Toward a Plan-Based Understanding Model for Mixed-Initiative Dialogues

Hiroaki Kitano* and Carol Van Ess-Dykema†
Center for Machine Translation
Carnegie Mellon University
Pittsburgh, PA 15213
hiroaki@cs.cmu.edu vaness@cs.cmu.edu

ABSTRACT

This paper presents an enhanced model of plan-based dialogue understanding. Most plan-based dialogue understanding models derived from [Litman and Allen, 1987] assume that the dialogue speakers have access to the same domain plan library, and that the active domain plans are shared by the two speakers. We call these features shared domain plan constraints. These assumptions, however, are too strict to account for mixed-initiative dialogues where each speaker has a different set of domain plans that are housed in his or her own plan library, and where an individual speaker's domain plans may be activated at any point in the dialogue. We propose an extension to the Litman and Allen model by relaxing the shared domain plan constraints. Our extension improves (1) the ability to track the currently active plan, (2) the ability to explain the planning behind speaker utterances, and (3) the ability to track which speaker controls the conversational initiative in the dialogue.

1. Introduction

In this paper, we present an enhanced plan-based model of dialogue understanding that provides a framework for computer processing of mixed-initiative dialogues. In mixed-initiative dialogues, each speaker brings to the conversation his or her own plans and goals based on his or her own domain knowledge, and which do not necessarily match those of the other speaker, even in cooperative situations. Thus, mixed-initiative dialogues exhibit a more complicated discourse structure than do dialogues in which a single speaker controls the conversational initiative.

The existing plan-based model of dialogue understanding (as represented by [Litman and Allen, 1987]) accounts for dialogues in which a single speaker controls the initiative. We call these dialogues *Single-Initiative Dialogues*. In modeling *single-initiative dialogues*, Litman and Allen assume a shared stack that represents a *joint plan* (joint domain plan). This *joint plan* is shared by the two speakers. We claim that this assumption is too restrictive to apply to mixed-initiative dialogues, because in mixed-initiative dialogues each speaker may have his or her own individual domain plans[1]. The assumption creates several functional problems in the Litman and Allen model, namely, its inability to process mixed-initiative dialogues and the need for a large amount of schema definition (domain knowledge representation) to handle complex conversational interactions.

The model we present builds on the framework of [Litman and Allen, 1987]. We hypothesize, however, that *speaker-specific plan libraries* are needed, instead of a single plan library storing *joint plans*, for a plan-based theory of discourse to account for mixed-initiative dialogues. In our framework, the understanding system activates the instantiated schemata (places them on the stack) from each speaker's individual plan library[2], thus creating two domain plan stacks. We also theorize that in addition to using the domain plans that are stored in a speaker's memory (plan library), speakers incrementally expand their domain plans in response to the current context of the dialogue. These extensions enable our model to:

- Provide a mechanism for tracking the currently active plan in mixed-initiative dialogues,
- Explain the planning behind speaker utterances,
- Provide a mechanism for tracking which speaker controls the conversational initiative, and for tracking the nesting of initiatives within a dialogue segment.
- Reduce the amount of schema definition required to process mixed-initiative dialogues.

Throughout this paper, we use two dialogue extrac-

*This author is supported, in part, by NEC Corporation, Japan.

†This author's research was made possible by a post-doctoral fellowship awarded her by the U.S. Department of Defense. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Defense or of the United States government.

[1] In this regard, we agree with [Grosz and Sidner, 1990]'s criticism of the *master-slave* model of plan recognition.

[2] Using the [Pollack, 1990] distinction, plans are *mental objects* when they are on the stack, and *recipes-for-action* when they are in the plan library.

tions from our data: 1) an extraction from a Japanese dialogue in the conference registration domain, and 2) an extraction from a Spanish dialogue in the travel agency domain.[3] SpA and SpB refer to Speaker A and Speaker B, respectively.

Dialogue I (Conference Registration, translated from Japanese):

SpA: I would like to attend the conference. (1)
SpA: What am I supposed to do? (2)
SpB: First, you must register for the conference. (3)
SpB: Do you have a registration form? (4)
SpA: No, not yet. (5)
SpB: Then we will send you one. (6)

Dialogue II (Travel Agency, translated from Spanish):
Prior to the following dialogue exchanges, the traveler (SpB) asks the travel agent (SpA) for a recommendation on how it is best to travel to Barcelona. They agree that travel by bus is best.

