

A monotonic model of denials in dialogue

Elena Karagjosova

Institut für Linguistik/Germanistik
Stuttgart University

elena.karagjosova@ling.uni-stuttgart.de

Abstract

The paper outlines a monotonic model of denial in dialogue that keeps a representation of the offensive material at the same time as it accounts for the tentative status of utterances with respect to the common ground (CG). It is cast in the Information state based approach to dialogue developed in the projects TRINDI and SIRIDUS (Cooper and Larsson, 1999; Larsson, 2002), and incorporates a notion of discourse commitments (DCs) that enables us to distinguish between information that is part of the CG and such that is merely proposed for consideration. The presented IS based model is meant as a first theoretical approximation towards an adequate DRT-based account of denial and correction.

1 Introduction

This paper deals with denials and their adequate modelling in terms of their effects in dialogue. It treats denial as a case of correction. Consider the dialogue in (1) where B denies the truth of the entire proposition expressed by A's utterance.

- (1) A: The earth is flat.
B: No, it isn't.

The analysis suggested here should be also applicable to other kinds of corrections like (2), where the objection concerns only a portion of the utterance of the previous discourse participant.

- (2) A: Anna ate spaghetti.
B: No, she ate salad.

Existing models of denial and correction are non-monotonic, i.e. they account of denial and correction as effectuating a removal of the corrected material from the context, see (Maier and

van der Sandt, 2003) and (van Leusen, 2004). Both accounts employ a notion of context that corresponds to Stalnaker's *common ground* (CG), i.e. the commitments the discourse participants (DPs) have agreed upon (Stalnaker, 1979). A denial in these models has the effect that the CG is revised by removing the offensive material from the representation of the discourse.

However, the existing accounts of denial and correction in terms of CG-revision do not do justice to the nature of these phenomena. The core of the problem seems to be that the dialogue models employed are not fine-grained enough to treat denial properly. Intuitively, the content of utterances that are rejected does not become part of the CG in the first place. Therefore, it cannot be removed from it by means of denial. Each utterance made is only a proposal on how to update the CG. Only if it is accepted by the other DP, is it added to the CG. A sequence of an assertion and a denial represents a dialogue segment where the DPs explicitly negotiate on how the CG should be updated. In other words, existing models of denial and correction do not have a way to account for the preliminary status of utterances with respect to the CG and more specifically, of the explicit CG-negotiation that denials represent.

Another objection to the non-monotonic models of denial, and especially to the one in (Maier and van der Sandt, 2003) is that it is counterintuitive that the corrected material should completely disappear from the dialogue representation. For comparison, (van Leusen, 2004) deals with corrections in a more fine-grained model that distinguishes the discourse record (a record of all utterances contributed during the discourse, discourse history) from the discourse meaning. However, the same criticism holds for this model as well when it comes to accounting of corrections as context revisions.

The present paper outlines an alternative model of denial in dialogue. The model is monotonic since it keeps a representation of the offensive material at the same time as it accounts for the tentative status of utterances with respect to the CG. It is cast in the Information state based approach to dialogue developed in the projects TRINDI and SIRIDUS (Cooper and Larsson, 1999; Larsson, 2002), and incorporates a notion of discourse commitments (DCs) that enables us to distinguish between information that is part of the CG and such that is merely proposed for consideration. The IS based approach to dialogue provides a framework that is flexible enough to implement the more fine-grained dialogue model needed. The presented IS based model is meant as a first theoretical approximation towards an adequate DRT-based account of denial and correction. As the discussion of the existing DRT approaches will show, in order to model correction adequately in DRT, more fundamental, non-trivial modifications to the theory have to be made. This is a large scale enterprise that will be addressed in future work.

The paper is structured as follows. Section 2 provides an overview over the two existing elaborate models of denial and correction. In section 3, I briefly clarify my understanding of the relation between denials and the notion of context in terms of CG. I implement this relation within the framework of the Information state based approach to dialogue in section 4, and section 5 presents a refined account that also implements the notion of dialogue history. Finally, a summary and outlook are presented in section 6.

