Logical Structure and Discourse Anaphora Resolution*

Livia Polanyi
FX Palo Alto Laboratory
3400 Hillview Ave Bldg 4
Palo Alto CA 94304
polanyi@pal.xerox.com

Martin van den Berg
FX Palo Alto Laboratory
3400 Hillview Ave Bldg 4
Palo Alto CA 94304
vdberg@pal.xerox.com

Abstract

Working within the Dynamic Quantifier Logic (DQL) framework (van den Berg 1992, 1996a,b), we claim in this paper that in every language the translation into a logical language will be such that the preference ordering of possible discourse referents for an anaphor in a sentence can be explained in terms of the scopal order of the expressions in the antecedent that introduce the discourse referents. Since the scope of terms is derived from arguments independent of any discourse theory, our account explains discourse anaphora resolution in terms of general principles of utterance semantics, from which the predictions of centering theory follow. When combined with the powerful discourse structural framework of the Linguistic Discourse Model (Polanyi (1985, 1986, 1988, 1996) Polanyi and Scha (1984), Scha and Polanyi (1988), Prüst, H., R. Scha and M. H. van den Berg, 1994; Polanyi, L. and M. H. van den Berg 1996; van den Berg, M. H. 1996b), we provide a unified account of discourse anaphora resolution.

1 Introduction

In this paper, we use a semantic theory based on Dynamic Quantifier Logic (van den Berg 1992, 1996a,b) to present an approach to discourse anaphora resolution under the Linguistic Discourse Model (Polanyi (1985, 1986, 1988, 1996) Polanyi and Scha (1984), Scha and Polanyi (1988), Prüst, H., R. Scha and M. H. van den Berg, 1994; Polanyi, L. and M. H. van den Berg 1996; van den Berg, M. H. 1996b). Our treatment integrates the insights of the Centering framework (Joshi and Kuhn 1979, 1981; Grosz et.al. 1983, 1986, 1995; Gundel 1998; Walker et.al. 1998b) into a unified theory of discourse level structural and semantic relations. In our account, discourse level anaphora resolution effects fall out of a general theory of discourse quantification. Scope **orderings in the logical representation of the antecedent utterance result in differences in accessibility for potential referents in a target utterance.** No additional *centering mechanisms* are required, the centering predictions follow from this theory.

Our treatment is universal: explanations of relative coherence do not depend on conventions that might differ in different languages. Furthermore, we provide a treatment for the resolution of multiple anaphors, resulting from a range of possible antecedents including plurals and multiple antecedents.

The approach to discourse anaphora resolution we take in this paper integrates a rigorous formal semantic machinery within a theory of discourse structure. Before giving a detailed account of our treatment of discourse reference resolution, we would like to address explicitly some of the positions towards reference resolution and discourse structure which inform our work.

1.1 Theoretical and Methodological Considerations

To begin with, we should state explicitly that our enterprise is a semantic one: we are interested in developing and implementing a formalization capable of assigning a correct interpretation to each utterance in a discourse. In this, we are fully committed to the Dynamic Semantics enterprise (Kamp 1981, Heim 1982, Groenendijk and Stokhof, 1990, 1991, Chierchia 1992, van den Berg 1991, Kamp and Reyle 1993, Asher 1993, van den Berg 1996). Except in so far as it is provably necessary, we are not concerned with psychological issues of how human language users process discourse nor with what human beings intend when they use language to communicate with one another.

Our aim is to build machinery applicable to all genres and all modes of communication. Thus we can not assume that a discourse is necessarily "coherent" and that our goal is to provide an account of why that is so, nor can we assume that all discourse is written or spoken or occurs in a task context where the demands or reasonable expectations of an external activity are available to guide parsing and interpretation.

* The authors dedicate this paper to the memory of Megumi Kameyama (1953-1999), a dedicated researcher and a very dear friend.

Our theory is a formal one, therefore we can rely on well-known, rule-driven, parsing methods developed for sentences which allows us to parse discourse incrementally as it unfolds. In order to do so, our framework formalizes the relationship among constituent units in the discourse by specifying how antecedent units provide context for the interpretation of later units. In all cases, our method involves computing the resulting meaning of the combination of the meanings of the combined units, rather than identifying appropriate labels under which to characterize the relationship obtaining between the units. Our units of analysis are well defined semantic units. These units are usually encoded as single simple sentences or clauses but may also be realized by words, phrases or gestures which communicate exactly one elementary predication.

