

Action at a distance: the difference between dialogue and multilogue

Jonathan Ginzburg and Raquel Fernández

Dept of Computer Science

King's College, London

The Strand, London WC2R 2LS

UK

{ginzburg, raquel}@dcs.kcl.ac.uk

Abstract

The paper considers how to scale up dialogue protocols to multilogue, settings with multiple conversationalists. We extract two benchmarks to evaluate scaled up protocols based on the long distance resolution possibilities of nonsentential utterances in dialogue and multilogue in the British National Corpus. In light of these benchmarks, we then consider three possible transformations to dialogue protocols, inspired by Goffman's audience taxonomy and formulated within an issue-based approach to dialogue management. We show that one such transformation yields protocols for querying and assertion that fulfill these benchmarks. We indicate how these protocols can be implemented in terms of conversational update rules.

1 Introduction

Dialogue—two person conversation—is by now a topic with an ever increasing theoretical, corpus-based, and implementational literature. In contrast, the study of *multilogue*—conversation with 3 or more participants—is

still in its early stages. *The* fundamental issue in tackling multilogue is: to what extent do mechanisms motivated for dialogue (e.g. information states, protocols, conversational rules etc) scale up directly to multilogue?

There are of course various plausible views of the relation between dialogue and multilogue. One possible approach to take is to view multilogue as a sequence of dialogues. Something like this approach seems to be adopted in the literature on communication between autonomous software agents. However, even though many situations considered in multiagent systems do involve more than two agents, most interaction protocols are designed only for two participants at a time, perhaps in parallel. See e.g. the protocol specifications provided by FIPA (FIPA, 2003). Modelling of obligations and grounding becomes more complex when considering multilogue situations. The model of grounding implemented in the Mission Rehearsal Exercise (MRE) Project (Traum and Rickel, 2002), one of the largest multilogue systems developed hitherto, derives from the one designed by (Matheson, Poesio, and Traum, 2000) for dialogue and can only be used in cases where there is a single initiator and responder. It is not clear what the model should be for multiple addressees: should the contents be considered grounded when any of the addressees has acknowledged them? Should evidence

of understanding be required from every addressee?

Since their resolution is almost wholly reliant on context, non sentential utterances (NSUs) provide a large testbed concerning the structure of both dialogue and multilogue. In section 2 we present data from the British National Corpus (BNC) concerning the resolution of NSUs in dialogue and multilogue. The main focus of this data is with the distance between antecedent and fragment. We use this to extract certain benchmarks concerning multilogue interaction. In section 3 we sketch the basic principles of issue based dialogue management which we use as a basis for our subsequent investigations of multilogue interaction. This will include information states and formulation of protocols for querying and assertion in dialogue. In section 4 we consider three possible transformations on dialogue protocols into multilogue protocols. These transformations are entirely general in nature and could be applied to protocols stated in whatever specification language. We evaluate the protocols that are generated by these transformations with reference to the benchmarks extracted in section 2. Finally, in section 5 we discuss how these protocols can be implemented in terms of conversational update rules.

2 Long Distance Resolution of NSUs in Dialogue and Multilogue: some benchmarks

The work we present in this paper is based on data extracted from the British National Corpus (BNC). Our current corpus is a sub-portion of the BNC conversational transcripts consisting of 14,315 sentences. The corpus was created by randomly excerpting a 200-speaker-turn section from 54 BNC files. Of these files, 29 are transcripts of conversations between two dialogue participants, and 25 files are multilogue transcripts. A total of 1285 NSUs were found in our sub-corpus, 709

in dialogue and 576 in multilogue. All NSUs encountered within the corpus were classified according to the NSU typology presented in (Fernández and Ginzburg, 2002). Additionally, the distance from their antecedent was measured. Although the proportion of NSUs found in dialogue and multilogue is roughly the same, when taking into account the distance of NSUs from their antecedent, the proportion of long distance NSUs in multilogue increases radically: the longer the distance, the higher the proportion of NSUs that were found in multilogue. These differences are significant ($\chi^2 = 62.24$, $p \leq 0.001$). In fact, as Table 1 shows, NSUs that have a distance of 7 sentences or more appear exclusively in multilogue transcripts:

| | 1 | 2 | 3 | 4 | 5 | ≥ 6 |
|--------|--------------|------------|--------------|-------------|------------|------------|
| Dia. | 658 (93%) | 37 (5%) | 11 (1.5%) | 1 (1.5%) | 1 (.1%) | 1 (.1%) |
| Multi. | 467 (81%) | 45 (8%) | 15 (3%) | 8 (1.5%) | 6 (1%) | 35 (6%) |