2 Non-monotonic models of denial

2.1 (Maier and van der Sandt, 2003)

(Maier and van der Sandt, 2003) account of denial in terms of its discourse effects, which are claimed to be “a non-monotonic correction operation on contextual information”. More closely, they argue that the primary function of denial is “to object to information which has been entered before and to remove it from the discourse record.” They model denials in an extension of DRT. In DRT, discourse is modelled in terms of abstract structures, DRSs, which represent the meaning of the discourse as it evolves. Each new sentence is interpreted on the background of the representation of the preceding discourse, the background-DRS. Thus in DRT, the notion of context is modelled by

the DRS. In order to be able to model dialogue, (Maier and van der Sandt, 2003) propose an extension to standard DRT that allows keeping track of who said what and when in a dialogue.¹ In this extension, it is assumed that a DRS represents the CG of the DPs, i.e. the propositions that the DPs have agreed upon as being true. In this framework, it is not possible to model denial monotonically. As (Maier and van der Sandt, 2003, p.12) argue, “with respect to an incoming context that contains the offensive material the sentence cannot even be processed in view of the fact that this would result in a plain contradiction.” Therefore they model the effect of denial by means of the so-called “reversed anaphora”: the denial is not further processed but leaves a simple negated condition in the DRS. Then a preliminary sentence representation is constructed and merged with the background DRS. After that, a presupposition resolution mechanism with reversed anaphora collects the offensive material from the preceding utterance and moves it to the position of the negated condition. As a result of this process, the contribution of the corrected utterance is removed from the main DRS and the material it originally introduced ends up under the scope of the negation introduced by the denial. In other words, the offensive material is removed first from the dialogue representation (the CG), and then the content of the denial is added to it.

For instance, as a result of this process, the final representation of the dialogue in (3) is a DRS containing the representation of σ_2 , which is the negated sentence σ_1 , and the representation of σ_3 .

- (3) σ_1 A: The King of France walks in the park.
 σ_2 B: No, he doesn't,
 σ_3 France doesn't have a king.

However, the final representation of the dialogue in (3) does not seem satisfactory as a dialogue representation of an assertion-denial sequence since it only contains a representation of the negative assertion in σ_2 and contains no trace of the denied utterance σ_1 . I.e. the representation of the dialogue in (3) hardly differs from the representation of the negative statement *The King of France does not walk in the park*. It is unsatisfactory to let the contribution of the offensive

¹It is further refined to Layered DRT in order to be able to account for cases where only parts of the utterance are rejected while others are acknowledged.

material disappear from the dialogue representation. One reason for wanting to keep this material in a dialogue representation may be that a speaker may want to refer to it at some later stage of the dialogue. As already mentioned, in this approach non-monotonicity is necessary since the DRS represents the CG and must be kept consistent. On the other hand, the inadequacy of the proposed solution suggests that a more radical modification of DRT is needed to capture adequately the nature of dialogue and of phenomena like corrections.²

Another objection to the account presented in (Maier and van der Sandt, 2003) is that if a DRS reflects the CG, it is inadequate to treat the offensive material as being part of the CG, since the other DP hasn't acknowledged it yet. The same holds for the content of the denial itself, since the other DP may disagree and stick to his opinion. In the present model, the representation of the dialogue in (3) contains a representation of the denial σ_2 , which means that the denial is part of the CG. In general, if a DRS reflects the CG, then it is impossible to add new sentence representations to it, since each utterance in a dialogue is only a proposal on how the CG should be updated. The CG changes only after the proposal is accepted, explicitly or implicitly. In other words, before each update of the CG, there is a grounding step (see (Traum, 1994) on grounding). In the case of denial and correction, we can speak of a dialogue segment that has the purpose of explicitly negotiating how the CG should be updated. Consequently, dialogue models should provide separate representations for the level at which the content of the CG is negotiated, and for the one that represents the result of this negotiation. As it stands, Maier and Sandt's DRT model reflects the former but is intended to represent the latter. When a speaker denies an utterance of the other DP, neither the

²Another inadequacy of this model is that it does not account for the fact that it is not always the previous utterance that contains the offensive material, cf. corrections with German accented adverbial *doch*, as in (i), where the correction occurs several turns away from the turn that introduces the offensive material, see (Karagjosova, 2006) on corrections with *doch*.

- (i) A₁: **es geht nicht**. ('it does not work')
 B₁: du musst die Schraube drehen, [...] ('you must turn the screw')
 A₂: [...] hast recht ('you are right')
 B₂: Na siehst du? **es geht** DOCH ('What did I tell you? It works.')

content of the preceding utterance nor that of the denial are part of the CG. Consequently, denials cannot be identified with revisions in the CG. Denials are part of a negotiation phase in dialogue, and this is not reflected in the model.