In our view, a formal theory of discourse structure should give well defined structures on which inferencing operates and on which world knowledge applies. We strive to limit the role of world knowledge in so far as possible to a specific moment in discourse processing —namely at the precise moment when a choice must be made about how a newly incoming unit must be integrated into the unfolding discourse. Just as in sentence grammar where world knowledge is used to decide between syntactically equivalent alternatives in the case of PP attachment, for example, in discourse grammar the relationships between elements are purely grammatical, and world knowledge is only used to decide between syntactically equally reasonable alternatives.

Similarly, in calculating the structure of discourse, we do not rely on the use of cue words such as *so*, *anyway* or *therefore*, because these terms are never obligatory. The relationship of one unit to another is always calculated on the relationship between the meanings of the constituent utterances which may then be reinforced by the presence of terms which specify the nature of the intended relationship.

In the framework developed below, there is a close relationship between discourse referents and discourse structure. **We deal both with how anaphors are resolved to particular antecedents using the structure of the discourse, and how an antecedent gives meaning to an anaphor.** The problem of identifying the antecedent to which an anaphor refers is dealt with in Centering Theory, discussed in Section 2 below. After reviewing Centering, we will discuss Dynamic Quantifier Logic (Section 3) and then show how the insights of Centering can be integrated into a general theory of discourse syntactic and semantic structure (Section 5.1), We shall point out how our approach accounts for multiple anaphors to different antecedents as well as accounting for anaphoric reference to multiple antecedents, a problem which remain unsolved within that framework (5.2).

2 Centering Theory

Centering Theory first described in detail in Grosz, Joshi and Weinstein (1983, 1986 [1995]) is designed to provide an assignment of a preference order among discourse entities in a sentence for the purpose of anaphora resolution. Centering Theory, which built upon earlier work by Joshi and Kuhn (1979) and Joshi and Weinstein (1981, 1998), proposed that (1) is perceived to be more coherent than (2) because in (1)

(1) (a) *Jeff helped Dick wash the car.* (b) *He_a washed the windows as Dick washed the car* (c) *He_b soaped a pane.*

He_a and *he_b* are both co-referential with *Jeff*, while in (2)

(2) (a) *Jeff helped Dick wash the car.* (b) *He_a washed the windows as Dick waxed the car* (c) *He_b buffed the hood.*

the referent for *He_b* in (c) is *Dick* while *he_a* in (b) refers to *Jeff*.

We quote here from the concise description of Centering given in (Walker et.al., 1998b):

> The centering model is very simple. Discourses consist of constituent segments and each segment is represented as part of a discourse model. Centers are semantic entities that are part of the discourse model for each utterance in a discourse segment. The set of FORWARD-LOOKING CENTERS, C_f(Ui,D) represents discourse entities evoked by an utterance Ui in a discourse segment D (Webber 1978; Prince 1981). The [unique] BACKWARD-LOOKING CENTER, C_b(Ui,D) is a special member of the C_f, which represents the discourse entity that the utterance U most centrally concerns. ...The C_b entity links the current utterance to the previous discourse. ...(or not more than one) ...The set of FORWARD-LOOKING CENTERS, C_f, is ranked according to discourse salience. This ranking is a partial order. The highest ranking member of the set of forward-looking centers ...represents a prediction about the $C_b$of the following utterance. Walker, Joshi, Prince (1998b) in Walker, Joshi, Prince 1998a henceforth WJP) p. 3.

From a linguistic perspective (cf. papers and references in Walker, Joshi and Prince 1998; Strube 1998), Centering theorists have explained the choice of C_b in a sentence in terms of a large number of potential factors. In particular: the grammatical hierarchy with subjects ranking higher than objects (Grosz, Joshi, Weinstein 1983), topic or empathy marking (Kameyama 1985), surface order position

(Rambow, 1993) or grammatical function (Brennan, Friedman and Pollard 1987) of the encoding of discourse entities in the immediately preceding segment.