Table 1: NSUs in dialogue and multilogue sorted by distance

Table 2 shows the distribution of NSU categories and their antecedent separation distance.¹ The classes of NSU which feature in our discussion below are boldfaced.

The last row in Table 2 shows the distribution of NSU-antecedent separation distances as percentages of the total of NSUs found. This allows us to see that about 87% of NSUs have a distance of 1 sentence (i.e. the antecedent was the immediately preceding sentence), and that the vast majority (about 96%) have a distance of 3 sentences or less.

¹The distance we report is measured in terms of sentence numbers. It should however be noted that taking into account synchronous speech would not change the data reported in Table 2 in any significant way, as manual examination of all NSUs at more than distance 3 reveals that the transcription portion between antecedent and NSU does not contain any completely synchronous sentences in such cases.

| NSU Class | Total | Distance | | | | | | |
|---|-------|----------|-----|----|-----|-----|-----|-----|
| | | 1 | 2 | 3 | 4 | 5 | 6 | >6 |
| Acknowledgment <i>Mmm.</i> | 595 | 578 | 15 | 2 | | | | |
| Short Answer <i>Ballet shoes.</i> | 188 | 104 | 21 | 17 | 5 | 5 | 8 | 28 |
| Affirmative Answer <i>Yes.</i> | 109 | 104 | 4 | | | 1 | | |
| Clarification Ellipsis <i>John?</i> | 92 | 76 | 13 | 2 | 1 | | | |
| Repeated Ack. <i>His boss, right.</i> | 86 | 81 | 2 | 3 | | | | |
| Rejection <i>No.</i> | 50 | 49 | 1 | | | | | |
| Factual Modifier <i>Brilliant!</i> | 27 | 23 | 2 | 1 | 1 | | | |
| Repeated Aff. Ans. <i>Very far, yes.</i> | 26 | 25 | 1 | | | | | |
| Helpful Rejection <i>No, my aunt.</i> | 24 | 18 | 5 | | 1 | | | |
| Check Question <i>Okay?</i> | 22 | 15 | 7 | | | | | |
| Filler <i>... a cough.</i> | 18 | 16 | 1 | | 1 | | | |
| Bare Mod. Phrase <i>On the desk.</i> | 16 | 11 | 4 | | | 1 | | |
| Sluice <i>When?</i> | 11 | 10 | 1 | | | | | |
| Prop. Modifier <i>Probably.</i> | 11 | 10 | 1 | | | | | |
| Conjunction Phrase <i>Or a mirror.</i> | 10 | 5 | 4 | 1 | | | | |
| Total | 1285 | 1125 | 82 | 26 | 9 | 7 | 8 | 28 |
| Percentage | 100 | 87.6 | 6.3 | 2 | 0.6 | 0.5 | 0.6 | 2.1 |

Table 2: NSUs sorted by Class and Distance

The data in table 2 highlights two significant generalizations about multilogue: the first concerns short answers. With a few exceptions, NSUs that have a distance of 3 sentences or more are exclusively short answers. Not only is the long distance phenomenon almost exclusively restricted to short answers, but the frequency of long distance short answers stands in strong contrast to the other NSUs classes; indeed, over 44% of short answers have more than distance 1, and over 24% have distance 4 or more, like the last answer in the following example:

- (1) Allan(1): How much do you think? Cynthia(2): Three hundred pounds. Sue(3): More. Cynthia(4): A thousand pounds. Allan(5): More. Unknown(6): <unclear> Allan(7): Eleven hundred quid apparently. [BNC, G4X]

