A Speech Act Model of Air Traffic Control Dialogue

by
Karen Ward
B. Sci., University of Oregon, 1978

A thesis presented to the faculty of the
Oregon Graduate Institute of Science & Technology
in partial fulfillment of the
requirements for the degree of
Master of Science
in
Computer Science and Engineering
July 1992

The thesis "A Speech Act Model of Air Traffic Control Dialogue" by Karen Ward has been examined and approved by the following Examination Committee:

David G. Novick
Assistant Professor
Thesis Advisor

Ronald A. Cole
Professor

James Hook
Assistant Professor

Acknowledgments

I would like to thank Jim Hook and Ron Cole for their helpful and encouraging feedback as members of my thesis committee.

I am deeply indebted and grateful to David Novick for his guidance, encouragement, and unfounded faith. Without him, this thesis would never have been started, much less finished.

Finally, I would like to express my appreciation to Buko, who has put up with me through all of this.

Table of Contents

List of Tables

List of Figures

Abstract

This thesis develops a computational representation for air traffic control dialogue. Such a model might be used in developing a spoken language understanding system to represent and reason about utterances. Currently, speaker-independent, continuous-speech systems rely primarily on constraints such as sharply limited vocabularies, simple grammars, and word-pair probabilities to limit the possibilities considered in mapping sounds to phonemes, words, and ultimately, meaning. When speech understanding systems are applied to unconstrained speech in real-world settings, though, the range and complexity of potential utterances increases dramatically. To overcome this complexity, additional knowledge sources are needed. Of particular interest are "higher-level" knowledge sources, which describe information above the phonemic level.

One class of higher-level knowledge that is beginning to prove useful in speech understanding is dialogue modeling. By predicting the form and content of the next utterance from the content of prior utterances, dialogue models allow the recognizer to consider only a subset of the application's full grammar and vocabulary. Dialogue context is also used after the fact to correct the output of the speech recognizer, to select among several possible interpretations of the utterance, to handle ellipses and anaphora, and to disambiguate meaning. This thesis will focus on the use of two dialogue models, speech acts and the collaborative view of conversation, to explain and predict the intended meaning of an utterance.

The domain selected for this analysis is air traffic control, which exhibits several characteristics that make it interesting for both dialogue modeling and speech recognition studies. For this analysis, radio exchanges between air traffic controllers and pilots were

taped and transcribed. A complete dialogue, consisting of all exchanges between the controller and the pilot of a commercial flight approaching the airport to land, was explicated at the speech act level in terms of the beliefs and intentions of the conversants.

Chapter 1

Introduction

The goal of this research is to develop a computational representation for air traffic control dialogue. Such a model is expected to be of use in developing a spoken language understanding system to represent and reason about utterances.

This research is part of a larger effort directed toward improving the performance of spoken language understanding systems. Speaker-independent, continuous-speech systems rely primarily on constraints such as sharply limited vocabularies, simple grammars, and word-pair probabilities to limit the possibilities considered in mapping sounds to phonemes, words, and ultimately, meaning. As speech understanding systems begin to tackle unconstrained speech in real-world settings, though, the range and complexity of potential utterances increases dramatically. To overcome this complexity, we would like to bring additional knowledge sources to bear on the problem. Of particular interest are the "higher-level" knowledge sources, which describe information above the phonemic level.

One class of higher-level knowledge that is beginning to prove useful in speech recognition is dialogue modeling. By predicting the form and content of the next utterance from the content of prior utterances, dialogue models may let the recognizer consider only a subset of the application's full grammar and vocabulary. Dialogue context is also used after the fact to correct the output of the speech recognizer, to select among several possible interpretations of the utterance, to handle ellipses and anaphora, and to disambiguate meaning. This thesis will focus on the role of speech act dialogue models and the collaborative view of conversation in explaining and predicting the intended meaning of an utterance.

Speech act theory was proposed by the language philosopher Austin and developed by his student Searle ([Austin 62], [Searle 69], [Searle 75], [Searle 85]). This theory explains the motivation behind an utterance by considering it as an act intended to bring about change in the world. The collaborative view of conversation offers an explanation for conversational coherence by viewing conversation as an ensemble work in which the conversants cooperatively build a model of shared belief [Clark 89]. Recently, increasing interest has been shown in applying dialogue models — particularly speech act theory — to text-based natural language understanding problems (for example, [Allen 89], [Perrault 80], [Stubbs 83]). Novick has used these theories in developing a theory of meta-locutionary acts to explain control acts in conversation [Novick 88]. To bring these theories to bear on the problem of understanding spoken language, a computational model of dialogue at the speech act level in a tractable real-world domain is needed.

The domain selected for this analysis is air traffic control (ATC). ATC dialogue exhibits several characteristics that make it interesting for both dialogue modeling and speech recognition studies. Although it is unconstrained speech — that is, the conversants can and will use any phraseology necessary to communicate their meaning — ATC communications are built around a small core vocabulary of phrases with documented meanings. The radio communications protocols make explicit certain aspects of conversational control (e.g., turn taking) that are often difficult to capture in face-to-face conversation. Many troublesome aspects of conversational context, such as power relationships and prior interaction between the conversants, are known or minimized in this domain.

For this analysis, radio exchanges between air traffic controllers and pilots were taped and transcribed. A complete dialogue, consisting of all exchanges between the controller and the pilot of a commercial flight approaching the airport to land, was explicated at the speech act level in terms of the beliefs and intentions of the conversants.

Chapter 2 of this thesis reviews related research in spoken language understanding and in dialogue modeling. Chapter 3 describes in more detail the characteristics of ATC communications that make it an attractive domain for studying dialogue and provides an introduction to ATC tasks and terminology. Chapter 4 presents a speech act model for ATC dialogue and Chapter 5 illustrates the use of the model in explaining a typical dialogue. Chapter 6 contains a summary and conclusions.

Chapter 2

Dialogue Models for Spoken Language Understanding Systems

Spoken language understanding technology is expanding from the highly constrained systems that marked the first successes in the field [Reddy 76] to more ambitious systems designed to function in real-world settings (for example, the systems participating in DARPA's ATIS project [Price 91]). The long-term goal, of course, is a system capable of matching human performance in unconstrained, speaker-independent dialogue. For this, acoustic models alone are insufficient. Humans rely on a wealth of non-acoustic language information to aid in interpreting the speech signal: syntax and semantics, prosody and pragmatics. One kind of pragmatic knowledge is our understanding of the structure of discourse above the sentence level, sometimes called the dialogue model. The first half of this chapter reviews the use of dialogue models in spoken language understanding systems. The second half describes the linguistic theory underlying the dialogue modeling techniques that will be used in this work.

2.1 Language Models in Spoken Language Understanding Systems

Research in spoken language understanding has traditionally emphasized modeling and interpreting the acoustic signal. Language models have been used primarily to reduce the number of possibilities that the recognizer must consider in searching for the most likely word string to match the given acoustics. Thus the emphasis has been on simple, fast models that predict words or phonemes.

2.1.1 Word and Phrase Models

The most widely-used language model is N-gram modeling. This technique assumes that the probability of a given word's occurrence depends on the previous N words, most often one word (bigram models or word-pair grammars). N-gram models offer several advantages: they have been successful in greatly reducing the search space; they are reasonably straightforward to construct when the input is understood well enough so that the word occurrence probabilities can be estimated; because they do not incorporate notions of grammar, semantics or pragmatics, N-gram models have proven robust in handling the ungrammatical constructions that typify unconstrained speech [Price 91].

The lack of higher-level language modeling leads to several serious problems, however. First, the recognizer frequently returns an ungrammatical or nonsensical result. Second, an N-gram model at best merely returns a plausible string of words; additional processing is required to determine the speaker's intentions so that the system may respond appropriately. When N-gram models are used, then, the recognizer's output must be interpreted and corrected after the fact. The most common approach has been to modify the recognizer to return several possible word strings instead of a single assignment. This ranked list of the "N-best" possibilities is then evaluated by a separate language component (see, for example, [Schwartz 90] or [Soong 90]).

Traditional language grammars fare poorly in spoken language systems. Unconstrained speech is very different from the grammatical prose that we were encouraged to use in English class. As journalist Janet Malcolm observes:

> "When we talk with somebody, we are not aware of the strangeness of the language we are speaking. Our ear takes it in as English, and only if we see it transcribed verbatim do we realize that it is a kind

> of foreign tongue... we all seem to be extremely reluctant to come right out and say what we mean — thus the bizarre syntax, the hesitations, the circumlocutions, the repetitions, the contradictions, the lacunae in almost every non-sentence we speak" ([Malcolm 90], pg. 20).

Parsing techniques built around traditional grammars of English tend to do poorly in the face of this "foreign tongue," not surprisingly.

Frame-based approaches to parsing and semantic analysis have been more successful in handling unconstrained speech. CMU's Phoenix system uses slot-level grammars built around frame-based semantics to implement a phrase-driven flexible parser ([Ward 91], [Young 91]). The output of the SPHINX speech recognition system [Lee 90] is passed to a parser which applies grammatical restraints at the phrase level. The phrases fill in slots in semantic frames, which are then analyzed further to construct a database query. For example, the utterance

"Show me I want to see all flights to Denver after two pm"

would be initially parsed into a frame like this:

[list]:	I want to see
[flights]:	all flights
[arrive_loc]:	to Denver
[depart_time_range]:	after two pm

(example taken from [Ward 91]). This strategy assumes that unparsable utterance fragments represent restarts or repeats and may be ignored, an assumption which may not generalize well. The Phoenix system has performed well on the DARPA ATIS task, however [Price 91], and this approach is clearly a useful one.