2.2 (van Leusen, 2004)

The second nonmonotonic account of denial and correction in dialogue is proposed in (van Leusen, 2004). It is based on a more sophisticated dialogue model than the one in (Maier and van der Sandt, 2003), called Logical Description Grammar (van Leusen and Muskens, 2003), that distinguishes between the utterances made, or the dialogue history, and the discourse meaning. The former is modelled by means of "discourse descriptions", which describe the syntactic, semantic and pragmatic characteristics of the discourse and represent "the language user's knowledge of the discourse processed so far" (van Leusen, 2004). A second level constitutes the discourse meaning or context, which is a model that fits or verifies a discourse description. The incrementation of the discourse description is monotonic, including cases of corrections and denials. Corrections have an update effect only on discourse meaning.

In this model, discourse meanings and sentence meanings are DRSs. Correspondingly, contexts are DRSs, as well as the semantic contents of discourse contributions. The discourse meaning is built up in this model from the contents of the utterances of the DPs. Since DPs may disagree on certain points, it is argued, it is not necessary that all of the contextual information is believed or supported by each of the participants. Therefore it is assumed that the context at any point of the conversation represents what has been proposed for acceptance as CG by the most recent speaker.

This view is so far consistent with the position advocated in this paper. However, it is not entirely clear how the notion of context is defined in van Leusen's account. Thus, if the context only contains proposals on how to alter the CG, why does it have to be kept consistent? There may be contradicting proposals. Also, the effect of corrections on the context is modeled in terms of Gärdenfors' revision of epistemic states (Gärdenfors, 1988), which suggests that the notion of context employed in this model coincides with the notion of the CG, and the same criticism as for (Maier and van der Sandt, 2003) holds for this approach too.

3 Denials and the notion of context

As mentioned in section 1, the notion of context employed in the existing models of denial and correction is based on Stalnaker's notion of the CG. Stalnaker defines context as follows: "Think of a state of a context at any given moment as defined by the presuppositions of the participants." Presupposition is defined as "what is taken by the speaker to be CG of the participants", and the CG represents the propositions the dialogue participants have agreed upon, or mutually believe.

Based on this notion of context as CG, Stalnaker defines assertion in terms of its effects on the context, namely, the content of an assertion changes the context (i.e. the CG) by reducing the context set (i.e. all possible situations/worlds incompatible with what is said are eliminated). Stalnaker specifies further that this effect is only achieved provided there are no objections from the other DP. "This effect is avoided only if the assertion is rejected." In a footnote (footnote 9 on p. 324), Stalnaker argues further that "to reject an assertion is not to assert or assent to the contradictory of the assertion, but only to refuse to accept the assertion. If an assertion is rejected, the context remains as it was." More exactly, rejection of an assertion blocks the effect assertions have on the context, namely adding its contents to the CG.

Thus we find support for our criticism of the existing models of denial in Stalnaker's definition of assertion and its effects on the CG. A denial is a rejection to add the contents of a previously made contribution to the CG. The corrected material, i.e. the previous commitment, is thus not yet part of the CG, since its content has not been agreed upon yet. The corrected material, as well as the correction itself, are just *proposals* on how to update CG.

In the next two sections I implement an alternative, monotonic account of denial in the framework of the Information State based approach to dialogue.

4 The Information State based approach to dialogue

In this section, I use the framework of the Information State based approach to dialogue to implement a model of denial that does not assume CG revision.

The information state (IS) is an abstract data structure that represents information available to the DPs at any given stage of the dialogue. It is

based on Ginzburg's (Ginzburg, 1998) notion of the dialogue gameboard which in turn is a modification and elaboration on Stalnaker's "common ground" and Lewis' "conversational scoreboard" (Lewis, 1979).

Ginzburg's dialogue gameboard is a structure that includes propositions, questions and dialogue moves.³ The IS is an adaption of Ginzburg's DGB. The IS is a flexible construct that allows adding more complexity depending on the requirements of the dialogue phenomena that are modelled. This is also the strategy that I will follow here. I will start with a basic IS model, namely the IS used in (Larsson, 2002) in an implementation of the dialogue system IBiS, and see how far this IS can get us. It will turn out that additional complexities have to be added.