SpA: You would leave at night. (1)
SpA: You would take a nap in the bus on your way to Barcelona. (2)
SpB: Couldn't we leave in the morning ... instead of at night? (3)
SpA: Well, it would be a little difficult. (4)
SpA: You would be traveling during the day which would be difficult because it's very hot. (5)
SpB: Really? (6)

2. Limitations of the Current Plan-Based Dialogue Understanding Model

The current plan-based model of dialogue understanding [Litman and Allen, 1987] assumes a single plan library that contains the domain plans of the two speakers, and a shared plan stack mechanism to track the current plan structure of the dialogue. The shared stack contains the domain plans and the discourse plans from the plan library that are activated by the inference module of the dialogue understanding system. The domain plan is a joint plan shared by the two dialogue speakers. Although this shared stack mechanism accounts for highly task-oriented and cooperative dialogues where one can assume that both speakers share the same domain plan, the model does not account for mixed-initiative dialogues.

In this section we examine three limitations of the current plan-based dialogue understanding model: 1) the inability to track the currently active plan, 2) the inability to explain a speaker's planning behind his or her utterances, and 3) the inability to track conversational initiative control transfer. A dialogue understanding system must be able to infer the dialogue participants' goals in order to arrive at an understanding of the speakers' actions. The inability to explain the planning behind speaker utterances is a serious flaw in the design of a plan-based dialogue processing model. Tracking the conversational control initiative provides the system with a mechanism to identify which of a speaker's plans is currently activated, and which goal is presently being persued. We believe that an understanding model for mixed-initiative dialogues must be able to account for these phenomena.

2.1. Tracking the Currently Active Plan

The Litman and Allen model lacks a mechanism to track which plan is the currently active plan in mixed-initiative dialogue where the two speakers have very different domain plan schemata in their individual plan libraries. The currently active plan is the plan or action that the dialogue processing system is currently considering. In Dialogue I, after utterance (2), *What am I supposed to do?*, by SpA, the stack should look like Figure 1[4]. Although the manner in which the conference registration domain plans may be expanded on the stack depends upon which domain plan schemata are available in a speaker's domain plan library, we assume that a rational agent would have a schema containing the plan to attend a conference, Attend-Conference. This plan is considered the currently active plan and thus marked *[Next]*. When processing the subsequent utterance, (3), *First, you must register for the conference.*, the currently active plan should be understood as registration, Register, since SpB clearly states that the action[5] of registration is necessary to carry out the plan to attend the conference. The Litman and Allen model lacks a mechanism for instantiating a new plan within the domain unless the currently ac-

[3] Dialogue 1 is extracted from a corpus of Japanese ATR (Advanced Telecommunication Research) recorded simulated conference registration telephone conversations. No visual information was exchanged between the telephone speakers. Dialogue 2 is extracted from a corpus of recorded Spanish dialogues in the travel agency domain, collected by the second author of this paper. These dialogues are simulated telephone conversations, where no visual information was exchanged.

[4] Notational conventions in this paper follow [Litman and Allen, 1987]. In their model, the currently active plan is labeled *[Next]*. ID-PARAM in Plan2 refers to IDENTIFY-PARAMETER. I1 in Plan2 and AC in Plan3 are abbreviated tags for INFORMREF (Inform with Reference to) and Attend-Conference, respectively. Proc in Plan2 stands for procedure.

[5] The words *plan* and *action* can be used interchangably. A sequence of *actions* as specified in the decomposition of a *plan* carry out a *plan*. Each *action* can also be a *plan* which has its own decomposition. *Actions* are not decomposed when they are primitive operators [Litman and Allen, 1987].

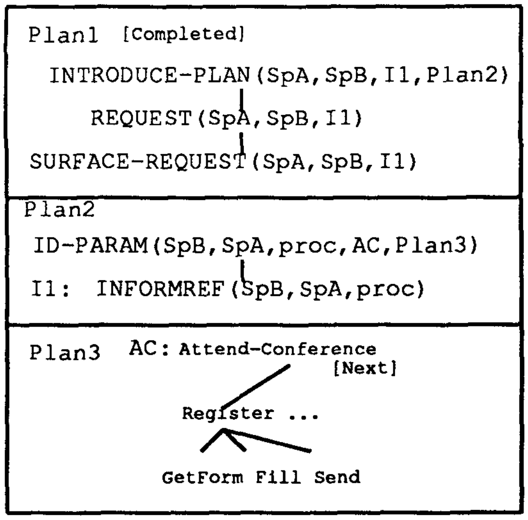
Plan1 [Completed]
INTRODUCE-PLAN(SpA,SpB,I1,Plan2)
REQUEST(SpA,SpB,I1)
SURFACE-REQUEST(SpA,SpB,I1)
Plan2
ID-PARAM(SpB,SpA,proc,AC,Plan3)
I1: INFORMREF(SpB,SpA,proc)
Plan3 AC: Attend-Conference
[Next]
Register ...
GetForm Fill Send