2.1.2 Integrating Language Models with Speech Recognition

Although frame-based language models represent an improvement over simple N-gram word models, the standard architecture still suffers from an inherent limitation: higher-level language knowledge is used only after the fact to correct and disambiguate the recognizer output. The lack of feedback between the language component and the recognizer means that the language component is limited to second-guessing the recognizer. Similarly, the speech recognizer cannot make use of higher-level constraints to further restrict and guide the search. Clearly, integrating the two should improve the accuracy of the system, but how best to accomplish this integration remains unclear.

One approach has been to develop "dynamic grammars," which change as the dialogue progresses. Fink and Biermann made an early attempt to implement a system that could recognize and exploit patterns in the discourse [Fink 86]. When a pattern was detected, the system formed expectations about what was likely to be said next. These expectations were used to bias the recognizer's grammar toward the expected next sentence. This design was noteworthy in that the system acquired these patterns dynamically by monitoring the discourse. The implementation, however, was limited by the decision to use off-the-shelf speech recognition equipment that had no provision for dynamically modified grammars. Instead, Fink and Biermann applied the expectations to the recognizer output as part of a post-processing error correction strategy.

The MINDS project was more successful at implementing a true dynamic grammar ([Young 89b], [Young 89a], [Hauptmann 88]). This system tracked the dialogue state and modified the grammar to reflect the current dialogue context and focus, user goals, and problem solving strategy. As an utterance was processed, the system constructed a grammar for the speech recognizer to use in interpreting the next utterance. If this small, specific grammar failed to produce an acceptable parse, the system relaxed its constraints somewhat to produce a less-focused grammar. This "layered grammar" technique was designed to permit the system to respond efficiently when the next utterance matched expectations while still exhibiting graceful degradation when an utterance was unexpected. Although intuitively appealing, this approach apparently was unacceptably inefficient; in their more recent work (e.g., [Ward 91],[Young 91]), Ward and Young have abandoned dynamic grammars in favor of a more traditional post-processing approach.

A more successful effort to integrate higher-level language models into the speech recognition process comes from MIT. The SUMMIT speech recognition system and the TINA language understanding system can be run in several configurations. In the most tightly coupled configuration, TINA's parser is called interactively during the recognizer's search phase to prune impossible theories from the search space [Goodine 91]. The TINA language understanding system currently includes only a fairly traditional context-free grammar and does not yet incorporate dialogue-level knowledge. Still, this architecture is promising because of its potential for incorporating multiple higher-level knowledge sources into the speech recognition process in a flexible and extensible manner.

2.2 Modeling Dialogue

The previous section discussed the language models currently used in spoken language understanding systems. This section looks at the theoretical basis for the dialogue models that will be used in this analysis.

2.2.1 Speech Acts

Traditional linguistics approaches the problem of understanding language from the bottom up, focusing on words and definitions, grammar and syntax. We cannot fully understand or predict language use by considering syntax and semantics alone, however. For instance, it seems intuitively correct to say that these sentences all request the same action:

Pass the pepper.
Would you please hand me the pepper?
Pepper, please.

In form, the first appears to be a command, the second a question, and the third isn't even a complete sentence. This suggests that there is a common underlying intention behind these sentences that is separate from their surface form.

The reverse of this problem is equally vexing. How do we account for the observation that the same words uttered in different circumstances are easily understood to mean very different things? For example, the question:

Do you know what time it is?

could represent, in the appropriate circumstances, either a simple request for information or a pointed suggestion that it's time to leave — and only rarely can it be interpreted as the simple "yes-no" question that its surface form would indicate. Meaning, then, must somehow depend on the context in which the utterance was made.

Furthermore, if communication is defined by the lexical meaning of words, how do we explain our understanding of nonverbal communications? As Austin notes, "in very many cases it is possible to perform an act of exactly the same kind not by uttering words, whether written or spoken, but in some other way." ([Austin 62], page 8). In the analysis which follows, we will see that what *isn't* said is often as significant as what *is* said. We can take this a step further by observing that silence is itself a potent carrier of meaning [Saville-Troike 85]. We communicate not only with speech and sound, but with gesture and pause.

We observe, then, a many-to-many mapping between the literal form of a communicatory action and our intuitive notion of what the speaker intended in performing it. Clearly the literal meanings of words can do no more than constrain the possible interpretations of the utterance. It is not enough to transcribe the words that may — or may not —

have been said. For reasoning about communication, we need a representation that captures the speaker's intention and the hearer's understanding.

Speech act theory suggests such a representation ([Austin 62], [Searle 69]). Speech act theory treats language as a tool the speaker uses to bring about changes in the world. Earlier language philosophers had considered communication to consist primarily of statements that were intrinsically true or false; Austin, however, realized that language use is fundamentally an action, something not conceptually all that different from, say, picking up a pencil.

Some utterances affect the world directly. For instance, saying "I bet you a quarter it will rain tomorrow," is not to make a report about the truth or falsehood of a bet; making the statement *creates* the bet. More commonly, an utterance may be intended to change the hearer's mental state in some way, or to motivate the hearer to perform some action. For example, the request "please pass the pepper" is generally designed to motivate the hearer to pass the pepper to the speaker. Casual conversation may have the more diffuse goal of building and affirming social relationships.

Speech act theory, then, emphasizes the intent of the speaker and the effect on the hearer, independent of the words — if any — actually used. The hearer draws on a combination of the literal meaning of the words actually uttered, the manner in which they were uttered (prosody and gesture), and the context in which the utterance occurred to infer the speaker's intentions in making the utterance. The interaction and redundancy among knowledge sources provide robustness to human communication in the face of noisy, inadequate, ambiguous, or even erroneous productions. A spoken language understanding system, then, needs to be able to derive and represent the speech acts that underlie the observed locutionary acts.

Searle proposed that speech acts could be recognized and defined by a set of rules [Searle 69]. In Searle's terminology, *propositional content* rules indicate the literal

Table 1. Properties of a *Request* Speech Act	
Rule	**Properties of *Request***
Propositional content	The literal meaning of the utterance refers to some future action of the hearer.
Preparatory conditions	The speaker believes that the hearer can do the action, and it is not obvious to both speaker and hearer that the hearer will perform the action in the normal course of events.
Sincerity conditions	The speaker wants the hearer to do the action.
Essential feature	The utterance represents an attempt to get the hearer to perform the action.

meaning of the utterance, *preparatory* and *sincerity* rules express the context in terms of the relevant beliefs and goals of the conversants, and the *essential feature* defines the intentions of the speaker in making the utterance[1]. For instance, an utterance that exhibits the properties summarized in Table 1 will generally be interpreted as a *request*. Searle's rules suggest several features that will be required for modeling language understanding, notably that such a model should incorporate some notion of conversants' beliefs about each other's wants, abilities, and expected actions.

2.2.2 Speech Act Taxonomies

In this thesis I will be presenting a set of speech acts useful for representing air traffic control dialogue. A question naturally arises from this: how universal are these

1. In his later work, Searle refined the rule categories somewhat [Searle 85]. In particular, his 1985 version includes additional categories designed to capture degree of strength (e.g., the intuitive difference between *request* and *beg*). For the purposes of this work, the original categories are sufficient.

speech acts? Is it possible — or even meaningful — to attempt to develop a general taxonomy or list of speech acts?

There have been efforts to develop a small list of basic, irreducible speech acts or to group speech acts into a small number of related families. Several people have suggested heuristics for recognizing verbs that can describe speech acts (e.g., [Austin 62], [Stubbs 83]). Austin estimated that there roughly one thousand speech act verbs in English, and he proposed a preliminary taxonomy based on an intuitive classification of related verbs. Searle later proposed a hierarchical taxonomy based on similarities among the speech act properties [Searle 85]. Allen [Allen 91a] presents an alternate taxonomy of intention-based speech acts, categorized as *understanding*, *information*, or *coordination* acts.

Whether any of these efforts succeed in capturing the expressive richness of the English language is a question perhaps best left to language philosophers. A better approach, I believe, is the one taken here: determine the speech acts relevant within a particular domain. There are some speech acts that crop up commonly in many domains, like *request* or *inform*, just as there are concepts that are pervasive in more traditional data modeling (*person* and *name*, for instance). Within a particular context, however, only a subset of possible concepts are likely to be relevant. For instance, a *request* differs from an *order* in that a successful *order* requires that the speaker has the authority to give the hearer an order. Note that a person with such authority may still make a *request*, implying that the hearer may choose to not comply. In representing an air traffic controller's communication with a pilot, however, the difference between *request* and *order* becomes vanishingly small; by law, a pilot must comply with controller directives if able to safely do so. The authority relationship is so extremely asymmetrical in this case that it becomes difficult — if not impossible — for the controller to make a *request* that would not functionally be an *order*, and so we will model only one act for the controller. Only in the context of a particular domain can one determine whether a particular shade of meaning is significant.

2.2.3 Using Speech Acts to Model Dialogue

A dialogue — a conversational exchange between two or more persons — exhibits structure above the utterance level. As Stubbs points out [Stubbs 83], we can readily distinguish between random sentences and actual dialogue, or grasp a joke that depends on faulty discourse order:

A: Yes, I can.
B: Can you see into the future?

It is this structure that lends a conversation its coherence.

Although neither Austin nor Searle directly concerned themselves with discourse above the sentence level, speech act theory contributes two important ideas to the understanding of dialogue structure. First, speech act theory provides the conceptual link between intention and utterance, thus providing a basis for modeling utterance meaning in terms of the speaker's goals. A person says something for a reason, and that reason is to effect change in the world. Second, speech act theory makes explicit the role of the hearer in understanding an utterance. People do not normally speak in a vacuum; they communicate with another person. The meaning of an utterance can only be considered in terms of its expected effect on some hearer in some context.

Thus, speech act theory suggests that we can motivate dialogue and explain conversational coherence by modeling conversation in terms of the conversants' goals and their plans for reaching those goals. A speaker plans utterances to accomplish certain goals; a hearer interprets those utterances in light of the inferred intentions — goals — of the speaker. This approach accords well with findings that the structure of discourse about a particular task closely follows the structure of the task itself (for example, [Oviatt 88], [Cohen 79], [Grosz 86]). Language, then, is viewed as just another tool to be used in accomplishing the goal, and utterance planning becomes incorporated into the larger task planning [Power 79].

