Developing a Model for the Representation and Function of Discourse Structure in a Natural Language Consulting System[1]

Elizabeth Garner

1992

[1]This research has been sponsored by the Austrian *Fonds zur Förderung der wissenschaftlichen Forschung*, Grant No. P7986-PHY.

Contents

Abstract

Discourse structure is concerned with how the parts of a text relate together to convey meaning. In task-oriented dialogues, such as those which occur in the domain of VIE-$\mathcal{DU}$, a NL consulting system designed to provide information on the granting of subsidies for house improvements, the structure of discourse corresponds strongly with the goals and intentions of the agents involved in the interaction, while its function is to provide sufficient information to allow agents to determine how their partner's actions contribute to the current task. In this paper a model is presented for a dynamic representation of discourse, in which the text is seen as playing a role as mediator between an agent's goals and intentions (the agent's *intentional structure*) and what she believes is the public view of these same goals and intentions (the *conversational record*). The process of mediation is performed by two kinds of inference rules: plan strategies and collaborative strategies, whereby the former are used to respond to a CP's *domain* goals and the latter to *discourse* goals by satisfying the principle of Informativity. The model is formulated in the situation semantics framework.

1 Introduction

Most work on plan recognition has followed what Grosz & Sidner (1990) have termed the *master-slave assumption*, i.e. that one agent (the user) has goals and produces utterances and the other agent (the system) attempts to infer the user's goals from her utterances and react to them. This assumption has also affected the discourse behaviour of the system so that utterances are generated simply in response to goals of the system to, for example, obtain information; goals which are themselves subordinate to the recognised goals of the user.

Wilensky & al. (1988) make use of the term *reasonable agent* to describe the desired behavior of the system in a consulting environment. By this they mean that as far as possible the system should act as a *participant* in the dialogue, with explicit goals of its own, rather than reacting blindly to the goals of the user. Attention is restricted, however, to certainx *domain* goals of the system, which may be in conflict with the (domain) goals of the user. Discourse once again plays a subordinate role to these domain goals.

According to Grosz & Sidner (1990), in order to account for extended sequences of utterances, it is necessary to realise that two agents may develop a plan *together* rather than merely execute the existing plan of one of them. While one of the discourse participants (the user) may initiate the plan, subsequent exchanges entail both system and user working together to satisfy it.

Collaboration between user and system occurs at the level of discourse. Indeed Grosz and Sidner claim that the role of discourse is precisely to provide sufficient information about the beliefs and intentions of the conversational participants (CPs) to allow their interlocutors to be able to determine how their actions contribute to their shared plan. Failure to provide such information leads to situations in which CPs may be unsure of one another's cooperativity.

Satisfying the demands of such collaborative discourse behaviour seems to require something akin to Grice (1975)'s principle of quantity, i.e. that a contribution to a dialogue should be as informative as is required for the current purposes of the exchange. We will refer to this as the Informativity principle. Such a view of discourse is clearly a step onwards from regarding it merely as a tool to effect certain domain goals.

In this paper I present a framework for ensuring that system utterances provide precisely such collaborative discourse behaviour. I achieve this by means of an approach to discourse structure which combines aspects of Rhetorical Structure Theory and the Intentional Structure suggested by Grosz & Sidner (1986) to provide a set of strategies aimed at avoiding potential discrepancies between the *conversational record* (a representation of those mutually held beliefs that play a role in the shared plan) and the system's own goals and intentions (which arise as a result of the plan). The approach is presented within a sit-

User: Wie muß ich vorgehen, um eine Förderung für eine Zentralheizung zu bekommen? (1)
How do I go about getting a subsudy for central heating?

System: Dazu müssen Sie einige Voraussetzungen erfüllen. (2)
To do that you have to fulfil a number of conditions.
Ich nehme an, Sie wohnen in einer Mietwohnung in Wien. Ist das richtig? (3)
I assume you live in a rented appartment in Vienna. Is that correct?

User: Ja. (4)
Yes.

System: Jetzt müssen wir feststellen, ob Ihr Haushaltseinkommen unter der zulässigen Höchstgrenze liegt. (5)
We now need to find out if your income lies under the permitted limit.
Wie hoch ist Ihr Haushaltseinkommen? (6)
What is your household income?

User: Ich verdiene 20.000 S im Monat. (7)
I earn 20,000 Schillings a month.

System: Wer lebt sonst noch bei Ihnen im Haushalt? (8)
Who else lives in your appartment?

User: Meine Frau, meine zwei Kinder und die Oma. (9)
My wife, my two children and their grandmother.

System: Hat jemand von ihnen ein Einkommen? (10)
Do any of them have a personal income?

User: Meine Frau und die Kinder nicht, die Oma hat ihre Pension. (11)
My wife and the children don't. Grandmother has a pension.

System: Wie hoch ist die? (12)
How much is that?

User: So ca. 8000,– S. (13)
About 8,000 Schillings.

System: Das heißt also, daß Ihr gemeinsames Haushaltseinkommen im Jahr 28.000 mal 14, das sind 392.000 S beträgt. (14)
That means, your total household income is 28,000 x 14, that's 392,000 Schillings.
Damit liegen Sie unter der zulässigen Höchstgrenze von 546.000 S für 5 Personen im Haushalt. (15)
In that case you are under the permitted limit of 546,000 Schillings for households with five members.

System: Liegt die Baubewilligung fuer die Wohnung weiter als 20 Jahre zurück? (16)
Was planning permission granted more than 20 years ago?

User: So genau weiß ich das nicht, das Haus ist aber schon ziemlich alt. (17)
I don't know precisely, but the house is quite old.

System: Na gut, nehmen wir das einmal an. (18)
OK, let's assume that it was.
Hat die Wohnung weniger als 150m^2? (19)
Is the appartment less than 150m^2?

User: Zählt da die Terrasse dazu? (20)
Is the balcony included in that?

System: Nein, außer sie ist überdacht und geheizt und wird für Wohnzwecke verwendet. (21)
No, unless it is covered, heated and used for living purposes.

User: Ja, also dann sind es sicher weniger. (22)
Yes, then it must be less.

System: Dann sind die Voraussetzungen alle erfüllt. (23)
You fulfil all the conditions.
Um die Förderung zu bekommen, gehen Sie zu Ihrer Bank und besorgen Sie sich ein Antragsformular X-17. (24)
To get the subsidy, go to your bank and ask for the X-17 application form.
Stellen Sie den Antrag bei der Magistratsabteilung 50. (25)
Hand in your application at the (name of office).

