegates intent to multiple planning agents.

3. MOTIVATION

While compiling intentional plans (as in [8]) dramatically improves the efficiency of their narrative generation, this still cannot be performed sufficiently quickly for online replanning, and it is not clear that the increase in the number of actions at each state will scale to more expressive domains.

Our goal is an alternative approach which preserves the satisficing of author and narrative goals of intentional approaches, but in which action intent is resolved for by multiple agents. Improvements in efficiency are realised by making these agents autonomous from the global search procedure that determines the final sequence of actions in the narrative. This decomposition is not only significantly more efficient at producing the same narratives as the monolithic planners, but also can be more expressive, generating alternative plausible narratives that may contain higher quality solutions.

Underpinning the intentional approach to narrative generation is an assumption that in order to generate believable narratives we require that every action in a narrative plan that is performed by a character has intent (in the sense used in [17]).

A simple illustration of this is given in [17] in which a princess, knight and king are present and the authored goals for the narrative are that by the end of the narrative the princess is jailed and the king dead. An example solution without perceptible intent is:

1) The princess kills the king;
2) The princess locks herself in the tower.

This is in contrast to a narrative with the appearance of intent, where:

1) The king locks the princess in the tower;
2) The knight kills the king.

This is a somewhat subjective constraint, but for the purposes of this paper we can use the same criteria as intentional planners: that we require that generated narratives feature some form of *causal link* from every narrative action to an intent of a relevant character and that characters have achieved their intent at some time prior to the point in time in which actions occur in a narrative. This is precisely defined in [17] but informally it refers to whether actions make some kind of sense to the user based on what they know about the story world – in the above examples, it makes no sense for the princess to lock herself in the tower, whereas its perfectly plausible for the king to lock up the princess.

3.1 Characteristics of Intentional Plans

We have identified three types of reasoning that contribute to character intent as produced by monolithic approaches: (1) cooperation of two or more characters; (2) characters predicting other characters' actions; and (3) the occurrence of chains of commands which propagate intent. Below we give an example to illustrate each of these based on the Aladdin fairy tale [17] and its' characters.

3.1.1 Example 1: Cooperation

Aladdin and the Genie both intend for Jasmine to love Jafar. Aladdin is incapable of love-spells, and the Genie is trapped in the lamp.

Neither character can fulfil their intent by acting alone, so if their reasoning is limited to considering only their own actions' effects on the world in isolation then no actions can be motivated. However, by co-operating, Aladdin can first free the Genie who can then cast the spell, fulfilling their shared intent. Thus both the freeing of the Genie and the casting of the spell can be assigned an intent via cooperation.

3.1.2 Example 2: Predicting Other Characters

The dragon has the lamp at the mountain. Aladdin intends to slay the dragon and Jafar intends to have the lamp.

From this state, Jafar reasoning in isolation cannot justify performing any action. However, by predicting the future actions of Aladdin, the narrative can have Jafar precede Aladdin to the mountain. His intent for this action is to await the anticipated arrival of Aladdin before pillaging the lamp from the slain dragon.

3.1.3 Example 3: Chains of Command

Jafar intends to marry Jasmine. Only the Genie can make Jasmine love Jafar. Only Aladdin can slay the dragon to get the lamp.

Jafar reasons that he can order Aladdin to help who will in turn command the Genie to help who will cast the love-spell, allowing Jafar to fulfil his intent. Jafar can only take the first step of ordering Aladdin if he can foresee the intent passing again from Aladdin to the Genie, thereby justifying it as working towards fulfilling his marriage. This encapsulates the version of Aladdin presented in [17], and we shall refer to it as the *canonical* Aladdin problem.

3.2 Extended Forms of Narrative

In addition there are further narratives that can be generated, where every action has plausible intent, but which cannot be generated by monolithic approaches. In particular those in which some agents have not succeeded in fulfilling their intents by the end of the story. Take the example of the previous section with a narrative goal of having Jafar succeed in marrying Jasmine. A second narrative solution is to first have Aladdin fall in love with Jasmine and intend to marry her. Aladdin can then be motivated to slay the dragon, and to fulfil the narrative goal Jafar pillages the lamp from the dragon before Aladdin. At the end of the narrative when Jafar marries, Aladdin's intent will remain unfulfilled. In the approach of [17] this narrative can not be found because in order to identify an association of intent to action, a chain of causal links is required from a state in which the intended fact is true, back to the action that is to be motivated by it. When no such state exists, as for Aladdin's intent in the example above, the causal chain back to his actions cannot exist and the intent will not be identified.