But task planning, although an important part of explaining dialogue structure, does not seem sufficient for explaining certain phenomena observed in actual dialogue. For example, how do subdialogues for correction or clarification fit into a planning model? What about "back-channel" responses, the "uh-huhs" and head nods that punctuate casual dialogue? How is conversational turn-taking coordinated? The plan-based analyses of Cohen [Cohen 84] or Litman [Litman 87], for instance, explain these phenomena only with difficulty.

2.2.4 The Collaborative View of Conversation

Speech act theory views communicative acts — verbal or nonverbal — as attempts to bring about change. But what is being changed by a conversation? In a series of studies, Clark and his colleagues proposed and developed a theory of conversation as collaborative process in which conversants work together to build up a mutual model of the conversation ([Clark 81], [Schober 89], [Clark 89], [Clark 86]).

In Clark's view, many of the characteristics of real-world dialogue can be explained in terms of mutuality of knowledge. Conversants build upon a basis of shared knowledge drawn from the information considered to be known by all members of a community, knowledge from prior interaction between the conversants, and information observed from the physical world around them. They add to this mutual model in an orderly way though collaboration; both conversants are responsible for ensuring that the speaker's contribution has been understood "to a criterion sufficient for current purposes" ([Clark 89], pg. 163). Clark modeled conversation, then, as a series of contribution-acceptance pairs. After each contribution by one conversant (A), the other conversant (B) accepts the contribution by displaying evidence of understanding. This evidence might consist of one of the following (taken from [Clark 89], pg. 267):

1. *Continued attention.* By continuing to listen, B indicates that A's presentation has been understood to B's satisfaction.

2. *Initiation of the relevant next contribution.* B shows that A's contribution has been understood by starting in on the next relevant contribution.
3. *Acknowledgment.* B nods, says "uh huh", or makes some other overt indication that A has been understood.
4. *Demonstration.* B demonstrates understanding, e.g., B performs the action that A has requested.
5. *Display.* B repeats verbatim all or part of A's presentation, e.g., B repeats back the address that A has dictated.

Notice that an acknowledgment is itself a contribution to the conversation, thus requiring acknowledgment. How do we keep from looping on acceptances of acceptances? Clark proposed that the types of evidence are ordered from weakest to strongest; we accept a contribution at one level by offering evidence of understanding at a weaker level.

Clark also suggested a hierarchy of evidence of trouble in understanding (from [Clark 89] pg. 268):

State 0: B didn't notice that A was attempting to communicate.

State 1: B noticed that A was attempting to communicate, but B wasn't in State 2. For instance, B may have heard A say something without catching all the words.

State 2: B correctly received A's communication, but wasn't in state 3. For instance, B may have understood all the words but doesn't understand what A meant.

State 3: B understood what A meant.

2.2.5 Example

The exchange reproduced in Figure 1 illustrates several types of evidence of understanding and of trouble in understanding. The exchange is drawn from a longer dialogue found in [Ward 90b].

As the conversation begins, the pilot (Ford 645) has attempted to contact the controller (Approach) for permission to make a sight-seeing flight over downtown Portland. The controller didn't quite catch the pilot's initial transmission — that is, he was in

(82)	Ford 645:	Portland Approach, Ford Trimotor niner six four five.
(83)	Approach:	((Garbled)) Portland Approach say again?
(84)	Ford 645:	Ford Trimotor niner six four five we're squawking twelve hundred. We're uh just off Troutdale southbound we'd like to: make one circuit over the city, we're uh one thousand three hundred, then we'll be southbound to McMinnville.
(85)	Approach:	Roger and uh say again the full call sign? I got the niner six and what was the last of it?
(86)	Ford 645:	Niner six, four five.
(87)	Approach:	Niner six four five is that it?
(88)	Ford 645:	That is correct.
(89)	Approach:	And what's your type aircraft niner six four five?
(90)	Ford 645:	Uh we're a Ford Trimotor.
(91)	Approach:	Ford uh, six four five squawk zero one zero five and ident remain outside the ARSA till radar identified.
(92)	Ford 645:	OK.

Figure 1. A Difficult Conversation

state 1 — so in utterance 83 Approach *acknowledges* that the transmission occurred while indicating that its content was not understood. Ford 645 *demonstrates* his understanding by repeating his identification as requested and then going on to explain what he wants (*relevant next contribution*). Approach, however, is still having difficulty with the call sign (*state 1*). Notice that in utterance 85 Approach again indicates that only part of the transmission was received and understood, this time using *display* to show that he has understood part of the call sign. In utterance 86, the pilot responds with a *demonstration* of understanding, repeating the call sign. The controller again uses *display*, the strongest form of evidence of understanding, to show that he has finally understood the call sign. The pilot accepts the controller's contribution with the relatively weaker *acknowledgment* form "that is correct." The controller then continues with the *relevant next contribution*, a request that the pilot supply his aircraft type. The pilot supplied this in utterance 84, but the controller did not catch this phrase either (*state 1*), possibly because the type was unusual; the Ford Trimotor is an antique. Only after this correction is completed does the controller respond to the pilot's original request (utterance 91). Notice that this exchange ends with the pilot's moderately-strong *acknowledgment;* the controller apparently has no *relevant next contribution*, the next-weaker form of evidence, and so he responds with the weakest evidence of understanding: *continued attention*, or silence.

2.2.6 Summary

The model proposed in this thesis is based on a synthesis of the principles supplied by the theories of dialogue summarized in the previous section. From speech act theory:

- Language use can be abstracted and represented in terms of speech acts.
- Speech acts can be motivated and described in terms of the speaker's beliefs, which include the speaker's beliefs about the hearer's beliefs.

From the collaborative view of conversation:

- Conversational state can be represented in terms of the conversants' beliefs and the mutuality of those beliefs.
- Mutuality of knowledge considerations can motivate and explain the information that conversants exchange.

Thus, in this model, conversation is viewed as an attempt to establish and build upon mutual knowledge using speech acts. This synthesis was first proposed by Novick [Novick 88] to explain conversational control acts. In this work, I will be using these principles to explain and motivate domain-level acts in a real-world task.

2.3 Domain Requirements

This research is part of a larger effort that is investigating the use of dialogue models in improving the performance of spoken language understanding systems. In choosing a domain in which to instantiate the investigation, therefore, one must consider the requirements and limitations of the state of the art in both dialogue modeling and speech recognition. This section discusses the factors that would characterize a tractable and interesting domain for this purpose.

2.3.1 Domain Requirements for Speech Recognition

The fundamental problem in speech recognition is coping with variability. For instance, the differences between the way two different speakers pronounce the same word may be far greater than the difference between two different words uttered by the same speaker, or even by one speaker saying the same word under different conditions [Doddington 85]. From a speech recognition standpoint, then, the biggest concern is controlling or compensating for acoustic variability. This has several implications for choosing a domain for speech understanding research.

2.3.1.1 Vocabulary

In simplified terms, a speech recognition system works by attempting to find the best match between its acoustic input and its vocabulary [Reddy 76]. As the number of alternatives grows, the search space quickly becomes unmanageably large. For speech recognition performance, then, a good domain is one with either a small vocabulary overall, or one with a vocabulary that has few legal alternatives at each point.

When presented with an out-of-vocabulary word or perhaps an unexpected noise, like a door slam, a speech recognizer will try to map that sound to some word in its vocabulary. At best, it may reject all matches. Thus, a good domain should have few out-of-vocabulary words.

We would also like the vocabulary to exhibit a low confusability. The greater the acoustic differences between alternatives, the lower the chance that the recognizer will confuse one word for another.

2.3.1.2 Signal Quality

Noise also presents multiple problems for speech recognition. A noisy channel distorts the acoustic signal and degrades performance. When background noise is constant, it is possible to compensate for it. When background noise is variable, however, it poses great problems for a recognizer. It may be difficult for the recognizer to determine where the noise ends and the speech signal begins, for instance, or to distinguish a door slam from a word. Also, people make distinct prosodic, acoustic, and phonetic changes in their speech when speaking in the presence of noise [Summers 88]. Although people instinctively change their speech production to increase their intelligibility, ironically these changes may be great enough to confuse a recognizer, especially if the noise levels — and thus the speakers' compensating productions — vary greatly. A good domain, then, is one where the noise levels are minimal, or at least consistent.

2.3.1.3 Data Collection

It will be necessary to gather and analyze many examples of conversations directed toward accomplishing the same task both to train the system and to test the robustness of the models. This suggests that a good domain should exhibit short, repetitive, well-defined, largely verbal tasks.

Another consideration in selecting the domain is the anticipated difficulty of capturing sample dialogues for analysis. One approach is to set up a controlled experiment in the lab, solicit subjects, and record protocols for study. This has the advantage of giving the researcher the greatest possible control over the data, but is time-consuming and expensive. Another possibility would be to draw protocols from sources that were produced for another purpose, either public sources such as television or radio, or other research projects. This option is likely to be cheaper and faster, but may require accepting more compromises in the domain characteristics.

2.3.2 Domain Requirements for Dialogue Modeling

If the primary concern from the speech recognition standpoint is acoustic variability, the corresponding concern from the dialogue modeling standpoint is contextual variability. As noted earlier, the same words in different contexts may take on very different meanings. "Context" is a broad term, though, encompassing a wide range of ill-defined factors: facial expression, gesture and body position, the physical surroundings, prior interaction between the conversants, the "background" or social knowledge that all members of a community are presumed to share, etc. For most of these, there is no accepted method of describing or representing the relevant features, nor even a theoretical basis for predicting when a given feature might be relevant [Cook 90]. For instance, how do we determine what "general knowledge" is relevant to a given conversation? It would therefore be desirable to eliminate or control as many of these contextual features as possible.