Figure 1: Sample Dialogue

Ds1 **Wie muß ich vorgehen, um eine Förderung für eine Zentralheizung zu bekommen?**

Ds2 Dazu müssen Sie einige Voraussetzungen erfüllen.

Ds3 Ich nehme an, Sie wohnen in einer Mietwohnung in Wien.
Ist das richtig?
Ja.

Ds4 Jetzt müssen wir feststellen, ob Ihr Haushaltseinkommen unter der zulässigen Höchstgrenze liegt.

Ds5 Wie hoch ist Ihr Haushaltsnettoeinkommen?
Ich verdiene 20.000 S im Monat.

Ds6

Ds7 Wer lebt sonst noch bei Ihnen im Haushalt?
Meine Frau, meine zwei Kinder und die Oma.

Ds8 Hat jemand von ihnen ein Einkommen?
Meine Frau und die Kinder nicht die Oma hat ihre Pension.

Ds9 Wie hoch ist die?
So ca. 8000,– S.

Das heißt also, daß Ihr gemeinsames Haushaltseinkommen im Jahr 28000 mal 14, das sind 392.000 S beträgt.

Damit liegen Sie unter der zulässigen Höchstgrenze von 546.000 für 5 Personen im Haushalt.

Ds10 Liegt die Baubewilligung für die Wohnung weiter als 20 Jahre zurück?
So genau weiß ich das nicht, das Haus ist aber schon ziemlich alt.
Na gut, nehmen wir das einmal an.

Ds11 Hat die Wohnung weniger als 150 m^2?

Ds12 **Zählt da die Terrasse dazu?**
Nein, außer sie ist überdacht und geheizt und wird für Wohnzwecke verwendet.

Ja, also dann sind es sicher weniger.

Dann sind die Voraussetzungen alle erfüllt.

Um die Förderung zu bekommen, gehen Sie zu Ihrer Bank und besorgen Sie sich ein Antragsformular X-17. Stellen Sie den Antrag bei der Magistratsabteilung 50.

Figure 2: Intentional Structure of Sample Dialogue

DS1	Know action required to obtain a subsidy for central heating installation
DS2	Know whether user fulfills the conditions for obtaining a subsidy
DS3	Know whether user lives in a rented appartment in Vienna
DS4	Know if user's household income falls under the limit for households of that size
DS5	Know the user's household income
DS6	Know if the user' household income has any other income besides the user's personal income
DS7	Know who lives in the user's household
DS8	Know if any of these household members have an income
DS9	Know the income of the user's grandmother
DS10	Know if planning permission for the user's appartment was granted more than 20 years ago
DS11	Know if the size of the user's appartment is less than 150m^2
DS12	Know if the balcony area is included in the area of the appartment

Figure 3: Primary Intentions of Sample Dialogue

uations semantics framework and is currently being implemented in the dialogue system VIE-$\mathcal{DU}$, a consulting system for providing information about subsidies for house improvements.

I begin with a look at the function of discourse structure in a sample dialogue taken from our domain and given in Figure 1. I then briefly present the two main approaches to discourse structure, Rhetorical Structure Theory (Mann & Thompson (1987), Mann & Thompson (1988)) and the Attentional, Intentional account of Grosz & Sidner (1986), pointing out which aspects of these systems are of use in VIE-$\mathcal{DU}$. An alternative approach is then presented, based on the concepts of Conversational Record, SharedPlan(*) and Intentional Structure which are updated by means of plan and collaborative strategies. Finally, I show how this approach can be used to analyse a part of the sample dialogue.

2 The Function of Discourse Structure

An important task for a natural language consulting system such as that we are building in VIE-$\mathcal{DU}$ is to determine a user's goals and beliefs from her utterances. In Garner (1991) an approach was presented in which a version of speech act theory is used to deduce a speaker's goals and beliefs from the type of her utterances. The approach is formulated in the traditional plan framework of preconditions, actions and effects, in which the action of making an utterance is seen as occuring in a certain environment (the preconditions) and leading to changes in that environment (the effects). So, for example, the precondition of a declarative sentence is that the speaker believes the content of the proposition expressed (assuming sincerity), while interrogatives represent goals of the speaker to be informed by the addressee. The effects of both sentence types is to make the addressee aware of the speaker's beliefs and/or goals, i.e. to raise their status to mutual beliefs. The approach is also able to deal with performative sentences.

It was suggested in Garner (1991) that dialogue control could be handled by a CP responding directly to a (mutually) believed goal of another CP by carrying out the goal. Alone, however, this approach is not capable of dealing with the complex interaction which occurs in a consulting dialogue within our domain. To illustrate this, we will examine a number of problems with such a simplistic account that arise in the sample dialogue given in Figure 1.

2.1 Determining Intentions from Utterances

The first point to be made is that while a user's goals may be explicitly available from the preconditions of her utterances this is not always the case. The user's query in line 1 of the dialogue contains the precondition that the user has a goal to be informed of an action which will enables him to obtain a subsidy for central heating installation. However, the response of the system clearly indicates that the system also believes the user also has a goal to obtain such a subsidy, otherwise the problem of ascertaining whether or not the user is eligible for such a subsidy (lines 3 to 23) would not arise. (The user's cooperation in this section of the dialogue sequence also represents a tacit acknowledgement of this goal).

2.2 Reacting to a User's Intention

Related to the first point is the fact that the system is required to do more than simply adopt the *stated* goals of the user. If the user does indeed have a goal to obtain a subsidy, and the system believes that obtaining a subsidy is dependent on the user's fulfilling a number of conditions, it would clearly be misleading to the user if the system merely informed him of the action he needed to take without first discovering if the user does

in fact meet the conditions. It would in fact give rise to an assumption by the user that carrying out the action will unproblematically bring about his goal. We need some means of determining what strategy the system, as a reasonable agent, should adopt in response to a user's goal in any given situation.