The basic structure of the IS is represented in figure 1 on page 5. The dialogue information is divided into two basic records: private and shared information. The record of information private to the DPs contains an agenda of actions the DP (by default the system) needs to perform in the near future (/PRIVATE/AGENDA), a dialogue plan for more long-term actions (/PRIVATE/PLAN), and a set of beliefs (/PRIVATE/BEL). Another record represents the shared information, information that is public to the DPs (system and user), containing a set of mutually agreed-upon propositions (/SHARED/COM), a stack of questions under discussion (/SHARED/QUD) and information about the latest utterance (/SHARED/LU): the speaker and the speech act/move realised by the utterance (assuming for simplicity that each utterance realizes only one move). The propositions in /SHARED/COM need not be actually believed by the DPs but they have committed to them for the purpose of the conversation.

Let us examine how the dialogue information is recorded in the IS in the case of denial. Consider the exchange in (4).

- (4) A: The earth is flat.
B: No, the earth is not flat.

After the first utterance, the IS looks like in figure 2. Here, the record /PRIVATE/BEL contains the information about the belief of the speaker that the earth is flat. In (Larsson, 2002), this slot is foreseen for utterances of the system, where for in-

³In this model, each DP has his own version of the DGB and there may be mismatches. This is also the model on which the IS in (Cooper and Larsson, 1999) is based.




Figure 1: IBiS information state type




Figure 2: A: *The earth is flat*




Figure 3: B: *The earth is not flat*

stance after a database search the system writes the result of the search into this field in order to communicate this result in a next utterance. The field /SHARED/COM is assumed to be empty at the start of the dialogue. It will not change at this stage since the content of the utterance does not become CG before it has been accepted by the other DP. In general, the update of the IS is governed by update rules defined for handling various dialogue moves, as well as for handling plans and actions. Concerning the CG, different grounding strategies may be assumed, such as optimistic (content of utterance automatically added to CG), cautious (content added but can be retracted if DP objects to it) and pessimistic grounding (content of utterance added to CG only after positive evidence for grounding) (Larsson, 2002). I will assume a pessimistic grounding strategy because it seems to reflect more adequately the nature of corrections. Following this strategy, the CG will not get updated after this first utterance until positive feedback is received. The utterance of the assertion has the effect that the corresponding yes/no-question is pushed on top of the QUD-stack, where it stays until it is resolved, i.e. gets accepted.⁴

I leave the field /PRIVATE/AGENDA empty since at this point it is not relevant for the current investigation. I also ignore completely the field /PRIVATE/PLAN for the same reason.

After the second utterance, the IS gets updated again and looks like in figure 3. The field /PRIVATE/BEL contains now the belief of the current speaker B that he communicates via utterance (4-

B). The CG is still empty. Ginzburg's utterance processing protocol foresees that when an assertion is rejected, its content is not added to the CG. The only effect it has is that the corresponding yes/no-question is pushed on /SHARED/QUD.

The record /SHARED/LU is updated with information about the current move. This record only keeps information about the latest move. The next move overrides it with its own information.

If A agrees with B, then the IS will look like the IS in figure 4:




Figure 4: A: *You are right, the earth is not flat*

Here, the speaker holds the belief that the earth is not flat. The CG will be updated with this proposition, and the question will be removed from the QUD-stack since it is resolved.

Implementing denial in the basic IS under the pessimistic grounding strategy captures in a way the preliminary status of utterances with respect to the CG. However, it does not provide means for keeping track of the actual commitments of the DPs in the course of the dialogue, but only reflects those that the DPs have agreed upon. This and several other points require adjustments to the simple IS and the rules for its update in order to model denial more adequately in this framework. This issue will be the subject of the next section.

⁴Otherwise it remains unresolved.

5 A modified IS for dealing with denials

In order to be able to deal more adequately with denials and corrections in this framework, I suggest some adjustments to the structure of the IS and its contents, as well as to the rules of its update. The modified IS is presented in figure 5 on page 7.

