COMMUNICATING WITH MULTIPLE AGENTS*

Elizabeth A. Hinkelman and Stephen P. Spackman

Deutsches Forschungszentrum für Künstliche Intelligenz
Stuhlsatzenhausweg 3, 66123 Germany
hinkelman@dfki.uni-sb.de, stephen@acm.org

Abstract

Previous dialogue systems have focussed on dialogues between two agents. Many applications, however, require conversations between several participants. This paper extends speech act definitions to handle multi-agent conversations, based on a model of multi-agent belief attribution with some unique properties. Our approach has the advantage of capturing a number of interesting phenomena in a straightforward way.

Motivation

The rise of spoken language NLP applications has led to increased interest in such issues as real time processing and on-line error recovery. But dialogue is an *inherently* online process; this manifests in such linguistic phenomena as turntaking [Sacks *et al.*, 1974], repair [Schegloff *et al.*, 1977], and content grounding [Clark and Schaefer, 1989]. Grounding is the phenomenon that establishes shared beliefs, for which simply making or hearing an utterance does not suffice. It makes hearers into full participants who actively signal success or failure of communication, as in this exchange:

Steph: that's friday at seven then. |
Lynn: | at seven.

Our long term goal is to show how, *from the perspective of a participant*, one plans and acts in an environment with other communicating individuals, even when those other individuals are not perfectly reliable, and even when the groups involved may be large enough that it is impractical to model all participants. For example, consider this familiar exchange, from the point of view of someone who remembers that the next group meeting is on tuesday:

Jan : so we should drop the mtn . cancel the
meeting on thursday.
Les : | tuesday
ME : | tuesday
Lou : | yeah.
Jan : | yes, right.

Here both our subject and another participant offer a correction, which is confirmed by Lou and by the original speaker. Other participants may be present.

In this paper, we focus on the effects of communicative actions on the participant's model of the situation. In contrast with previous dialogue work, we are

interested in cases where there are more than three agents. In a group of ten or more, it is hard to imagine how a participant can track the beliefs and disbeliefs of others accurately; it may not even be practical to track who they all are.

The advantages of analysing natural language utterances as communicative actions are by now well understood; they serve to summarise conversations for long-term storage [Schupeta, 1993], as a basis for generation [Moore and Paris, 1989], in top-down prediction of utterance function and structure [Alexandersson *et al.*, 1994], and most importantly, to provide a representation of natural language utterances that is uniform with that used in general facilities for planning and action [Allen, 1983].

We follow [Traum and Hinkelman, 1992] in regarding speech acts as fully joint actions between conversational participants. Not only are joint speech acts co-operatively undertaken, but they have at least nominally joint effects: if they complete but still fail to result in shared goals or shared beliefs, this should be attributable to politeness, dishonesty (cf. [Perrault, 1990]), or other social functions.

This perspective on speech act processing forces us to deal with issues of jointly held goals and beliefs at a very basic level. These matters are by now quite well-studied, but analytic solutions from logical first principles tend to be relatively complex, yielding neither perspicuous notations nor plausible computational or cognitive models. In short, normative analyses are not necessarily descriptive ones.

Aside from relatively involved calculations, there are several sources of difficulty:

- when multiple participants are involved, the number of 'belief spaces' (patterns of modal embedding) tends to blow up rapidly;
- when the *actual* state of affairs regarding the extent of others' knowledge is unknown (as is the case for an online conversational participant) the number of cases to be considered can become large;
- when dealing with large organisations, some form of aggregate modelling becomes an absolute necessity.

Consider, for instance, the case in which you believe that the government knows your income from last year. What you believe is not that each individual government employee knows it, but that anyone from the tax department, whose business it is to know, and who actually wants to, will. Thus we would typically assume that an employee of the tax department who, in a professional capacity, makes contact with your accountant, would actually have this information to hand. We want to abstract away from the communi-

*The work underlying this paper was supported by a research grant, FKZ ITW 9002 0, from the German Bundesministerium für Forschung und Technologie to the DFKI project DISCO.

cation channels that make this possible while retaining the availability of the conclusion; and we would ideally like to do so in a manner consistent with the needs of online dialogue processing.