Roberts (1998) argues that C_b is an unordered set of backward-looking centers in terms of classical Discourse Representation Theory notions of familiarity, compatibility and logical accessibility (Kamp 1981, Heim 1982, Kamp and Reyle 1993, Asher 1993), with an additional constraint that the set of discourse referents are attentionally accessible, a notion taken from Grosz and Sidner (1986). Under Roberts' treatment, the set of preferred centers, takes the place of the original C_b. Walker (1998) also replaces a unique C_b with a set of possible backward looking centers computed from a set of possible forward looking centers using agreement features, selection constraints of the verb and contra-indexing conditions.

The choice of segment also remains contested ground in Centering, with most linguists choosing for the sentence or clause while Walker (1998), argues for integrating Centering with a more global model of discourse focus. Within computational linguistics, several Centering Algorithms have been proposed, most notably by Brennan, S, M. Friedman and C. Pollard (1987), Walker, Iida and Cote (1990, 1994) and, more recently, by Strube and Hahn (1996), Strube (1998), and Walker (1998) which reflect these various perspectives.

Although the several variants of Centering can be argued to be better suited to one or another task or to account for phenomena in one or another language, they all fail to account for the interpretation of common examples such as (3)[1].

(3) (a) *Joan*^{1} *went to work at eight.* (b) *Bill*^{2} *arrived at nine.* (c) *They*$_{1+2}$ *met in the conference room.*

In (3), no entity in a single target clause or sentence resolves the plural pronoun in (3c). *They*$_{1+2}$ refers to a complex semantic entity created by combining entities in (3a) and (3b).

In the reformulation of Centering in terms of Dynamic Quantifier Logic presented in Section 3, below, we show how multiple anaphoric elements can be handled and each assigned its preferred resolution. DQL allows us to calculate a preference ordering on the discourse referents that can be used to account for multiple anaphors refering to different antecedents. When paired with the LDM, we also provide a means for one anaphor to refer back to multiple antecedents.

[1] Notational Convention: Introduced indices are written as superscripts; indices that are old (refer back) are written as subscripts.

3 Dynamic Quantifier Logic

DQL combines Generalized Quantifier Theory (GQT) (Barwise and Cooper 1991) and Plural Quantifier Logic (Scha 1981; van der Does 1992) with Dynamic Semantics. DQL was designed to handle phenomena such as plurals and complex relations between discourse referents often left unaddressed by other formal semantic frameworks (see van de Berg 1992,1996a,b).

Dynamic Quantifier Logic is based on the observation that NPs are generally anaphoric, quantificational and can be the antecedent of further anaphora, as illustrated by (4):

(4) (a) *The children*^{1} *arrived at the natural history museum early in the morning.* (b) *Three*$_1$ *boys*^{2} *disappeared in the gift shop.* (c) *They*$_2$ *had a great time touching almost everything.*

In (4b), *three boys* is anaphoric: its domain of quantification is given by *The children*. Within this domain, it is quantificational:, there are exactly *three boys* that *disappeared in the gift shop*. Finally, it is an antecedent: it introduces a referent picked up by *They*1 in (4c) to refer back to *the three boys*.

DQL, designed to explain examples like (4), was defined to preserve as far as possible the prediction of its precursors while inheriting most of their results. Under DQL well known, solid results and established procedures remain unchanged. As an illustration of a DQL representation of a sentence, take the simplified representation of (5b) below

(5) (a) *Some children*^{x} *were playing in the backyard.* (b) *Every*$_x$ *girl*^{y} *was wearing a hat*$_z$. (c) *They*$_y$ *had put them*$_z$ *on before they left the house.*

(5'b) $\forall y \subset x\,(\mathrm{girl}(y), \exists z \subseteq \bullet\,(\mathrm{hat}(z), \mathrm{wear}(y,z)))$

Formula (5'b) states that *for ever y entity that is a girl, taken from the domain given by the discourse referent x', it is the case that there is a hat such that she wears it.* This expression is very similar to classical translations into logic of (b). The only difference in the form of the expression is the explicit mention of the context set that sets the domain of quantification. These context sets are given by discourse referents. The universal quantification *Every girl* takes its range from the discourse referent x, and introduces a subset y, the indefinite *a hat* takes its domain from an as yet unspecified domain ($\bullet$).

