On Coordination and the Common Ground

Anton Benz

1 Introduction

In this talk we consider the role of common ground for some coordination problems in dialogue. There, we concentrate especially on the difference between explicit and implicit representations. A prototypical example for a coordination problem is the so-called *coordinated attack* problem[1]:

(1) Two divisions of an army are camped on two hilltops. In the valley awaits the enemy. It is clear that if both divisions attack the enemy simultaneously they will win the battle, while if only one division attacks it will be defeated. As a result, neither general will attack unless he is absolutely sure that the other will attack with him. In particular, a general will not attack if he receives no massage. The commanding general of the first division wishes to coordinate a simultaneous attack.

The two armies have to coordinate a joined project — simultaneous attack — with a common goal — defeating the enemy. There is a view on language use[2] which emphasises especially the role of joined projects and their relation to common ground and public goals. In this talk we want to consider the role of common ground more closely, especially its two aspects as *common information* which is provided by general rules of language interpretation, and as *explicitly represented common knowledge*.

If we consider only monologue, then we can identify the common ground at a certain point with the information which has been introduced in the previous text. There are principally two ways to represent the common ground: As a set of possible worlds or in a more syntactic way, e.g. as a Discourse Representation Structure (DRS)[3]. It is common ground that a sentence ψ holds iff ψ is true in all worlds, or iff it is a DRS–condition. The meaning of a DRS is a set of worlds, or world–assignment pairs. This allows to compare the two representations. The main difference is with respect to fine–grainedness and *logical omniscience*[4]. We call the possible worlds representation the *common information*, and the DRS-like representation a *common DRS*.

If we look at dialogue, then the distinction between these representations has more interesting aspects. As an example, we look at the use of *specific* indefinite NPs and reference to them by use of pronouns or definite descriptions[5].

(2) A: Yesterday, a man ran into my office, who inquired after the secretary's office.

B: Was *the man* wearing a purple jogging suit?

A: If it was Arnold he was, and if it was somebody else he was not.

Dekker observes that A's answer sounds strange, even if we assume that there was more than one person coming into the office, one of them Arnold. We assume that a definite description def φ picks out an object (1) which must be given in the common ground, and (2) which is the only one in the common ground which has the property φ. For dialogue there are again the two ways to represent the common ground: As a set of possible worlds or in a syntactic way. According to the first view, the common ground can be identified with the set of all worlds which are epistemically possible for one of the dialogue participants, or which are possible for one of the participants according to the beliefs of the other participants about his beliefs, etc. Intuitively, B's use of *the man* in the example refers back to a specific person introduced by A. But if we represent the common ground as a set of possible worlds, then there is no unique *real* object such that a definite description can pick it out as the only object which is a man. There is only a unique discourse referent, but discourse referents are not part of the common information. Hence, we can try a DRS–like representation for the common ground, but then we have to explain in which sense such

[1] See [4, pp. 176–183]. The example is taken from p. 176.

[2] Of course, it is *H.H. Clark*'s view we have in mind [2].

[3] [6], or e.g. as a sentence in Dynamic Predicate Logic [5].

[4] See [4, Ch. 7] and [7, Ch. 2] for a discussion of various notions of knowledge and belief in modal approaches and the problem of logical omniscience.

[5] The example is due to *P. Dekker* [3]. In Dekker's original version B's utterance contains a pronoun but our point becomes clearer if we replace it by a definite description.

a DRS is *common*, and how interlocutors manage to refer to the same *object*. The last point is critical especially if we assume that referents introduced by specific indefinites are free variables. We will argue that:

- Both representations have to be kept distinct.
- Success of coordination is defined relative to common information.
- There are fixed rules in a community for how to update the common DRS.

Our main example which shows that we have to make a distinction between common information and common DRSes is the use of specific indefinites and anaphoric reference to the objects introduced by them. Here, we assume that all participants use DRS-like representations for the common ground and that the interpretation of definite descriptions depend on them. That the description then refers to the same object is a coordination problem. We discuss this example in Section 3. We model it by use of *multi–agent* systems. This systems consist of a set of situations, joined acts, a transition operation, and a function P which specifies which joined acts can possibly be performed in which situation. Semantic and pragmatic conditions enter as restrictions on the function P. We introduce multi–agent systems in more detail in Section 2. Success has to be defined with respect to common information. This is motivated by Example **(1)**. Where *success* means that interlocutors manage to solve a coordination problem. In order to show that fixed update rules play a crucial role for coordination we discuss the *time imprecision problem*[6], and the *Muddy Children Puzzle*[7] in Section 4.