It should be emphasized that long distance short answers is primarily a multilogue effect. Table 3 shows the total number of short answers found in dialogue and multilogue respectively, and the proportions sorted by distance over those totals. Note that only

18% of short answers found in dialogue have a distance of more than 1 sentence, with all of them having a distance of at most 3. This dialogue/multilogue asymmetry argues against reductive views of multilogue as sequential dialogue.

| Short Answers | Total # | 1 | 2 | 3 | > 3 |
|---------------|---------|----|----|---|-----|
| Dialogue | 54 | 82 | 9 | 9 | 0 |
| Multilogue | 134 | 44 | 11 | 8 | 37 |

Table 3: % over the totals found in dialogue and multilogue

The other striking generalization is the adjacency to their antecedent utterance of the remaining majoritarian classes of NSUs, Ack(nowledgements), Affirmative Answer, CE (clarification ellipsis), Repeated Ack(nowledgements), and Rejection. These are used either in grounding interaction, or to affirm/reject propositions.² The overwhelming adjacency to their antecedent underlines the locality of these interactions.

These data suggest two benchmarks protocols for multilogue need to satisfy:

- (2) a. **Multilogue Long Distance short answers (MLDSA)**: querying protocols for multilogue must license short answers an unbounded number of turns from the original query.
- b. **Multilogue adjacency of grounding/acceptance (MAG)**: assertion and grounding protocols for multilogue should license grounding/clarification/acceptance moves only adjacently to their antecedent utterance.

MLDSA and MAG have a somewhat different status: whereas MLDSA is a direct generalization from the data, MAG is a negative

²Acknowledgements and acceptances are, in principle, distinct acts: the former involves indication that an utterance has been understood, whereas the latter that an assertion is accepted. In practice, though, acknowledgements in the form of NSUs commonly simultaneously signal acceptance. Given this, corpus studies of NSUs (e.g. (Fernández and Ginzburg, 2002)) often conflate the two.

constraint, posited given the paucity of positive instances. As such MAG is more open to doubt and we shall develop alternatives to it in the sequel.³

3 Dialogue Protocols and Conversational Rules

In this section we outline some of the basic principles of Issue-based Dialogue Management (Ginzburg (1996, forthcoming), Larsson, 2002) which we use as a basis for our subsequent investigations of multilogue interaction. Following (Larsson, 2002; Cooper, 2004), our dialogue theory is formulated in Type Theory with Records (TTR). This allows simple interfacing with the grammar, which is a Constraint-based Grammar closely modelled on HPSG but formulated in TTR, rather than using typed feature structures. See (Ginzburg, forthcoming) for details.

Within this approach, each dialogue participant's view of the common ground, the dialogue gameboard (DGB), are records of the type given in (3). We will frequently find it useful to talk directly of the first element of the Moves and QUD lists, referring to them respectively as **LatestMove** and **MaxQUD**.

$$(3) \quad \left[\begin{array}{l} \text{facts : Prop} \\ \text{Moves : list(IllocProp)} \\ \text{QUD : list(Question)} \end{array} \right]$$

The querying/assertion protocols (in their most basic form) we assume for dialogue are summarized in Table 4.⁴

These protocols arise from the composition of *conversational (update) rules* akin to those introduced by (Larsson, 2002). A conversational rule is a mapping that specifies how one DGB configuration (the *preconditions*) can be

modified into another (the *effects*). Two conversational rules *part1*, *part2* can be composed if they satisfy $\text{preconds}(\text{part2}) \sqsubseteq \text{effects}(\text{part1})$.

| querying | assertion |
|---|--|
| LatestMove = Ask(A,q) | LatestMove = Assert(A,p) |
| A: push q onto QUD; release turn; | A: push p? onto QUD; release turn |
| B: push q onto QUD; take turn; make max-qud-specific; utterance ⁵ take turn. | B: push p? onto QUD; take turn; Option 1: Discuss p? Option 2: Accept p |
| | LatestMove = Accept(B,p) |
| | B: increment FACTS with p; pop p? from QUD; |
| | A: increment FACTS with p; pop p? from QUD; |