3.1 Quantification and Reference Resolution

In DQL, all discourse anaphoric effects take place through discourse referents functioning as context sets to quantifiers. Variables that are quantified over[2] are introduced as discourse referents to function as context sets in subsequent sentences.

Although (5b) introduces both referents y for the *girls* and z for the *hats*, the referents do not have equivalent status. This is caused by the quantificational structure. The set of girls is given as a simple subset of the set of children, and as such is readily available. The set of hats, on the other hand, is only introduced relative to the set of girls. The hats are not introduced independently, but rather are introduced indirectly as belonging to the girls. Referring back to the set of hats is much more computationally expensive than referring back to the set of girls; to refer to the *hats* we must implicitly refer to the girls relative to which the set of hats is defined.

A consequence of the fact that the *hats* are introduced relative to the *girls*, is that there is an implied ordering of the discourse referents that we use in referring back to these sets. The discourse referent corresponding to the *girls* is much easier to pick up from the context than the discourse referent referring back to the *hats*[3]. Everything else being equal, the discourse referent referring to the *girls* will be preferred over the discourse referent referring to the *hats* because accessing it requires less computation.

This preference order corresponds closely to the forward-looking centers C_f. However, there is nothing in the construction of the preference ordering based on complexity of retrieval sketched above that would lead us to believe that there is at most one backward-looking center. In fact, our treatment gives the same predictions as Centering for the first pronoun resolved, but results in different predictions for embedded anaphors. The following diagram representing the scopes of (b) and (c) illustrates this:

(6)

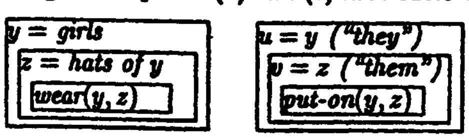

3.2 Anaphora Resolution Preference Order

It follows from the argument we have laid out above, that the referent y in (c), is preferred for anaphoric linking. However, once the girls are available as a set in (c) via u, the hats are also available, via discourse referent z, to serve as an antecedent. The set of girls, being already available no longer adds to the computational burden of calculating the set of hats. Within the scope of *they*, the referent z is much more accessible than outside that scope.

We can push this line of reasoning further. Consider example (7a). In this example, the subject, *Every woman*, has scope over the object, *a car*. As (7-8) shows, the preferred center of the C_f corresponding to this is the set of women because the cars are introduced as a function of the women. To refer correctly to the set of cars, we must also refer indirectly to the set of women since we are interested in retrieving only the cars owned by the women, not cars owned by men. On the other hand, to refer to the women, we need no information about their cars. This does not mean that we cannot refer to the cars in a subsequent sentence, as (9b) shows.

(7) (a) *Every woman in this town has a car.*
(b) *They park them in their garages.*

Where the set of women is referred to with *They*, *the cars* can be referred to directly. There is then no longer a hidden cost of retrieving the set of women in addition to *the cars*, since *cars* are already given in the sentence.

But now consider (8) and (9):

(8) (a) *Every woman in this town has a car.*
(b) *They use it to drive to work.*

(9) (a) *Every woman in this town has a car.*
(b) *They are parked in their garages.*

Note that (7–9) are decreasing in acceptability. (8) is more problematic than (7), because in (7) only the set of cars need be retrieved, while in (8) also the actual dependence of *the cars* on *the women* that own them is invoked by the use of the singular *it*[4]. (9) is much less acceptable than either (7) or (8), because in (9) *They* refers to the cars without the help of an explicitly given set of women.

The fact that once we have used a discourse referent, we can use other discourse referents that depend on it has important consequences as soon as we consider anaphora more complex than pronouns. Consider example (10).

(10) (a) *Seventeen people^1 in our lab have their own computers^2.* (b) *Three of them$_1$ are silly and switch them$_2$ off every night.*

In (10a), a discourse referent d_1 to a set of seventeen people is introduced, and as well as a discourse referent d_2 to a the set of computers they own, which depends on d_1. In (10b), *Three of them* quantifies over the domain given by d_1, and states that within d_1, there are exactly three people who switch their

[2]Like y and z in (5′b).