As we shall show in subsequent sections, it is possible to generate this class of narratives using our heuristic search with multiple planning agents, IMPRACTical.

4. INTEGRATING AUTONOMOUS AGENTS WITH STATE-BASED SEARCH

The fundamental concept underlying our proposed approach to generating intentional narratives is to delegate the problem of reasoning about intent to a collection of agents. Each agent represents a character in the narrative and its responsibility is, given a current world state, to provide the narrative generator with all actions that it perceives as being relevant to its intents. This differs from standard planning agents that select a single action to apply, based on a single plan they have committed to in order to fulfil their intent. By providing *all* actions that can be part of *any* reasonable plan with the same goal, the narrative search can explore many potential branches without having to make an early commitment to a specific plan for an action to belong to.

Where a monolithic planner expands the search space to consider the product of intent and action, IMPRACTical reduces the search space to only those actions deemed relevant. The trade-off is made in the additional computation for calculating action relevance, but so long as this is less expensive than the cost of evaluating the larger number of partial narratives this approach will be an improvement. In general, exchanging greater computation per search node with a reduction in branching factor is a good one to make as problem sizes increase. Many planning heuristics make this choice be creating and solving abstract versions of their problem at every node evaluation [9]. In this paper we perform similar reasoning with character agents and achieve similar gains in overall efficiency.

Mitigating the need for agents to reason at every node in the search is the fact that their reasoning need not be optimal. Agents represent imperfect characters, and an action they deem relevant can be part of a longer, more costly or even an ultimately invalid plan, so long as their mistake is not 'unbelievable'. However an agent can not permit all actions or the search will revert to that of a generic planner without intent. Conversely, even if the agents could produce optimal plans in reasonable time, this may preclude optimal narratives as many plausible yet sub-optimal actors' actions would be ignored that could have formed part of the optimal solution. The reasoning about causality should be apparent to a narrative's audience, but it is preferable for agents to be optimistic with their choice of actions.

For an agent to determine actions relevant to its intents it must be provided with sufficient information. As well as domain information such as which actions it has access to, it must also have knowledge of the current world state so that it can determine which of these actions can be applied. Plan-space planners like IPOCL can only provide a stateless partial plan at any given point in the planning process. If a complete state is to be provided whenever a decision on adding an action is to be made we are restricted to forward search through the state-space. Fortunately this form of search has proven effective for many modern planners such as HSP [3] and FF [10] and our approach is in this tradition.

In terms of the information available to an agent at any given decision point, we are consistent with existing intentional extensions to planning domains. This means that an agent has available:

The set of all possible facts F.
The world state $S \subseteq F$, containing all true facts.
All agents A, including self.
Open intents I of the form $\langle b, j\rangle$ with $b \in A$ and $j \in F$.

and a set of actions O, where each $a \in O$ is defined by

$pre(a) \subseteq F$. a is applicable when $pre(a) \subseteq S$.
$add(a) \subseteq F$. When a is applied $S = S \cup pre(a)$.
$del(a) \subseteq F$. When a is applied $S = S \setminus del(a)$.
$actor(a)$, one or zero agents from A requiring motivation.
$add_i(a)$, a set of $\langle b, j\rangle$. When a is applied $I = I \cup add_i(a)$.

Furthermore, when an action a is applied all fulfilled intents are removed,

$$I = I \setminus \{\langle b, j\rangle | b \in A, j \in add(a)\}$$

For the sake of efficiency we assume the agents consider fulfilling single intents in isolation, in a completely known world state, with actions only performed by themselves and agents explicitly cooperating with them. From this basis in the next 3 sub-sections, we will describe 3 simple procedures for combining agents' resulting relevant actions that will integrate them into the search such that the properties of cooperation, prediction and chains of command that were identified in Section 3 are enabled. We represent the basic reasoning with a function $f(B, i, S) = R$ that maps a set of cooperating agents $B \subseteq A$, a target intent fact $i \in F$, and the world state S, to a set of relevant actions R, and with a function $g(B, S)$ that returns all reachable actions from S by the agents B.

