An Intent-Driven Planner for Multi-Agent Story Generation

Mark Owen Riedl
Liquid Narrative Group
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
moriedl@ncsu.edu

R. Michael Young
Liquid Narrative Group
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
young@csc.ncsu.edu

Abstract

The ability to generate narrative is of importance to computer systems that wish to use story effectively for entertainment, training, or education. We identify two properties of story – plot coherence and character believability – which play a role in the success of a story. Plot coherence is the perception by audience members that character actions have relevance to the outcome of the story. Character believability is the perception that character actions are motivated by agents' internal beliefs and desires. Unlike conventional planning in which plan goals represent an agent's intended world state, multi-agent story planning involves goals that represent the outcome of a story. In order for the plans' actions to appear believable, multi-agent story planners must determine not only how agents' actions achieve a story's goal state, but must also ensure that each agent appears to be acting intentionally. We present a narrative generation planning system for multi-agent stories that is capable of generating narratives with both strong plot coherence and strong character believability. The planning algorithm uses causal reasoning and a simulated intention recognition process to drive plan creation.

1. Introduction

Narrative as entertainment, in the form of oral, written, or visual stories, plays a central role in our social and leisure lives. There is evidence that suggests that we, as humans, build cognitive structures that represent the real events in our lives using models similar to the ones used for narrative in order to better understand the world around us [4]. This *narrative intelligence* [12] is central in the cognitive processes that we employ across a range of experiences, from entertainment contexts to active learning.

The ability to generate narrative, that is, to create and structure novel event sequences so that they can be understood as elements of a story, is of importance to systems that wish to effectively use story for entertainment, training, or education. Most existing narrative-oriented virtual worlds are built using pre-scripted action sequences; story world characters play out the same elements of the same story each time the system is run. In contrast, a system that generates a novel narrative structure for each user session could tailor its narratives to the individual preferences or needs of the user instead of relying on scripted sequences prepared in advance.

A story is, by definition, narrative. A narrative, in its simplest form, is a temporally ordered sequence of events. Typically, the events that make up the narrative represent the actions of one or more actors – characters – that exist in the story-world. To distinguish a narrative from a random series of occurrences, narratologists argue that a narrative must have "a continuant subject and constitute a whole" [17]. One way to enforce a continuant subject is to insist on some form of causal relationship between events that an audience can understand [19]. Unfortunately, the definition of narrative provided by narratologists does not specify qualitative measures with which to characterize a given narrative. We propose two principles which we believe to be integral to the success of a story: *plot coherence* and *character believability*. Plot coherence is the perception that the events of the narrative have meaning and relevance to the outcome of the story. Character believability is the perception that a character's actions are performed due to his or her internal traits and desires. Character believability refers to the numerous elements that allow a character to achieve the "illusion of life" [2]. In this work, we narrow our focus to intentionality and personality. Both plot coherence and character believability help determine how compelling a story is: a story without strong plot coherence will seem pointless and meandering; a story without strong character believability will be unable to persuade an audience to suspend its disbelief.

2. Related Work

Tale-spin [13] generates stories through inference about character goals. As characters are introduced, they are given goals to achieve. An inference engine determines how each character reacts to the changing world state in order to achieve their individual needs. Narrative emerges from the interaction between characters chosen by the inference engine. Careful consideration must be given when defining the initial state of the world and the character goals or the ensuing narrative will be poorly structured [13].

The Oz project [2; 10; 20] situates a user in a virtual environment populated by autonomous, animated agents. Each agent has a set of goals and behaviors and

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

autonomously works toward achieving its personal goals by applying personality-rich behavior descriptions. In order to achieve some aspects of plot coherence, an external module – a drama manager – discreetly manipulates the agents' goals in order to force coherent narrative to emerge [20]. Kelso and Bates observe that when the user is immersed within the story world, the user is less likely to notice inconsistencies in character behavior or plot coherence [8].

The I-Storytelling system [6] implements autonomous agents that use HTN planners to achieve their individual goals. The initial configuration of the story world (and how the user manipulates the world) affects how the plot emerges. As the world is changed by the agents or by user intervention, other agents' actions fail and are re-planned, affecting the global emergence of plot[1].

The Universe system [9] uses a hierarchical planner to select plot fragments and piece together a narrative involving many story world characters. The planner in Universe incorporates character actions into the narrative sequence that contribute to the systems storytelling goals. Storytelling goals are high-level descriptions of plot such as "keep lovers apart" (Universe operates in the domain of soap-operas). Since the goal is a set of author intentions instead of character intentions, Universe selects character actions based on whether they solve the goal instead of considering whether it makes sense of a character to perform that action.