2.3 Relations between Utterances

Thirdly, not all the information contained in the dialogue is revealed by the preconditions and effects of the individual utterances of which it consists. Sequences of utterances are frequently related to one another by certain types of relations, and even if such relations are not made explicit they are in many cases intended to be recognised by the addressee. For example, in line 2 the user is made aware of the system's belief that obtaining a subsidy is dependent on a number of conditions being met. The subsequent query by the system is designed to elicit whether or not the user meets the first condition, i.e. whether he lives in a rented appartment in Vienna. The user is clearly expected to draw the inference that he is being asked about one of the conditions, although it is never explicitly stated that the query has anything to do with the said conditions.[1] We need some explanation of how the user is able to recognise such unstated relations between utterances.

2.4 The Use of Metadialogue to Convey Discourse Information

Finally, the system utterances contained in the dialogue can be classified as performing two different functions. On the one hand they serve to provide or secure information in order to satisfy the current *domain* goals of the system. This is true, for example, of the queries in lines 3, 6, 16 and 19, designed to discover the user's eligibility. Secondly, the system's utterances serve to provide information aimed to keep the user informed as to how the system is proceeding in the satisfaction of domain goals. So, for example, in line 5 the user is informed of the system's current goal to find out if the user's household income falls under the permitted limit. While in lines 14, 15 and 23 the user is offered information as to the results of the system's goals, which serves to inform the user that the system's goal has now been satisfied. These two types of information are not relevant to the satisfaction of domain goals but seem to satisfy certain *discourse* goals of the system.

[1] Note that it would be perfectly possible to make this fact explicit by means of an utterance such as: "The first condition requires that you be resident in rented accommodation in Vienna".

2.5 Conclusion

These four problems point to a need for some kind of analysis above the level of the utterance. Discourse structure, which is concerned with the segments into which texts may be divided and the relations which exist between these segments, seems to offer a means of attempting to deal with these problems. In the next section I will take a look at two of the best-known approaches to discourse structure and discuss to what extent they are able to deal with these problems.

3 Existing Approaches to Discourse Structure

There have been two leading approaches to the study of discourse structure: Attentional and Intentional structure (Grosz & Sidner (1986)) and Rhetorical Structure Theory (RST) (Mann & Thompson (1987),Mann & Thompson (1988)).

3.1 The Grosz-Sidner Approach

The Grosz-Sidner approach sees discourse structure as consisting of three interacting components: linguistic structure, intentional structure and attentional structure. The *linguistic structure* is the structure of the sequence of utterances that comprise the text. These utterances can be grouped into *discourse segments*, connected to each other by an embedding relation. *Intentional structure* is concerned with the *purpose* of each individual discourse segment, related to each other by the two structural relations dominance and satisfaction-precedence. In the former case this means that a discourse segment purpose (DSP) A contributes to a DSP B, and in the latter case that the DSP A precedes the DSP B. *Attentional structure* concerns itself with the *focus of attention* of the discourse, i.e. which discourse entities are salient in a discourse at any particular time. Of these three structures, it is the intentional structure which is seen as central, the other two structures deriving parasitically from it.

Figure 2 presents an attempt to provide an intentional structure for the sample dialogue given in Figure 1. Here the importance of the purpose of a segment as a means of defining its boundaries is apparent. However, let us consider how the approach is able to deal with the problems mentioned in the previous section.

Firstly, the DSP is not necessarily, according to Grosz and Sidner, reflected in the text itself, but can only be reconstructed by looking at the meaning carried by the text. However, it is unclear from this account exactly how this is achieved, no method of reconstructing DSPs from the text is offered. For this reason the problems mentioned in 2.1 and 2.2 cannot even be tackled.

Secondly, the model offers no analysis of how the system reacts to user utterances, i.e. what types of responses are appropriate in what situations. Yet clearly such responses are restricted in a coherent text, a fact that we would do well to capture.

Many of the segments contained in the dialogue can be seen as related by the relations of dominance (e.g. DS4 and DS5) and the discourse segments DS4, DS5, DS10 and DS11 are related in terms of satisfaction precedence. However, there are other relations between segments which cannot so easily be described by one or the other of these relations, e.g. the relation between the user's goal to know an action which will enable him to obtain a subsidy and the system's goal to find out if the user is eligible for such a subsidy (DS1

and DS2). The problem seems to be that individual CPs may at any point in time seem to be working towards different goal; only in the text as a whole does their collaboration towards a single task become apparent.

The use of dialogue to indicate the system's current goal and its satisfaction, as described in 2.4, fits neatly with the demarcation of segment boundaries given in the intentional structure. However, there is nothing in the description which accounts for the role performed by such utterances, or which explains when they are needed. (Not all discourse segment boundaries are demarcated by such utterances.)

3.2 Rhetorical Structure Theory

In RST a text is broken down into clauses, the smallest unit of RST, and then these units are linked together by means of the rhetorical relations in order to form schemas, units of text structure spanning a particular portion of text. A selection of the rhetorical relations suggested by RST are Enablement, Motivation, Concession, Circumstance, Elaboration, Solutionhood, Antithesis, Purpose and Condition. A schema decomposes into a "Nucleus-Satellite" relation. The nucleus is the central assumption of the schema, the satellite ancillary information connected to the nucleus by the specified relation. RST works then under the assumption that within a text certain portions are more central to the text as a whole, while others provide supplementary, elaborating material, it thus also provides something like a dominance relation à la Grosz-Sidner but defines the relation in greater detail.

To each relation is attached a definition, a component of which is the *effect* of the satellite, e.g. in the Condition relation the effect is that the reader of the text recognises how the realisation of the situation presented in the nucleus depends on the realisation of the situation presented in the satellite. Hence RST contains a method of explaining the plausible purpose of the writer in including a particular relation in the text.

Looking now at how RST is able to deal with the problems mentioned in Section 2, we see that firstly, according to Mann & Thompson (1988), in judging the functions of a particular portion of text the analyst must often go beyond the literal readings of the text. Once again then we have a statement of the problem referred to in 2.1 but no suggestion as to its solution — in RST the burden of identifying such functions is placed on the interpretive abilities of the analyst/reader. We are offered no model as to how clause-level purposes may be derived and as with the Grosz-Sidner approach no analysis of how the system reacts to user utterances.

RST has thus concerned itself more with the *variety* of relations which may hold between utterances in a piece of discourse and thus might seem to offer a solution to the problem mentioned in 2.3. However, the RST relation descriptions are rather complex, designed as

they are to deal with many different types of, primarily, written text. It is questionable to what degree such complexity is required to deal with task-specific problem-solving dialogues as occur in VIE-*DU*. Grosz & Sidner (1986) have argued, furthermore, that while such descriptions may provide a meta-level description of discourse, their precise role in discourse interpretation is unclear. While a CP may well have such relations in mind when she produces an utterance it is less clear whether her partner in the action needs to infer them. Grosz and Sidner also claim that such an approach fails to capture generalisations which exist between certain types of relations, generalisations which may be important for, for example, reference resolution.