4.1 Cooperation

Agents are assumed to be cooperating if and only if they share an intent. Given a set of cooperating agents the actions relevant to their intent is defined by f, so determining the agent set for each intent and applying f, as shown below, is sufficient for finding all relevant actions with cooperation.

```
procedure RELEVANTACTIONS-BASIC(S, I)
    R ← ∅                               ▷ Relevant actions set
    H ← {i|⟨a, i⟩ ∈ I}                  ▷ Open intent facts
    for i ∈ H do
        B ← {a|⟨a, i⟩ ∈ I}              ▷ Cooperating agents
        R ← R ∪ f(B, i, S)
    end for
    return R
end procedure
```

This procedure is called in every state S, and the applicable actions in that state are all $x \in$ RELEVANTACTIONS-BASIC(S, A) such that $pre(x) \subseteq S$.

4.2 Prediction

As with other intentional narrative generators, the characters, and therefore their agents, have full world knowledge at all times. This includes the intent of other agents and the reasoning process they will use to fulfil these intents. Agents can therefore plan to make use of other actors' future actions without requiring explicit cooperation.

Marking actions as relevant based on the modelling of expected future actions of other agents can be integrated into the forward search using an iterative wrapper. In each iteration the agents' relevant actions are added to the current state. In following iterations other agents are then able to use the effects of these expected actions as preconditions to their own future actions.

```
procedure RELEVANTACTIONS-PREDICT(S, I)
    R ← ∅                               ▷ Relevant actions set
    C ← S                               ▷ Current state
    repeat
        T ←RELEVANTACTIONS-BASIC(C, I)
        T ← T \ R                       ▷ Keep the new relevant actions
        for a ∈ T do
            C ← C ∪ add(a)              ▷ Update the state
        end for
        R ← R ∪ T                       ▷ Store this iteration
    until T = ∅
    return R
end procedure
```

4.3 Chaining Commands

For narratives in the form of the example in Section 3.1.3, integration can be enabled by replacing calls to the function f with a procedure that recurses down possible chains of command, calling f at each step. The idea is to have agents use commands or orders – actions a that act on characters c which can be grounded to have any fact j in $\langle j, c\rangle \in add_i(a)$ – to pass their current goal intent to other agents. They will then meet the requirement to be considered 'cooperating' and a new call to f with the additional agent can be made. The final set of relevant actions is all those relevant to the intent fact, plus all those relevant to reaching commands that act on agents whose cooperation was required.

procedure F-COMMANDS(B, i, S)

$R \leftarrow f(B, i, S)$ ▷ Cooperating action set
if $R = \emptyset$ **then** ▷ If i is not achievable
  $U \leftarrow g(B, i, S)$ ▷ Reachable actions
  $V \leftarrow \emptyset$
  $C \leftarrow \emptyset$
  for $u \in U$ **do**
    if $\exists b$ s.t. $\langle b, i \rangle \in add_i(u) \wedge b \notin B$ **then**
      $V \leftarrow V \cup \{u\}$ ▷ Useful commands
      $C \leftarrow C \cup \{b\}$ ▷ Commandable agents
    end if
  end for
  $B \leftarrow B \cup C$ ▷ Commandable agents join in
  $R \leftarrow$ F-COMMANDS(B, i, S) ▷ Recurse
  $A \leftarrow A \setminus B$ ▷ Commandable agents are done
  for $v \in V$ **do** ▷ Actions to reach used commands
    if $\exists a \in R$ s.t. $actor(a) \in C$ **then**
      $R \leftarrow R \cup \bigcup\{f(B, j, S) | j \in pre(v)\}$
    end if
  end for
end if
return R
end procedure

5. IMPLEMENTATION

In the previous section we outlined three procedures that can be used to determine actions based on cooperation, prediction and chaining commands, but the basis for all of these is the relevant actions function f. In this section we give an example implementation of this function that can be performed efficiently and maintains the causal relationship between relevant actions and their target intent.