Façade [11], unlike other related work described above, explicitly addresses the balance of character and plot. Façade implements a reactive behavior planner that selects, orders, and executes fine-grain plot elements called *beats* that describe action/reaction behaviors that story world characters will perform. At any given time, Façade chooses the beat from a pool of eligible (based on applicability constraints) beats that is most believable and most likely to match the desired plot structure. The success of Façade, as defined by the discrepancy between emergent action and the desired plot structure, is dependent on the ability of the human author to provide, at design time, beats that satisfy every possible situation [11].

3. An Intent-Driven Planner for Narrative Generation

Effective dynamic story generation involves two seemingly conflicting requirements: character believability and plot coherence. Our approach merges partial-order planning with techniques derived from the BDI framework.

Because a planner reasons about the entire solution, a planner can reason about the overall coherence of a narrative. Young suggests that planning has many benefits as a model of narrative [21]. One of which is that partial-order plans rely on causal relationships between plan steps. Ensuring strong causal relationships between actions in a narrative is one way of achieving a "continuant whole" [19] and thus increasing plot coherence. The type of planning that is done for a single agent acting believably in a virtual world, however, is different from the type of planning that must be done for many different agents to act believability in a virtual world. Intentionality in a single agent plan is implicitly captured in the plan's goals. In a multi-agent story plan, the goals express the outcome of the narrative and do not necessarily describe a desired world state of any agent. Agents must therefore have other intentions that, when achieved, surreptitiously bring about the story outcome.

Much work in believable agents considers story world characters independently of plot. Plot is assumed to emerge from the believable interactions between characters [1]. One approach to building intelligent, autonomous agents that are capable of interacting with other autonomous agents is to build agents according to the BDI framework. The BDI framework [3] models an agent's behavior as the interaction between its beliefs, desires, and intentions. Beliefs are propositions that the agent believes to be true about the world in which it is situated. Desires are states of the world that the agent would like to make true. Intentions are the desires that an agent commits to make true. The act of committing to an intention is the process of choosing a set of consistent desires and forming a goal to make those desires true. Once a goal is established, the agent constructs a plan to achieve it [3]. The plan is relevant only to a single agent attempting to transform the world state it believes to hold into an intended world state.

We have built a special-purpose story planner that plans actions for many story world characters. The causal relationships between character actions support plot coherence. Apparent intentionality of character actions supports character believability. Causality is an important aspect of the planning process because it ensures logical necessity of actions that are performed in the story and thus ensures plot coherence. In addition, the planner has a second mechanism used when constructing a plan that ensures that all character actions added to the plan appear to be intentional. The planner simulates the intention recognition performed by the audience to determine whether character actions will be perceived by the audience as intentional. This process is integrated into the planning process: if a character's action does not appear intentional, the plan is considered flawed and additional plan structure is instantiated to repair the plan and thus give that character the appearance of intentionality. Our algorithm is an intent-driven partial order causal link (IPOCL) planner that plans actions for many story world characters based on two mechanisms: causality and intentionality.

[1] I-Storytelling operates in the domain of situation comedies and the situations in which agent plans fail are desirable [6].

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

3.1 Simulating Intention Recognition in Narrative Generation

Plan recognition, and its cousins, goal recognition and intention recognition (see [5] for a comprehensive overview), is the process of taking a partially or wholly observed sequence of agent actions and inferring the plan (or goal or intention) that the executing agent is pursuing. We believe that intention recognition is an integral part of narrative planning. Gerrig and Bernardo [7] suggest that an audience is not a passive observer of a narrative. Instead, people actively perform problem-solving in order to predict the fate of favorable story world characters. This problem-solving involves interpreting character actions, inferring future events and the probability of favorable outcomes (e.g. the low probability of a favorable outcome invariably leads to feelings of suspense in the audience [7]). Since narrative structure is transparent to the audience, the audience is only able to draw upon the visible actions of the story world characters as a source for their problem-solving conclusions. It makes sense, therefore, for a narrative planner to emulate this problem-solving process as the narrative plan is generated as a check of story "goodness". The intent-driven narrative planner performs intention recognition on character actions as a way of guiding the plan construction process. Unlike a heuristic that orders alternatives, we use simulated intention recognition process to generate plan structure.

Our intent-driven story planner utilizes two levels, the domain level and the intention level. The domain level contains physical and mental character actions of which all causal and temporal relationships are captured while the intention level records character's goals and commitments. In the domain level, a causal link [14] connects two plan steps s_1 and s_2 via condition e, written $s_1 \rightarrow^{e} s_2$, when s_1 establishes the condition e in the story world needed by subsequent action s_2 in order for step s_2 to execute. We define a character action, s_i, in the domain level as *intentional* if there is a path from s_i to some internal character goal, in the directed acyclical graph (DAG) comprised of plan steps. Internal character goals differ from the goals present in the goal state of the plan in that internal character goals reside in the intentional level and are "caused" to exist by domain-level plan steps that transform beliefs and desires into intention.