2 Dialogue and Multi–Agent Systems

We describe dialogue fragments as *Multi–Agent Systems*. Here, we follow the theory developed in [4]. Within this framework we can describe dialogue and the update of the knowledge bases of the participants as a game. This has the advantage that we can exploit standard techniques to define the information of agents in a possible worlds framework, and we get the usual definition of mutual knowledge. A multi–agent system consists of the following components:

1. A set $\mathcal{S}$ of *global states*.
2. A set ACT of possible dialogue acts.
3. A function P[8] which tells us which dialogue acts can be performed in which dialogue situations.
4. A transition operation $\tau :$ ACT $\times \mathcal{S} \longrightarrow \mathcal{S}$ which models the effect of performing a dialogue act.
5. A set of initial dialogue situations S_0.

We identify dialogues with the set $\mathcal{D}$ of all sequences

$$\langle s_0, \mathtt{act}_0, \ldots, \mathtt{act}_{n-1}, s_n \rangle$$

where s_0 is an initial dialogue situation, and:

- $\mathtt{act}_i \in P(s_i)$, i.e. $\mathtt{act}_i$ is a possible joined action in s_i.
- $s_{i+1} = \tau(\mathtt{act}_i, s_i)$.

[6][4, pp. 395-397]
[7][4, pp. 3-7]
[8]In [4] this function is called a *protocol*.

Furthermore, we assume that a context s divides into three components: Two for the interlocutors, and one for the whole *environment* including the situation talked about. This means that a context is of the form $\langle e, s_S, s_H \rangle$, where e denotes the state of the environment, s_S the state of the speaker, and s_H the state of the hearer. We assume that assertions don't change the environment.

It is usual to identify the information of a participant X in a situation s_i with the set of all situations s' where the local state of X is identical with its local state in s_i. Hence, we represent the *information of agent X at time i in dialogue* $D = \langle s_{D,0}, \mathtt{act}_{D,0}, \ldots, \mathtt{act}_{D,n-1}, s_{D,n} \rangle$ by:

$$I(X, i, D) := \{D' \in \mathcal{D} \mid s_{D,i,X} = s_{D',i,X}\}.$$

Then we can represent the *common information* $\mathrm{CI}(i, D)$ at time i in Dialogue D by:

$$M^0 := \{D\},$$

$$M^{n+1} := \bigcup\{I(i, D', X) \mid X = S, H \;\&\; D' \in M^n\},$$

$$\mathrm{CI}(i, D) := \bigcup_{n \in \mathbf{N}} M^n.$$

If M characterises some property of dialogues, i.e. if $M \subseteq \mathcal{D}$, then it will be *common information* at time i that the actual Dialogue D has this property, iff $\mathrm{CI}(i, D) \subseteq M$.

We characterise a *joined goal* of interlocutors by a set G of dialogue situations. The goal is reached, if the actual situation is an element of G. Hence, their coordination problem is to choose their joined actions in such a way that it is common information that the resulting state belongs to G.

We reconsider Example **(1)**. The commanding general of the first division wishes to coordinate a simultaneous attack. Therefore, he has to decide at which time t the attack should take place. We assume for the example that he has to inform the general of the second division by sending an appropriate message. Hence, G is the set of all situations where both generals know that they have to attack at time t. The *dialogue moves* in this example are messages (and acknowledgements) which can be sent from one division to the other. In the end, it has to be common information that the attack takes place simultaneously at the specified time. If there remains a possibility where one of the two divisions does not attack, then the other one should also not attack. If the transmission of the messages takes no time, and if it is certain that they reach the addressee, then this goal can be reached by a simple message where the first general states the time of attack. This follows by induction over the M_ns[9]. In the actual dialogue D it is guaranteed that the addressee gets the message. The speaker knows this, hence, it implies that in all dialogues which are indiscernible for him it is also guaranteed that the message will be delivered immediately. And the addressee knows for all his possible dialogues that he received a message. But this can only be because the speaker sent one. The reasoning can be repeated in each induction step.

This example shows also that a common goal needs not to be known to all interlocutors. In the initial situation G is only known to the general of the first division.

3 Example I: Referential Chains

We describe a dialogue fragment with assertions which allows for the interpretation of specific indefinite NPs, and reference to objects introduced by them[10].