Finally, while the meta-dialogue information provided in our sample dialogue receives a description in RST, (e.g. the statement of goals acts as satellite in a solutionhood relation) we once again have no explanation as to when such information needs to be provided by the system.

3.3 Conclusion

While aspects of both theories seem to have something to offer, neither provides exactly what is required in response to the problems mentioned in Section 2. Moreover, both approaches fall short on the issue of *collaboration*. They provide a method of analysing a static text from the point of view of a single reader/analyst. However, texts such as that in Figure 1 are not static objects but evolve in the course of the interaction. Individual utterances contain interpretations of a single CP towards what she perceives of as the current situation, an interpretation which may need to be explained to her partner in the context as a whole, due to its seemingly being at odds with the current task. No attempt is made to show how discourse functions to counter such a problem. Moreover, at any single point in the dialogue CPs may analyse a text differently, depending on what each of them sees as the current task.

In the account which follows I will attempt to remedy this situation by separating the notion of intentional structure from the text itself, and regarding the text as a record of the information which the CPs mutually believe on the basis of the exchange, a record which enables them to remain convinced of each other's cooperativity towards the goals being pursued.

4 The Conversational Record, SharedPlans and the Intentional Structure

The core to the approach presented here is the *conversational record*, a term taken from Thomason (1990). According to Thomason, the conversational record (CR) is an evolving representation of the state of a conversation. It contains information which is public, i.e. that can be supposed to be available to all CPs. The CR is, however, more than a linear record of the conversation itself. In addition to the information carried by utterances it also contains such information as all CPs believe to be *mutually inferable* from these utterances. Such information might include, Thomason suggests, the intentional and attentional structures proposed by Grosz & Sidner (1986), since discourse purposes, focus management and the appropriate use of referring expressions[2] are things which need to be public if a conversation is to function smoothly. We might represent the CR then as containing different categories of information corresponding to these structures.

Providing an actual processing model for the construction of the intentional structure is, according to Grosz & Sidner (1990) dependent on underlying theories of intention, actions and plans which deal adequately with collaboration. Their contribution to such a theory is the notion of *SharedPlan*, a schema based upon the plan description of Pollack (1986) in which plans are looked at not in terms of *individual* beliefs and intentions but in terms of *mutual* beliefs of the CPs. A SharedPlan is a construct consisting of a bundle of mutual beliefs about the intentions of each CP and how these intentions relate to the SharedPlan.

SharedPlans are constructed from a combination of those beliefs and intentions explicitly mentioned by the CPs, and on what Grosz & Sidner (1990) term *prior* mutual beliefs, i.e. beliefs needed for the construction of the SharedPlan but not explicitly mentioned. Such prior mutual beliefs must be inferable on the basis of what has been made explicit in the discourse.

The full collection of the beliefs and intentions which make up the SharedPlan only becomes available upon its completion; until that point what is mutually believed by the CPs is a *SharedPlan**, the collection of those mutual beliefs available up to the current state of the discourse. The SharedPlan* may thus be compared to the evolving intentional structure of the CR.

Since not all the intentions in the CR are necessarily explicit, it may seem that mutual belief is an untenable assumption for the intentions of the SharedPlans(*). CPs may be mistaken about exactly which inferences are mutually believed (in the same way as the meaning intended by an utterance itself may be misunderstood by a CP). Since CPs must, however, have good reason for including inferable information in the CR, we will assume that CPs treat such information as mutually believed unless evidence to the contrary arises.

[2] Clark & Wilkes-Gibbs (1990) provide evidence for the colloborative nature of referring expressions.

This seems to be a natural assumption in discourse.

The role of discourse in a collaborative framework is, as we have stated, to provide sufficient information to allow CPs to be able to determine how their various actions contribute to the shared plan. This means making sure that all intentions in an individual CP's model of the SharedPlan* are made explicit, either directly in an utterance or by means of an inference which an interlocutor can be expected to draw. How can this be achieved?

Besides the mutually believed information carried by the CR a CP in a consulting system also has available an intentional structure (IS) of her own, which contains a record of privately known goals and intentions, and the relations which hold between them, as formulated in response to the goals adopted in a dialogue exchange. Collaboration towards the SharedPlan requires that precisely this information be known to all CPs. In some cases the CP's plan itself will entail that such information be made explicit. This is so, for instance, with intentions to inform a CP which allow the CP to infer the goal behind the inform action. In other cases, however, the goals of the CP are not recoverable from her utterances.

Not all information, however, that is not made explicit by plan goals needs to be so. Much information is inferable on the basis of standard relations between utterances. An intention to be informed, for example, allows a partner to reconstruct a goal to know. In many cases the relations which exist between clauses (e.g. cause and effect) can be identified by a CP without any need for specific cues. In other cases a CP's competence may allow her to reconstruct quite complex relations between the mutually known goals and intentions already present in the SharedPlan*. Here the SharedPlan* provides a crucial element of the background against which utterances are interpreted. Only in cases where a CP believes her partner will be unable to reconstruct the intentions behind her utterances will such relations need to be made explicit. What is required, then, is a set of constraints describing situations in which this is so, as well as strategies that can then be applied to successfully update the CR in such a way as to ensure that the principle of informativity will be satisfied.

5 Basic Framework

We require then two mechanisms, the Shared Plan* (a component of the CR) and the system's personal intentional structure, IS, which are updated during the course of the interaction. Changes in these structures are effected by means of plan strategies and collaborative strategies which apply according to the context of the discourse. Here context refers to the current SharedPlan* and the IS as well as the beliefs of the system (in particular with regard to the user's competence).

To model such a framework we have chosen a semantic theory that allows for the integration of context - Situation Semantics (cf. Barwise & Perry (1983), Barwise & Cooper (1991)). In situation semantics meaning is considered relational, as constraints linking utterance situations and described situations. Conditional constraints are used to model the relation holding between situations relative to a given context (Barwise (1989)). So, for example, the situation S *involves* (signalled by $\Rightarrow$ the situation S' relative to a context C:

(1) $$S \Rightarrow S' \mid C$$

We make use of such conditional constraints to represent the strategies for updating the IS and the SharedPlan*.