5.1 Relaxed Agent Reasoning

Our implementation relies on the concept of a relaxed planning domain (as described for example in [10]). To relax the domain, every actions' delete effects are removed. This means that no pair of actions can be mutually exclusive – one action can never delete required preconditions of another – making it possible to efficiently find solutions to the relaxed version of a planning problem.

The relaxed domain used here is a reasonable approximation in many cases and has proven suitable for use in heuristics for planning to optimise plan length [9]. Plans in this abstracted space are sufficiently close to plans based on the original, unrelaxed actions to produce a view of plausible, though at times imperfect, characters.

Solving problems in the relaxed domain usually involves two steps: firstly finding all reachable actions and facts by working forward from the current state and storing them in a relaxed *planning graph* (RPG); and secondly searching back through the graph from goal facts to extract the solution. It is the second step that is most time consuming in this process, as there can still be many possible combinations of relaxed actions to consider.

The construction of the RPG is shown in Figure 1 performed by applying relaxed actions to a set of facts, beginning with the current state S, until the agents' intent fact is found. The graph is represented with levels L_1 to L_n, alternating between levels with nodes representing facts and

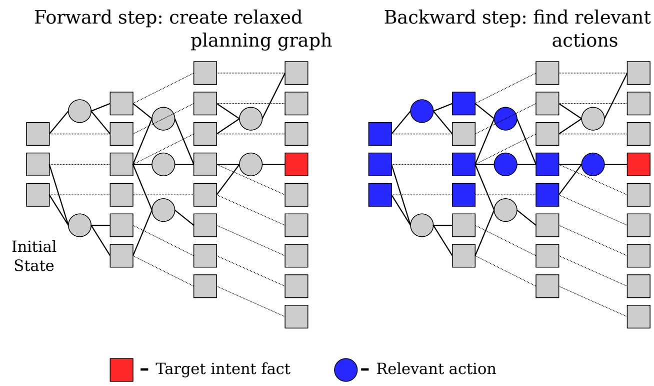

Figure 1: Determining relevance for a single agent's actions. Solid lines indicate effect to fact or fact to precondition relations, dotted lines indicate a no-op. After the backwards step, of the two actions with their preconditions met in the current state, one is found to also be relevant to the actor's intents.

levels with nodes representing actions. The set facts in a level $L_i$is a superset of the previous fact level L_{i-2}, with each new fact $s \in L_i \setminus L_{i-2}$ being linked to each action $a \in L_{i-1}$ with $s \in add(a)$. Facts that also appear in earlier levels are not only connected to any actions that add them, but also to a 'no-op' action in the previous level. Each action in the graph is associated with a single node, so an action a_1 is causally linked to an action a_2 in a later level if and only if $add(a_1) \cap pre(a_2) \neq \emptyset$ in which case they are connected through a path containing fact and no-op nodes. What this means for agents that reason in the relaxed domain is that following a chain of actions back through the graph is guaranteed to maintain causal links at every step: precisely the property required to maintain the plausibility of actions in intentional planning.

Halting the RPG at the goals means any plan extracted minimises makespan: the number of plan steps when parallel execution is permitted. The relevant actions for an agent is the union of these plans, as any one of these solutions are suitable for narrative believability. All these actions are causally linked to the goals, and this is guaranteed to be at least as efficient as constructing the graph. This is our implementation of the relevant actions function f for each agent.

The reachable actions g only require calculating in the case that relevant actions cannot be found, i.e. the intent was not reached in the RPG. In this case the actions in the RPG are the reachable actions, and these can be quickly extracted.

5.2 Heuristic

Modern heuristic search-based planning relies on solving in simplified, abstract versions of the domain [9]. However, simply solving a standard planning problem with delete effects removed will greatly underestimate the distance to the narrative goals for our approach. Not only will this abstraction ignore mutex relationships due to deleting facts, but it will also ignore any requirement of agent intent.