(3) At 7:00 am Anna observes how a Doberman bites the young girl Melanie. Some minutes later the Doberman attacks also another young girl, Stefanie. This time, it is Debra who observes it.

[9] We will here and also later present only sketches (!) of proofs. Of course, for a real proof we would have to provide detailed definitions of the multi–agent systems, the possible actions, their update effects, and the local states. We can refer to [4] for a precise treatment of coordinated attack problems.

[10] This is a very much simplified version of [1]. We refer to this paper for the precise definitions of the multi–agent systems.

a) Then Anna meets Bob and Chris and tells them that she has seen how a Doberman attacked a young girl.

b) The next day, Debra meets Bob, and she tells him that the dog attacked also another young girl.

c) Later, she meets also Chris and tells him the same.

d) Chris, who does not know that Bob knows already the whole story, meets Bob again and says to him: "*The young girl* was not the only one who was attacked by the dangerous Doberman."

It is a standard assumption that a definite description $\text{def}\, x.\varphi(x)$ picks out an object a which must be an element in the common ground such that $\varphi(a)$ holds, and there must be no other object b in the common ground such that $\varphi(b)$ holds. If we represent the common ground as common information, then we normally can't find such an object. The use of *the young girl* by Chris is felicitous although both of them know that there have been two young girls who were attacked by the Doberman. Moreover, there is no real object a such that $\varphi(a)$ is mutually known. Intuitively, the description *the young girl* refers to Melanie because she is the only person who is available through a common referent.

The following picture shows (a part of) the situation of Example **(3)** after a) where Anna told Bob and Chris that a young girl was bitten by a Doberman.

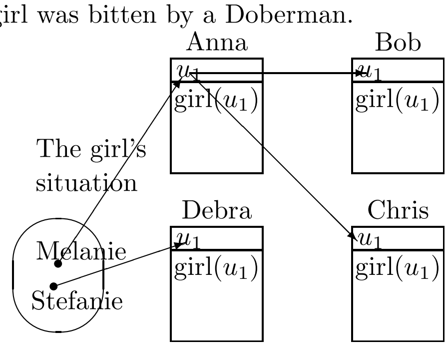

There are parallels between the interpretation of referential NPs and the coordinated attack problem. The task to coordinate simultaneous attack at a certain time can be seen in parallel with the task to coordinate identical reference to a certain object. The speaker chooses the object, and we assume that he can introduce it into the common ground with an indefinite NP. In the end, it has to be common information that an appropriate definite description refers for all interlocutors to the same object.

We describe the use of specific indefinites by a multi–agent system. Each use of an indefinite NP introduces a new discourse referent into the knowledge base of the hearer. We represent the local states of interlocutors as DRSes, and the update triggered by an assertion by a DRT mechanism. If a new object is introduced by a specific indefinite NP, then it is chained to an object which serves as its *source*. A global dialogue state is of the form

$$\langle \mathcal{M}, \rightarrow, D_1, \ldots, D_m \rangle$$

$\mathcal{M}$: A first order model which describes the situation talked about.

$\rightarrow$: A relation between objects and subjects, or subjects and subjects, where a subject is a pair $\langle a, u \rangle$ of a participant a and a discourse referent u.

D_a: A simple DRS which represents the discourse which has taken place so far.

Hence, a DRS D_a will have the form:

Participant a
$\mathcal{U}_{D_a}$
Con_{D_a}
Previous Dialogue Acts

We write u^a for $\langle a, u \rangle$. A (partial) representation for the situation in Example **(3)** after a) then looks as follows:

$\mathcal{M}$	$\to$	Anna *An*	Debra *De*	Bob *Bo*	Chris *Ch*
Mel	$Mel \to u_1^{An}$	u_1	u_1	u_1	u_1
Ste	$Ste \to u_1^{De}$	girl(u_1)	girl(u_1)	girl(u_1)	girl(u_1)
	$u_1^{An} \to u_1^{Bo}$				
	$u_1^{An} \to u_1^{Ch}$				

But these structures do not allow to represent what agents know about other agents, especially, there is no explicit representation of the common ground. For subsequent dialogues, it will be necessary for the involved persons to keep track with whom they share which referent. If an agent a interprets a definite description def φ, then he has to search a DRS D_a for a discourse referent x such that it is the only one such that $\varphi(x)$ is an entry in Con_{D_a}. There needs to be no real object such that it is common information that this object has the property φ. This example makes clear where there is a coordination problem. All participants can search only their own DRS. Hence, they will pick out the same object, only if their DRSes contain the same referents with the same properties, and if these referents are chained to the same objects. But this means that they have to be more or less identical. Therefore, we assume that all participants maintain in addition for each group H a DRS where they store the information which they share with this group. I.e. we can represent a local state by a function with arguments $H \subseteq$ DP and DRSes as values. In addition, the local state must relate the discourse referents in the shared DRSes to the referents in the DRS representing the *private* knowledge.