In the following sections we offer an overview of the basic structure of the model, followed by detailed examples which show how the collaborative strategies are used to increase the informativity of the CR.

5.1 Representation of Beliefs

At the outset of the interaction the system has a set of beliefs. Most basic of these are beliefs about general principles relating information, e.g. that an action achieves a goal state. Such general principles can be represented by means of *types* in situation semantics, e.g:

(2) $$[\mathrm{A,G}|\ (\mathrm{s}\ |\!=\!\langle\!\langle \mathrm{achieve,A,G;1} \rangle\!\rangle\ )]^{3}$$

i.e. A and G are parameters which may be anchored to actions and goals in the situation s, in which the condition that A achieves G obtains. Domain and discourse-specific beliefs are then formulated by means of application of such types. For example, a piece of discourse knowledge available to both system and user (i.e. mutually believed) is that an informing action achieves a goal to know something:

(3) $$[\mathrm{A}{:}\langle\!\langle \mathrm{inform,X,Y,P} \rangle\!\rangle,\mathrm{G}{:}\langle\!\langle \mathrm{know,Y,P} \rangle\!\rangle|\ (\mathrm{s}\ |\!=\!\langle\!\langle \mathrm{achieve,A,G;1} \rangle\!\rangle\ )]$$

[3]The situation semantic formulas were produced using macros developed by J. Barwise and R. Cooper

which denotes the same object as:

(4) $(s \models \langle\langle achieve, \langle\langle inform, X, Y, P \rangle\rangle, \langle\langle know, Y, P \rangle\rangle; 1 \rangle\rangle)$[4]

Relations other than the *involves* relation may be sensitive to a given context. We can show this by means of a restriction. So, for example, the system believes at the outset of the dialogue that:

(5) $(s \models \langle\langle enable, \left(\begin{array}{l} \langle\langle go, user, bank \rangle\rangle \wedge \\ \langle\langle obtain, user, X\text{-}17 \rangle\rangle \wedge \\ \langle\langle hand\text{-}in, user, X\text{-}17, MA\text{-}50 \rangle\rangle \end{array} \right), \langle\langle get, user, subsidy \rangle\rangle \rangle\rangle \mid Conds)$

where Conds refers to the set of eligibility conditions:

(6)
1. living in a rented appartment in Vienna
2. having a household income under the given limit
3. the size of the user's appartment being less than $150m^2$
4. planning permissing having been granted more than 20 years ago

This is a piece of domain knowledge about which the user is not assumed to have competence. The system also has beliefs about such competence. At the outset of the interaction, for example, the system will not expect the user to have competence regarding domain-specific knowledge, but will expect the user to have competence relating to the user herself. Such information can be represented as restricted types, e.g:

(7) $[TU: \langle\langle concern, TU, household; 1 \rangle\rangle \mid (r \models \langle\langle competent, U, TU; 1 \rangle\rangle)]$

(8) $[TD: \langle\langle concern, TD, domain; 1 \rangle\rangle \mid (r \models \langle\langle competent, U, TD; 0 \rangle\rangle)]$

i.e. in (7), if TU is a parameter for a topic object concerning the user's household then in the discourse situation (DS) the user is competent regarding TU. In addition, the system expects the general information principles and the discourse knowledge to be mutually believed by both user and system.

During the course of the interaction system beliefs are updated on the basis of information contained in the IS and SharedPlan*. This process will be explained in the next section.

[4]For simplicity, I will in future omit the polarity value of the infon unless it is negative or for some reason not clear.

5.2 Plan and Collaborative Strategies

Goals and intentions[5] lead to strategies to update the IS in the course of the interaction. We distinguish between two kinds of strategies, plan and collaborative.

5.2.1 Updating the IS — Plan Strategies

Plan strategies apply in response to the goals and intentions present in the SharedPlan* and the IS, i.e. in response to *domain* goals. In the former case such strategies are reactive, i.e. they are triggered by believed goals and intentions of the user, as derived from the user's utterances. In the latter case the strategies are active, i.e. they are triggered by goals and intentions of the system (themselves subordinate to the goals contained in the Shared Plan*). The strategies differ then according to under whose initiative they are launched. An example of a reactive plan strategy is:

(9) $$(CR\models\langle\langle goal,X,(s\models\langle\langle inform,Y,X,P\rangle\rangle)\rangle\rangle \mid (r\models\langle\langle know,Y,P\rangle\rangle)) \Rightarrow (IS\models\langle\langle intend,Y,(s\models\langle\langle inform,Y,X,P\rangle\rangle)\rangle\rangle)$$

i.e. when an agent X has a (mutually believed) intention to be informed of P, in a context where an agent Y knows P, Y adopts an intention (in the IS) to carry out the inform. 'CR $\models$' is used to represent the fact that a *mutual belief* holds in the current resource situation, r (where r contains the beliefs and goals of the system). 'IS $\models$' represents the fact that the situation described currently holds in the system's IS.

A second reactive discourse strategy is used to deal with restricted contexts:

(10) $$(CR\models\langle\langle goal,X,(s\models\langle\langle inform,Y,X,P\rangle\rangle)\rangle\rangle \mid (r\models\langle\langle believe,Y,P\rangle\rangle \mid C)\ )\ ^{6} \Rightarrow (IS\models\langle\langle intend,Y,(s\models\langle\langle inform,Y,X,P\rangle\rangle)\rangle\rangle \mid C)$$

i.e. the CP adopts the intention to inform only if the restriction holds. Again, this can be assumed to be a mutually believed strategy.

Some examples of active plan strategies are:

(11) $$(IS\models\langle\langle goal,X,G\rangle\rangle \mid (r\models\langle\langle believe,X,(s\models\langle\langle achieve,A,G\rangle\rangle)\rangle\rangle)) \Rightarrow (IS\models\langle\langle intend,X,A\rangle\rangle)$$

i.e. if a CP has a goal to achieve a state G and believes that an action A will bring about this state the CP adopts an intention to perform A.

[5] We make a distinction here between goals which aim to achieve a certain *state*, and intentions to carry out an *action*.