Figure 1: State of the Stack after Utterance (2) in Dialogue I

tive plan (or an action of the domain plan) marked by *[Next]*, is executed. Thus, in this example, only if the plan `Attend-Conference` marked as *[Next]*, is executed, can the system process the prerequisite plan, `Register`. Looking at this constraint from the point of view of an event timeline, the Litman and Allen model can process only temporally sequential actions, i.e., the `Attend-Conference` event must be completed before the `Register` event can begin.

This problem can be clearly illustrated when we look at the state of the stack after utterance (4), *Do you have a registration form?*, shown in Figure 2. Utterance (4) stems from the action `GetForm` (GF) which is a plan for the conference office secretary to send a registration form to the participant. It is an action of the `Register` plan. Since the `Attend-Conference` plan has not been executed, the system has two active plans, `Attend-Conference` and `GetForm`, both marked *[Next]*, in the stack where only `GetForm` should be labeled the active plan.

2.2. Explaining Speaker Planning Behind Utterances

A second limitation of the Litman and Allen model is that it cannot explain the planning behind speaker utterances in certain situations. The system cannot process utterances stemming from *speaker-specific domain plans* that are enacted because they are an active response to the previous speaker's utterance. This is because the model assumes a *joint plan* to account for utterances spoken in the dialogue. But utterances that stem from an active response stem from neither shared domain plans currently on the stack nor from a plan which concurrently exists in the plan libraries of the two speakers.

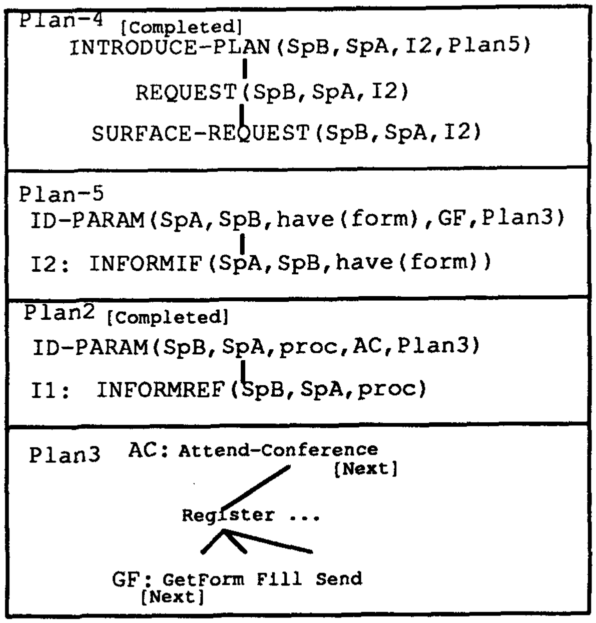
Plan-4 [Completed]
INTRODUCE-PLAN(SpB,SpA,I2,Plan5)
REQUEST(SpB,SpA,I2)
SURFACE-REQUEST(SpB,SpA,I2)
Plan-5
ID-PARAM(SpA,SpB,have(form),GF,Plan3)
I2: INFORMIF(SpA,SpB,have(form))
Plan2 [Completed]
ID-PARAM(SpB,SpA,proc,AC,Plan3)
I1: INFORMREF(SpB,SpA,proc)
Plan3 AC: Attend-Conference
[Next]
Register ...
GF: GetForm Fill Send
[Next]

Figure 2: State of the Stack after Utterance (4) in Dialogue I

In Figure 1, the `Attend-Conference` domain plan from Dialogue I is expanded with the `Register` plan after the first utterance because utterance (4), *Do you have a registrationform?*, and the subsequent conversation cannot be understood without having domain plans entailing the `Register` plan in the stack. If this were a joint domain plan, SpA's utterance *What am I supposed to do?*, could not be explained. It can be inferred that SpA does not have a domain plan for attending a conference, or at least that the system did not activate it in the stack. The fact that SpA asks SpB *What am I supposed to do?* gives evidence that SpA and SpB do not share the `Register` domain plan at that point in the dialogue.