2.3.2.1 Physical Context

One method of controlling the context of expression, gesture, and physical surroundings is to examine dialogue produced under limited modality conditions. The term *modality* refers to the communications modes available to the conversants. For example, face-to-face conversation is an interactive, audio-plus-visual modality; voice-mail is a non-interactive, audio-only modality.

The audio-only modality offers some advantages for dialogue research. Face-to-face communications rely heavily on gaze, gesture, and facial expression to control and coordinate the conversation. These phenomena are rapid and subtle, making them difficult to capture and analyze. Furthermore, there is no general agreement on how these phenomena should be recognized or categorized: was that eye-blink a significant part of the conversation, or was it an involuntary action that passed unnoticed by the conversants? Finally, there are few tools available to automate the task of capturing and classifying gesture, much less make the physical context available to a speech understanding system in real time. By studying communications carried out without benefit of visual modalities, i.e., through telephone or radio, the researcher avoids the theoretical and practical difficulties of capturing and quantifying nonverbal communication. Also, the lack of visual interaction forces the conversants to convert nonverbal feedback into more explicit verbalizations, making them easier to detect and classify.

2.3.2.2 Prior Interaction

The context of prior interaction can present many problems in understanding and analyzing dialogue. How are we to understand, much less anticipate for speech recognition purposes, a conversation that begins: "Did you get that done?" Similarly, it may be important to be able to anticipate and allow for the effects of roles, power relationships, and authority structures. These affect both the words we choose [Hovy 91] and the interpretation we place on what we hear [Stubbs 83]. For instance, the primary difference between a suggestion and an order lies in whether the speaker has the authority to order the hearer to perform the action [Searle 69]. As Stubbs notes, if the authority relationship is sufficiently lopsided - as with teacher and pupil, for instance - it may be difficult for the more powerful speaker to successfully make a suggestion without being very explicit ("Now, you don't have to do this if you don't want to, OK?"). To minimize these effects, it is common to design experiments involving strangers or to observe people interacting in situations where their roles are formalized or well-defined (for example, Stubbs' analysis of classroom dialogue [Stubbs 83]).

2.3.3 Summary

An ideal domain, then, would have these characteristics:

Vocabulary:
- perplexity low
- confusability low
- small core vocabulary
- few out-of-vocabulary words
- letters and numbers
- semantic ambiguity minimized

Noise:
- signal quality high
- background noise low and consistent
- few nonverbal noises
- little overlapping talk

Data Collection:
- short, repetitive, well-defined tasks
- easy to capture

Physical Context:
- audio modality emphasized

Prior Interaction:
- prior interaction minimized
- roles well-defined

The domain selected for this study, that of air traffic control, exhibits many of the characteristics desired for dialogue modelling: the vocabulary is relatively small and formal; domain tasks are short, repetitive, and well-documented; communications take place via radio, an advantage both in terms of data collection and modality; roles and context are unusually well-defined for a real-world task. For speech recognition, however, the picture is not so rosy: signal quality is poor, background noise is high. Because this study focuses on dialogue modelling, it was decided that the advantages outweighed the disadvantages. These factors will be discussed more thoroughly in the following chapter.

Chapter 3

The Air Traffic Control Domain

The first part of this chapter provides an introduction to the air traffic control (ATC) tasks and terminology that will be used in this study. The second discusses the characteristics of ATC-pilot communications that make this domain particularly attractive for studying dialogue. The chapter concludes with a summary of other modeling efforts in the ATC domain.

3.1 Introduction to Air Traffic Control

This represents a simplified version of approach control procedures from the standpoint of an Instrument Flight Rules (IFR) pilot flying into a moderately busy airport such as Portland International. The information in this section is drawn primarily from three sources: *Air Traffic Control* [FAA 89], the air traffic controller's primary manual; *Airman's Information Manual* [FAA 91], the pilot's major source of information on air traffic control procedures; *A User's View of the Air Traffic Control (ATC) System* [Rosenbaum 88], an unusually thorough explanation of the air traffic control system written for the non-pilot.

3.1.1 Controllers

An air traffic controller's primary responsibility is to ensure the separation of aircraft. Each controller is responsible for coordinating traffic within a particular piece of the airspace. During the course of a flight, a pilot communicates with several controller positions:

Clearance Delivery. A pilot flying under Instrument Flight Rules (IFR) must file a flight plan showing the proposed route of the flight. Clearance Delivery is responsible for approving or amending the flight plan and for entering it into the ATC system. A pilot normally communicates with this controller immediately before departure.

Ground Control. The ground controller has responsibility for traffic on the taxiways within an airport. This controller directs the pilot from the ramp area (next to the gates) to the runway.

Control Tower. The tower controller is responsible for traffic on the runways; a pilot communicates with the tower during takeoff and landing. At larger airports (e.g., Portland International) the tower controller has the use of a radar display; at smaller airports such as Hillsboro, OR, the tower controller relies on binoculars and radio communications to locate and coordinate traffic.

Departure/Approach Control. This controller (Departure/Approach Control is one position, not two) is responsible for a section of the airspace immediately around the airport. Specifically, this controller is particularly concerned with the coordination of arriving and departing aircraft as they enter the standard approach and departure routes established around an airport. Pilots that are taking off will generally address this controller as "Departure," while pilots that are landing or overflying the airport tend to use the term "Approach." The approach controller relies exclusively on radar and radio communications to coordinate traffic. The airspace that this controller is responsible for is called an Airport Radar Service Area (ARSA), and the radar facility from which this controller operates is called a Terminal Radar Approach Control (TRACON).

Enroute Controller. The enroute controller operates from a facility called an Air Route Traffic Control Center (ARTCC). The enroute controller controls the airspace between and above airports using radar and radio.

Each controller communicates with pilots on a separate, published radio frequency. As an aircraft moves from one airspace to another, the flight is handed off to the next controller, that is, the controller transfers responsibility for the aircraft to the next controller before instructing the pilot to contact the next controller. Although the frequencies are published and the pilot is responsible for looking them up in advance, the handoff procedure includes a reminder to the pilot of which frequency to use next.

Because each controller position is responsible for a different kind of airspace and operates on a separate frequency, a controller working a given position expects to encounter only a subset of the range of ATC tasks. For instance, ground control would not expect to hear a request for lower altitude, and an enroute controller would be very surprised to be contacted for takeoff clearance.

The approach controller's radar display shows the airplane's position and a data block with the aircraft's current status. The controller has a printed flight progress strip that shows the flight plan filed by the pilot and information added by controllers along the aircraft's route. The console also displays the current airport conditions being broadcast on the Automatic Terminal Information System (ATIS), a recording of current airport conditions that is continuously broadcast on a separate frequency.

As of December 30, 1990, an aircraft is required to have a Mode C transponder when operating within an ARSA. A Mode C transponder provides altitude information to the controller's radar display. Although the exchange analyzed here was recorded a few months before this requirement took effect, it is reasonable to assume that all aircraft involved were equipped with Mode C transponders.

3.1.2 Pilots

A pilot is required to comply with ATC instructions unless doing so would jeopardize the safe operation of the aircraft.

Portland International (PDX) is surrounded by a type of controlled airspace called an Airport Radar Service Area (ARSA). Federal regulations require two-way communication between pilot and the controlling facility (in this case, Portland Approach) before an aircraft may enter an ARSA (AIM, par. 3-101).

Before landing, a pilot needs to know the airport conditions, particularly wind, altimeter setting, and the runway and approach procedure in use. A pilot may get this information verbally from the ARTCC controller just before handoff, or from the approach controller. Frequently, though, the pilot listens to the Automatic Terminal Information Service (ATIS) broadcast. Because the ATIS is updated frequently, each update includes an identifying letter (e.g., "alfa[1]," "bravo," "charlie"). A pilot who has listened to the ATIS broadcast will include this letter in the initial callup so that controller and pilot may verify that the pilot has indeed heard the current version.

According to the procedures recommended in the Air Traffic Control handbook and the Airman's Information Manual, the pilot and the approach controller should confirm that the pilot is aware of current airport conditions and of the approach procedure to expect. They often do not do so. A local pilot reports that Seattle Center (ARTCC) often confirms ATIS information just before handoff [Schwartz 92]; if this is common procedure, then Portland Approach may not expect to confirm ATIS with an incoming IFR flight. For the purposes of this analysis, then we will assume that this is the case.

3.1.3 IFR Approach Procedures

The dialogue to be analyzed in this study is between an approach controller and the pilot of a commercial flight making an IFR approach to Portland International Airport. Approach procedures differ for aircraft flying under Instrument Flight Rules and Visual Flight Rules. This section briefly describes the pilot-controller interactions expected when the flight is IFR.

An IFR flight is monitored by the ATC system from start to finish. The pilot files a flight plan showing details of the flight's expected route, and the pilot is expected to adhere to that plan. The flight plan is entered into the ATC system and is made available (by computer, generally) to each controller along the flight's route. While the aircraft is within a particular controller's airspace, the controller may update the flight plan to reflect

1. FAA documentation uses the spelling "alfa" instead of "alpha."

changes that occur during the flight (changes in estimated arrival time, for instance, or changes in routing).

As the flight leaves one controller's airspace and enters the next, responsibility for the flight must be explicitly transferred from one controller to the next. This procedure is called a handoff, and it may be coordinated either verbally (by phone) or electronically (though the controllers' consoles). In either case, the receiving controller is required to take explicit action to accept responsibility for the aircraft, and the pilot should not be instructed to contact the next controller until after the controllers have completed their handoff. When the pilot first contacts the controller, then, the controller is expecting the call; the explicit handoff should guarantee that the controller knows the aircraft is there. Furthermore, the controller knows quite a lot about the history of the flight and about the pilot's intentions from the flight plan.