(12) $$(IS \models \langle\langle goal, X, G(Pa) \rangle\rangle \mid (r \models \langle\langle know, X, Pa; 0 \rangle\rangle)) \Rightarrow (IS \models \langle\langle goal, X, (s \models \langle\langle know, X, Pa \rangle\rangle) \rangle\rangle)$$

i.e. if a CP has an intention to perform an action A which contains a parameter Pa and currently does not know the value of Pa, she adopts a goal to know the value of the parameter.

(13) $$(IS \models \langle\langle intend, X, A \rangle\rangle \mid C \wedge (r \models \langle\langle know, X, C; 0 \rangle\rangle)) \Rightarrow (IS \models \langle\langle goal, X, (s \models \langle\langle know, X, C \rangle\rangle) \rangle\rangle)$$

i.e. if a CP has an intention to perform an action conditional upon a restriction C, and currently does not know if C holds, the CP adopts a goal to know C.

(14) $$(IS \models \langle\langle goal, X, (s \models \langle\langle know, X, P \rangle\rangle) \rangle\rangle \mid (r \models \langle\langle believe, X, (r \models \langle\langle know, Y, P; 0 \rangle\rangle); 0 \rangle\rangle)) \Rightarrow (IS \models \langle\langle goal, X, (s_1 \models \langle\langle inform, Y, X, P \rangle\rangle) \rangle\rangle)$$

i.e. if a CP X has a goal to know P in a context where she believes another CP Y knows P, she adopts a goal that Y inform her of P.

These strategies are also available to both system and user.

5.2.2 Updating the CR — Collaborative Strategies

Intentions to inform the user lead to utterances, and a subsequent update of the SharedPlan*. The amount of information carried in the update depends on the prior mutual beliefs recoverable from the context. The SharedPlan* should contain the same information as that section of the IS which represents already launched goals.

Certain information is available from the preconditions of the utterance, as already described in section 2.

Additionally, certain other information is available on the basis of the domain of the utterance. A request to be informed of an action which enables the user to gain a subsidy, for example, involves the user not only having a goal to be informed (made explicit by the form of the utterance) but also allows the system to deduce a goal to obtain a subsidy. We may express this in the following manner:

(15) $(CR \models \langle\langle goal, X, (s \models \langle\langle inform, Y, X, [A \mid (s \models \langle\langle achieve, A, G \rangle\rangle)] \rangle\rangle) \rangle\rangle) \Rightarrow (CR \models \langle\langle goal, X, G \rangle\rangle)$

i.e. if a CP X has a goal that another CP Y inform her of an action which will bring about a particular goal state, G, then X also has the goal to achieve G. By means of such inferences the problem mention in section 2.2. can be countered.

Utterances, however, are usually subordinate to goals and intentions present in the IS, but not necessarily overtly reflected in the SharedPlan*. In the plan strategy given in (14), for example, the goal to be informed related back to a goal to know, a goal present in the IS. (14) represents, however, mutually believed information. We may therefore assume that whenever a goal to be informed is made public by an interrogative the goal to know to which it is subordinate in the IS is also mutually believed. This us true whenever mutually believed strategies are employed.

In some situations the system utterances, as directed by the IS, may not enable the user to reconstruct the underlying goals and intentions present in the IS. This would lead to a breakdown in informativity. To avoid this the system makes use of collaborative strategies. I examine here two types of situations where a collaborative strategy is necessary for dealing with the example dialogue given in Figure 1. In other environments further strategies, not here discussed, may be required.

The first strategy is applied in situations where a planned next utterance of the system is not in accord with a user's expectation, as contained in the SharedPlan*. The plan strategy given in (9), for example, gives rise to an expectation that the system will respond to the user's query by immediately providing an answer to the query. However, there may be environments where the system does not react in this way, for some reason. This could occur, for example, if the system believes the correctness of an answer is dependent on some context, as is the case in the plan strategy given in (10), where the system first attempts to find out if the context holds. If this attempt involves launching a sub-dialogue with the user, it may lead to utterances whose relevance is not immediately available to the user on the basis of the user's expectation.[7] To ensure that the user is kept informed of what is happening a collaborative strategy is applied which explains the relation of the planned subdialogue to the initial user query:

(16) $(IS \models \langle\langle goal, sys, (s \models \langle\langle inform, user, sys, C \rangle\rangle) \rangle\rangle \mid (CR \models \langle\langle goal, user, (s_1 \models \langle\langle inform, sys, user, Q \rangle\rangle) \rangle\rangle) (r \models \langle\langle believe, sys, Q \rangle\rangle \mid C))$
$\Rightarrow (IS \models \langle\langle intend, sys, (s_2 \models \langle\langle inform, sys, user, Q \rangle\rangle \mid C) \rangle\rangle)$

The utterance to which this strategy leads causes an update in the CR to the effect that

[7] Another way of looking at this, is that in a certain situation there exist two possible mutually believed strategies, one of which applies by default. The user must be informed if, however, the non-default strategy applies.

the queried information is dependent on a context. The user will no longer expect an immediate response to her query, but will realise that strategy (10) is now operative. Subsequent utterances will be interpreted in this light.

The second and third types of collaborative strategy I will look at both occur in a similar environment. The second serves to demarcate the opening of a subgoal, and the third to denote the satisfaction of a subgoal.

The second is used in an environment where the system launches a subgoal that would not, on the basis of plan strategies to inform, be reflected in the SharedPlan*, about a subject regarding which the user is not expected to have competence. The subgoal itself leads via plan strategies to further subgoals, and subsequent utterances by the system. However, the system does not believe that the information thus made available in the CR will enable the user to reconstruct the initial subgoal. The collaborative strategy is designed to establish this link:

(17) $(\mathrm{IS}\models\langle\langle \mathrm{goal,sys,}(s\models\langle\langle \mathrm{inform,user,sys,G2}\rangle\rangle)\rangle\rangle \mid (\mathrm{CR}\models\langle\langle \mathrm{goal,sys,G0}\rangle\rangle) \wedge \exists\ \mathrm{G1}\ \langle\langle \mathrm{subgoal,G1,G0}\rangle\rangle \wedge \langle\langle \mathrm{subgoal,}\langle\langle \mathrm{know,sys,G2}\rangle\rangle\mathrm{,G1}\rangle\rangle \wedge \langle\langle \mathrm{concern,G1,domain}\rangle\rangle)$

$\Rightarrow (\mathrm{IS}\models\langle\langle \mathrm{intend,sys,}(s_1\models\langle\langle \mathrm{inform,sys,user,}\ (\mathrm{IS}\models\langle\langle \mathrm{goal,sys,G1}\rangle\rangle)\rangle\rangle)\rangle\rangle)$

The user is able to infer, on the basis of the plan strategy given in (9) the goal of the system to know G2. The user is also aware from the SharedPlan* of the superordinate goal G0. There is however no mutually believed plan strategy which links G0 to G1 — it rests on a belief which lies outside the user's competence. The collaborative strategy supplies the missing subgoal G1, allowing the user to recreate the relation between the four goal present in the system's IS.