An internal character goal is the end result of a sequence of intentional character actions meant to transform the world state to match the internal character goal. We refer to this sequence of actions as an *interval of intentionality*.

Definition 1 (Interval of intentionality). *An interval of intentionality is tuple,* <S, c, g_c, s_f> *such that* $s_f \in S$, *where* S *is a set of plan steps,* c *is a symbolic reference to a character,* g_c *is an internal character goal held by* c, *and* s_f *– referred to as the final step of the interval – has* g_c *for an effect and does not temporally precede any other step in* S.

The interval of intentionality, I, is the set of actions that character c performs to achieve his internal goal, g_c. An interval of intentionality can contain more than one step with g_c as an effect. This is necessary in the case where another action undoes g_c in the world and the condition must be reestablished.

Structurally, within the narrative plan, an interval of intentionality is represented as part of a *frame of commitment*, which is a data structure recording the commitment that a story world character has to achieving some internal character goal. A frame of commitment indicates that character c is committed to achieving an internal character goal, g_c, and if g_c is undone, character c will strive to reestablish the desired world state [3].

Definition 2 (Frame of commitment). *A frame of commitment is a tuple,* <c, g_c, I>, *where* c *is a symbolic reference to a character,* g_c *is the internal character goal that* c *is committed to, and* I *is an interval of intentionality which shares the same character and internal character goal with the frame of commitment.*

While the internal character goal, g_c, is the raison d'etre for the frame of commitment to exist, the character, c, must be observed by the audience to adopt (and thus commit to) g_c. A frame of commitment is associated with a condition, e_g of the form `(intends c gc)`, which indicates that for a character to commit to an internal character goal, c must intend to bring about that world state. The condition, e_g, is established in the same way that an open condition on another plan step is established: by the effect of some preceding domain-level plan step. In this case, we refer to the step that establishes e_g as the *motivating action* for the frame of commitment. The motivating action is perceived to convert a character's beliefs and desires into intention.

Domain-level actions cause characters to commit to goals (e.g. they establish conditions for frames of commitment to exist). Domain-level actions must also be intentional.

Definition 3 (Intentionality). *An action is intentional if it belongs to some interval of intentionality that is part of a frame of commitment. The intended purpose of that action is the partial fulfillment of the internal character goal of that frame of commitment.*

If the final step of an interval of intentionality uses its effect, g_c, to causally establish a precondition of another step in another frame of commitment, we say that the frame of commitment that the final step belongs to is *in service of* that other step and, consequently, in service of that other step's frame of commitment.

In the remainder of this paper, we refer to a plan step as being part of a frame of commitment if it is a member of the frame of commitment's interval of intentionality. We now define what a plan is with respect to the intent-driven planner.

Definition 4 (Plan). *A plan is a tuple,* <S, B, O, L, C>, *where* S *is a set of steps,* B *is a set of binding constraints* *on the free variables in* S, O *is the set of ordering constraints on steps in* S, L *is a set of causal links between steps in* S, *and* C *is a set of frames of commitment.*

The sets, *S*, *B*, *O*, and *L* are defined in the standard way [14]. *C* is described in Definition 2. We must modify the definition of completeness as follows.

Definition 5 (Plan completeness). *A plan is compete if and only if (1) all preconditions for all steps are established, (2) all causal threats are resolved, and (3) all steps are intentional.*

Conditions 1 and 2 together make up the standard definition of plan completeness, which we refer to as *causally complete*. A plan can be causally complete without being fully complete under the third condition. When this situation occurs, there is no plan repair possible and the planner must backtrack to find another solution.

3.2 Integrating Intention Recognition into Least-Commitment Planning

Frames of commitment are products of the simulation of a hypothetical audience's process of intention recognition which is performed routinely on the incomplete story plan. In conventional, causal link planners, a causally incomplete plan has at least one flaw. For example, an open condition [14] is a plan flaw in which a precondition of a plan step is not causally established by a preceding step or the initial state. New plan steps are instantiated in a backward-chaining fashion with the sole purpose of establishing open conditions. However, for character actions to appear intentional, every character action in the story plan must be part of the interval of some frame of commitment. Thus, when a plan step is newly instantiated, it must be declared part of the interval of an existing frame of commitment or a new frame of commitment must be created that describes a possible intention that the character has for performing that action. Either way, every plan step in a complete plan is linked to a frame of commitment and is thus declared intentional.

Intention recognition simulation is applied to the incomplete plan when a character action is newly instantiated or when a character action is reused. The purpose of intention recognition is to realize that a new character action may be intended as part of an interval of an existing frame of commitment, or to recognize the character action as part of a new intention. Figure 1 shows the IPOCL algorithm. The algorithm is broken up into three parts: causal planning, motivation planning, and intent planning.