The following figure describes (part of) the local state of Bob in Example 3 after b) his talk with Debra. The first column represents his total knowledge about the biting situation, the second his protocol for what he heard in common with Anne and Chris, and the third for what he has in common with Debra.

Bo	$\{Bo, An, Ch\}$	$\{Bo, De\}$
u_1, u_2, u_3, u_4	u_1, u_2	u_1, u_2
Dob(u_1)	Dob(u_1)	Dob(u_1)
girl(u_2)	girl(u_2)	girl(u_2)
bit(u_1, u_2)	bit(u_1, u_2)	bit(u_1, u_2)
Dob(u_3)		
girl(u_4)		
bit(u_3, u_4)		

It has to be clear how the discourse referents in the second and third column are connected to the referents in the first one. u_1 and u_2 in the first have been introduced by the same assertion as u_1 and u_2 in the second one, whereas u_1 and u_2 in the third one have been introduced together with u_3 and u_4. Hence, in a full representation we have to add functions $l_{Bo,H} : \mathcal{U}_{D_{Bo,H}} \longrightarrow \mathcal{U}_{D_{Bo,\{Bo\}}}$ to each column. We can assume that the group–DRSes are updated in the same way as the private DRSes except for the fact that only shared information is added. Then, if all group members do it in the same way, then we can show in general that the DRSes $D_{a,H}$, which are internal representations of agent a for the referents and conditions which are common for the group H, are identical for all members $a, b \in H$[11]. Moreover, identical discourse referents are chained to the same objects. If Bob meets Chris, then he can apply the uniqueness condition which is connected to the definite description *the girl* to his DRS in the second column. He finds there u_2 as the only referent which represents a girl. As Chris will have the same representation for the common DRS, they both will interpret the description as relating to a subject which is *chained* to Melanie. But if Bob meets Debra, his use of *the girl* will single out the referent u_2 in the third column,

[11]How this can be spelled out in detail can bee found in [1, Section 7]. Of course, there is some freedom for how to do it. E.g. identity requires that all participant choose the same referent for the same object. This can be guaranteed by suitable update rules, but one can imagine other rules where strict identity is lost.

which is chained to the girl Stephanie. The identity of the $D_{a,H}$'s explains why the group can coordinate their interpretation of referentially used definites. It makes also clear in which sense the group–DRS is a *common* DRS. If the overall update rules, which must be specified in the multi–agent system, guarantee that all members maintain the same DRS, then it follows that it is also common information that these DRSes are identical. And, moreover, it allows to identify the *common ground for the group H in a dialogue D* with this shared DRS. Hence, these observations suggest that:

- Success of coordination is defined relative to common information.
- There are fixed rules in a community for how to update the common DRS.

The last condition is at least useful to guarantee that all members of a certain group maintain the same DRS for the common knowledge[12]. Now what is the relation between common information and a common DRS? Of course, this depends very much on the update rules. But in general a DRS will be a more fine grained representation for the common ground as it is not closed under logical consequences. On the other side, the common information contains all pragmatic information, and context information — like the referential chains. Again, it depends on the actual update rules whether this information is also represented in a group–DRS. We may expected that the explicitly represented information is at least weaker than the common information, or, that the rules allow to introduce a condition φ in the common DRS only if it is common information that φ is true. We introduce the Muddy Children Puzzle[13] as an example which supports this view, and the Time Imprecision Problem[14] as an example which shows that it is problematic.

4 Examples II: Muddy Children and Time Imprecision

The following example shows that it is unrealistic to assume that dialogue participants explicitly represent all common information. This holds, even if we assume that the DRSes are closed under logical consequences (for the non–modal part). The problem is the reasoning about each other.