[3]This is related to the discussion in Joshi and Weinstein 1998, which motivates Centering from the perspective of complexity of inference in discourse.

[4]For some people (8) is totally impossible, because they demand a plural here as in (7), seemingly preferring semantic number agreement over syntactic number agreement. However, syntactic agreement does occur, as the following example illustrates:

Every soldier is responsible for his own gun. He has to clean it and will be reprimanded if any dirt is found on it.

own computers off every night. If the discourse referent introduced by their own computers would simply refer to the set of computers owned by people in the company, and not be dependent on the people, *them*$_2$ would refer to this set, rather than only to the set of computers owned by the three people. The meaning of (10b) would then be that these three people switch off all computers in the company, not just their own. This, of course, in not the correct reading.

4 Quantifier Scope and Anaphora Resolution

Under our analysis, the preferred antecedent for a pronoun is based on computational complexity arising from universal facts of scope ordering in the logical representation of the antecedent utterance. Different approaches to centering will be better or worse at predicting ordering relations depending on the match between the ordering scheme decided upon and the underlying scopal ordering.

We argue as follows.

If the discourse referent A is introduced by a term that has scope over a term introducing discourse referent B, and discourse referent B is introduced by a term that has scope over discourse referent C, A will be preferred over B and B will be preferred over C. Since this explanation is not dependent on conventions that might be different in different languages our treatment is universal. This is not the case for explanations based on linear ordering of syntactic constituents or arguments based on grammatical function, for example. Because in English the subject has scope over the objects, and the objects have scope over more deeply embedded terms, the ordering of discourse referents familiar to us from the literature will result in the well known C_f predictions.

Rejecting a preferred ordering for a less preferred ordering is a computationally complex operation. First the preferred order is computed, then this analysis is rejected —perhaps on pragmatic grounds. The calculations must then be re-done and the resulting less preferred ordering checked to see if it fits the pragmatic facts of the situation described in the target utterance. Differences in computational complexity arising from rejecting more preferred interpretations for less preferred thus result in the judgments of relative coherence which have been noted in the literature. Our account thus explains how Centering effects originate and why some anaphoric choices may involve more attention to the referent retrieval process than others[5].

4.1 Acceptability Predictions

To return then to examples (1) and (2), reproduced here as (11) and (12)

(11) (a) *Jeff helped Dick wash the car.* (b) *He*$_a$ *washed the windows as Dick washed the car* (c) *He*$_b$ *soaped a pane.*

(12) (a) *Jeff helped Dick wash the car.* (b) *He*$_a$ *washed the windows as Dick waxed the car* (c) *He*$_b$ *buffed the hood*

Since the discourse referent *Jeff* is introduced by a term that has scope over a term introducing discourse referent *Dick*, *Jeff* will be preferred over *Dick*. The difference in perceived coherence between (1/11) and (2/12) falls out of the more general fact that wide scope quantifiers are preferred over narrow scope quantifiers.

We will now turn to discussing how discourse structure and Anaphora Resolution interact to produce different acceptability predictions for different structures of discourse.

5 Discourse Structure and Anaphora Resolution

Although Centering Theory is associated with the Discourse Structures Theory of Grosz and Sidner (1986) which considers speaker intention and hearer attention as the critical dimensions to be modeled in discourse understanding, there are alternative models for understanding the relations among utterances in a discourse which are based on other principles. In particular, Dynamic Quantifier Logic, the anaphora resolution mechanism based on quantifier scope we are working with here, has been designed to provide the semantic machinery for the Linguistic Discourse Model (LDM). The LDM provides an account for discourse interpretation in terms of structural and semantic relations among the linguistic constituents making up a discourse[6].