In the next section we describe a method of representing the information necessary for processing multiparticipant speech acts, one which treats groups as agents in their own right, and does so in such a way that speech acts can operate on them naturally and directly.

Corporate Agents and Attitude Propagation

The basis of our model is the *corporate agent*. These ‘agents’ represent groups, but, like individual agents, they may have beliefs (which we write B_{agentp}) and goals (G_{agentp}) ascribed to them directly. Thus, they can be thought of as intermediate in status between sets of component agents and prototypical group members. They differ from simple sets in three striking ways: first, they are intensional structures that may be distinct even when co-extensive (as, for example, when the members of the marketing department form a volleyball team on tuesday nights); second, attitudes are ascribed to them *directly*, and potentially at variance with any attitudes we might ascribe to their members, or other subagents (see the discussion section); and third, that (other than perhaps in the case of a ‘real’, singleton agent) subsethood is not a sufficient condition for subagency—some intramural volleyball teams are clearly, as social entities, agents of the company, and others rather the reverse.

While not in a position to make detailed psychological claims, we believe that structures of this kind are compatible with what we know about the linguistic and cognitive handling of aggregates in other contexts.

These corporate agents will be used to represent both long-term social groups and transient groups of conversational participants.

Figure 1: Agents

In the remainder of this paper we illustrate relationships between agents with diagrams such as that in figure 1. Here the playing-card shapes represent agents and the heavy lines connecting them the subagent relation ($\sqsubseteq$): the agents include the system itself (sec.Jan), the system’s boss (Jan), Jan’s office, their coworker

Les, their common corporate employer (WidgetCorp), and another ‘random’ person, Steph, who does not belong to WidgetCorp. Later we will represent attitudes ascribed to agents by small shapes within the playing cards and their propagation by thin curved arrows between them.

Note that since we are discussing the system’s own representation of the world, the double-bordered playing card really represents the system’s self-model and the whole diagram the system; but we shall not belabour the point here, assuming that information passes freely between the two.

The model we use to compute the transfer of attitudes between agents is approximate, partly to simplify computation and partly because it is in any case unusual to have enough empirical data to compute an exact solution. The same model (with parametric variation in the domain knowledge) is applied to all the attitudes the system ascribes to agents; in our present implementation, these may be either beliefs or goals.¹ Unlike representations based on conventional logics of belief, it does not introduce nested contexts unless they are explicitly represented in the content of utterances (as would be the case with a sentence like “But Lynn thinks we should have the meeting anyway.”), though extended reasoning processes may compatibly do so.

In this simplified model, the propagation of attitudes is performed lazily, as determined by their semantic relevance to the various agents involved. Ideally, this judgment would derive from social world-knowledge and information about the purposes of groups; in our current implementation it is approximated using a simple classification of corporate agents, participant roles and message topics into the domain ontology. Delays and chance are not modelled; all relevant attitudes are presumed to propagate between agents, subject to the following constraints:

Figure 2: Common context constraint

¹Since our model is not analytic we do not want or need a notion of ‘knowledge’: the system lacks direct access to empirical truth and to social consensus, and does not have the cognitive sophistication to validate arguments against some independent notion of rationality. In short, none of the usual theories of truth can in principle be made to apply.

1. Attitudes propagate *only* between superagent and subagent (or vice-versa). This stipulation amounts to saying that communication only occurs between agents in the presence of *some* common social context. Of course, new corporations can be introduced when new social groups form.

Thus in figure 2, beliefs ascribed to WidgetCorp can propagate to the subagents, Jan and Jan's electronic secretary; but they do not reach Steph, who is not a subagent of WidgetCorp.² We use the convention that attitudes are drawn filled in if they are known by direct evidence, hollow otherwise; and that dotted structures are 'negated'—they do *not* arise as drawn because they violate some constraint under discussion.