(4) Imagine n children playing together. The mother of these children has told them if they get dirty there will be severe consequences. So, of course, each child wants to keep clean, but each would love to see the others get dirty. Now it happens during their play that some of the children, say k of them, get mud on their foreheads. Each can see the mud on others but not on his own forehead. So, of course, no one says a thing. Along comes the father, who says, "At least one of you has mud on your forehead," thus expressing a fact known to each of them before he spoke (if $k > 1$). The father then asks the following question, over and over: "Does any of you know whether you have mud on your own forehead?" Assuming that all the children are perceptive, intelligent, truthful, and that they answer simultaneously, what will happen?

There is a "proof" that the first $k-1$ times he asks the question, they will all say "No," but then the k^{th} time the children with muddy foreheads will answer "Yes."[15]

If the father states that at least one child has a muddy forehead (φ_1), then this becomes common knowledge for the whole group. If all children maintain a DRS which represents the common ground of the group, then the update rules should be formulated such that this information becomes part of the DRS. Before, it was only part of the private beliefs. If then the father asks his question and all children answer with "No," then it becomes common information that there are at least two children with muddy foreheads (φ_2). This presupposes, of course, that all children are able to conclude as follows: *If there are at least $n+1$ children with muddy foreheads, and I can perceive only n children with mud on their foreheads, then I must have a muddy forehead.* If we

[12] Of course, this is not a necessary condition.
[13] [4, Sec. 1.1]
[14] [4, Sec. 11.3]
[15] [4, p. 4]

furthermore assume that all common information becomes explicit common knowledge, then we can formulate a rule which allows to introduce (φ_2) into the common DRS. By induction it follows that every time the children answer with "No" they have to add a condition for *there are at least* $n+1$ *children with muddy foreheads*. But, of course, children do not update in this way. Hence, they can conclude that they have mud on their own forehead if they can see only one other child with mud and all children have answered with "No." But as there is no guarantee that others are equally intelligent they can't infer that φ_2 becomes part of the common ground.

It is not surprising to find an example where the explicitly represented information is really weaker than the common information. But is there an example where the update rules introduce a condition φ in the common DRS although it is not common information that φ is true?

The following example is a version of the coordinated attack problem, Example **(1)**. We assume that the general of the first division can send simple messages of the form *Simultaneous attack starts at time t*! We further assume that it is certain that the message will be delivered to the general of the second division within a certain interval of time. Hence, there is an $\epsilon > 0$ such that it is certain that for all sending times t_0 the message reaches the addressee at some time t_1 where $0 < t_1 < t_0 + \epsilon$. Furthermore, we assume that neither the general of the first division can know when exactly the message reaches the other general, nor does the general of the second division know when it was sent. It is possible to prove that it will never be common information that the message was successfully delivered, and therefore, that no general should attack because he can not be certain that the other one attacks simultaneously. The *proof* works as follows: As the general (A) of the first division does not know when the message reaches the general (B) of the second division, he has to wait until $t_0 + \epsilon$ until he can be certain that it is delivered. But then the general B must wait until $t_1 + \epsilon$ until he can be sure that general A knows that the message is delivered. Hence, general A has to wait for $t_0 + 2\epsilon$ until he can know that B knows that A knows that the message is delivered, etc. It follows that it will take infinite time until it is common knowledge that the message is delivered[16]. It is crucial that the sending time, and the time of delivery are not commonly known. But one would expect that both generals can be sure that they simultaneously attack at the specified time. This may be the case because we naturally expect a constraint which guarantees that the first general sends a message only if he is sure that it must be delivered before time t, i.e. $t_0 + \epsilon < t$. In this case they have at time t common information that the message is delivered, and they can simultaneously attack. But we can strengthen the example.

Assume that it is not a time of attack which is delivered but that the speaker introduces just a discourse referent by use of an indefinite NP (lets say *a girl*) with specific reference, and that he subsequently wants to refer to this object by use of a definite description (*the girl*). If this should be successful, it must be common knowledge that the referent has really been introduced. If one of the two interlocutors does not know that there is a girl, he will not be able to interpret *the girl*. But if we assume, as it is natural, that the delivery of the speakers assertion takes some time $\epsilon > 0$, and if the time of asserting and delivery are not commonly known (e.g. in a telephone call), then the same reasoning as in the Time Imprecision Problem shows that it can never become common knowledge that both interlocutors know of a *girl*. And, there is no time t such that it is common knowledge that the message was sent before $t - \epsilon$.