Another example of speaker planning that the Litman and Allen model cannot explain, occurs in Dialogue II. After a series of interactions between SpA and SpB, SpB says in utterance (3), *Couldn't we leave in the morning ... instead of at night?*, as an active response to SpA. In order to explain the speaker planning behind these utterances, the current model would include the schemata shown in Figure 3[6]. Utterance (3), however, does not stem from speaker action. One way to correct this situation within the current model would be to allow for the ad hoc addition of the schema,

[6] This is a simplified list of schemata, excluding prerequisite conditions and effects. Like the Litman and Allen model, our schema definition follows that of NOAH [Sacerdoti, 1977] and STRIPS [Fikes and Nilsson, 1971].

State-Preference. The consequence, however, of this approach is that too large a number of schemata are required, and stored in the plan library. This large number of schemata will explode exponentially as the size of the domain increases.

2.3. Tracking Conversational Initiative Control

A third problem in the Litman and Allen model is that it cannot track which speaker controls the conversational initiative at a specific point in the dialogue, nor how initiatives are nested within a dialogue segment, e.g., within a clarification subdialogue. This is self-evident since the model accounts only for *single-initiative dialogues*. Since the model calls for a *joint plan*, it does not track which of the two speakers maintains or initiates the transfer of the conversational initiative within the dialogue. Thus, that the conversational initiative is transferred from SpA to SpB at utterance (3) in Dialogue II, *Couldn't we leave in the morning ... instead of at night?*, or that SpA maintains the initiative during SpB's request for clarification about the weather, utterance (6), *Really?*, cannot be explained by the Litman and Allen model.

3. An Enhanced Model

In order to overcome these limitations, we propose an enhanced plan-based model of dialogue understanding, building on the framework described in [Litman and Allen, 1987]. Our model inherits the basic flow of processing in [Litman and Allen, 1987], such as a constraint-based search to activate the domain plan schemata in the plan library, and the stack operation. However, we incorporate two modifications that enable our model to account for mixed-initiative dialogues, which the current model cannot. These modifications include:

- Speaker-Specific Domain Plan Libraries and the Individual Placement of Speaker-Specific Plans on the Stack.
- Incremental Domain Plan Expansion.

First, our model assumes a domain plan library for each speaker and the individual placement of the *speaker-specific domain plans* on the stack. Figure 4 shows how the stack is organized in our model. The domain plan, previously considered a joint plan, is separated into two domain plans, each representing a domain plan of a specific speaker. Each speaker can only be represented on the stack by his or her own domain plans. Progression from one domain plan to another can only be accomplished through the system's recognition of speaker utterances in the dialogue.

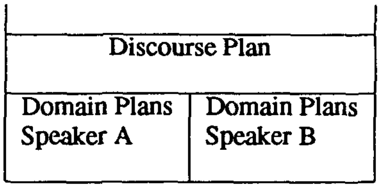

Figure 4: New Stack Structure

Second, our model includes an incremental expansion of domain plans. Dialogue speakers use domain plans stored in their individual plan library in response to the content of the previous speaker's utterance. The domain plans can be further expanded when they activate additional domain plans in the plan library of the current speaker. For example, if a domain plan is marked *[Next]* (currently active), the system decomposes the plan into its component plan sequence. Then the first element in the component plan sequence (which is an *action*) is marked *[Next]* and the previous plan is no longer marked. Figure 5 illustrates how the domain plans in Dialogue I can be incrementally expanded. In Figure 5(a), Attend-Conference is the only plan activated, and it is marked *[Next]*. As the plan is expanded, *[Next]* is moved to the first action of the decomposition sequence (Figure 5(b)). This expansion is attributed to information provided by the previous speaker, for example, *First, you must register for the conference.* (If such an utterance is not made, no expansion takes place.) Then, if the subsequent speaker has a plan for the registration procedure, the domain plan for Register is expanded under Register. Again, *[Next]* is moved to the first element of the component plan sequence, GetForm (Figure 5(c)).

We are implementing this model using the Spanish travel agency domain corpus and the Japanese ATR conference registration corpus. The implementation is in CMU CommonLisp, and uses the CMU FrameKit frame-based knowledge representation system. The module accepts output from the Generalized LR Parsers developed at Carnegie Mellon University [Tomita, 1985].