First of all, we need to distinguish between different layers of information in dialogue, CG and discourse history. (Gunlogson, 2003) proposes a dialogue model that allows to keep track of the discourse commitments (DCs) of the DPs. DCs are beliefs publicly attributed to each participant in the conversation. I.e. if A says *p* then it becomes CG that A believes *p*. A public belief of a DP does not have to be mutual. I.e. if it is CG that A believes *p*, from this does not follow that it is CG that *p*. Thus in a way, DCs capture the notion of dialogue history. In this model, dialogue history is part of the CG: it is in the CG that A has committed himself to *p*. Gunlogson adopts Stalnaker's definition of the CG as the set of propositions representing what the participants in a conversation take to be mutually believed, or at least mutually assumed for the purpose of the discourse.

However, implementing directly the DCs as part of the CG will not lead us far in the case of denial and correction, since in the case the DP accepts the denial, the CG must be revised or else it will become inconsistent. However, this would mean that we remove a commitment made by a speaker from the dialogue record, which is not satisfactory: intuitively, even if the DP makes a contradictory commitment, the fact that he has made the earlier commitment remains. A more adequate solution will be therefore to separate the DC from the CG. I therefore implement Gunlogson's discourse commitments as a separate field of the IS. The DCs represent the propositions that the DPs have committed to in the course of the conversation. Each utterance (at least each assertion) leads to updating the DC (the field /SHARED/DC in figure 5 on page 7) with the information that the speaker believes the proposition expressed by the utterance. E.g., after A's utterance in (4), /SHARED/DC is updated with the information $B_A flat(earth)$. By means of implementing the DCs, we can keep track of the dialogue history, a record of all utterances contributed by the DPs in the course of the entire dialogue, independently of the CG status of their contents. In the field /SHARED/DC, mutually contra-

dicting utterances of the DPs can coexist. The data type is assumed to be an ordered set (although it may be inconsistent) of beliefs.

Second, since we model interaction between human DPs, we need a way to keep track of both DP's beliefs and commitments. I.e. we need to be able to represent the beliefs of the DPs separately. This can be done by using a belief operator *B* indexed with the speaker of the utterance and holder of the respective belief. E.g., $B_A flat(earth)$. In other words, the field /PRIVATE/BEL is a set of beliefs. The update will not overwrite the information in this field, but augment it with the beliefs of the next DP.⁵

Thus the field /SHARED/DC will partly contain the same information as the field /PRIVATE/BEL. The difference will be that while we cannot retract commitments, we can revise belief states, i.e. delete certain beliefs from /PRIVATE/BEL. Thus the information in /PRIVATE/BEL is not redundant but can be used to model the dynamics of the belief states of the DPs during the exchange. As already said, DPs need not actually hold these beliefs, but it suffices that they act as if they were.

The CG is represented by a separate field. In order to avoid confusion with the Discourse Commitments, I call the field that records the CG /SHARED/BEL (instead of "shared commitments"), since it concerns the propositions that the DPs mutually believe.⁶ The shared believes correspond to the notion of CG, i.e. commitments the DPs have agreed upon. The data type is a set of propositions.⁷ Note that the CG does not include information that is merely public, or manifest, to the DPs, such as the information captured by the other subfields in the SHARED-field, but rather concerns only the content of the utterances.

By separating the DC from the CG we can capture the CG-negotiating effect of denial⁸ and

⁵Another possibility is to have different copies of the IS for each DP, as in (Cooper and Larsson, 1999). However, this solution will unnecessarily complicate matters and will not be further pursued for the time being.

⁶Note that in (Cooper and Larsson, 1999) /SHARED/COM is /COMMON/BEL, which reflects more adequately the intended purpose of this field as a set of agreed upon, or commonly believed, propositions.

⁷Actually, the CG may contain not only propositions, but also beliefs attributed to other DPs, or introspective belief, which means that it is a set of propositions and beliefs. I will ignore this potential complication for the time being.

⁸There exist IS-based dialogue models that provide means for modelling the process of negotiating content, such as the PENDING field in (Ginzburg and Cooper, 2004) and the list of ungrounded discourse units in (Poesio and Traum, 1998)

$$\left[\begin{array}{l} \text{PRIVATE} : \\ \text{SHARED} : \end{array} \left[\begin{array}{l} \text{BEL} : \text{Set}(\text{Bel}) \\ \text{DC} : \text{Set}(\text{Bel}) \\ \text{BEL} : \text{Set}(\text{Prop}) \\ \text{QUD} : \text{Stack}(\text{Quest}) \\ \text{MOVES} : \text{Set}(\text{Moves}) \end{array} \right] \right]$$

Figure 5: IS modified

need not assume that denial updates nonmonotonically the CG.⁹ A revision becomes only necessary within the private beliefs of the DPs (the field /PRIVATE/BEL) in case the corrected DP accepts the correction. Note that assuming that the private beliefs of the DP can be revised does not make our model of denial less monotonic. The monotonicity concerns only the DC field, which accumulates all utterances made during the conversation independently of whether their contents are accepted or rejected by the other DP.