Table 4: Protocols for querying and assertion

Specifically, the conversational rules that give rise to the protocols in Table 4 are the following TTR formulated rules from (Ginzburg, forthcoming), which for reasons of space are stated here informally in English:

- (4) a. **QUD-Specificity (QSPEC)**: given $\text{MaxQUD} = q$, one can make an utterance which is q-specific.
- b. **Ask QUD Update**: given $\text{LatestMove} = \text{Ask}(A,B,q)$, q becomes QUD maximal
- c. **Assert QUD Update**: given $\text{LatestMove} = \text{Assert}(A,B,p)$, p? becomes QUD maximal
- d. **Accept**: given $\text{LatestMove} = \text{Assert}(A,B,p)$, B can make utterance such that $\text{LatestMove} = \text{Accept}(B,A,p)$.
- e. **UpdateFacts + DowndateQUD**: Given $\text{LatestMove} = \text{Accept}(B,p)$,

³We would like to thank an anonymous reviewer for Dialogor for strengthening our open mindedness regarding MAG.

⁴For reasons of space we do not formulate an explicit protocol for grounding here—the structure of such a protocol resembles the assertion protocol. Our subsequent discussion of assertion can be modified *mutatis mutandis* to grounding.

⁵An utterance whose content is either a proposition *p* about max-qud or a question *q*₁ on which max-qud Depends. For the latter see footnote 8. If one assumes QUD to be a stack, then ‘max-qud-specific’ will in this case reduce to ‘q-specific’. But the more general formulation will be important below.

conjoin p with FACTS, downdate p?
and all other qs from QUD resolved
by FACTS

NSU Resolution We assume the account of NSU resolution developed in (Ginzburg and Sag, 2000). The essential idea they develop is that NSUs get their main predicates from context, specifically via unification with the question that is currently *under discussion*, an entity dubbed the *maximal question under discussion* (MAX-QUD). NSU resolution is, consequently, tied to conversational topic, viz. the MAX-QUD.⁶

Dialogue short answers The QUD-based resolution strategy affords the potential for non adjacent short answers in *dialogue*, given the assumption that QUD is a stack. These, as discussed in section 2, are relatively infrequent. Two commonly observed *dialogue* conditions will jointly enforce adjacency between short answers and their interrogative antecedents: (a) Questions have a simple, one phrase answer. (b) Questions can be answered immediately, without preparatory or subsequent discussion. For multilogue (or at least certain genres thereof), both these conditions are less likely to be maintained: different CPs can supply different answers, even assuming that relative to each CP there is a simple, one phrase answer. The more CPs there are in a conversation, the smaller their common ground and the more likely the need for clarificatory interaction. A pragmatic account of this type of the frequency of adjacency in dialogue short answers seems clearly preferable to any actual mechanism that would *rule out* long distance short answers. These can be perfectly felicitous—see

⁶The resolution of NSUs, on the approach of (Ginzburg and Sag, 2000), involves one other parameter, an antecedent sub-utterance they dub the *salient-utterance* (SAL-UTT). This plays a role similar to the role played by the *parallel element* in higher order unification-based approaches to ellipsis resolution (see e.g. (Pulman, 1997)). For current purposes, we limit attention to the MAX-QUD as the nucleus of NSU resolution.

e.g. example (1) above which would work fine if the turn uttered by Sue had been uttered by Allan instead.

4 Scaling up Protocols

(Goffman, 1981) introduced the distinction between *ratified participants* and *overhearers* in a conversation. Within the former are located the speaker and participants whom she takes into account in her utterance design—the intended addressee(s) of a given utterance, as well as *side participants*. In this section we consider three possible principles of protocol extension, each of which can be viewed as adding roles for participants from one of Goffman's categories. The final principle we consider, Add Side Participants (ASP), seems to yield the best results, relative to the benchmarks we introduced in section 2. We state the principles informally as transformations on operational construals of the protocols and then in section 5 we indicate how such protocols could be implemented in terms of conversational update rules.