We instead propose to solve in a domain with relaxed delete effects and a close approximation of the selection of relevant actions that will be available at each search state. The approximation of relevant actions relies on the fact that agents are reasoning in the same relaxed domain, and so the narrative RPG can be constructed from the relevant actions and commanded actions found in the process of RELEVANTACTIONS-PREDICT and F-COMMANDS.

Two steps complete the heuristic calculation, illustrated in Figure 2. The addition of arcs in the RPG between relevant command actions adding $\langle b, j \rangle$ and the relevant actions of b, so that causality from intent is explicitly represented in the graph (see **B**). A minor modification to RELEVANTACTIONS-PREDICT that includes the addition of all applicable happenings (see **C** and **D**) in addition to each agents' relevant actions. From this point any procedure that extracts a solution from the RPG can be used. We use the well-established h_{FF} [10] as this final step in our heuristic: h_R.

6. EVALUATION

6.1 Experimental Setup

Existing intentional approaches have been evaluated based on the production of a single expected narrative using a single domain, namely, the Aladdin domain and problem presented in Section 3.1.3. The first objective of our evaluation was to show that IMPRACTical can replicate these results. In addition a further objective was to demonstrate the practical application of the approach for IS where narrative generation must be possible with multiple character actions and motivations. Hence, we evaluated our IMPRACTical approach with a new domain that is a combination of actions taken from a number of narratives that have appeared as examples in narrative generation including *The Vengeful Princess*[20], a Mexica[21] story, and *Aladdin*[14, 17].[1]

We compared IMPRACTical against problems compiled using the approach of Haslum [8]. The planner was an A-star forward search in the manner of HSP [3] using either the h_{max} admissible heuristic, or the more informative but inadmissible h_{FF}. This is compared against the same search procedure with the applicable actions restricted to happenings or those from RELEVANTACTIONS-PREDICT, with each agent using F-COMMANDS to determine relevant actions. In addition to the two heuristics used by the monolithic planner, the h_R heuristic was assessed. While we shall measure both run time and the number of search-states evaluated, we also wish to compare the overhead of the agents' relevance computation with the saving from reducing search space. As both the heuristic calculation and agents' reasoning is based on the time-consuming production of RPGs, we will use the number and size of these as a measure of the complexity. For run time calculation, a single core on a consumer desktop 3GHz AMD CPU was used to run the experiments.

6.2 Results

6.2.1 Efficiency

Table 1 shows the results of running the canonical Aladdin domain and problem using the compiled domain compared to our approach. Of particular note is the average branching factor which is approximately ten times greater when searching the space of compiled intent-action pairs. In any

[1]The domain model is available to download from http://www.scm.tees.ac.uk/j.porteous/aladdin-files.zip