Figure 3: Propagation as default

2. Attitude propagation is only a *default* (the particular choice of default logic need not concern us here). If there is direct evidence for ascribing a contrary attitude to an agent, propagation from an external source is inhibited. This property is crucial to modelling dishonesty, negotiation, compromise, and atypical group members in general.

Such blocking is illustrated in figure 3. In this case our model of Jan *fails* to inherit a goal from office.Jan because it conflicts with another goal (the square box) for the ascription of which to Jan we have prior independent evidence.

3. The system may never assume that its *own* attitudes automatically propagate upwards to a superagent. The upward attitude propagation path models the effect of external agents independently attending to their various communicative goals, but the system must still plan—and execute—its own actions.

Thus, in figure 4 the system—sec.Jan—is prohibited from simply assuming that its own beliefs are shared by its employer, though those of fellow employees would be propagated when otherwise consistent. (Some humans seem to relax this constraint.)

²In order to place some limit on the promiscuity of attitude propagation, it seems best to insist that indirect transfer must occur through a single agent that is a transitive super- or sub- agent of both terminal agents. Thus, even if Jan and Steph both belonged to some peer of WidgetCorp with a similar semantic domain, propagation would still be not permitted along the resulting N-shaped path. Common membership in Everyone will not transmit beliefs either, because its relevance filter is maximally restrictive.

Figure 4: The exceptional nature of self

Figure 5: The need for speech

4. *Nonce corporations*, introduced dynamically to represent the temporary grouping of participants in an active conversation, never *inherit* attitudes from their subagents, but must acquire them as the effects of observable actions. The idea here is that while participating in (or directly observing) a conversation, the system is in a position to observe the construction of the public record of that conversation [Lewis, 1983] directly, and this record consists exactly to the attitudes we wish to ascribe to the conversational group itself. In conversation even a new 'unspoken understanding' should be based on inference from observed communication, and not just the system's private beliefs about other participants' views.

The fact that we still permit conversational groups to inherit from *superagents* allows us to place a discussion within a social context that supplies shared background assumptions. The fact that we permit their *subagents* to inherit from them models the actual adoption of information from the public record by individual participants, including the system itself, without additional mechanism.

Figure 5 depicts this situation: nonce1, the conversational grouping, represents a shared social construct distinct from our understanding of Jan's private views. This allows us to deal gracefully with the situation in which we, sec.Jan, catch (or perhaps even conspire with) Jan in telling a lie.

The most important property of this model of attitude ascription is that the only belief spaces it introduces are those that are independently motivated

by identified social groupings or the records of the actual conversations in which the system participates. This reduces the chance that the system will become mired in irrelevant structural detail, and specifically avoids the 'spy novel' style of belief space nesting that is characteristic of classical normative models. Attribution by default inference allows an individual to be represented as a member of several different groups holding conflicting beliefs, and inheriting only those beliefs consistent with those represented as being held privately.

The results are thus substantially different from those obtained in classical logics [Allen, 1983; Kraus and Lehmann, 1988; Appelt, 1985; Cohen and Levesque, 1990]. They differ from other path-based algorithms [Ballim and Wilks, 1991] in the provision of semantic relevance conditions and in addressing the need for shared attitudes by ascribing them directly to groups, rather than by maintaining complex accounts of which agents believe what. This allows us to describe and process conversational mechanics without recourse to nested (x believes that y believes that ...) belief spaces, though such structures may remain necessary for other, less routine feats of cognition.

In the next section we show how our model of attitude ascription can be used to implement multiparticant speech act processing.

Multiparticant Speech Acts

As in [Traum and Hinkelman, 1992], we assume that a core speech act is ultimately realised by a *sequence* of utterances that embody the grounding process. The model requires that the definitions of the speech acts themselves abstract away from grounding and provide high level actions that can be integrated with non-linguistic domain actions in planning. Using our multi-agent attitude attribution mechanism, we can simplify matters further, defining speech acts as joint actions whose effects apply directly to the conversational group being modelled.