5.1 The Linguistic Discourse Model

The LDM is designed as a discourse parser designed to construct a meaning representation of the input discourse icrementally. The LDM treats a discourse as a sequence of *basic discourse units* (BDUs)

[5]The DQL formalism has been explicitly designed to look as similar as possible to well-known, standard logics. To argue about issues of accessibility of the referents, a logical system that is less natural, but externalizes the dependencies between ranges of values for referents might be more suitable, such as van der Does' E-type logic (1993), a logical system very similar to DQL. we thank one anonymous reviewer for pointing out the work of Ranta (1991), who's use of Martin-löf's type theory may also be suitable as an analysis tool.

[6]In Prüst, Scha and van den Berg 1991, a resolution mechanism for unification based discourse grammar for verb phrase anaphora is defined, in terms of the Linguistic Discourse Model (LDM; Polanyi and Scha 1984; Polanyi 1987, 1988, 1996), which takes semantic representations as input. This treatment was later extended to a unification based discourse grammar acting on dynamic quantifier logic in Polanyi 1996, van den Berg and Polanyi 1996. The current paper extends that work.

each of which encodes formal semantic, syntactic and phonological properties of either an elementary predication or a discourse function. Using rules of discourse wellformedness which specify how to compute the relationship between a BDU and the previous discourse, the LDM constructs a parse tree by successively attaching the BDUs to a node at the right of edge of the emerging tree. The nodes of the tree are called *Discourse Constituent Units* (DCUs)[7]. DCUs encode formal semantic, syntactic and phonological properties that are calculated by following construction rules corresponding to the relationship computed as a result of the attachment process.

The discourse parse tree represents the structural relations obtaining among the DCUs. There are three basic types of relations among DCUs: *Coordination*, *Subordination* and *Binary Relation*. Corresponding to these relations, a DCU can be attached at a node on the right edge of a tree in one of three ways[8]:

1. The input DCU will be Coordinated with a node present on the right-edge of the tree if it continues a discourse activity (such as topic chaining or narrating) underway at that node.
2. The input DCU will be Subordinated to a node on the right-edge of the tree if it elaborates on material expressed at that node or if it interrupts the flow of the discourse completely.
3. The input DCU will be Binary-attached to a node if it is related to that node in a logical, rhetorical or interactional pattern specified explicitly by the grammar.

The LDM is a compositional framework. Simultaneous with the incremental construction of the structural representation of the discourse by attaching incoming DCUs, a semantic representation of the meaning of the discourse is constructed by incorporating the interpretation of an incoming DCU in the semantic representation on the discourse.

The LDM accounts for both structural and semantic aspects of discourse parsing using logical and structural notions analogous to units and processes constituting lower levels of the linguistic hierarchy. It is an ideal framework for understanding the relations between sentential syntax and semantics, on the one hand, and on the other hand, the texts and interactions that are constructed using sentential linguistic structures.

[7] BDUs once attached to the tree are DCUs.

[8] Besides these three basic composition relations between DCUs, a complex DCU can also be constructed by an operator having a DCU as an argument and within sentences, a DCU can occur embedded in another DCU. These two cases will not be discussed here.

5.2 Reference Resolution in the Linguistic Discourse Model

Let us now look at several short example of the interaction of anaphora resolution with discourse structure using the Dynamic Quantifier Logic framework above.

(13) (a) *Susan came home late yesterday.* (b) *Doris had held her up at work.* (c) *She needed help with the copier.*

In (13) the relationship between DCU (13a) and DCU (13b) is a Subordination relation because (13b) supplies more detailed information about why Susan came home late. As is shown in (13a), the S node inherits all information about the dominating DCU. In this case (a). A representation of Susan is therefore available at this constructed node. (13c) gives more explanation about what went on when Doris held Susan up at work and is therefore Subordinated to (b). Susan and Doris available for reference at that node. In (14) the situation is different.

(14) (a) *Susan came home late yesterday.* (b) *Doris had held her up at work.* (c) *She didn't even have time for dinner.*

(13′)
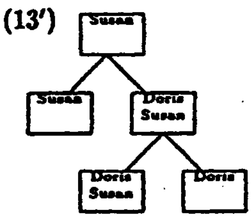

(14′)
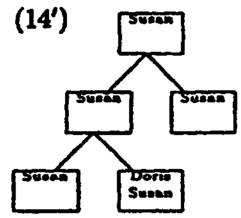

Although the relationship between DCU (14a) and DCU (14b) is a Subordination relation, as shown in (14a), as the discourse continues with (14c), the state of the discourse POPS from the embedded explanation to continue describing the state of affairs of Susan's evening. (14c) is therefore in a Coordination relation with (14a) as shown. Only Susan is now available as a potential referent in the current context.