Now, we assume that both interlocutors maintain DRSes which represent the linguistic common ground, and that the update rules are such that (1) the speaker simultaneously with the use of *a girl* at t_0 introduces a condition of the form $girl(u_n)$ into his common DRS, and that (2) the addressee introduces the same condition into his common DRS as soon as he recognises the use of *a girl* at a time of delivery t_1. t_1 varies for all possible courses of events but the speaker has the information that there will be a time when his message reaches the addressee. Hence, he has information that the addressee will introduce a condition of the form $girl(u_n)$. On the other side, the hearer has at time t_1 the information that the speaker introduced at some sending time t_0 a condition of the form $girl(u_n)$ in his representation of the linguistic common ground. It follows

[16]This is the so-called *Time Imprecision Problem*, see [4, Ch. 11]. They discuss there also various solutions to this problem within modal approaches.

that it becomes at time t_1 common information that in all possible courses of events there are at the time of delivery[17] discourse referents in the common DRSes of both interlocutors. Hence, at the same time it becomes common information that the speaker can refer to this referent with the description *the girl*[18]. If ϵ is very small, then such a update rule is also natural[19]. The important point for our discussion lies in the fact that neither update is justified by common information. It only works because both participants adhere to these update rules, and they guarantee that the coordination problem for the use of a subsequent definite description is solved. But this means that there are natural examples where the update rules introduce conditions φ into the common DRSes where it is *not* common information that φ holds.

5 Conclusions

It is often assumed that the linguistic common ground can be defined as the set of all worlds which some participant believes to be possible, or where some participant believes that another participant beliefs it to be possible, or where someone believes that somebody else believes that some participant believes that the world is possible, etc. But if we look at examples like specificity in dialogue then we can see that we have to acknowledge for contextual information about the dialogue itself like referential chains. Hence, we defined *common information* as the set of all *dialogues* which are possible according to the beliefs of participants, or according to beliefs about each other.

On the other side, for the interpretation of anaphoric definite descriptions, where it is assumed that their interpretation depends on a uniqueness condition, we have argued that we need an *explicit* representation for the common ground with discourse referents. We've introduced a fragment where this is a DRS–like syntactic structure.

We have seen that the contents of the two representations for the common ground are not definable by one another. Especially, there are examples where the explicit representation is informationally weaker (Muddy Children), and examples where it is stronger (in some versions of the Time Imprecision Problem) than the common information. Update rules play here a crucial role.

We have seen examples which show that we need both, common information and common DRSes in order to explain coordination. E.g. in case of the anaphoric use of a definite description we saw that the uniqueness condition applies to *discourse referents*, not to real objects. Hence we need common DRSes for their interpretation. But in order to explain why it is guaranteed that the interlocutors refer to the same object we had to look at the common information. We have seen that the last problem is essentially a coordination problem.

[17]This *time of delivery* varies form dialogue to dialogue. But in every possible dialogue there is a time of delivery, and it is known to the interlocutors that at this time they both have introduced a suitable discourse referent which can serve as antecedent for a definite description. This becomes *common information* at the actual time of delivery because then the interpreter learns that the speaker has introduced an object with *a girl*.

[18]If, of course, it is also the only girl which was introduced.

[19]If delivery takes some time, e.g. if we send a letter, then the speaker must wait until he can be sure that the hearer has updated.

Bibliography

[1] A. Benz (December 2000): *Referential Chains and the Common Ground;* Ms., available from http://www.anton-benz.de.

[2] H.H. Clark (1996): *Using Language;* Cambridge.

[3] P. Dekker (1997): *On First Order Information Exchange;* in: A. Benz & G. Jäger (ed.), *Proceedings of the Munich Workshop on Formal Semantics and Pragmatics of Dialogue MunDial'97*, CIS Lecture Notes 106, München, pp. 21–39 .

[4] R. Fagin, J.Y. Halpern, Y. Moses, M.Y. Vardi (1995): *Reasoning About Knowledge;* MIT–Press, Cambridge, Massachusetts.

[5] J. Groenendijk, M. Stokhof (1991): *Dynamic Predicate Logic;* Linguistics and Philosophy 14, pp. 39–100.

[6] H. Kamp, U. Reyle (1993): *From Discourse to Logic;* Dordrecht.

[7] J.–J. CH. Meyer, W. van der Hoek (1995): *Epistemic Logic for AI and Computer Science;* Cambridge University Press, Cambridge.