Consider the generalised action operator representing one simple core speech act:

Inform_{ap}
conditions : $B_{bp} \wedge b \sqsubseteq a \wedge \text{live } a$
effects : B_{ap}

This Inform is a true joint action. Agent a is the nonce corporation representing all the participants taken jointly (the *live* predicate requires that this nonce correspond to an³ ongoing conversation). Though the singleton subagent b is the source of the information, the action has its effect directly on our model of the group. From that point propagation downwards to the individual participants is a function of the attitude ascription model, and is subject to the constraints given above. (The system effectively assumes that corresponding updates actually take place in the minds of the conversational participants.)

The correctness of this formulation relies on two facts. The first is that the grounding structure realis-

³Our current implementation actually deals with email rather than live speech, and must cope with multiple active dialogues.

ing the core speech act operator ensures the content is successfully conveyed. The second is that *if* a speech act that has an effect on the conversational group is fully realised and properly grounded, then any hearer who *dissents* from those effects must explicitly act to cancel them. That is, acknowledgement of receipt of a message establishes a presumption of assent to the content. Note, however, that when the speech act remains unchallenged this means only that the conversational participants will let it stand as part of the public record; it does not mean that they are truly persuaded of its content, and the rules we have given only predict that they adopt it if there is no evidence to the contrary.

Successful *requests* have effects on the goals rather than the beliefs of the group. It is crucial that both communicative and noncommunicative are introduced. The first goal below is noncommunicative and represents simply that the requested action be performed. Note that although the requested action's (possibly corporate) agent participates in the dialogue, there is no restriction that it not include the requester. Writing Bif_{ap} for $B_{ap} \vee B_{a\neg p}$ and $\Diamond$ for *eventually*, we have

Request_{ae}
conditions : $\text{agent } e \sqsubseteq a \wedge \text{live } a$
effects : $G_a \Diamond e \wedge G_a \text{Bif}_a \Diamond e$

The second goal in the effects is a communicative one; the group acquires the goal of finding out *whether* the requested action will be performed. The consequence of this is that the requester gets an indication of whether the request was successful: even under the assumption of co-operativity, goal conflicts and plan constraints sometimes lead to the rejection of a successfully communicated request.

In the next section we describe how COSMA, our implemented calendar management system, processes an actual exchange.

Processing an N-Way Speech Act

Speech acts like the above can now figure in the plan- and inference- based algorithms of communicative intelligent agents. Since dialogue may include unpredicted events, such agents must be able to react to changing circumstances rather than relying completely on advance planning. As each incoming speech act arrives, the agent updates its beliefs and goals; these beliefs and goals are the basis for subsequent action. This is not only appropriate for the interface between task and dialogue, but absolutely crucial for the grounding process.

A typical application task for N-way speech acts in a multiagent environment is appointment scheduling, with dialogue systems serving as personal appointment secretaries to human agents. Our implemented system, COSMA, operates in this domain. We model a human/secretarial pair as a kind of corporate agent in which beliefs about appointments propagate up and down from both members, and in which goals about appointments propagate from the human up to the pair and from there down to the secretary.

When this example begins, the dialogue system

Discussion

An important property of the corporate agent model presented in this paper is its scaling behaviour. Although the number of 'agents' in a nontrivial world model may be large, we only introduce belief spaces corresponding to 'actual' objects about which the system has knowledge. In particular, the corporate agents that are used correspond to either durable social groupings or records of actual conversations. Individual speech act definitions, though they account for all the agents in the dialogue, need make reference to at most two agents.

In contrast with normative models, *our speech act processing model at no point requires that individual addressees be modelled*. Of course, dialogue is typically motivated by the desire to modify the addressees' mental states, but our system is free to make these updates on demand. Thus, so long as constructing detailed partner models is not independently necessary, the effort required to plan and respond to speech acts remains almost constant as the number of conversational participants grows.