Add Overhearers (AOV) This involves adding participants who merely observe the interaction. They keep track of facts concerning a particular interaction, but their context is not facilitated for them to participate:

- (5) Given a dialogue protocol π , add roles C_1, \dots, C_n where each C_i is a silent participant: given an utterance u_0 classified as being of type T_0 , C_i updates C_i .DGB.FACTS with the proposition $u_0 : T_0$.

Applying AOV yields essentially multilogues which are sequences of dialogues. A special case of this are moderated multilogues, where all dialogues involve a designated individual (who is also responsible for turn assignment.). AOV will not allow for long distance short answers across more than two participants, as in e.g. (1), so will fail the MLDSA benchmark.

Duplicate Responders(DR)

- (6) Given a dialogue protocol π , add roles C_1, \dots, C_n which duplicate the responder role

Applying DR to the querying protocol in Table 4 yields a protocol in which each responder to A's query q gets to provide their input concerning q (i.e. a q -specific utterance). This yields interactions such as (7) from the BNC:

- (7) Anon (1) How about finance then? <pause> Unknown1 (2): Corruption. Unknown2(3): Risk <pause dur=30> Unknown3(4): Wage claims <pause dur=18>

Such a querying protocol licenses long distance short answers, so satisfies the MLDSA benchmark. On the other hand, the contextual updates it enforces will not enable it to deal with the following (constructed) variant on (7), in other words does not afford responders to comment on previous responders, as opposed to the original querier:

- (8) A(1): Who should we invite for the conference? B(2): Svetlanov. C(3): No (=Not Svetlanov), Zhdanov. D(4): No (= Not Zhdanov, \neq Not Svetlanov), Gergev

Applying DR to the assertion protocol will yield a protocol in which multiple responders get to sequentially accept an assertion. This will licence long distance acceptance and thus is inconsistent with the MAG benchmark. On the other hand, it is potentially useful for interactions where there is explicitly more than one direct addressee.

Add Side Participants (ASP) This is a principle intermediate between AOV and DR:

- (9) Given a dialogue protocol π , add roles C_1, \dots, C_n , which affect the same contextual update as the interaction initiator.

In terms of the protocols introduced in section 3, ASP involves the same protocols modified such that (a) the audience is a non-singleton, (b) one member of this audience instantiates the addressee role and responds, the others update their DGBs in similar fashion to the original speaker.

This will yield a protocol for assertion that satisfies the MAG benchmark in that acceptance is strictly local. This is because it enforces *communal acceptance*—acceptance by one CP can count as acceptance by all other addressees of an assertion. There is an obvious rational motivation for this, given the difficulty of a CP constantly monitoring an entire audience (when this consists of more than one addressee) for acceptance signals—it is well known that the effect of visual access on turn taking is highly significant (Dabbs and Ruback, 1987). It also enforces quick reaction to an assertion—anyone wishing to accept p or dissent from p must get their reaction in early i.e. immediately following the assertion since further discussion of p ? is not countenanced if acceptance takes place. The latter can happen of course as a consequence of a dissenter not being quick on their feet; on this protocol to accommodate such cases would require some type of backtracking.⁷

Applying ASP to the dialogue querying protocol yields a protocol that improves on the DR generated protocol because it does allow responders to comment on previous responders—the context is modified as in the dialogue protocol. Nonetheless, as it stands,

⁷In this respect an example pointed out by an anonymous Dialor reviewer is relevant; the reviewer suggests that ‘that a disagreement by one respondent need not precede acknowledgement by another. E.g. I don't think there's anything wrong with this dialogue:

(i) A: We're inviting Svetlanov. B: Right. C: No we're not.