4. Examples

4.1. Tracking the Currently Active Plan

In our model, we provide a mechanism for consistently tracking the individual speaker's currently active plans. First, we show how the model keeps track of a speaker's plans within mixed-initiative dialogue. The state of the stack after utterance (2), *What am I supposed to do?*, in Dialogue I, should look like Figure 6. Plan 3 represents a domain plan of SpA,

```
((HEADER: Set-Itinerary)
     (Decomposition:  Set-Destination Decide-Transportation ...)

((HEADER: Decide-Transportation)
     (Decomposition:  Tell-Depart-Times Tell-Outcomes Establish-Agreement))
```

Figure 3: Domain Plan Schemata for Dialogue II (Partial Listing)

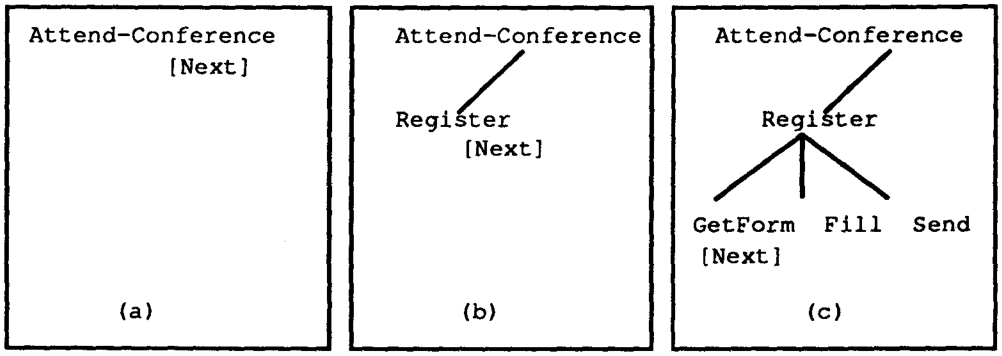

Figure 5: Incremental Domain Plan Expansion for Dialogue I

and Plan 4 represents a domain plan of SpB. Since SpA does not know what he or she is supposed to do to attend the conference, the only plan in the stack is Attend-Conference. SpB knows the registration procedure details, so his or her domain plan is expanded to include Register, and then its decomposition into the GetForm Fill Send action sequence. The first element of the decomposition is further expanded, and an action sequence notHave GetAdrs Send is created under GetForm. The action sequence notHave GetAdrs Send is a sequence where the secretary's plan is to ask whether SpA already has a registration form (notHave), and if not, to ask his or her name and address (GetAdrs), and to send him or her a form (Send).

Figure 7 shows the state of the stack in Dialogue I after SpB's question, utterance (4), *Do you have a registration form?*. From the information given in his or her previous utterance, (3), *First, you must register for the conference.*, SpA's domain plan (Plan3) was expanded downward. Thus, Plan3 has a Register plan, and it is marked *[Next]*. For SpB, notHave is marked *[Next]*, indicating that it is his or her plan currently under consideration. Although SpB's currently active plan is notHave, SpA considers the Register plan to be the current plan because SpA does not have the schema that includes the decomposition of the Register plan.

4.2. Explaining Speaker Planning Behind Utterances

Second, our model explains a speaker's active planning behind an utterance. In the Litman and Allen model, SpA's utterance (2) in Dialogue I, *What am I supposed to do?*, cannot be explained if the domain plan Attend-Conference is shared by the two speakers. In such a joint plan both speakers would know that a conference participant needs to register for a conference. However, the rational agent will not ask *What am I supposed to do?* if he or she already knows the details of the registration procedure. But, if such an expansion is not made on the stack, the system cannot process SpB's reply, *First, you must register for the conference.*, because there would be no domain plan on the stack for Register. This dilemma cannot be solved with a *joint plan*. It, however, can be resolved by assuming individual domain plan libraries and an active domain plan for each speaker. As shown in Figure 6, when SpA asks *What am I supposed to do?*, the active domain plan is solely Attend-Conference, with no decomposition. SpB's domain plan, on the other hand, contains the full details of the conference registration procedure. This enables SpB to say *First, you must register for the conference.* It also enables SpB to ask *Do you have a registration form?*, because the action to ask whether SpA has a form or not (notHave) is already on the stack due to action decomposition.

Our model also explains speaker planning in Dialogue II. In this dialogue, the traveler (SpB)'s utterance (3), *Couldn't we leave in the morning ... instead of at*