The causal and motivation planning portions of the algorithm are an implementation of partial-order planning [14] with the addition of a frame of commitment discovery phase. Causal planning repairs open condition flaws and also simulates the intention recognition process. Suppose a newly instantiated action, s_1, is a behavior to be performed by character c. Possible frames of commitment for s_1 are "discovered" from the following.

- All existing frames of commitment, F=<c, g_c, I>, such that there is a path from s_1 to s_2 in the DAG comprised of plan steps and causal links, and where s_2 is the final step of I.

IPOCL (<*S*, *B*, *O*, *L*, *C*>, *F*, *Λ*)

On the initial call to IPOCL, there are only two steps in *S* – the dummy initial and final steps – and a single ordering constraint between them in *O*. *F* contains an open condition flaw for each precondition of the dummy final step. $B = L = C = \{\}$. *Λ* is the set of action schemata.

I. Termination. If *O* or *B* is inconsistent, fail. If *F* is empty and $\forall s \in S, \exists c \in C \mid s$ is part of *c*, return <*S*, *B*, *O*, *L*, *C*>. Otherwise, if *F* is empty, fail.

II. Plan Refinement. Non-deterministically do one of the following.

- **Causal Planning**
  - **Goal Selection.** Select an open condition flaw f = <s_{need}, p> from *F*. Let $F' = F - f$.
  - **Operator selection.** Let s_{add} be a step that adds an effect *e* that can be unified with *p* (to create s_{add}, non-deterministically choose a step s_{old} already in *S* or instantiate an action schema in *Λ*). If no such step exists, backtrack. Otherwise, let $S' = S \cup \{s_{add}\}$, $O' = O \cup \{s_{add} < s_{need}\}$, $B' = B \cup$ bindings needed to make s_{add} add *e*, including the bindings of S_{add} itself, and $L' = L \cup \{$<s_{add}, e, p, s_{need}>$\}$. If $s_{add} \neq s_{old}$, add new open condition flaws to *F'* for every precondition of s_{add}.
  - **Frame discovery.** Let $C' = C$.
    - If $s_{add} \neq s_{old}$, non-deterministically choose an effect *e* of s_{add} or e = nil. If $e \neq$ nil, construct a new frame of commitment *c* with internal character goal *e*, let s_{add} be part of *c*, let $C' = C' \cup \{c\}$, create a new open motivation flaw, f = <c>, and let $F' = F' \cup \{f\}$.
    - Let C" be the set of existing frames of commitment that can be used to explain s_{add}. For all $d \in C''$, create an intent flaw f = <s_{add}, d> and let $F' = F' \cup \{f\}$.
  - **Threat resolution.** Performed in the standard way.
  - **Recursive invocation.** Call IPOCL(<*S'*,*B'*,*O'*,*L'*,*C'*>, *F'*, *Λ*).
- **Motivation Planning**
  - **Goal selection.** Select an open motivation flaw f = <c>. Let *p* be the condition of *c*. Let $F' = F - f$.
  - **Operator selection.** Same as causal planning above, except $O' = O \cup \{s_{add} < s_i \mid s_i$ is part of $c\}$.
  - **Frame discovery.** Same as for causal planning.
  - **Threat resolution.** Performed in the standard way.
  - **Recursive Invocation.** Call IPOCL(<*S'*,*B'*,*O'*,*L'*,*C'*>, *F'*, *Λ*).
- **Intent Planning**
  - **Goal selection.** Select an intent flaw f = <s, c> from *F*. Let $F' = F - f$.
  - **Frame selection.** Let $O' = O$. Non-deterministically choose to do one of the following.
    - Make *s* part of *c*. Let s_m be the motivating step of *c*. $O' = O' \cup \{s_m < s\}$. For each $s_{pred} \in S$ such that <s_{pred}, p, q, s> $\in L$, create an intent flaw, f = <s_{pred}, c>. Let $F' = F' \cup \{f\}$.
    - Do not make *s* part of *c*.
  - **Recursive invocation.** Call IPOCL(<*S*, *B*, *O'*, *L*, *C*>, *F'*, *Λ*).

Figure 1. The IPOCL algorithm.