We agree that the dialogue is fine. However, intuitively, it seems to us, (and indeed on any protocol in which acceptance does not itself require acceptance,) that C's move will potentially give rise to some sort of backtracking, at least from B. See below though for a version of acceptance, *distributed acceptance* that can accommodate such cases.

this protocol won't fully deal with examples such as (7)—the issue introduced by each successive participant takes precedence given that QUD is assumed to be a stack. This can be remedied by slightly modifying this latter assumption: we will assume that when a question q gets added to QUD it doesn't subsume *all* existing questions in QUD, but rather only those on which q does not depend:⁸

- (10) q is $\text{QUD}^{\text{mod}(\text{dependence})}$ maximal iff for any q_0 in QUD such that $\neg \text{Depend}(q, q_0)$: $q \succ q_0$.

This is conceptually attractive because it reinforces that the order in QUD has an intuitive semantic basis. The effect of this will be to ensure that any polar question $p?$ introduced into QUD, whether by an assertion or by a query, subsequent to a wh-question q on which $p?$ depends does not subsume q . Hence, q will remain accessible as an antecedent for NSUs, as long as no new unrelated topic has been introduced. Assuming this modification to QUD is implemented in the above ASP-generated protocols, both MLDSA and MAG benchmarks are fulfilled.

5 Conversational Rules for Multilogue

In this section we consider how the protocols scaled up according to the principles ASP and DR discussed in section 4 can be compositionally decomposed from conversational rules akin to those in (4).⁹ QSPEC does not require any modification—once a question q is pushed on QUD, licensing a q -specific utterance is characteristic of both querying and assertion protocols.

⁸ The notion of dependence we assume here is one common in work on questions, e.g. (Ginzburg and Sag, 2000), intuitively corresponding to the notion of 'is a subquestion of'. q_1 depends on q_2 iff any proposition p such that p resolves q_2 also satisfies p is about q_1 .

⁹ Adding overhearers (AOV) involves no substantive change to the previously discussed protocols: AOV is already in the form of an update rule, which concerns solely the overhearers.

Adding Side Participants (ASP) involves one rather minor modification to the rules in (4). The illocutionary propositions that constitute the values of LatestMove in the various rules now need to have a plural set of individuals as their type. For instance:

- (11) **Ask QUD Update (plural audience):** given LatestMove = Ask(A, { C_1, \dots, C_n }, q), q becomes QUD maximal for { A, C_1, \dots, C_n }

- (12) **UpdateFacts + DowndateQUD (plural audience):** Given LatestMove = Accept(B, { A, C_1, \dots, C_n }, p), { B, A, C_1, \dots, C_n } conjoin p with FACTS, downdate $p?$ and all other q s from QUD resolved by FACTS

Pluralized QUD update rules are also components of DRed querying and assertion rules. Given the modification to QUD proposed in the previous section, a reasonably direct treatment of DRed querying follows: following a query q by A, Ask QUD update enables the next speaker to provide a q -specific answer. By the ordering in QUD, q will remain maximal for any subsequent speaker who has not downdated it.

The main additional modification seems to concern acceptance. Consider first the preconditions for an acceptance move—the difference from the dialogue case is that they no longer involve adjacency of the assertion in question. They now involve the combination of the existence of a prior assertion of p and the maximality of $p?$ in QUD:

- (13) **Distributed Accept:** given Moves = $\langle \dots \text{Assert}(A, \{C_1, \dots, C_n\}, p) \dots \rangle$, and Max-QUD = $p?$, C_i can make utterance such that LatestMove = Accept($C_i, \{A, C_1, \dots, C_n\}, p$).

It seems like fact-incrementation/QUD downdate needs to be divided into two sub-cases: one that concerns the addressees, the

other that concerns the original asserter. To take these in order: for the addressees, given the distributed nature of acceptance here, the precondition for fact-incrementation/QUD-downdate has to be an acceptance by that particular individual. For the original asserter the precondition for fact-incrementation/QUD-downdate is the existence of acceptance acts of p by all addressees:

- (14) **Distributed UpdateFacts + Downdate-QUD (audience version):** Given Latest-Move = Accept($C_i, \{A, C_1, \dots, C_n\}, p$), C_i conjoin p with FACTS, downdate p ? and all other q s from QUD resolved by FACTS
- (15) **Distributed UpdateFacts + Downdate-QUD (asserter version):** Given Moves = $\langle \dots \text{Assert}(A, \{C_1, \dots, C_n\}, p) \dots, \text{Accept}(C_1, \{A, \dots, C_n\}, p), \dots, \text{Accept}(C_n, \{A, C_1, \dots\}, p) \rangle$, A conjoin p with FACTS, downdate p ? and all other q s from QUD resolved by FACTS

6 Conclusions and Future Work

In this paper we have considered how to scale up dialogue protocols to multilogue. We have extracted two benchmarks, MLDSA and MAG, to evaluate scaled up protocols based on the long distance resolution possibilities of NSUs in dialogue and multilogue. In light of these benchmarks, we consider three possible transformations to protocols, which can be intuited as adding roles that correspond to different categories of an audience originally suggested by Goffman. We then indicate how such protocols could be implemented in terms of conversational update rules.

In the future we intend to implement multilogue protocols in CLARIE so it can simulate multilogue. We will then evaluate its ability to process NSUs from the BNC.

Acknowledgements

We would like to thank three anonymous Dialor reviewers for some very useful comments and challenges, which reshaped

our thinking on a number of significant issues. We would also like to thank Pat Healey, Shalom Lappin, Richard Power, and Matt Purver for discussion. Earlier versions of this work were presented at colloquia at ITRI, Brighton, and at the Université Paris, 7. The research described here is funded by grant number RES-000-23-0065 from the Economic and Social Research Council of the United Kingdom.

References

- H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge.
- R. Cooper. 2004. A type theoretic approach to information state update in issue based dialogue management. Invited paper, *Catalog'04, the 8th Workshop on the Semantics and Pragmatics of Dialogue*, UPF, Barcelona.
- J. Dabbs and R. B. Ruback. 1987. Dimensions of group process: amount and structure of vocal interaction. *Advances in Experimental Social Psychology* 20, pp. 123–169.
- N. Fay, S. Garrod, and J. Carletta. 2000. Group discussion as interactive dialogue or serial monologue. *Psychological Science*, pp. 481–486.
- R. Fernández and J. Ginzburg. 2002. Non-sentential utterances: A corpus study. *Traitement automatique des langues. Dialogue*, 43(2):13–42.
- FIPA. 2003. The foundation for intelligent physical agents. Interaction protocol specifications. <http://www.fipa.org>.
- J. Ginzburg and R. Cooper. 2001. Resolving ellipsis in clarification. In *Proc. of the 39th Meeting of the ACL*, Toulouse.
- J. Ginzburg and I. Sag. 2000. *Interrogative Investigations: the form, meaning and use of English Interrogatives*. CSLI Lecture Notes, 123. CSLI Publications.
- J. Ginzburg. (forthcoming). *Semantics and Interaction in Dialogue* CSLI Publications and U. of Chicago Press.
- J. Ginzburg. 1996. Interrogatives: Questions, facts, and dialogue. In Shalom Lappin, editor, *Handbook of Contemporary Semantic Theory*. Blackwell, Oxford.
- E. Goffman 1981 *Forms of Talk*. University of Pennsylvania Press, Philadelphia.
- S. Larsson. 2002. *Issue based Dialogue Management*. Ph.D. thesis, Gothenburg University.
- C. Matheson and M. Poesio and D. Traum. 2000. Modelling Grounding and Discourse Obligations Using Update Rules. *Proc. of NAACL 2000*, Seattle.
- S. Pulman. 1997. Focus and higher order unification. *Linguistics and Philosophy*, 20.
- M. Purver. 2004. *The Theory and Use of Clarification in Dialogue*. Ph.D. thesis, King's College, London.
- D. Traum and J. Rickel. 2002. Embodied agents for multi-party dialogue in immersive virtual world. In *Proc. of AAMAS 2002*, pp. 766–773.
- D. Traum. 2003. Semantics and pragmatics of questions and answers for dialogue agents. In *Proc. of the 5th IWCS*, pp. 380–394. University of Tilburg.