Within the ARSA, the approach controller directs the aircraft from its current path to the approach gate, a point where it can begin to execute one of several standard procedure for approaching the airport and landing. These directions, termed vectors, will include instructions to change heading, speed, and altitude. A large airport like PDX will have several approach procedures for various weather conditions. The ATIS broadcast informs the pilot which procedures are in use at that particular time.

The exact vectors given will depend on many factors: the approach procedure in use, the current position of the aircraft, other traffic in the area, weather conditions, the capabilities of the aircraft and pilot. In general, the pilot does not know exactly what vectors to expect. Controllers tend to use a handful of standard routes, however, so pilots that fly into a particular airport frequently may learn those routes and come to anticipate, or at least recognize, the route they're being directed though.

Instrument Flight Rules flights into PDX use a navigation system called the instrument landing system (ILS). This system includes several components; the two that will concern us in this analysis are the marker beacon and the localizer. A marker beacon triggers a light in the cockpit when the airplane flies over it, thus providing the pilot with an accurate point of reference. Marker beacons are used to define legs of an ILS approach procedure; the aircraft should be at a certain altitude and heading when passing over a particular marker. The localizer is a directional radio beacon that defines the path that the aircraft must follow in descending to the runway. Once the pilot has intercepted the localizer signal, he or she can follow it directly to the runway without additional instructions from the controller. When the pilot has begun flying along the localizer beam, the aircraft is said to be *established* on the localizer.

The controller guides the aircraft to the localizer. After the aircraft is established on the localizer, the controller authorizes the pilot to follow one of the standard approach procedures. The approach procedure specifies the altitudes, speeds, and directions that the aircraft is to use, so the controller no longer needs to provide vectors. As the aircraft enters the airspace controlled by the tower, the approach controller hands the flight off to the tower controller.

The normal procedure for an IFR flight entering an ARSA for landing, then, is this:

- The pilot establishes communications with the approach controller,
- The controller directs the aircraft to the approach gate by issuing a series of vectors,
- The aircraft intercepts the localizer,
- The controller authorizes, or *clears*, the pilot to fly a particular approach procedure,
- The controller transfers authority for the flight to the tower controller.

As can be seen from the description above, communications between the approach controller and the IFR pilot are highly formalized, with strong expectations on both sides about what should be happening at every point in the conversation.

3.2 Characteristics of ATC Communications

To a naive listener, the dialogue between pilot and controller is almost unintelligible. The exchanges are short, rapid-fire, and full of jargon. In fact, ATC communications pose difficulties even for experienced participants. Communications take place over noisy, crowded radio channels by people engaged in the complex tasks of flying airplanes or coordinating the movements of multiple aircraft. ATC dialogue exhibits several adaptations to these challenges that make it particularly interesting for studying dialogue and speech understanding. This section discusses the factors that make ATC communications unusual.

3.2.1 Vocabulary

Recall from the previous chapter that a good domain for speech recognition applications should have a small core vocabulary and few out-of-vocabulary words. Additionally, the core vocabulary should include relatively few confusable words, and the words used should have clearly-defined meanings.

The conditions under which ATC communications are performed have had a strong effect on the vocabulary. The quality of the acoustic signal is generally poor; radio is a low-bandwidth medium made worse by the noise of the aircraft engines and by uncertain reception quality. To compensate for this, considerable human-factors work has gone into increasing the intelligibility of ATC messages. The FAA suggests, and conversants generally use, a standard lexicon and phraseology. These phrases are built around a small core vocabulary of words known to exhibit a low degree of phonetic confusion and to remain differentiable when spoken by non-native speakers of English.[1] In some cases, words are given alternate pronunciations to increase their intelligibility. For instance, "five" and "nine" are notoriously confusable over the radio, so they are pronounced "fife" and "niner" ([FAA 89], [FAA 91]).

1. English is the international language of aviation and is used at major airports throughout much of the world.

The number of words in the core vocabulary is quite small. Hall, for instance, found that the grammar of the Local Controller's prescribed vocabulary is built around just four verbs [Hall 88]. The core vocabulary also incorporates a high proportion of letters and numbers, a speech recognition task that has been particularly well-studied (e.g., [Fanty 92]).

The prescribed vocabulary is somewhat misleading, however; the range of language actually used in pilot-controller communications is much more extensive and complex than one might think from reading the ATC manual (see, for example, the dialogues recorded in [Ward 90b]). This happens for several reasons. First, although controllers are rigorously trained in the standard phraseology, pilots — especially general aviation pilots — are not. We see both a greater variation in their phraseology and a greater degree of disfluency in their speech. Recall this transmission from the example in the previous chapter (Figure 2):

(84) Ford 645: Ford Trimotor niner six four five we're squawking twelve hundred. We're uh just off Troutdale southbound we'd like to: make one circuit over the city, we're uh one thousand three hundred, then we'll be southbound to McMinnville.

The transmission is far from fluent; it contains filled pauses ("**uh**") and lengthened syllables ("**to:**"), and the information is poorly organized (the pilot intersperses position information with his request to fly over the city).

Second, the ATC handbook prescribes exact phrases only for the most common situations. It cannot of course include local place names and landmarks. "Lloyd Center" and "the volcano" are mentioned regularly in the Portland area, for example. Also, standard phrases are augmented as necessary to accomplish the task at hand, resulting in a more "English-like" vocabulary when unusual situations arise. This example is taken from a conversation that took place at Detroit Metro Wayne County Airport [Ward 90b]. The controller is asking the pilot of Northwest 752 to allow a smaller commuter flight to take off before him. Apparently, Northwest's company policy requires a 15 mile minimum spacing between company aircraft, so the controller can permit the commuter to take-off immediately but must delay the Northwest flight:

(281) Local Control: Northwest seven fifty two roger I got to put you fifteen miles in trail of your company. What I'd like to have you do if we can work it out i:s take position and hold three center, but pull off to your right because there is a Jetstream commuter behind you that I can get out right now.

A third type of variation from the "standard" phraseology is slang, as illustrated by this controller's use of "three holer" to indicate a 727 (from [Ward 90b]):

(239) Local Control: Continental two oh one, follow the Northwest seven twenty seven approaching from your left, he'll be the second three holer.

Finally, there is a certain amount of chitchat that takes place on occasion. It is not unheard of for enroute controllers to provide pilots with score updates during the World Series, for example [Dunham 91]. Although the volume of such "social" discourse is still far smaller than in most other domains, it does occur.

Despite these exceptions, ATC communications represent an unusually constrained and focused subset of English. The advantages to a speech understanding system are obvious: relatively small vocabulary, words selected to be easily identifiable, and a simple phrase-oriented grammar. These are exactly the characteristics that one looks for in a speech-understanding application. Unfortunately, ATC communication exhibits these characteristics precisely because it takes place under conditions that present difficulty even for people; the poor signal quality puts ATC discourse — especially the pilots' utterances — beyond the capability of current speech recognition technology. If spoken languages understanding systems are to succeed in this challenging domain, they will have to make use of higher-level knowledge sources such as dialogue context.

3.2.2 Radio Communications

Communication between pilot and controller is largely limited to radio and radar. This forces most communication to be carried out without the use of non-verbal channels. The nonverbal communications that do occur in this domain are rare and extremely limited. When radio communications are impossible, a tower may communicate with an aircraft via a light gun, for instance. There are also reserved transponder codes that the pilot can use to alert the controller to certain situations, e.g., "7600" to indicate radio failure. The only significant avenue of nonverbal information exchange is the radar image displayed on the controller's workstation, however. The controller can see the aircraft position and thereby draw some conclusions about the pilot's actions and intentions. Also, most aircraft have a Mode C transponder, which supplies approximate aircraft altitude to the controller.

Radio has another important modality limitation; it is a half-duplex communications medium. A person cannot simultaneously listen and transmit on a single channel. This has several implications for ATC communications. The most obvious one is that conversants cannot speak at the same time (at least, not understandably), which allows the researcher to avoid the difficulties of interpreting and representing the sort of simultaneous conversation that is common in face-to-face modalities.

The half-duplex modality also affects the structure of dialogue. For instance, it tends to discourage the installment-type exchanges noted by Clark [Clark 89], where complex information — say, an address — is presented in several parts with the speaker waiting for acknowledgment after each piece of the contribution. In ATC communications, however, we see the opposite pattern: multiple items of information are usually packaged into one transmission. This appears to be a response to both the overhead of establishing two-way communications (an installment inform would be painfully slow) and the crowded bandwidth around major airports; a pilot is encouraged to say everything immediately, because it may be difficult to get another chance to transmit.

Congested frequencies are a particular problem around heavily used airports [Fischetti 86]. Pilots sometimes have difficulty getting through to the controller in a timely fashion. As the channel becomes saturated with calls, the controller may begin issuing the next instruction without waiting for an acknowledgment of the previous one. The problem, of course, is that the controller assumes that the instructions are understood and accepted by the pilot unless the pilot objects. This accords well with Clark's observation that we accept silence as evidence of understanding and acceptance [Clark 89]. On a crowded radio channel, however, the lack of response could signify that the pilot was unable to gain control of the frequency.

The emphasis on verbal communications in ATC is changing, however. Although general aviation and smaller commercial carriers will continue to rely on radio for the foreseeable future, larger carriers have already begun to use the Mode S transponder system. This system provides a data link between the ATC system and the cockpit. Pilots and controllers will be able to exchange routine information such as weather data and altitude confirmations without using the radio [Heppenheim 90]. This is expected to reduce the frequency congestion, but at the cost of a reduction in the pilots' situational awareness.

3.2.3 Task, Context, and Role

Task, context, and the relative roles of the conversants have tremendous effect on the structure and content of dialogue. In most "real-world" situations, however, these factors are extremely difficult to control or even understand. ATC communications take place in an unusually structured environment.