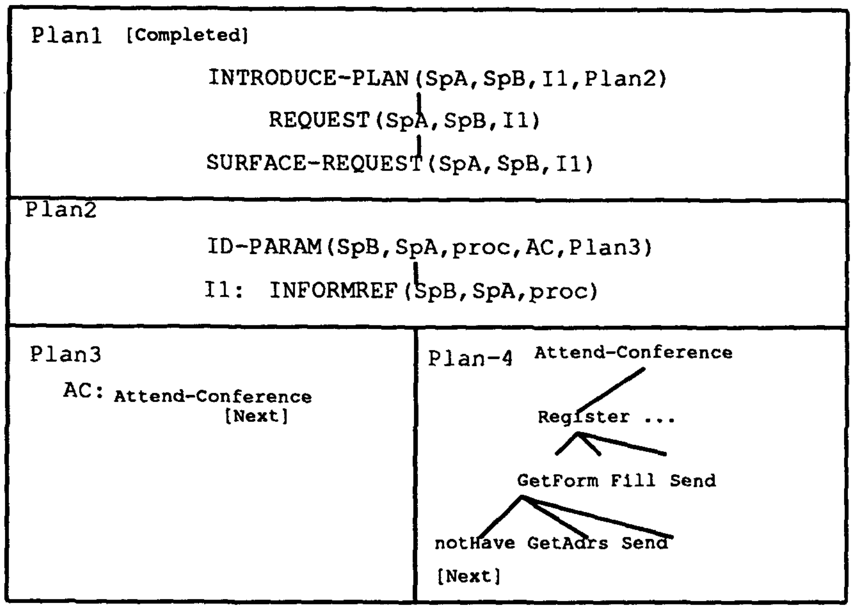

Figure 6: State of the Stack after Utterance (2) in Dialogue I

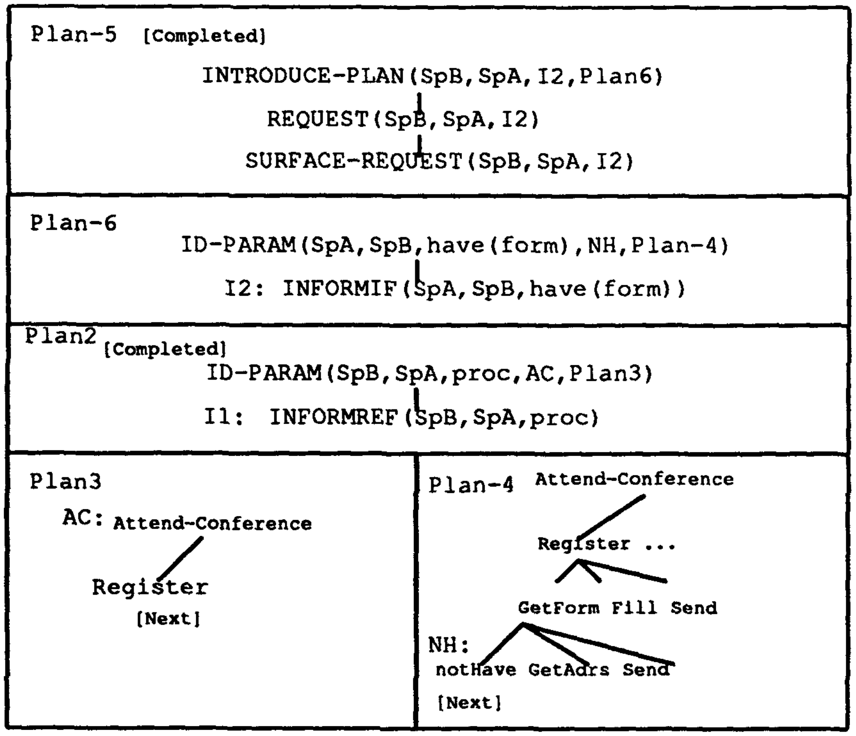
Plan-5 [Completed]
INTRODUCE-PLAN(SpB,SpA,I2,Plan6)
REQUEST(SpB,SpA,I2)
SURFACE-REQUEST(SpB,SpA,I2)
Plan-6
ID-PARAM(SpA,SpB,have(form),NH,Plan-4)
I2: INFORMIF(SpA,SpB,have(form))
Plan2 [Completed]
ID-PARAM(SpB,SpA,proc,AC,Plan3)
I1: INFORMREF(SpB,SpA,proc)
Plan3
AC: Attend-Conference
Register
[Next]
Plan-4
Attend-Conference
Register ...
GetForm Fill Send
NH:
notHave GetAdrs Send
[Next]

Figure 7: State of the Stack after Utterance (4) in Dialogue I

night?, can be explained by the plan specific to SpB which is to State-Depart-Preference. In our model, we assign plans to a specific speaker, depending upon his or her role in the dialogue, e.g., traveler or travel agent. This eliminates the potential combinatorial explosion of the number of schemata required in the current model.

4.3. Tracking Conversational Initiative Control

Third, our model provides a consistent mechanism to track who controls the conversational initiative at any given utterance in the dialogue. This mechanism provides an explanation for the initiative control rules proposed by [Walker and Whittaker, 1990], within the plan-based model of dialogue understanding. Our data allow us to state the following rule:

- When Sp-X makes an utterance that instantiates a discourse plan based on his or her domain plan, then Sp-X controls the conversational initiative.