Further, as already said, each assertive utterance results in raising the respective yes/no-question in QUD, i.e. pushing it on top of the QUD-stack.

Another change concerns the LU-record. In standard IS there is no relation between the utterance and its utterer after the IS gets updated - after the next utterance, the speaker is a different person, and we do not have a way to relate the contents of the utterance with its originator beyond the respective turn. LU only shows who the last speaker was. Having represented the DCs, we do not need information about the latest speaker. We keep however the information about the move realized by the utterance, where in order to keep track of who realized which move, we index the moves with the respective DP, e.g. $assert_A(\text{flat}(\text{earth}))$. Also, in order to have a more complete record of the course of the dialogue, we do not let the dialogue move-field be overwritten after each update, but make sure that it gets augmented with the next moves.¹⁰ It would also be useful that the information in the move-field is ordered, i.e. we assume that it has the data type ordered set.¹¹

I also ignore for the time being the fields /PRIVATE/AGENDA and /PRIVATE/PLAN, since they

⁹Of course, there may be situations in a dialogue where the CG has to be revised, e.g. when both DPs adopt a belief that contradicts their earlier common knowledge.

¹⁰A similar strategy is also applied in other dialogue models, see e.g. (Ginzburg and Fernandez, 2005).

¹¹It may actually be more reasonable to assume a sequence or a stack, since there could be multiple entries. The same holds for the DC field.

are not immediately relevant for my purpose.

All updates must be handled by respective update rules, whose definition however I have to ignore for the time being.

Let us go through an example to see how this model works. Figure 6 reflects the IS after the first utterance in (4). It contains the private belief of the speaker (under the assumption of cooperativity) that he communicates with his utterance. In the shared record, the CG is empty, for the sake of simplicity, i.e. this is the first utterance in a dialogue. The QUD is whether the earth is flat, and the communicated belief is recorded as a DC of A in /SHARED/DC.

$$\left[\begin{array}{l} \text{PRIV} : \\ \text{SH} : \end{array} \left[\begin{array}{l} \text{BEL} : \{B_A \text{ flat}(\text{earth})\} \\ \text{DC} : \{B_A \text{ flat}(\text{earth})\} \\ \text{BEL} : \{\} \\ \text{QUD} : \langle ?\text{flat}(\text{earth}) \rangle \\ \text{MV} : \{ \text{assert}_A(\text{flat}(\text{earth})) \} \end{array} \right] \right]$$

Figure 6: A: *The earth is flat.*

The IS after the second utterance is represented in Figure 7. It reflects in addition the private belief of the speaker B, which is just opposite to what A asserts. In the shared record, the CG is still empty, since the DPs have not yet agreed on a proposition. The field /SHARED/DC is updated by the DC of B. Topmost on QUD is still the question whether the earth is flat.

$$\left[\begin{array}{l} \text{PRIV} : \\ \text{SH} : \end{array} \left[\begin{array}{l} \text{BEL} : \{B_A \text{ flat}(\text{earth}), \\ B_B \neg \text{flat}(\text{earth})\} \\ \text{DC} : \{B_A \text{ flat}(\text{earth}), \\ B_B \neg \text{flat}(\text{earth})\} \\ \text{BEL} : \{\} \\ \text{QUD} : \langle ?\text{flat}(\text{earth}) \rangle \\ \text{MV} : \{ \text{assert}_A(\text{flat}(\text{earth})), \\ \text{deny}_B(\text{flat}(\text{earth})) \} \end{array} \right] \right]$$