In the ATC domain, contexts are limited and largely available to the researcher. Pilots and controllers do not in general know each other, so the context of prior interaction is minimized. Furthermore, roles are formally defined and quite limited. Because an unusual effort has gone into specifying and documenting procedures in this domain, many of the conversants' expectations are available to the researcher in the form of charts, regulations, and recommended practices. Thus, conversants are tightly focussed on a small set of well-defined tasks that are completed over a short period of time. This structure makes it possible to collect and analyze many examples of brief conversations with nearly identical context and goals.

3.2.4 Multi-Person Dialogues and Overhearers

Another interesting property of ATC dialogue is that it represents a sort of midway point between the two-person dialogues that are usually examined by researchers and a full multi-party conversation. ATC dialogue can be viewed as many interleaved yet fairly independent conversations. The controller converses with first one pilot, then another. These conversations are not entirely separate, though; pilots routinely monitor the frequency in order to build up a situational awareness of the activities around them. Thus they may, on occasion, respond or refer to dialogue that was not directed to them, as in this exchange from [Ward 90b]:

(66)	Ground Control:	Northwest two oh nine, uh, you still got your radar on?
(67)	NW 209:	That's a negative, Northwest two zero nine.
(68)	Ground Control:	OK.
		(.)
(69)	NW 255:	Two fifty five does.

To understand such exchanges, it will be necessary to model a utterance's effect on all hearers. The effect will not in general be the same for all hearers, of course. A pilot and retired controller comments [Dunham 91]:

> With no visual display of traffic in the cockpit, the only way to keep track of what's going on is via the radio. For instance, I once had a

> pilot question me after I cleared an aircraft to hold at the same place and altitude as he was. They were still 120 miles apart, and the plan was that Aircraft A was going to be long gone by the time B got there, but Pilot A couldn't tell that from just listening. He spoke up when he heard what sounded like a dangerous situation developing.

This accords well with Schober and Clark's findings that it is more difficult for overhearers to understand a conversation [Schober 89].

3.2.5 Safety and Sincerity

Safety issues have a strong effect on ATC communications. This is seen particularly clearly in the emphasis on explicit confirmation of information and instructions, and in the relative lack of insincere and non-literal speech.

Some of the information exchanged between controller and pilot is expected by the hearer. Rather than representing information that the other does not know, it represents confirmation of information that the other may well already have. This is seen most clearly in the emphasis on explicit acknowledgment and readback of critical instructions. For example, pilots routinely report their altimeter reading when entering a controller's airspace even though the Mode C transponder also reports the aircraft altitude automatically. This is done because altimeters are frequently imprecise ([FAA 89], [FAA 91]) and the aircraft altitude is critical in maintaining aircraft separation. Radar or computer failures can cause erroneous or missing data on the controller's display, making confirmation of critical information even more important [Fischetti 86].

Searle and other writers (e.g. [Grimshaw 80]) have wrestled with the issue of insincerity in speech. Searle's definitions of speech acts include sincerity requirements [Searle 69]. For instance, to recognize a *request* the hearer must believe that the speaker honestly wants the proposed action to occur. This requirement explains how we easily understand sarcastic comments that appear on the surface to be requests ("Go ahead, kick me when I'm down!") and suggests a mechanism for explaining both politeness constructs

("Oh, you shouldn't have!") and deliberate falsehood. It vastly complicates dialogue modelling, however, so the researcher would prefer a domain in which the conversants can be assumed to be telling the truth.

In ATC, safety concerns greatly lessen questions of sincerity; it is simply too dangerous to hide one's true intentions. In this domain, we can generally assume that people truly want what they ask for, although they may request indirectly. One exception to this may be in the handling of emergencies. These is anecdotal evidence that pilots are reluctant to declare an emergency, especially in marginal cases. Controllers seem to be aware of this and will treat situations as emergencies even though the pilot does not formally declare one.

3.3 Other Modeling Efforts in the ATC Domain

The unusual nature of ATC communications has made it an attractive domain for other language and speech modelling.

Hall used a PC-based system to build a speaker-dependent speech recognition application for the tower positions: local controller, ground controller, and clearance delivery [Hall 88]. His efforts were constrained by the decision to use an off-the-shelf PC-based speech recognition system, however. He developed a static case-based grammar implemented in the proprietary software of the speech recognizer that he was using. A slightly simplified version of the ground controller portion of the grammar was tested in an operational setting. Hall made no effort to model dialogue; his application recognized and captured only the controller's side of the conversation. Thus, his system did not attempt to build up an understanding of the context, nor did it have any way of capturing the pilots' contributions to the conversation.

Wesson modelled problem-solving for enroute controllers and achieved performance rivalling that of trained controllers for the given domain [Wesson 77]. His system successfully addressed the problem of modelling change and planning in a dynamic system.

He did not address communications, though, and recognized the lack of speech recognition as a barrier to his system (he tested his system with hand-entered data). Wesson's original system has since been expanded into the successful TRACON game and has formed the basis of more complex simulators used by the FAA in training air traffic controllers.

The FAA and the US Navy have recently installed a system called the Tower Simulator for training controllers in tower operations. The Tower Simulator is a simulated ATC tower inside a huge building. The trainees can look out the "window" and see a projected representation of the area around a tower. The computers provide a simulation of the airspace around the tower — including such factors as wind, traffic, pilots' responses to ATC, etc. — and alters the projection accordingly. An instructor at the FAA Academy describes this life-sized voice-controlled video game as being "very realistic" [Harold 91]. The Tower Simulator uses a speaker-dependent speech recognition system with a vocabulary of 260 "words" (some of these are actually short phrases), and 2.5 to 4 hours are required to train the system to recognize a new person.

Chapter 4

Representation

Chapter 2 described the theoretical basis for motivating dialogue and explaining conversational coherence in terms of the beliefs and intentions of the conversants. This chapter presents and discusses a computational representation designed to support language understanding by modeling the beliefs and intentions of conversants in the air traffic control domain. The first section describes the method that will be used in this study. The second section discusses representation issues and establishes the goals for the model. The final section describes the representation itself.

4.1 Method

This thesis develops a computational model for representing air traffic control dialogue at the speech act level. This model is intended to be incorporated into a spoken language understanding system. Unfortunately, such a system is not currently available. For this study, then, a more modest approach must be used.

Because the model should represent actual dialogue and not just encode the formalisms suggested by the FAA, the model was developed after studying protocols of actual pilot-controller conversation. As indicated in the previous chapter, it was decided to focus on Terminal Radar Approach Control (TRACON) operations for this study. TRACON operations — approach and departure control around a moderately busy airport — are attractive for several reasons. First, the TRACON controllers rely completely on radar and radio for interacting with the pilots. Unlike tower controllers, TRACON controllers cannot see the aircraft directly; TRACON controllers work in a room with no windows. Thus, the only non-verbal channel of communication available to the controller is radar, which could in principle be made available to a spoken language understanding system. Another advantage of studying TRACON operations is that there is typically a steady stream of both commercial and general aviation traffic into and out of the airspace, so it is particularly easy to collect many examples of controllers interacting with pilots of various skill levels. Finally, Portland International Airport, which has a TRACON facility, is nearby so it was possible to visit the TRACON to gain a clearer understanding of the controller task.

The protocols used in this study were made from a recording taped from the radio at a site approximately two nautical miles from the airport on July 1, 1990, between 10:45 and 11:15 AM [Ward 90b]. From these a suitable dialogue was selected for explication. From this sample dialogue I developed a model sufficient to represent and motivate the exchanges that were observed.

In selecting a dialogue for this study, I looked for one that might be considered "typical" so that I could examine the base requirements of the task and domain with as few complications as possible. The initial representation could then be expanded in future work to handle a wider range of dialogues. Also, the dialogue selected had to be complete; many of the dialogues in the corpus were in progress as taping began or were not finished at the time taping ended. The dialogue that was selected is described in detail in Section 5.2.

4.2 Representation Issues

This section discusses the goals to be met by the representation, then briefly discusses other representation issues that arose in the course of this study.

4.2.1 Goals of Representation

To represent an utterance in terms of speech acts, the representation must be able to model the beliefs and intentions of the conversants in a form that facilitates inferential analysis. The representation should also support reasoning about conversational coherence in terms of mutuality of belief. As will be seen, the safety-critical nature of the ATC domain leads to many exchanges in which conversants confirm information that each believes the other to already have. By modeling conversants' beliefs about the mutuality of this information, we can motivate these confirming exchanges and model their effects.

Conversants may likewise have inconsistent beliefs. Misunderstanding can be modeled in terms of inconsistent beliefs; conversants' mental models of the conversation may have diverged without the conversants realizing it. A person may then take some action intended to convey some speech act, only to be surprised when the other's response indicates that a completely different speech act was understood. Similarly, disagreements can be modeled in terms of the conversants' beliefs about the other conversants' differing beliefs. For instance, I may believe that A is true while recognizing that you believe that A is not true. Beliefs of this sort can also be used to account for certain types of correction subdialogues ("No, I said 'three', not 'free.'"). The model, then, should support both recognized and unrecognized belief inconsistency.

The representation should capture sufficient contextual information to model and explain the observed exchanges at the speech act level. Because this study focuses on the approach controller's position, it will be sufficient to represent only those domain elements that apply to the tasks of communicating with aircraft flying instrument approaches into Portland International Airport. In this domain one can generally assume that controller and pilot do not know one another; thus, the context of prior interaction will be ignored in this representation. The representation should, however, capture the context of the dialogue itself, that is, it should model the effect of the earlier part of the conversation on the conversants' current beliefs and intentions.

As was noted earlier, conversations in the ATC domain represent an interesting midway point between two-party conversations and full multi-party conversations. Although most utterances are clearly directed at just one other conversant, some exchanges cannot be explained without modeling the effect that an utterance has on the other pilots listening on the radio channel. The representation should therefore explicitly support multiple conversants by accounting for the effect that an utterance has on overhearers, including mutuality of belief among multiple agents. The issue of multiple conversants has been largely ignored in dialogue modelling efforts to date; a notable exception is Novick's saso system, which simulates multi-agent conversations at the speech act level [Novick 90]. Novick does not define the belief structures necessary to represent mutuality of belief among multiple agents, however.