In fact, the antecedent of an anaphora need not be one specific earlier utterance, but may be a constructed higher node in the parse tree as in (15):

(15) (a) *Joan went to work at eight.* (b) *Bill arrived at nine.* (c) *They met in the conference room.*

Q^{1+2}(ARRIVE(AT-TIME)) and They_{1+2}(meet-in-C)

Q^{1+2}(ARRIVE(AT-TIME)) They_X(meet-in-C)

Joan^1(got-to-work(at eight)) Bill^1(arrive(at eight))

In this case, the antecedent of (15c) is not (15a) or (15b), but the discourse node that constitutes the list (15a+b). In this higher node, there is a constructed schematic representation of what (15a) and (15b) share, and *They* is resolved to this. Very schematically, it amounts to resolving the anaphor X to the outer quantifier of its antecedent, Q^{1+2}.

6 Conclusions

Within our unified framework we are able to provide a detailed account of how anaphora resolution works across stretches of discourse, Because the LDM requires specific calculation of the information available at intermediary nodes. Computationally, during parsing, a rich data structure is created representing the meaning of the discourse. This, we would argue, is a distinct advantage of Dynamic Semantic approaches such as the LDM/DQL system over current computational alternatives such as Discourse Structures Theory (Grosz and Sidner 1989) and Rhetorical Structure Theory (Mann and Thompson 1987) which rely upon inferring the attentional and intentional states of language users, in one case, and on labeling the coherence relations among clauses, in the other. Looking towards formal discourse systems, we believe that while it would be possible to integrate the insights of DQL into a DRT approach such as that taken by Asher (1993), the approach taken here is computationally more tractable than more standard implementation of DRT for discourse parsing. The increased tractability results from the separation of discourse syntax and semantics which our approach imposes, taken together with the restriction of appeals to inference and world knowledge to specific moments in interpretation. In the case of the LDM, appeals to external knowledge are made only at the moment of DCU attachment to the parse tree.

References

Nicholas Asher. 1993. Reference to Abstract Objects in Discourse. Dordrecht, Kluwer.

John Barwise, and R. Cooper. 1981. Generalized Quantifiers and Natural Language *Linguistics and Philosophy* 4:159–219.

Susan E. Brennan, Marilyn E. Friedman and Carl J. Pollard. 1987. A Centering Approach to Pronouns. In: *Proceedings of the 25st Annual Meeting of the Association for Computational Linguistics*, Stanford CA 155–62.

Gennaro Chierchia, 1992. Anaphora and Dynamic Binding. *Linguistics and Philosophy* 15:111–183.

Jeroen Groenendijk, Theo M. V. Janssen and Martin Stokhof (ed) 1981. Formal Methods in the Study of Language. Amsterdam, Mathematical Centrum.

Barbara Grosz, Aravind Joshi and Scott Weinstein. 1983. Providing a Unified Account of Definite Noun Phrases in Discourse. Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics, Cambridge, MA 44–50.

Barbara Grosz, Aravind Joshi and Scott Weinstein. 1995. Towards a Computational Theory of Discourse Interpretation. Computational Linguistics, 21/2:203–25.

Barbara Grosz and Candace Sidner. 1986. Attention, intention, and the Structure of Discourse. Computational Linguistics, 12/3:175–204.

Jeanette K. Gundel. 1998. Centering Theory and the Giverness Hierarchy. In (Walker et.al., 1998a) 183–198.

Irene Heim. 1982. The Semantics of Definite and Indefinite Noun Phrases. Ph.d. Thesis. University of Massachusetts at Amherst.

Aravind Joshi and Steve Kuhn. 1979. Centered Logic: The Role of Entity Centered Sentence Representations in Natural Language Inferencing. In *Proceedings of the 6th International Joint Conference on Artificial Intelligence*. Tokyo, 435–9.