We have thus achieved the extension of speech acts to multiagent environments, a step beyond other speech act based models [Dols and van der Sloot, 1992; Meyer, 1992; Bunt, 1989; Traum and Hinkelman, 1992]. In the process, we have reduced the complexity of the task/dialogue interface.

Figure 9: WidgetCorp expands

An interesting consequence of not needing to model all members of a conversational group is that it becomes unnecessary to identify them. While in some circumstances this may be an advantage, it does leave the door open to an interesting glitch: without an independent check, the system's model of who it is addressing may turn out to be inaccurate. A related thing can happen when the system plans on the basis of attitude propagation: it can perform an action that 'ought' to result in a given agent's coming to hold some view through social processes, but since social channels are quite imperfect, the message never gets through. Human agents are at times more cautious, and may model delays in the grapevine, but this 'lost sheep' phenomenon occurs sufficiently often in real life to make the utterance "Oh, I'm sorry, I guess you weren't at the meeting." sound very familiar.

Generalisation remains a hard problem, of course. Our system has no special advantages when faced with

a question like "Are conservatives porcupine-lovers?" Vague questions about large groups require extensive search or some independent theory of generalisation, and seem to be difficult even for humans.

Related to this is the *Nixon diamond* anomaly faced by many default inference systems. In our case, when we find propagation paths that will support the ascription of contradictory attitudes to a single agent, how should we choose between them? It turns out that selecting whichever result we first discover we would like to use is a surprisingly good solution. Such 'arbitrary' judgments tend to facilitate the conversation by using the inference mechanism not to seek a reliable proof but to find the most convenient supportable argument, regardless of actual truth.

Perhaps the most interesting deviation of our model from the behaviour of systems founded on mutual belief is the *social fiction* anomaly: one can fairly easily reach a state in which an attitude is ascribed to a corporation which is held by *none* of its members. Incredibly, this also corresponds to a familiar situation in everyday life. Three examples should serve to illustrate the point. In the first place, consider this exchange, in which Jan asks Les to compile a tedious report:

Jan: I'll need that on my desk by friday.

Les: friday. no problem.

Such a dialogue may occur even when Jan will not have time to look at the report until the middle of the following week, *and* Les knows that the work cannot possibly be completed before the weekend.⁵

Secondly, we propose the example of a couple who are jointly *but not severally* on a diet. When together, neither partner ever takes dessert, and this policy is verbally reinforced. Yet either of them will happily join you in eating a large slice of strawberry cheesecake, if they are apart.

Finally, imagine that you are a lone bicyclist approaching a Very Large Hill. You might now say to yourself—a conversational nonce of one—"It's not far to the top!" Processing this speech act results in another unsupported belief. You can now try to be convinced.

In light of all of the above anomalies, it begins to appear that human agents may be struggling with limitations similar to those of our own model.

It may still be objected that our model falls short in failing to support the detailed 'spy novel' reasoning used in conventional logics of belief, keeping track of whether Lou believes that Lee believes that *p*, and whether or not common beliefs are truly mutual. Our response is threefold:

- Reflective problem solving is always an option, but in humans appears to be an independent process. Responsiveness demands make it unsuitable for mandatory online use in dialogue processing, though it may be important to use models (like ours) with which it can be integrated simply.
- Conversational mechanisms have evolved to cope with the cognitive shortcomings of humans. To the

⁵The authors disagree as to whether this particular pattern is more likely to arise through malice or optimism.

extent that the performance errors of a dialogue agent mirror human failings, co-operative recovery performance with human communicative tools should be enhanced.

- Even given access to an ideal normative dialogue model, a full system would benefit from running a less precise and more descriptive model in parallel. This would assist in isolating those parts of a communicative plan where confusion on the part of *other* agents is predictable.