A suitable representation, then, should meet the following goals:

- It should support reasoning about the beliefs and intentions of the conversants.
- It should permit agents to hold beliefs about other agents' beliefs.
- It should permit different agents to hold different, possibly inconsistent, beliefs, and it should support reasoning about the mutuality or inconsistency of conversants' beliefs.
- It should capture sufficient domain and physical context to model the exchanges, particularly the evolving context of the dialogue itself.
- It should support multiple (more than two) conversants.

4.2.2 Other Representation Issues

The emphasis of this representation is on modeling the role played by beliefs and intentions in motivating dialogue and in explaining conversational coherence in the air traffic control domain. ATC dialogue takes place as part of a larger cooperative task, however, that of trained individuals coordinating the safe movements of aircraft. The conversants' expectations and inferences include detailed mental models of the events taking place around them. If this were an attempt to model all aspects of air traffic control dialogue — for instance, to build a working system to understand ATC dialogue — there would be other factors that would require consideration. Although these are beyond the scope of this thesis, they will be mentioned briefly here.

4.2.2.1 Time

The representation presented here includes no notion of the passage of time. In the world, however, things change with time, and the passage of time thus has a profound effect on conversants' beliefs and expectations. Expectations for both pilot and controller are heavily dependent on time and on the conversants' mental models of how long something is likely to take or of when something should occur. For instance, consider this exchange [Ward 90b]:

```
(57) Horizon 64:  Horizon sixty four cleared for the approach?

(58) Approach:    Horizon sixty four you're: one zero miles
                  from Laker, maintain four thousand five
                  hundred till established on the localizer
                  cleared ILS two eight right approach,
                  maintain one seven zero knots until Laker.
```

The pilot expected to receive a clearance by a certain point in the approach. When no such clearance was forthcoming, the pilot finally prompted the controller. By modelling the pilot's unfulfilled expectation, the question can be motivated — but only by including in the model some notion of the passage of time and of the changing position of the aircraft.

4.2.2.2 Change

The model presented here is based on a fairly traditional notion of world state and events; at any given point, the world is in a particular state, and the world can change its state only when an event or action (an event performed by some agent) occurs. Problems arise in representing the effects of continuous and asynchronous change.

In many applications, events can be successfully modelled as if they were instantaneous. Events in the aviation domain — and in the world in general — are not so neat. Most predicates represent facts that are neither static nor changed only in response to a particular event. Instead, things change more-or-less continuously through time. For instance, a plane is at a particular position only for an instant, and then only approximately. It takes time to for a controller to transmit and confirm instructions to a pilot and more time for the pilot to implement the requested action, and during that time the situation may change, as in this example [Ward 90b]:

(33) Approach: **Alaska two oh five, nine from Laker, maintain three thousand till established on the localizer cleared the ILS two eight right approach, maintain speed (of) one eight zero until Laker.**

(34) Alaska 205: **Roger Alaska two oh five, uh, cleared approach three thousand till established on the localizer at three thousand, and uh hundred 'n uh ninety till Laker?**

(35) Approach: **Alaska two zero five, uh turn left heading two five zero now to join the localizer and maintain three thousand till established on the localizer, cleared ILS two eight right approach, maintain speed one eight zero until Laker.**

To make sense of the above exchange, one must realize that the exchange took some amount of time to complete; during that time, the aircraft travelled some distance on its current heading, so that the controller had to modify the initial instructions when issuing the clarification.

4.2.2.3 Inference

The model presented here is envisioned as the state portion of a production system. Throughout the discussion of the application of this model to a real-world dialogue (Appendix E) I indicate in general terms the inference rules that will be required to deduce speech acts and beliefs from the text of an utterance and to update the agents' belief states based on the acts recognized. A detailed specification of the rule set is beyond the scope of this work, though.

4.2.2.4 Natural Language Understanding

There are many linguistic issues that are not addressed here. While traditional linguistic notions of syntax and semantics are not sufficient for determining the intent of an utterance, they are certainly necessary. Furthermore, there are many other pragmatic issues that are not addressed here, such as common-sense knowledge, intonation and pause as modifiers of meaning [Steedman 90], inference generation. Each of these problems are difficult in isolation; to attain human performance in understanding unconstrained dialogue, however, we need to have at least rudimentary methods of handling all of these linguistic phenomena. Allen has proposed a system architecture designed to integrate models of reference and focus, discourse structure, and speech acts [Allen 89]; his Discourse System is designed for text, however, not spoken language.

4.3 Representation

4.3.1 Model of Agent Interaction

Figure 2 summarizes the conceptual model behind this representation. Conversants are modelled as autonomous agents, each having a separate belief space. An agent's

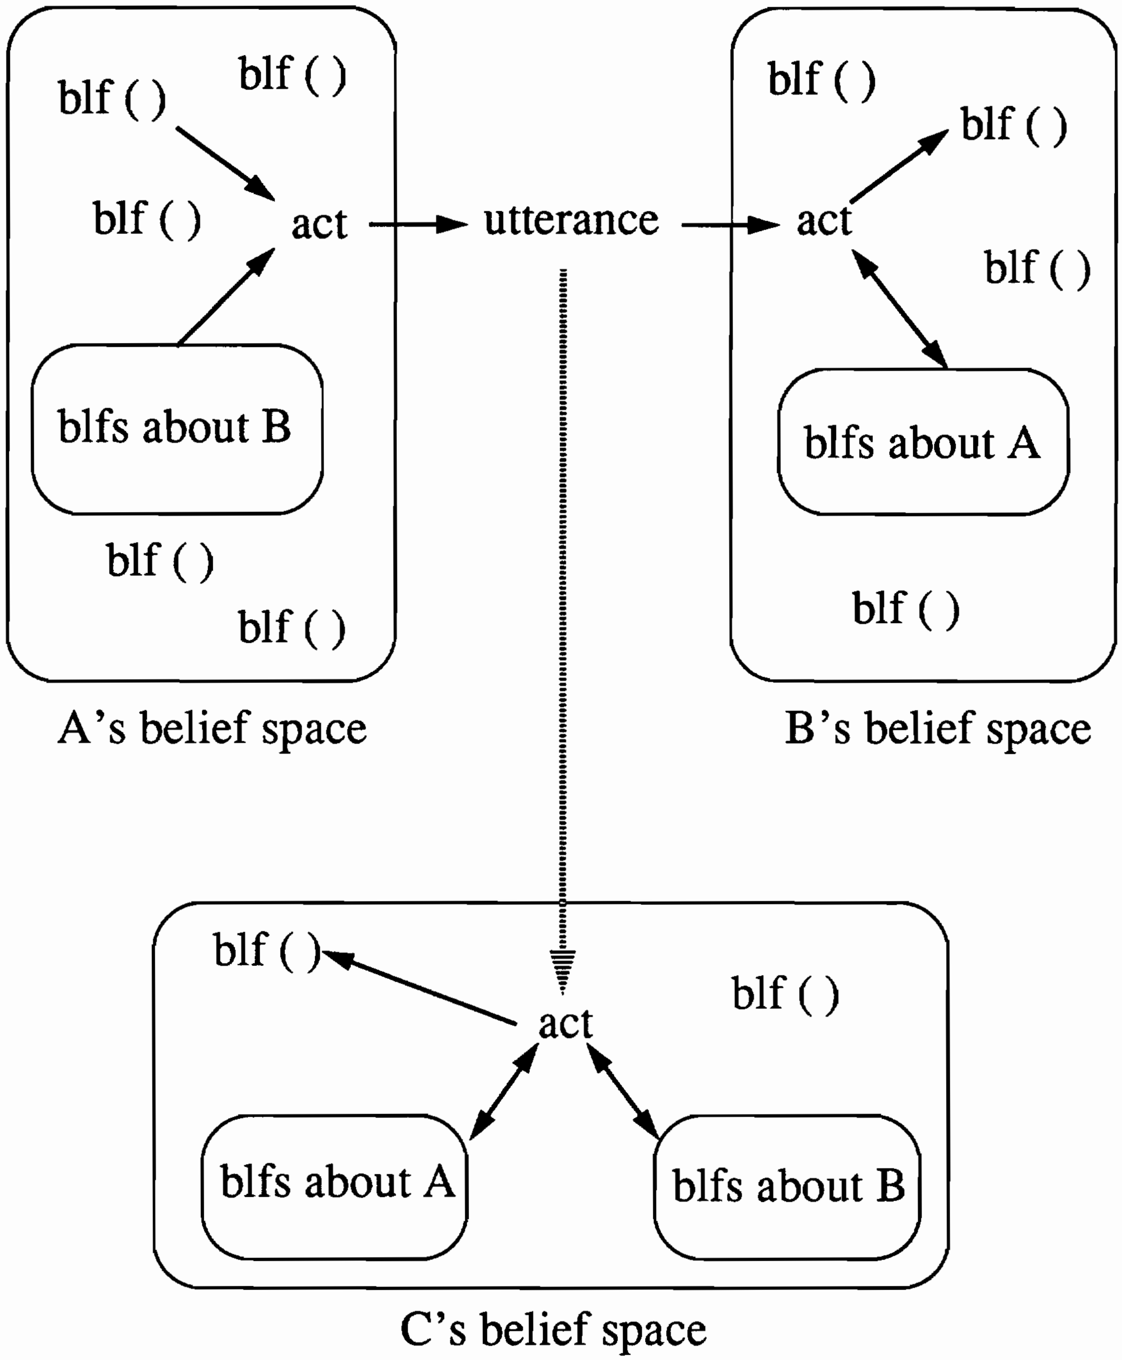

Figure 2. Model of Agent Interaction

beliefs may include beliefs about another agent's beliefs, shown in Figure 2 as smaller areas within an agent's belief space. Agents communicate in a multi-step process:

- Agent A forms the intention to perform an act directed toward Agent B. This intention is based on A's beliefs about B's beliefs along with other beliefs that A holds.
- A's intended act is expressed as an utterance, and it is this utterance that is transmitted to other agents in the system.
- B interprets this utterance as an act based on B's own beliefs and B's beliefs about A's beliefs. Note that if B's beliefs about A are in error or incomplete, B may infer a different act than A intended.
- Agent B's belief space is updated to reflect the effects that A's act had on B's beliefs.