Figure 7: B: *The earth is not flat.*

Suppose that after some convincing argumentation of B, A finally accepts B's counterproposal on how to update the CG, namely with the proposition that the earth is not flat. This situation is reflected in figure 8. Then, this proposition will

be added to the CG, here the field /SHARED/BEL. The QUD will be empty - the question whether the earth is flat is resolved and can be popped out from the QUD-stack. The beliefs of A will be revised: the old abandoned belief of A that the earth is flat will be deleted. This can be reflected in the field /PRIVATE/BEL by either simply removing the respective belief from the set, or by marking it somehow as not held anymore (e.g. by crossing the respective belief out), if we want to be able to capture the dynamics of the DPs' beliefs. In the example, I choose the first option for simplicity.

$$\left[\begin{array}{l} \text{PRIV} \\ \text{SH} \end{array} : \left[\begin{array}{l} \text{BEL} : \{ B_B \neg \text{flat}(\text{earth}), \\ \quad : B_A \neg \text{flat}(\text{earth}) \} \\ \text{DC} : \{ B_A \text{flat}(\text{earth}), \\ \quad : B_B \neg \text{flat}(\text{earth}), \\ \quad : B_A \neg \text{flat}(\text{earth}) \} \\ \text{BEL} : \{ \neg \text{flat}(\text{earth}) \} \\ \text{QUD} : \langle \rangle \\ \text{MV} : \{ \text{assert}_A(\text{flat}(\text{earth})), \\ \quad : \text{deny}_B(\text{flat}(\text{earth})), \\ \quad : \text{accept}_A(\neg(\text{flat}(\text{earth}))) \} \end{array} \right] \right]$$

Figure 8: A: *You are right, the earth is not flat*

6 Summary and outlook

In this paper I present a model of denial in dialogue that assumes that denial does not revise the CG but represents a phase in a dialogue with the purpose to negotiate the contents of the CG. I implement this idea in the IS based approach to dialogue and argue that it is important to be able to keep track of the dialogue history in order to deal adequately with denials. An obvious drawback of the proposed implementation is that the IS and especially the fields private beliefs, DC, and moves can become extremely long for realistic applications. But since the purpose of this investigation is a theoretical one, this fact is irrelevant for the time being. The ultimate goal of the present investigation that will be pursued in future work, is the development of a DRT-based model of denial and correction in dialogue that distinguishes between the CG and the dialogue history, and takes into account the private beliefs of the DPs. A DRT-based model should also be able to provide a proper semantics for these notions, an issue that was neglected in the present paper.

Acknowledgement I would like to thank three anonymous reviewers for detailed comments and suggestions.

References

- Robin Cooper and Staffan Larsson. 1999. Dialogue moves and information states. In *Proceedings of the Third IWCS*, Tilburg.
- Peter Gärdenfors. 1988. *Knowledge in flux. Modelling the dynamics of epistemic states*. MIT Press, Cambridge, Massachusetts.
- Jonathan Ginzburg and Robin Cooper. 2004. Clarification, ellipsis, and the nature of contextual updates in dialogue. *Linguistics and philosophy*, (27):297–365.
- Jonathan Ginzburg and Raquel Fernandez. 2005. Action at a distance: the difference between dialogue and multilogue. In *Proceedings of DiaLor*.
- Jonathan Ginzburg. 1998. Clarifying utterances. In *Proceedings of TwenDial*.
- Christine Gunlogson. 2003. *True to form. Rising and falling declaratives as questions in English*. Routledge, New York.
- Elena Karagjosova. 2006. Correction and acceptance by contrastive focus. In *Proceedings of BranDial*.
- Staffan Larsson. 2002. *Issue-based dialogue management*. Ph.D. thesis, Gteborg University.
- D. K. Lewis. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–359.
- Emar Maier and Rob van der Sandt. 2003. Denial and correction in layered DRT. In *Proceedings of DiaBruck*.
- Massimo Poesio and David Traum. 1998. Towards an axiomatization of dialogue acts. In *Proceedings of TwenDial*.
- Robert Stalnaker. 1979. Assertion. In P. Cole, editor, *Syntax and Semantics*, volume 9. Academic Press, London.
- David Traum. 1994. *A Computational Theory of Grounding in Natural Language Conversation*. Ph.D. thesis, University of Rochester.
- Noor van Leusen and Reinhard Muskens. 2003. Construction by description in discourse representation. In J. Peregrin, editor, *Meaning, the Dynamic turn*. Elsevier.
- Noor van Leusen. 2004. Incompatibility in context. A diagnosis for correction. *Journal of Semantics*, 21(4).