Aravind Joshi and Scott Weinstein. 1981. Control of Inference: Role of Some Aspects of Discourse Structure—Centering. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*. Vancouver 385–7.

Aravind Joshi and Scott Weinstein. 1998. Complexity and Control of Inference. In (Walker et.al., 1998a), 31–9.

Megumi Kameyama. 1985. Zero Anaphora: The Case of Japanese. Ph.D. Thesis. Stanford University.

Hans Kamp. 1981. A Theory of Truth and Semantic Representation. In (Groenendijk et.al., 1981)

William C. Mann and Sandra A. Thompson. 1987. Rhetorical Structure Theory. In G. Kempen (ed.) Natural Language Generation. The Hague, Mouton.

Per Martin-Löf. 1984. Intuitionistic type Theory Bibliopolis, Naples.

Livia Polanyi and Martin H. van den Berg. 1996. Discourse Structure and Discourse Interpretation. In *Proceedings of the Tenth Amsterdam Colloquium*. ILLC, Amsterdam.

Livia Polanyi and Remko Scha. 1984. A Syntactic Approach to Discourse Semantics. In *Proceedings of the 6th International Conference on Computational Linguistics*. Stanford CA.

Livia Polanyi. 1987. Keeping it all Straight: Interpreting Narrative Time in Real Discourse. WCCFL 6: 229–245.

Livia Polanyi. 1988. A Formal Model of Discourse Structure. In *Journal of Pragmatics* 12:601–638.

Livia Polanyi. 1996. The Linguistic Structure of Discourse. Stanford CA: CSLI Technical Report.

Hub Prüst, Remko Scha and Martin H. van den Berg. 1994. Discourse Grammar and Verb Phrase

Anaphora. Linguistics and Philosophy, 17:261–327.

Owen Rambow. 1993. Pragmatic Aspects of Scrambling and Topicalization in German. Paper presented at the Workshop in Centering Theory, Istitute for Research in Cognitive Science, University of Pennsylvania, Philadelphia. PA.

Aarne Ranta. 1991. Intuitionistic categorial Grammar. Linguistics and Philosophy, 14:203–239.

Craige Roberts. 1998. The Place of Centering in a General Theory of Anaphora Resolution. In (Walker et.al., 1998a), 359–400.

Remko Scha. 1981. Distributive, Collective and Cumulative Quantification. In (Groenendijk et.al., 1981).

Michael Strube. 1998. Never Look Back: An Alternative to Centering. In: Coling-ACL '98: Proceedings of the 17th International Conference on Computational Linguistics and the 36th, Annual Meeting of the Association for Computational Linguistics. Montreal, Quebec, Canada, Aug 10–14, pp.1251–1257.

Michael Strube and Udo Hahn. 1996). Functional Centering. Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics, Santa Cruz, CA 270–7.

Martin H. van den Berg. 1992. Dynamic Generalized Quantifiers. In Does J. M. v. d. and Eijck J. v.: Quantifiers, logic and Language. CSLI Lecture Notes 54, Stanford CA.

Martin H. van den Berg. 1996a. The Internal Structure of Discourse. Ph.D. Dissertation. ILLC, University of Amsterdam.

Martin H. van den Berg. 1996b. Discourse Grammar and Dynamic Logic. Proceedings of the Tenth Amsterdam Colloquium, ILLC, Amsterdam.

Jaap van der Does. 1991. Applied Quantifier Logics collectives and Naked Infinitives. Ph.D. thesis. University of Amsterdam.

Jaap van der Does. 1993. Formalizing E-type Anaphora. Proceedings of the Ninth Amsterdam Colloquium, ILLC, Amsterdam.

Marilyn Walker, Aravind Joshi and Ellen Prince (ed). 1998a. Centering Theory in Discourse. Oxford. Clarendon Press.

Marilyn Walker, Aravind Joshi and Ellen Prince. 1998b. Centering in Naturally Occurring Discourse: An Overview. In (Walker et.al., 1998a) 1–29.

Marilyn Walker. 1998. Centering, Anaphora Resolution, and Discourse Structure. In (Walker et.al., 1998a), 359–400.