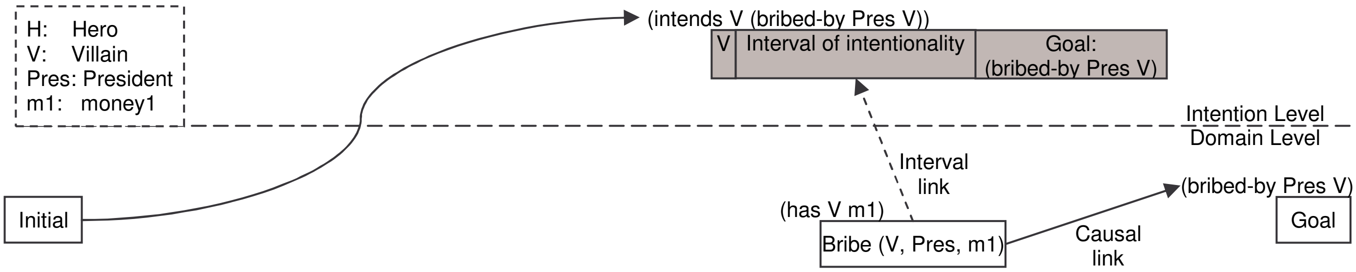

Figure 2: A story plan with a single frame of commitment.

- All new frames of commitment, F=<c, g_c, I>, where g_c is an effect of s_1, and s_1 is the only member of I. An open condition flaw is recorded for F.

The intent-driven planner considers the possibility that s_1 is a member of any of the possible frames of commitment or more than one frame of commitment (or none). A branch in the plan space is created for each of these possibilities. Heuristic functions are used that penalize against prolific construction of new frames of commitment, preferring instead plans with a minimal number of frames. If the frame of commitment, F, is newly instantiated, s_1 is added to F's interval of intentionality as the final step of the interval. If the frame of commitment, F, is an existing frame, then an intent flaw is noted, indicating that the planner must determine whether or not to associate the step with F. If s_1 is part of no frame of commitment, it is an *orphan* and the planner assumes that it will be used to causally satisfy some other (possibly yet-to-be-discovered) frame of commitment.

Motivation planning repairs an open motivation flaw in which a condition on a frame of commitment must be satisfied by a new or existing plan step. Motivating steps are explicitly ordered before all other steps in the frame's interval of intentionality.

The intent planning portion of the algorithm resolves an intent flaw, f = <s, c> where s is a plan step and c is an existing frame of commitment. The planner non-deterministically determines whether to add s to the interval of intentionality of c or to leave the structure of the plan untouched. If s is added to the interval of intentionality of c, then s is explicitly ordered after the motivating step of c. Let s_{pred} be a step that precedes s and is causally link to s – referred to as an establishing step. The inclusion of s in the interval of c makes it possible for establishing actions to be included in the interval of intentionality of c if the following cases hold.

- s_{pred} is performed by the same character as s.
- s_{pred} is not a part of the interval of intentionality of c.
- The intent flaw, f_{pred} = <s_{pred} , c> has not already been proposed and/or solved for (to preserve the systematicity of the algorithm).

New intent flaws are noted for all establishing steps of s for which the three cases hold.

3.2.1. An Example. We illustrate the algorithm with a story about an arch-villain who bribes the President of the United States with a large sum of money. We trace one possible path through the branching plan search space generated by the intent-driven planning algorithm. The domain level of the story plan contains, to begin with, the goal state `(bribed-by President villain)`. As an open condition, the goal state is established by the character action, `Bribe(villain, President, money1)`, in which the villain achieves the goal state of the story by bribing the President with some money. From the planner's perspective, the `Bribe` action is causally motivated by the open condition of the goal state. The audience, as active problem-solvers, will be considering a different reason for this action to take place having to do with some commitment the villain has to some internal character goal, (as yet unspecified due to the backward-chaining nature of the planner). Upon instantiation of the new character action, intention recognition is invoked. No previous frame of commitment exists, so a new frame of commitment is non-deterministically chosen with the internal character goal of `(bribed-by President villain)`.

The plan, even with the new frame of commitment, is still flawed because there is no reason for the villain character to have the internal goal of bribing the President. That is, the villain character needs to *form* the intention. This is denoted by the fact that the frame of commitment has an open motivation condition, `(intends villain (bribed-by President villain))`. For simplicity sake, the villain's intention is declared as part of the initial state of the world. The planner non-deterministically establishes the open motivation condition with a causal link from the plan's initial state. The story plan, as generated so far, is presented in Figure 2.