This rule also holds in the nesting of initiatives, such as in a clarification dialogue segment:

- When Sp-X makes an utterance that instantiates a discourse plan based on his or her domain plans and Sp-Y replies with an utterance that instantiates a discourse plan, then Sp-X maintains control of the conversational initiative.

In Dialogue II, illustrated in Figure 8, SpB's question, utterance (3), *Couldn't we leave in the morning ... instead of at night?*, instantiates discourse Plan 5. It stems from SpB's domain plan State-Depart-Preference. In this case, the first conversational initiative tracking rule applies, and the initiative is transferred to SpB.

In contrast, SpB's response of *Really?* to SpA's utterance (5), *You would be traveling during the day which would be difficult because it's very hot.*, is a request for clarification. This time, the second rule cited above for nested initiatives applies, and the initiative remains with SpA.

5. Related Works

The ATR dialogue understanding system [Yamaoka and Iida, 1990] incorporates a plan hierarchy comprising three kinds of universal pragmatic and domain plans to process cooperative and goal-oriented dialogues. They simulated the processing of such dialogues using the following plans: 1) *Interaction plans* - plans characterized by dialogue turn-taking that describes a sequence of communicative acts. Turn-taking allows other embedded turn-takings. 2) *Communication plans* - plans that determine how to execute or achieve an utterance goal or dialogue goals. 3) *Dialogue plans* - plans for establishing a dialogue construction. 4) *Domain plans.* The ATR model attempts to capture complex conversational interaction by using a hierarchy of plans whereas our model tries to capture the same phenomena by speaker-specific domain plans and discourse plans. Their interaction, communication, and dialogue plans operate at a level above our *speaker-specific domain plans.* Their plans serve as a type of meta-planning to their and our domain plans. An extension enabling their plan hierarchy to operate orthogonally to our model would be possible.

Our model is consistent with the initiative control rules presented in [Walker and Whittaker, 1990]. In their control rules scheme, however, the speaker controls the initiative when the dialogue utterance type (surface structure analysis) is an assertion (unless the utterance is a response to a question), a command, or a question (unless the utterance is a response to a question or command). In our model, the conversational initiative control is explained by the speaker's planning. In our model, control is transferred from the INITIATING CONVERSATIONAL PARTICIPANT (ICP) to the OTHER CONVERSATIONAL PARTICIPANT (OCP) when the utterance by the OCP is made based on the OCP's domain plan, not as a reply to the utterance made by the ICP based on the ICP's domain plan. Cases where no initiative control transfer takes place despite the utterance type (assertion, command or question) substantiate that these utterances are (1) an assertion which is a response by the ICP through ID-PARAM to answer a question, and (2) a question to clarify the command or question uttered by the ICP, and which includes a question functioning as a clarification discourse plan. Our model provides an explanation for the initiative control rules proposed by [Walker and Whittaker, 1990] within the framework of the plan-based model of dialogue understanding. [Walker and Whittaker, 1990] only provide a descriptive explanation of this phenomenon.

Carberry [Carberry, 1990] discusses plan disparity in which the plan inferred by the user modeling program differs from the actual plan of the user. However, her work does not address mixed-initiative dialogue understanding where either of the speakers can control the conversational initaitive.

6. Conclusion

In this paper we present an enhanced model of plan-based dialogue understanding. Our analysis demonstrates that the *joint-plan* assumption employed in the [Litman and Allen, 1987] model is too restrictive to track an individual speaker's instantiated plans, ac-