Corporate agents are an alternative to normative logics of belief which capture a number of interesting social and communicative phenomena straightforwardly. With their help, we can reformulate core speech act definitions cleanly and scalably for the case of many agents. The level of planning abstraction that results seems well-suited to the needs of intelligent communicative agents operating in an environment that includes many human agents.

References

- [Alexandersson *et al.*, 1994] Jan Alexandersson, Elisabeth Maier, and Norbert Reithinger. A robust and efficient three-layered dialog component for a speech-to-speech translation system. subm. Conference on Applied Natural Language Processing, Stuttgart, 1994.
- [Allen, 1983] James Allen. Recognizing intentions from natural language utterances. In Michael Brady and Robert C. Berwick, editors, *Computational Models of Discourse*. MIT Press, 1983.
- [Appelt, 1985] D. E. Appelt. *Planning English Sentences*. Cambridge University Press, Cambridge, 1985.
- [Ballin and Wilks, 1991] Afzal Ballin and Yorick Wilks. *Artificial Believers: The Ascription of Belief*. Lawrence Erlbaum Associates, 1991.
- [Bunt, 1989] H. C. Bunt. Information dialogues as communicative action in relation to partner modelling and information processing. In M.M Taylor, F. Neel, and D. G. Bouwhuis, editors, *The Structure of Multimodal Dialogue*. Elsevier Science Publishers B.V., 1989.
- [Clark and Schaefer, 1989] Herbert H. Clark and Edward F. Schaefer. Contributing to discourse. *Cognitive Science*, 13:259–294, 1989.
- [Cohen and Levesque, 1990] Phillip R. Cohen and Hector J. Levesque. Persistence, intention, and commitment. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.
- [Dols and van der Sloot, 1992] F. J. H. Dols and K. van der Sloot. Modelling mutual effects in belief-based interactive systems. In W. J. Black, G. Sabah, and T. J. Wachtel, editors, *Abduction, Beliefs and Context: Proceedings of the second ESPRIT PLUS workshop in computational pragmatics*, 1992.
- [Hinkelman, 1992] Elizabeth A. Hinkelman. Intonation and the Request/Question Distinction. In *Proceedings of the International Conference on Spoken Language Processing*, Banff, Canada, 1992.
- [Kraus and Lehmann, 1988] Sarit Kraus and Daniel Lehmann. Knowledge, belief and time. In *Theoretical Computer Science*, volume 58, pages 155–174, 1988.
- [Lewis, 1983] D. Lewis. *Scorekeeping in a Language Game*, volume 1, chapter 13. Oxford University Press, Oxford, 1983. Philosophical Papers.
- [Meyer, 1992] Ralph Meyer. Towards a conceptual model of belief and intention in dialogue situations. In W. J. Black, G. Sabah, and T. J. Wachtel, editors, *Abduction, Beliefs and Context: Proceedings of the second ESPRIT PLUS workshop in computational pragmatics*, 1992.
- [Moore and Paris, 1989] J.D. Moore and C.L. Paris. Planning text for advisory dialogues. In *Proc. Conf. of the 27th Annual Meeting of the ACL*, pages 203–211, Vancouver, 1989.
- [Perrault, 1990] C. R. Perrault. An application of default logic to speech act theory. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.
- [Sacks *et al.*, 1974] H. Sacks, E.A. Schegloff, and G.A. Jefferson. A simplest systematics for the organization of turn-taking in conversation. *Language*, 50:696–735, 1974.
- [Schegloff *et al.*, 1977] E. A. Schegloff, G. Jefferson, and H. Sacks. The preference for self correction in the organization of repair in conversation. *Language*, 53:361–382, 1977.
- [Schupeta, 1993] Achim W. Schupeta. Cosma: Ein verteilter terminplaner als fallstudie der verteilten ki. In Juergen Mueller and Donald Steiner, editors, *Kooperierende Agenten*, Saarbrücken, Germany, 1993.
- [Traum and Hinkelman, 1992] David R. Traum and Elizabeth A. Hinkelman. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3):575–599, 1992. Special Issue on Non-literal language.