Continuing the bribery example, we see from Figure 2 that the bribe action has a single open precondition, `(has villain money1)`. The planner non-deterministically chooses to satisfy this with a new action, `Give(hero, villain, money1)`, in which the hero character gives the money to the villain. The only existing frame of commitment belongs to the villain character; there is no appropriate frame of commitment for the hero's action. The plan recognition process infers that the possible goals for the hero character are `(has villain money1)` and `(not`

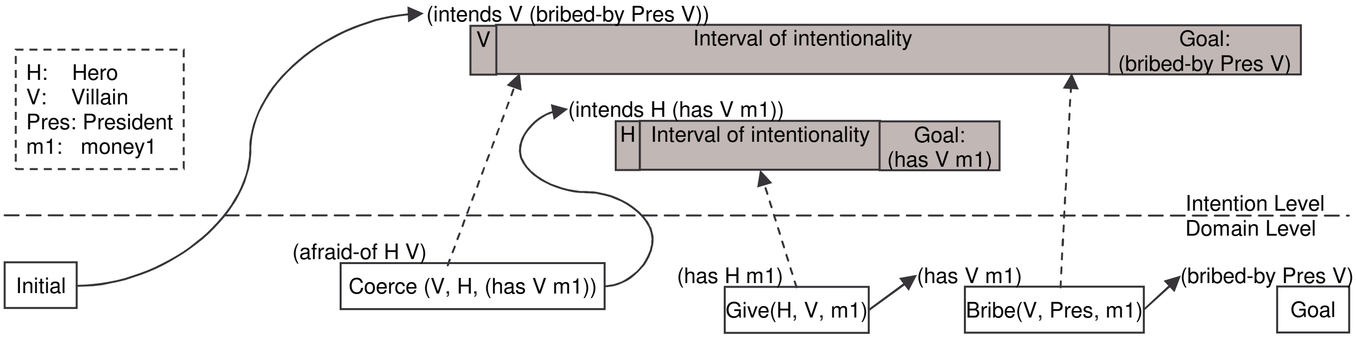

Figure 3: A story plan with several interrelated frames of commitment

`(has hero money))` by inspecting the effects of the given action. Suppose the former internal goal is non-deterministically chosen. A frame of commitment is created for that internal character goal which has an open motivation condition, `(intends hero (has villain money1))`. There are many possible actions that will establish this intention; the villain might persuade the hero to give him the money if they are friends, or the villain might coerce the hero. The latter, `Coerce(villain, hero, (has villain money1)`, is chosen as the cause of the hero's intention that the villain has the money.

At this point, the planner now has to "discover" a frame of commitment to explain the `Coerce` action. The system is able to determine that the villain's previous frame of commitment is acceptable along with any possible newly created frames of commitment. Although it is not obvious from the Figures, the hero's frame of commitment records the fact that it was created *in service of* the villain's frame of commitment. Therefore, there is a path from the `Coerce` action to the hero's frame of commitment to the `Bribe` action to the villain's original frame of commitment. Had the intent-driven story planner instead chosen `(not (has hero money1))` as the hero's internal goal, such a path would not exist because the hero's frame of commitment would not be created in direct fulfillment of a precondition of the villain's `Bribe` action. The planner non-deterministically chooses the branch that uses the existing frame of commitment for `Coerce` and the result is shown in Figure 3.

As the planning process continues, the open condition, `(has hero money1)`, on `Give(hero, villain, money1)` is established by a new action, `Steal(hero, money1, bank)`, in which the hero steals the money from the bank. The `Steal` action establishes a precondition of `Give` and can therefore be considered part of the hero's existing frame of commitment. Likewise, the open condition, `(afraid-of hero villain)`, on `Coerce(villain, hero, …)` is established by a new action, `threaten(villain, hero)`, in which the villain threatens harm onto the hero. This action is declared part of the villain's commitment to bribing the President.

3.2.2. Supporting Actions for Multiple Commitments. The example only considers the case where an action is part of a single frame of commitment. Every action in the plan must be part of some frame of commitment. However, there may be situations where a single action is intended to fulfill more than one internal character goal. That is, a single action is performed by some character with the intention of satisfying more than one internal character goal. This corresponds to Pollack's notion of *overloading* [16].

Every time an action – belonging to one frame of commitment – is used to satisfy an open condition of a successor action that belongs to a different frame of commitment, the system must non-deterministically decide whether the condition-establishing action belongs to both frames of commitment or remains only a member of its original frame. This decision corresponds to the intent planning portion of the algorithm in Figure 1. The decision about interval membership affects the possible orderings of motivating actions. Motivating actions are temporally ordered before all actions in the interval of the frame of commitment that the motivational action establishes. Any motivational action can be placed temporally in the plan at any point before the interval of its frame of commitment begins. Heuristically, we suggest that a motivating action be positioned as close to, but before, the earliest action in the interval. When an action is a member of more than one frame of commitment, the possible placement of motivational actions is constrained as in Figure 4.

For example, if a character has an internal goal of killing a deer and an internal goal of robbing a bank, then that character may pick up and load a gun as part of his commitment for one or both of those internal goals. If the act of loading the gun is part of both frames of commitment, then the character must form intentions to kill a deer and rob the bank before loading the gun. However, if loading the gun is only part of the commitment to kill a deer, then the character can form the intention of robbing a bank after the gun is loaded.