The third collaborative strategy is used when a complex subgoal such as that decribed above is launched, and serves to indicate the satisfaction of the subgoal, as soon as this information becomes available. Once again, it is triggered in an environment where a planned utterance of the system relates to a subgoal of the satisfied goal:

(18) $(\mathrm{IS}\models\langle\langle \mathrm{goal,sys,}(s\models\langle\langle \mathrm{inform,user,sys,G2}\rangle\rangle)\rangle\rangle \mid (\mathrm{CR}\models\mathrm{goal,sys,G0}) \wedge (\mathrm{CR}\models\mathrm{goal,sys,G1}) \wedge \langle\langle \mathrm{subgoal,G1,G0}\rangle\rangle \wedge \langle\langle \mathrm{subgoal,}\langle\langle \mathrm{know,sys,G2}\rangle\rangle\mathrm{,G1}\rangle\rangle)$

$\Rightarrow (\mathrm{IS}\models\langle\langle \mathrm{intend,sys,}(s\models\langle\langle \mathrm{inform,sys,user,G1}\rangle\rangle)\rangle\rangle)$

6 Examples of Plan and Collaborative Strategies in Use in the Sample Dialogue

In this section I offer an analysis of part of the sample dialogue given in Figure 1, showing how and where the plan and collaborative strategies outlined in the previous section are applied.

The sample dialogue in Figure 1 begins with a question by the user to be informed about the action necessary to obtain a subsidy. The question is represented by an abstract:

(19) $$Q = [A|\ (s_1 \models \langle\langle \text{enable,A,}(s_2 \models \text{get,user,subsidy})\rangle\rangle\ )]$$

for which an assignment Asg = [A $\rightarrow$ action] is sought. The user has then a (public) intention with the theme:

(20) $$I = (s \models \langle\langle \text{inform,sys,user,Q,Asg}\rangle\rangle)$$

and thus the SharedPlan* is instantiated by inserting the mutual belief of this intention:

(21) $$(CR \models \langle\langle \text{intend,user,I}\rangle\rangle)$$

The plan strategy given in (9) leads to the expectation that the system adopt an intention to inform the user of the value of the assignment.

As stated in (4) the system believes that an action which will achieve the desired goal is dependent on the four conditions given in (6) being met. The plan strategy given in (10) applies, leading to the system's IS being updated with an intention to inform the user of I if the restriction holds.

The four conditions form an I-sum. The polarity of the conditions as a whole depends on the polarity of each one. Subgoals are therefore launched to discover this.

The first condition requires that the user live in a rented appartment in Vienna. Since the system believes that this is information about which the user is this competent it gives rise, via the mutually believed plan strategy (14) to an intention:

(22) $$(IS \models \langle\langle \text{goal,sys,}(s \models \langle\langle \text{inform,sys,user,I}\rangle\rangle)\rangle\rangle)$$

where I is the assignment for:

(23) POL | ⟨⟨live-in,user,Vienna;POL⟩⟩

By itself, this intention would lead to the question given in line 3. However, the user is expecting, according to the current state of the SharedPlan*, to be informed of the value of an assignment for A. In this context the collaborative strategy given in (16) is triggered, which leads to a statement of the relation between the planned query (line3) and the user's goal, i.e. the utterance in line 2.[8] The SharedPlan* can now be updated with this information.

The mutually believed plan strategy (10) now leads the user to review his expectation for an assignment for Q subject to the conditions holding. Moreover, the mutually believed plan strategy given in (13) enables the user to reconstruct the system's goal to know if the conditions hold and to interprete the system's subsequent queries in this light. Hence the query in line 3 is recognised as referring to the first condition.

The second condition requires that the user's household income not exceed a limit set for households of that size. Since the system currently does not know the user's household income, this leads, via plan strategy (12) to a subgoal to find out the user's household income. Again, this falls within the user's area of competence, leading to a further subgoal to be informed by the user.

Taken alone, this subgoal would lead to the query in line 6, which would allow the user to reconstruct the system's goal to know the user's household income, but would not allow the reconstruction of the superordinate goal, to know if the user's household income exceeds the given limit, since this is part of domain knowledge, outside the user's field of competence. What we have is the environment required to trigger the collaborative strategy given in (17) where G0 is the superordinate goal to know if the conditions hold (present in the SharedPlan* on the basis of the utterance in line 2 and the mutually-believed plan strategy given in (13)) and G1 the missing subgoal. Application of (17) leads to the following intention being added to the IS:

(24) (IS⊨intend,sys,(s⊨inform,sys,user,(IS⊨goal,sys,know-if-condition2-holds)))

This leads to the utterance in line 5, and a subsequent update of the SharedPlan*. The goal to know the user's household income can now be recognised as an application of the plan strategy given in (12) to know the value of a parameter in a superordinate goal.

The situation also triggers the collaborative strategy given in (18) adding to the IS a goal to inform the user the outcome of the subgoal when this is known, i.e:

[8] (16) must be adapted to deal with the fact that we are here dealing with a goal to be informed of an I-part of C, rather than C itself.

(25) (IS|=intend,sys,(s|=inform,sys,user,I))

where I is the assignment for:

(26) POL | ⟪condition2;POL⟫

This results, upon completion of the subdialogue in the utterance in line 15.

The user reacts to the system's query to know his household income by giving his personal income. While the two concepts could be equivalent, the system checks this by a strategy to find out if the user's household possesses any other income (lines 8-13). This includes the three goals to know who else lives in the user's household, whether they have an income, and if so, the level of the income, all of which are subordinate to this goal. Each of these subgoals is reconstructable from an utterance by strategy (9), and the superordinate goal to know the user's household income is already present in the SharedPlan*. Hence only collaborative strategy (18) is here triggered, leading to the utterance in line 14.