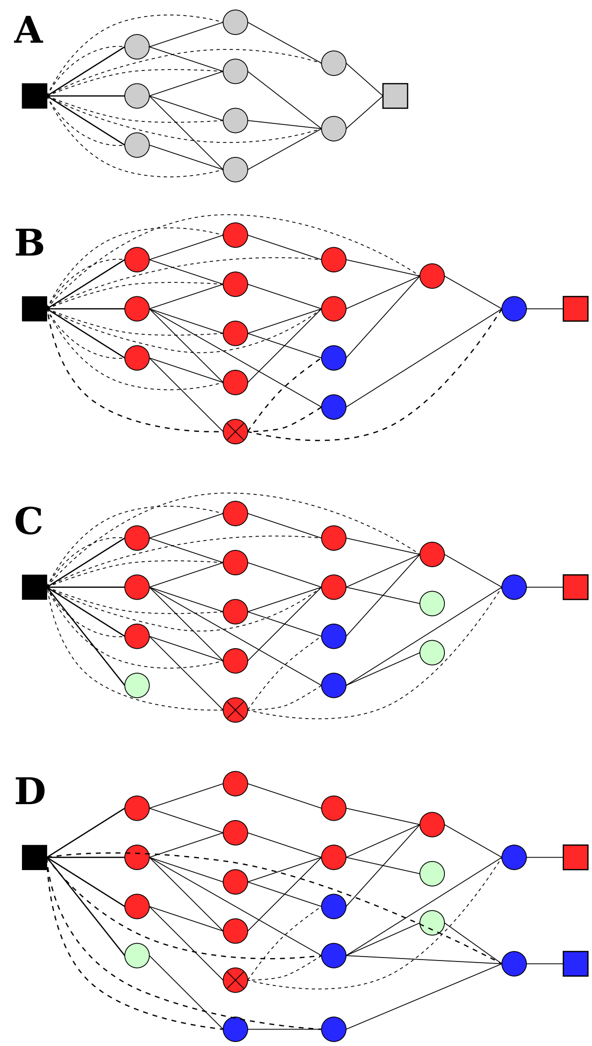

Figure 2: Construction of the RPG for heuristic evaluation. Fact nodes other than a combined initial state (left) and goal facts (right) are excluded. Solid arcs indicate causal relationships between actions. **A**) A single agent's RPG with additional dashed arcs showing that each action is self-motivated.**B**) An RPG in which a command (crossed action) was used to command a second (blue) agent to assist. The blue agent's actions are motivated by the command. **C**) Here the red agent's motivations are not shown for clarity. Another iteration begins from (B) adding actions that require no actor and that require not additional motivating arcs. Effects of these enable the blue agent to fulfil an intent. Note that one of the blue agent's actions now has two potential sources of motivation. **D**) In the next iteration after (C) one of the agent-less actions added an intent for the red agent. An arc from this action to the red actions leading to the intent is added.

Table 1: Comparison of Haslum's compiled intentional problem (Mono.) and IMPRACTical on the canonical Aladdin domain and problem.

	State count	RPG count	RPG size	Branch factor	Time (s)
Mono. (h_{max})	>1M			44.1	
Mono. (h_{FF})	2946	2447	2600	32.4	20.4
IMP. (h_{max})	1816	2723	32.0	3.8	1.0
IMP. (h_{FF})	215	357	50.3	3.3	0.5
IMP. (h_R)	46	274	29.5	2.4	0.3

Table 2: Comparison of the extended domain with Aladdin problem.

	State count	RPG count	RPG size	Branch factor	Time (s)
Mono. (h_{FF})	> 1M			41.4	
IMP. (h_{max})	291368	340321	130.3	9.9	24.3
IMP. (h_{FF})	29561	33959	156.3	9.2	3.9
IMP. (h_R)	165	12345	60.2	8.7	1.8

given search state the agents presented only three to four reasonable options on average, indicating how restricted this example domain is.

The compiled problem completed in 20.4s when using h_{FF} which is at the upper limit of acceptable running times. The lack of immediacy in response to interactions would greatly restrict IS systems on this problem. The less informed h_{max} heuristic was halted prior to completion after it had evaluated a million search states having run for several minutes. The curse of dimensionality claimed another victim due to the high branching factor of the compiled domain.

IMPRACTical finds a solution sufficiently quickly using any heuristic. Using h_{FF}, the number of states evaluated is under 10% of that evaluated by the compilation approach. The additional RPG evaluations have little impact – there are still fewer than 15% as many constructed over the course of the search.

As expected, the more informed heuristic h_R gives a significant reduction in the number of states evaluated over h_{FF}. Fewer RPGs are constructed and on average these contain fewer actions as more are agent plans rather than the larger narrative plans.

Table 2 repeats the comparison of Table 1 using the same problem description but with the extended domain with a greater number of available actions. Even using the h_{FF} heuristic that was able to direct the solution in the smaller domain, the compiled problem was unable to find a narrative solution in any reasonable time. For IMPRACTical, the branching factor was almost 3 times greater which resulted in around 50 to 100 times more RPGs and a run time of more than 6 times as long as the simpler domain.

To evaluate the scalability of IMPRACTical, another more complex problem with different initial state and narrative goals was evaluated. This problem took 24.2s to solve–longer than the desired response time for an IS system. This problem is an example near the upper limit in terms of complexity that our new approach can manage. The solution narrative contained 18 actions, had an average branching factor of 12.1, and required 200946 RPGs to be constructed.