Figure 4 illustrates two overlapping frames of commitment for a single character, *c*. Suppose that the steps in the plan are added as numbered. By solving for

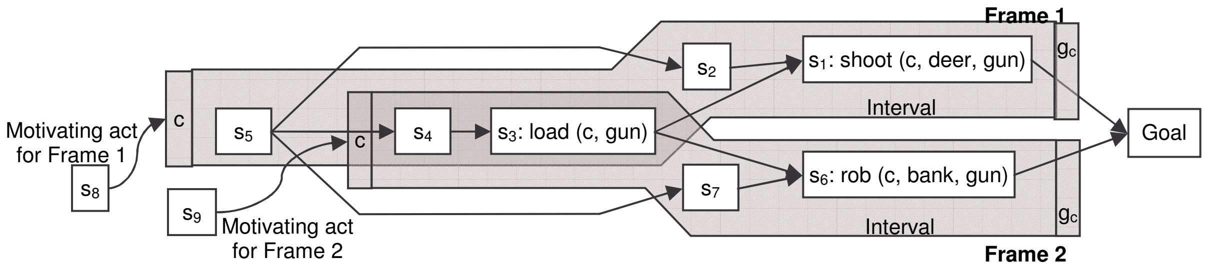

Figure 4: Frames of commitment that share sequences of actions.

intent flaws, the interval of intentionality of Frame 1 spreads to incorporate steps s_1 through s_5. Similarly, the interval of intentionality of Frame 2 spreads to incorporate s_6 and s_7. Given that step s_3 is an establishing step for actions in Frame 2, the planner has two choices: s_3 can remain solely in Frame 1 or it can assume joint membership in Frame 2. If the latter option is selected, the interval membership status of step s_4 comes into question. The planner decides whether s_4 remains part of the first frame or whether it assumes joint membership as well. Had s_3 been originally left as an orphan, the planner would have had different choices to consider: s_3 could remain an orphan, or s_3 could join the second frame of commitment. Thus, through two decision points, all four membership possibilities exist for step s_3: orphan, member of Frame 1, member of Frame 2, or joint member of both frames.

3.2.3. Orphans. The planner allows actions to be orphans (i.e. to not belong to any frame of commitment) in order to avoid making an overly strong commitment to the frame membership of an action. We recognize the fact that, for completeness, an action might need to be part of a frame of commitment that has not been discovered yet. Orphaned actions represent flaws in the plan. Unlike other flaws (e.g. open conditions, causal threats, etc.), orphans are not explicitly repaired. Instead, orphans are repaired surreptitiously when they are adopted into new intervals of intentionality because they causally establish other, intentional actions. Orphan flaws cannot be repaired directly because frames of commitment are discovered opportunistically instead of created in a least-commitment approach (as plan steps are).

This strategy of leaving orphans with the hope that they will be adopted eventually is not without some risk. It is possible that an orphan is never used to establish another open condition and thus be adopted. In this case, the plan can be causally complete but not complete with respect to intentionality. If this occurs, the plan is simply pruned and the planner backtracks to find another solution plan.

4. The Role of Intention in Narrative

The work presented here is consistent with previous work on intention [3] and integrates the notion of intentionality into a multi-agent planning framework. In a narrative, intention is one mechanism for creating believability in character actions. One of the central features of a story world that gives an audience insight into the nature of a story's character is the actions that the character chooses to perform. From this information, the audience must infer character and plot, both of which are known to the hypothetical author but initially unknown to the audience [7]. It is essential that story world characters *appear* to be acting intentionally. To facilitate this appearance, the IPOCL algorithm ensures that domain-level actions that motivate an agent's adoption of internal character goals are inserted into the narrative plan. This differs from more conventional BDI treatments that handle a single agent with a single task where goal commitment occurs separately from planning. To integrate intentional structure with causal structure, our planning approach accounts for the type of intention recognition that a story's audience might perform; this process eliminates plans that contain actions that cannot be accounted for based on intentional structure.

The representation of intentional structure within a multi-agent plan facilitates additional plan features that can be used to constrain the problem of narrative generation, achieving more intricate plots and more believable characters. For example, personality – the characteristic of a person that account for consistent patterns of behavior [15] – can be used to prune the narrative plan search space by eliminating possible narratives in which story world characters perform contradictory actions [18]. Within the computational model of intentional planning, character commitments to internal goals are represented by domain-level actions (e.g. the act of deciding to adopt an intention). Therefore, personality plays the double role of also eliminating plans in which characters adopt inconsistent intentions. However, to strictly exclude actions that contradict personality affects planning completeness [18]. There are situations when it is desirable for an agent to act "out of character". For instance, the example in Section 3.2.1 demonstrates one situation where acting out of character is appropriate. The hero character, whom we shall presume has a lawful personality, steals money from a bank because he is coerced to adopt a goal that is not of his own desire. How personality is represented computationally and the exact mechanism for applying agent personality to narrative plan search space management is beyond the scope of this paper.