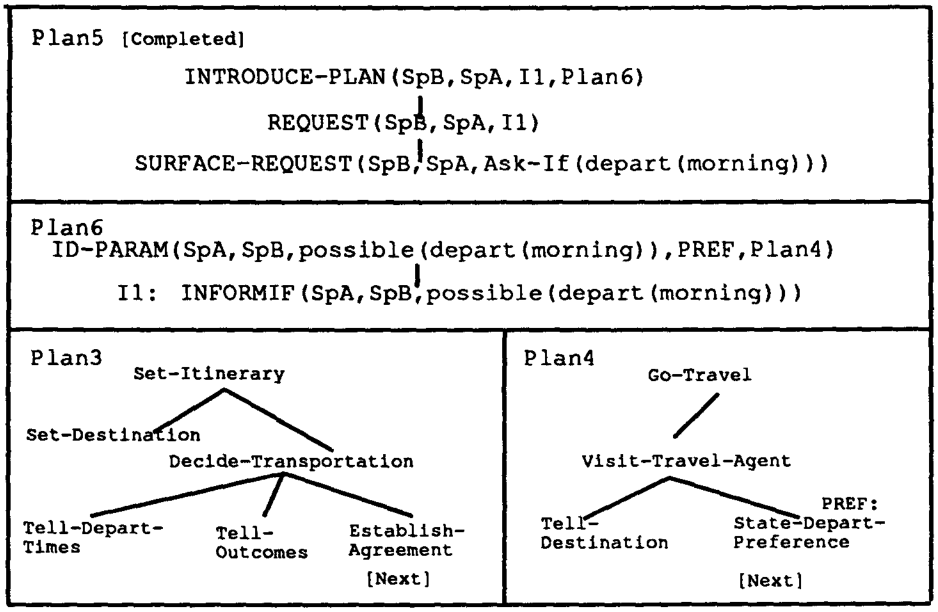
Plan5 [Completed]
INTRODUCE-PLAN(SpB,SpA,I1,Plan6)
REQUEST(SpB,SpA,I1)
SURFACE-REQUEST(SpB,SpA,Ask-If(depart(morning)))
Plan6
ID-PARAM(SpA,SpB,possible(depart(morning)),PREF,Plan4)
I1: INFORMIF(SpA,SpB,possible(depart(morning)))
Plan3
Set-Itinerary
Set-Destination
Decide-Transportation
Tell-Depart-Times
Tell-Outcomes
Establish-Agreement
[Next]
Plan4
Go-Travel
Visit-Travel-Agent
Tell-Destination
PREF: State-Depart-Preference
[Next]

Figure 8: State of the Stack after Utterance (3) in Dialogue II

count for active planning behind speaker utterances and track the transfer of conversational initiative control in dialogues, all of which characterize mixed-initiative dialogues. Our model employs *speaker-specific domain plan libraries* and the incremental expansion of domain plans to account for these mixed-initiative dialogue phenomena. We have used representative dialogues in two languages to demonstrate how our model accounts for these phenomena.

7. Acknowledgements

We would like to thank Dr. John Fought, Linguistics Department, University of Pennsylvania, for his help in collecting the Spanish travel agency domain corpus, and Mr. Hitoshi Iida and Dr. Akira Kurematsu for providing us with their Japanese ATR conference registration domain corpus. We also thank Mr. Ikuto Ishizuka, Hitachi, Japan and Dr. Michael Mauldin, Center for Machine Translation, Carnegie Mellon University for implementation support.

References

[Carberry, 1990] Carberry, S., *Plan Recognition in Natural Language Dialogue,* The MIT Press, 1990.

[Fikes and Nilsson, 1971] Fikes, R., and Nilsson, N., "STRIPS: A new apporach to the application of theorem proving to problem solving," *Artificial Intelligence,* 2, 189-208, 1971.

[Grosz and Sidner, 1990] Grosz, B. and Sidner, C., "Plans for Discourse," In Cohen, Morgan and Pollack, eds. *Intentions in Communication,* MIT Press, Cambridge, MA., 1990.

[Litman and Allen, 1987] Litman, D. and Allen, J., "A Plan Recognition Model for Subdialogues in Conversation", *Cognitive Science 11* (1987), 163-200.

[Pollack, 1990] Pollack, M., "Plans as Complex Mental Attitudes," In Cohen, Morgan and Pollack, eds. *Intentions in Communication,* MIT Press, Cambridge, MA., 1990.

[Sacerdoti, 1977] Sacerdoti, E. D., *A Structure for Plans and Behavior,* New York: American Elsevier, 1977.

[Tomita, 1985] Tomita, M., *Efficient Algorithms for Parsing Natural Language,* Kluwer Academic, 1985.

[Van Ess-Dykema and Kitano, Forthcoming] Van Ess-Dykema, C. and Kitano, H., *Toward a Computational Understanding Model for Mixed-Initiative Telephone Dialogues,* Carnegie Mellon University: Technical Report, (Forthcoming).

[Walker and Whittaker, 1990] Walker, M. and Whittaker, S., "Mixed Initiative in Dialogue: An Investigation into Discourse Segmentation," *Proceedings of ACL-90,* Pittsburgh, 1990.

[Yamaoka and Iida, 1990] Yamaoka, T. and Iida, H., "A Method to Predict the Next Utterance Using a Four-layered Plan Recognition Model," *Proceedings of the European Conference on Artificial Intelligence,* Stockholm, 1990.