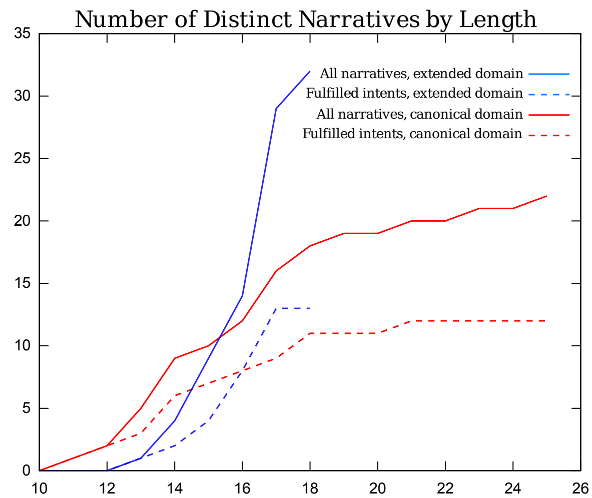

Figure 3: Number of distinct narratives for the Aladdin problem using the canonical domain (red), the extended domain (blue), permitting characters with unfulfilled intents (solid lines) and only fulfilled intents (dashed).

6.2.2 *Expressiveness*

The results presented in the previous section confirm that our IMPRACTical approach is capable of generating the same narratives as intentional planners, but much more efficiently. In this section, to quantitatively evaluate the scope of other narrative possibilities we have extended the search to find distinct sets of narratives. We define 'distinct' narratives based on their sets of actions. Two action sets O_1 and O_2 are distinct when $|O_1 \setminus O_2| \geq 2$ and $|O_2 \setminus O_1| \geq 2$. The search builds the set of narratives in increasing length for which this constraint holds between all pairs. Some of the narratives excluded by these constraints may be unique and interesting, and some that are included may have other interesting variants that can arise through action re-ordering. However, we believe these omissions are an acceptable compromise to avoid considering redundant narratives.

Using the canonical Aladdin domain and problem a total of 22 distinct solutions are found. All of these feature agent cooperation as in example 3.1.1; 9 of these feature agents acting on predictions of others future actions as in example 3.1.2; and 8 feature chains of two or more commands as in example 3.1.3. None of the narratives found by the planner using the compiled domain were distinct from these, however 10 of the 22 contained intents that could not be fulfilled and therefore were outside the scope of the compiled approach. An example of a new narrative found that demonstrates cooperation, prediction, a chaining of commands and unfulfilled intents is given in Figure 4. This clearly illustrates IMPRACTical can generate new forms of narratives that still display the properties identified in intentional plans.

Figure 3 shows the effect of moving to the extended domain. Within 5 actions of the optimal solution the extended domain's search space already contains over 30 distinct narratives, more than in the entire space of the canonical domain. The dashed lines show that only with the addition of narratives in which some agents' intents remain unfulfilled can this rapid expansion of narrative possibilities occur.

```
1.  Jafar falls in love with Jasmine
2.  Aladdin falls in love with Jasmine
3.  Aladdin travels to the mountain
4.  Jafar travels to the mountain (predicting)
5.  Aladdin slays Dragon
6.  Aladdin pillages the magic lamp from Dragon
7.  Jafar orders Aladdin to help marry (1st)
8.  Aladdin summons Genie from the magic lamp
9.  Aladdin commands Genie to help marry (2nd)
10.  Genie appears threatening to Jafar
11.  Jafar orders Aladdin to slay the Genie
12.  Jafar travels to the Castle
13.  Genie casts a spell on Jasmine (cooperating)
14.  Jafar marries Jasmine
15.  Aladdin slays Genie
```

Figure 4: Generated IMPRACTical narrative illustrating: character co-operation; prediction of other character intent; chaining of commands and intent; and unfulfilled intents. For further detail see text.

7. CONCLUSIONS

We have shown that extensive search for high-quality narratives is not mutually exclusive with run time efficiency. Our IMPRACTical approach begins to bridge the gap from monolithic intentional planners to multi-agent simulation, successfully decomposing intent and action causality into global search and agents' actions' relevance.

We have detailed an implementation for agents based on relaxing the domain that is able to determine the union of an agent's reasonable plans efficiently. This was shown to be 40 times faster than solving the current state-of-the-art's example compiled narrative problem.