5. Limitations and Future Work

Plot coherence and character believability are desirable properties for narrative. The IPOCL planner is capable of generating story plans that have strong plot coherence and character believability. The planner is implemented and an empirical evaluation is underway. However, the IPOCL algorithm does not computationally define what makes a "good" or "interesting" story. IPOCL extends least-commitment planning and is, consequently, complete; given the initial conditions, if a coherent, believable plan exists then it is in the planner's reachable search space. The primary limitation of IPOCL is that, to date, we have not defined the search space heuristic functions that inform the search. That is, IPOCL is unable to determine whether one branch through the search space is better than another, either in terms of plan completeness or in terms of "goodness" or "interestingness." One of the goals of our future work is to develop computational definitions for "goodness" and create corresponding heuristic evaluation functions.

6. Conclusions

An audience that watches a story unfold within a virtual world are active observers that continuously perform cognitive problem-solving tasks in order to predict the fate of the story world's characters [7]. Since a plot is comprised of actions that story world characters perform, audience problem-solving involves inferring character beliefs, desires, and intentions from observable character actions. Because plot coherence and character believability relate directly to the active problem-solving processes of a story's audience, narrative generation systems must account for both of these features in the stories they create.

The intent-driven planning algorithm presented in this paper demonstrates an approach to narrative generation that simultaneously solves for plot coherence and character believability. Planning invokes causal reasoning to produce action sequences that are connected by causal necessity. The BDI agent framework lends the formulation of intent and intention recognition to the problem of narrative generation. Even though the intent-driven planning algorithm is based on conventional planning approaches, both causality and intentionality are used as mechanisms for building plan structure so that character actions are part of a coherent plot and character actions are intentional and thus believable.

7. Acknowledgements

The work has been supported by National Science Foundation CAREER award 0092586.

8. References

[1] Aylett, R. (1999). Narrative in virtual environments – towards emergent narrative. In the *AAAI Fall Symposium on Narrative Intelligence*, 1999.

[2] Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7), 122-125.

[3] Bratman, M.E. (1987). *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.

[4] Bruner, J. (1990). *Acts of Meaning*. Cambridge, MA: Harvard University Press.

[5] Carberry, S. (2001). Techniques for plan recognition. *User-Modeling and User-Adapted Interaction*, 11(1-2).

[6] Cavazza, M., Charles, F., & Mead, S.J. (2001). Characters in search of an author: AI-based virtual storytelling. In *Proceedings of the 1st International Conference on Virtual Storytelling*, Avignon, France, 2001.

[7] Gerrig, R. & Bernardo, D. (1994). Readers as problem-solvers in the experience of suspense. *Poetics*, 22.

[8] Kelso, M.T. & Bates, J. (1992). Dramatic presence. Presence: *The Journal of Teleoperators and Virtual Environments*, 2(1).

[9] Lebowitz, M. (1985). Story-telling as planning and learning. *Poetics*, 14.

[10] Mateas, M. (1997). *An Oz-centric review of interactive drama and believable agents*. Technical Report CMU-CS-97-156, Carnegie Mellon University.

[11] Mateas, M. (2002). *Interactive Drama, Art and Artificial Intelligence*. Ph.D. Dissertation, Carnegie Mellon University.

[12] Mateas, M. & Sengers, P. (2003). *Narrative Intelligence* Philadelphia: John Benjamins.

[13] Meehan, J.R. (1976). *The Metanovel: Writing Stories by Computer*. Ph.D. Dissertation, Yale University.

[14] Penberthy, J.S. & Weld, D. (1992). UCPOP: A sound, complete, partial-order planner for ADL. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*.

[15] Pervin, L.A. (1993). *Personality: Theory and Research*. 6th Edition. New York: Wiley & Sons.

[16] Pollack, M.E. (1992). The uses of plans. *Artificial Intelligence*, 57(1), 43-68.

[17] Prince, G. (1987). *A Dictionary of Narratology*. Lincoln, NE: University of Nebraska Press.

[18] Rizzo, P., Veloso, M., Miceli, M., & Cesta, A. (1999). Goal-based personalities and social behaviors in believable agents. *Applied Artificial Intelligence*, 13.

[19] Trabasso, T. & Sperry, L.L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, 24, 595-611.

[20] Weyhrauch, P. (1997). *Guiding Interactive Fiction*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA.

[21] Young, R.M. (1999). Notes on the use of plan structures in the creation of interactive plot. In the *AAAI Fall Symposium on Narrative Intelligence*, 1999.