In lines 23-25 the answer to the user's original query in line 1 is given, the restrictions on the plan strategy given in (10) having been validated.

7 Future Work and Implementation

I have concentrated on how to update the IS of the system (and susbsequently the system's model of the CR) by means of collaborative strategies to increase informativity. Such strategies are based, as we have explained, on beliefs of the system about the user's competence. However, clearly such beliefs may be misplaced, leading to discrepancies between the CR of the system and the CR of the user. A further problem in our framework consists of identifying and remedying such mismatches.

In this paper I have focussed on an intentional structure component of the CR. However, as we stated in Section 2, the CR also contains other public information. A future task is to construct a means of representation of the attentional structure for the CR, building on the ideas for attentional structure given in Grosz & Sidner (1986) in order to represent mutually believed information regarding focussing and referring expressions. Such information is obviously important both for recognising the identity of anaphora used by the user and for deciding what type of referring expression to use in generating system utterances.

One of the functions of cue words (Grosz & Sidner (1986)) is to assist in the identification of discourse segment boundaries. They may thus bear a close relationship to the colloborative strategies considered in this paper, a relationship which needs to be made explicit in the generation component of VIE-$\mathcal{DU}$.[9]

This approach is currently being implemented in the consulting system VIE-$\mathcal{DU}$ which is being developed at the Austrian Research Institute for Artificial Intelligence. An overview of the system is available in Buchberger et al. (1991).

[9] *jetzt*, 'now', *dann*, 'then' and *das heißt also*, 'that means then' as they are used in Figure 1 can be seen as playing such a role.

8 Conclusion

In extended sequences of advisory dialogue there emerge phenomena which point to a need for a treatment of text above the level of the utterance. Popular approaches to discourse structure such as the Attentional-Intentional Structure of Grosz & Sidner (1986) and RST (Mann & Thompson (1987), Mann & Thompson (1988)) are, however, unable to deal satisfactorily with such phenomena due to their failing to take into consideration the function that discourse plays in revealing how CPs are collaborating towards the given problem.

In this paper an alternative approach to discourse structure has been presented which sees text not as a static but as a dynamic structure, mediating between a CPs goals and intentions (the CP's *intentional structure*) and what she believes is the public view of these same goals and intentions (the *conversational record*). The process of mediation is performed by two kinds of inference rules: plan strategies and collaborative strategies, whereby the former are used to respond to a CP's *domain* goals and the latter to *discourse* goals by satisfying the principle of Informativity. A complete 'representation' of discourse structure is captured only by the relation between these elements.

By means of this model we are able to deal with the textual phenomena mentioned in Sectin 2 that occur in a typical dialogue in our domain. Goals not directly reflected in utterances can be captured by means of a plan strategy such as (15). A CP's reaction to a user goal is explained by plan strategies. Metadialogue functions of text emerge as a result of the application of collaborative strategies, serving to point out the relation between the CP's behaviour and the publicly stated goals of the interaction. Such collaborative strategies function to predict when informativity may be violated on the basis of not fulfilling a partner's expectations, or making use of plan strategies not mutually believed by a partner.

The model shows the relevance of both the dominance relation of Grosz & Sidner (1986) and certain rhetorical relations of RST to a dynamic representation of discourse structure, both acting as triggers to the application of collaborative strategies. The approach is modelled in the situation semantics framework.

References

Barwise, J. (1989) *The Situation in Logic*, CSLI Lecture Notes 17, CSLI, Stanford

Barwise, J. and R. Cooper (1991) *Simple Situation Theory and its Graphical Representation Working Version*, to appear in *Proceedings of the Third Conference on Situation Theory and its Applications*

Barwise, J. and J. Perry (1983) *Situations and Attitudes*, MIT Press, Cambridge, Mass.

Buchberger, E., E. Garner, W. Heinz, J. Matiasek and B. Pfahringer, (1991) 'VIE-DU — Dialogue by Unification', in Kaindl, H. (ed.) *7. Österreichische Artificial-Intelligence-Tagung, Proceedings*, Springer, Berlin. Austrian Research Institute for Artificial Intelligence, Vienna

Clark, H. and D. Wilkes-Gibbs (1990) *Referring as a Collaborative Process*, in P.R. Cohen, J. Morgan and M.E. Pollack (eds.) *Intentions in Communication*, MIT Press, Cambridge, Mass, 463-93

Garner, E. (1991) 'A Model for the Representation of Speech Acts', Technical Report TR-91-15, Austrian Research Institute for Artificial Intelligence

Garner, E. and W. Heinz (1991) 'On the Representation of Speech Acts in Situation Semantics', in Christaller, Th. (ed.), *Proceedings of the 15th German Workshop on Artificial Intelligence*, Springer, Berlin

Grice, P. (1975) 'Logic and Conversation', in P. Cole and H.L. Morgan (ed.) *Speech Acts*, Syntax and Semantics 3, Academic Press, New York

Grosz, B.J. and C.L. Sidner (1986) 'Attention, Intention, and the Structure of Discourse', *Computational Linguistics* **12**, 175-204

Grosz, B.J. and C.L. Sidner (1990) 'Plans for Discourse', in P.R. Cohen, J. Morgan and M.E. Pollack (eds.) *Intentions in Communication*, MIT Press, Cambridge, Mass, 417-45

Mann, W.C. and S.A. Thompson (1987) 'Rhetorical Structure Theory: A Theory of Text Organization', in L. Polanyi (ed.) *The Structure of Discourse*, Ablex Publishing Corporation, Norwood, N.J.

Mann, W.C. and S.A. Thompson (1988) 'Rhetorical Structure Theory: Toward a functional theory of text organization', in *Text* **8(3), 243-81**

Pollack, M.E. (1986) 'A Model of Plan Inference that Distinguishes between the Beliefs of Actors and Observers', *Proceedings of the 24th Annual Meeting of the Association of Computational Linguistics*, 207-14

Thomason, R. (1990) 'Accommodation, Meaning, and Implicature', in P.R. Cohen, J. Morgan and M.E. Pollack (eds.) *Intentions in Communication*, MIT Press, Cambridge, Mass, 335-63.

Wilensky, R., Chin, R., Luria, M., Martin, J., Mayfield, J. and Wu, D. (1988) 'The Berkeley UNIX Consultant Project', in *Computational Linguistics* **14.4**