With a new extended domain we have shown that monolithic intentional planners do not scale when action options are increased. With our approach the heuristic h_{FF} is shown to still produce narratives in real-time, and we have presented a new more informed heuristic for IMPRACTical search that makes further significant reductions in the number of states evaluated.

The combination of global search and multiagent planning presented here is a practical solution to the generation of believable narratives for IS that guarantee author and narrative goals.

8. REFERENCES

[1] R. Aylett, J. Dias, and A. Paiva. An affectively-driven planner for synthetic characters. In *Proc. of 16th Int. Conf. on Automated Planning and Scheduling (ICAPS-06)*, 2006.

[2] J. Bates. The role of emotion in believable agents. *Communications of the ACM*, 37:122–125, 1994.

[3] B. Bonet and H. Geffner. Planning as heuristic search: New results. In *Proc. of the European Conference on Planning (ECP 1999)*, pages 360–372. Springer, 1999.

[4] M. Brenner. Creating Dynamic Story Plots with Continual Multiagent Planning. In *Proc. of the 24th AAAI Conf. on Artificial Intelligence (AAAI-10)*, 2010.

[5] M. Cavazza, F. Charles, and S. J. Mead. Character-based interactive storytelling. *IEEE Intelligent Systems*, 17(4):17–24, 2002.

[6] H.-M. Chang and V.-W. Soo. Planning to Influence Other Characters in Agent-Based Narratives. In *Integrating Technologies for Interactive Stories Workshop, Int. Con. on Intelligent Technologies for Interactive Entertainment*, 2008.

[7] D. Dennett. *The Intentional Stance*. MIT Press, Cambridge, MA, 1989.

[8] P. Haslum. Narrative Planning: Compilations to Classical Planning. *Journal of AI Research*, 44:383–395, 2012.

[9] M. Helmert and C. Domshlak. Landmarks, Critical Paths and Abstractions: What's the Difference Anyway? In *Proc. of the 19th Int. Conf. on Automated Planning and Scheduling*, 2009.

[10] J. Hoffmann and B. Nebel. The FF Planning System: Fast Plan Generation throughHeuristic Search. *Journal of AI Research*, 14:253–302, 2001.

[11] B. Magerko, J. Laird, M. Assanie, and A. K. D. Stokes. AI Characters andDirectors for Interactive Computer Games. In *Proceedings of AAAI*, 2004.

[12] M. J. Nelson, D. L. Roberts, C. L. Isbell, Jr., and M. Mateas. Reinforcement learning for declarative optimization-based drama management. In *Proc. of 2nd Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-06)*, 2006.

[13] J. Porteous, J. Teutenberg, F. Charles, and M. Cavazza. Controlling Narrative Time in Interactive Storytelling. In *Proc. of 10th Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS-11)*, 2011.

[14] M. Riedl. *Narrative Generation: Balancing Plot and Character*. PhD thesis, Department of Computer Science, North Carolina State University, 2004.

[15] M. Riedl. Incorporating Authorial Intent into GenerativeNarrative Systems. In *Proc. of AAAI Spring Symp. Intelligent Narrative Technologies*, 2009.

[16] M. Riedl and R. Young. Open-world planning for story generation. In *Proc. of the Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, 2005.

[17] M. O. Riedl and R. M. Young. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research*, 39:217–267, 2010.

[18] M. Si, S. Marsella, and D. Pynadath. Evaluating Directorial Control in a Character-Centric Interactive Narrative Framework. In *Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-10)*, 2010.

[19] M. Si, S. C. Marsella, and D. V. Pynadath. Thespian: using multi-agent fitting to craft interactive drama. In *the 4th international joint conference on Autonomous Agents and Multi-Agent Systems, (AAMAS-05)*, 2005.

[20] S. R. Turner. Minstrel: A computer model of creativity and storytelling. Technical report, Technical Report UCLA-AI-92-04, Computer Science Department, University of California, 1992.

[21] R. P. y Pérez. *MEXICA: A computer model of creativity in writing*. Ph.d. thesis, The University of Sussex, Brighton, UK, 1999.