Agent C represents an overhearer, an agent who hears A's utterance but is not the intended target (shown by a grey line from utterance to agent). Agent C interprets the overheard utterance in light of C's beliefs, including C's beliefs about the beliefs of A and B. Of course, C may arrive at an interpretation of the utterance that differs from that of either A or B.

4.3.2 Belief Representation

An agent's understanding of the world is represented as a set of beliefs of the form:

```
blf(belief-item,
    truth-value,
    [belief-group])
```

Each agent has a separate belief space. In determining the interpretation that an agent will place upon a particular utterance, only the beliefs within that agent's belief space are considered. For example,

```
approach:
        blf(airspeed([sun512],190),
            true,
            [[approach]])
```

```
blf(altitude([sun512],5000),
    true,
    [[approach]])
```

represents a portion of the belief space of an agent named Approach. Two beliefs are shown: that Approach believes Sundance 512 is travelling at an airspeed of 190 knots, and that Approach believes that the altitude of Sundance 512 is 5000 feet.

In this presentation it will be assumed that new beliefs are added to an agent's belief space. Old beliefs are retracted or replaced only if specified. In an implementation it would probably be desirable to retain a history of old beliefs with some indication of when and why they were retracted, so that the agents would be able to "remember" that they used to believe something.

4.3.3 Belief Items

The *belief-item* clause represents the item about which a belief is held. An agent may hold beliefs about the *state* of the world,.e.g., that United flight 101 is flying at an altitude of 5000 feet. An agent may also hold beliefs about acts that have occurred and about acts that agents intend, representing an agent's expectations of future actions and memory of past action. A *belief-item* clause, then, takes one of three forms:

intend ([*agent*], *act*)

occurred ([*agent*], *act*)

state (*state-data*)

An **intend** clause represents an agent's expectation of future action, while the **occurred** clause is used to represent an agent's memory of past events. For example:

```
approach:
      blf(intend([sun512],
                 contact ([approach])),
          true,
          [[approach]])
```

```
        blf(occurred([delta745],
                     contact ([approach])),
            true,
            [[approach]])
```

represents the approach controller's beliefs that Sundance 512 intends to contact the approach controller, and that Delta 745 has already done so.

The list notation for the agent in the **intend** and **occurred** clauses is intended to accommodate coordinated actions by groups of agents.

An *act* represents an agent's interpretation of observed events. Because this interpretation depends upon an agent's current beliefs, different agents may interpret the same event as different acts.

Belief-items of the form *state* (*state-data*) represent the agent's beliefs about the current state of the world. These *state* beliefs are assumed to be current and ongoing. For example,

```
approach:
        blf(airspeed([sun512],190),
            true,
            [[approach]])
```

would represent the approach controller's belief that Sundance 512 was flying at an air-speed of 190 knots. More complex beliefs are possible, e.g.:

```
approach:
        blf(intend( [sun512],
                   report( [approach],
                           blf(altitude([sun512],5000),
                               true,
                               [[sun512],[approach]])))),
        true,
        [[sun512],[approach]])
```

which would be read as: "The approach controller believes that the pilot of Sundance 512 intends to report that the airplane is at 5000 feet."

This representation does not accommodate conditions that are in the process of changing from one state to another, e.g., Sundance 512 is slowing from 190 knots to 170 knots. As mentioned earlier, it does not capture notions of time; a more general representation would add a time interval argument to indicate when the agent believed the action would or did occur (see, for example, [Allen 91b]). It also does not capture events that take place without being triggered by the actions of some agent. This limitation is acceptable for this representation because its purpose is to model dialogue and not to model general events.

The complete list of acts and states used in analyzing the example dialogue (Appendix E) are defined in Appendix D.

4.3.4 Truth Value

The *truth-value* clause may take the values `true`, `false`, or `unknown`. Clauses which do not appear in the agent's belief space are considered to have a truth value of `unknown`.

The *truth-value* clause does not directly indicate the agent's opinion about the truth the *belief-item* clause; instead, it reflects the agent's understanding of the beliefs of the agents in the *belief-group*. This indirection allows an agent to hold contradictory beliefs about other agents' beliefs. For example, assume the following clause appeared in the approach controller's belief space:

```
approach:
        blf(altitude (sun512,5000),
            true,
            [[sun512]])
```

This would represent a belief that the approach controller had about the belief of the pilot of Sundance 512, and might be read something like this: "Approach believes that Sundance

512 thinks that the aircraft is at an altitude of 5000 feet." Note that this belief would not conflict at all with the following:

```
approach:
        blf(altitude(sun512,4500),
            true,
            [[approach]])
```

and both beliefs could easily occur together in the approach controller's belief space:

```
approach:
        blf(altitude(sun512,4500),
            true,
            [[approach]])

        blf(altitude(sun512,5000),
            true,
            [[sun512]])
```

This might be read as: "The approach controller believes that Sundance 512 is at an altitude of 4500 feet, although the controller realizes that the pilot thinks that the airplane is at 5000 feet." Time for an altimeter check, clearly.

The *truth-value* representation does not capture the infinite gradations of uncertainty that people are certainly capable of expressing. In this domain, however, agents have unusually strong expectations about each others' beliefs and intentions, so the simplified range of truth values (`true`, `false`, `unknown`) is sufficient for the current purpose.

4.3.5 Belief Groups

The *belief-group* represents sets of agents who mutually hold a belief. A mutually-held belief is one that has been explicitly confirmed by the participating agents. For example:

conversational state is represented in terms of the agents' separate and possibly conflicting beliefs about the beliefs of other agents. An agent's belief associates a *belief-item* and *truth-value* to a *belief-group* of agents who, in the opinion of the believing agent, assign that *truth-value* to the *belief-item*. This indirection allows an agent to hold contradictory beliefs about other agents' beliefs. The use of a flexible set representation makes this model able to represent mutuality and inconsistency of belief in conversations involving more than two agents.

Chapter 5

Analysis

5.1 Introduction

The previous chapter described a computational representation designed to support language understanding by modeling dialogue in terms of beliefs and intentions. In this chapter, the use of this representation is illustrated by showing its application to a typical dialogue in the air traffic control domain. A complete actual dialogue between an air traffic controller and a pilot was explicated at the speech act level in terms of the conversants' evolving beliefs and expectations. The full analysis and a detailed discussion may be found in Appendix E. This chapter discusses the analysis at a more general level, focussing on the characteristic patterns seen in the dialogue and on the representation issues that arose during the analysis.

5.2 The Dialogue

The dialogue used in this example was selected from a larger corpus of air traffic control dialogue [Ward 90b]; the dialogue is shown in its entirety in Figure 3. It represents a complete conversation between an approach controller at Portland International Airport and the pilot of a commercial flight approaching the airport to land. Controller and pilot are cooperatively performing the task of guiding the aircraft through the controller's airspace to the approach gate, a point from which the pilot can begin flying a standard ILS landing procedure. As discussed in Chapter 3, the ILS approach procedure consists of the following steps:

(108)	Sun 512:	Approach Sundance five twelve's with you at five thousand, one hundred on the heading one ninety on the speed.
(109)	Approach:	Sundance five twelve Portland Approach roger.
(110)	Approach:	Sundance five twelve maintain speed one seven zero.
(111)	Sun 512:	One seven zero, Sundance five twelve.
(114)	Approach:	Sundance five twelve turn left heading zero five zero.
(115)	Sun 512:	Left zero five zero Sundance five twelve.
(125)	Approach:	Sundance five twelve turn left heading zero one zero.
(126)	Sun 512:	Left zero one zero Sundance five twelve.
(127)	Approach:	Sundance five twelve descend and maintain four thousand five hundred.
(128)	Sun 512:	Four thousand five hundred, Sundance five twelve.
(131)	Approach:	Sundance five twelve eight miles from Laker turn left heading three one zero. Maintain four thousand five hundred till established localizer cleared ILS two eight right approach, maintain one seven zero knots until Laker.
(132)	Sun 512:	Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.
(157)	Approach:	Sundance five twelve contact tower one one eight point seven.
(158)	Sun 512:	Five twelve good-bye.

Utterance numbers are from the original transcript [Ward 90b]. See Appendix C for context.

Figure 3. Dialogue between Sundance 512 and Portland Approach

- The pilot establishes communications with the approach controller. In this dialogue, this takes place in utterances 108 and 109.
- The controller directs the aircraft to the approach gate by issuing a series of vectors, seen here in utterances 110 through 128.
- The controller authorizes the pilot to begin the approach procedure (utterances 131 and 132).
- The controller hands off the flight to the tower controller (utterances 157 and 158).

The ILS approach task was selected because it represents a common task and a common form of interaction between pilot and controller. The 30 minutes of dialogue transcribed in [Ward 90b] include seven complete examples of the ILS approach task and portions of five more. Thus there are several samples available for analysis, and more can be readily obtained. Also, the task is short; in the seven complete examples, the task is accomplished in just 13 to 29 utterances, with all but one being completed in fewer than 20 utterances.

The Sundance 512 dialogue was selected for use in developing and illustrating the initial representation because it is a good baseline dialogue for this task. It depicts each step of the ILS approach procedure without including any unusual occurrences that would unnecessarily complicate the initial analysis. Also, it does not include exchanges that are not central to the ILS approach task, such as reports of other traffic or references to exchanges with other pilots.

The absence of references to other dialogues is particularly important because for this analysis the dialogue was removed from its original context. As was mentioned earlier, ATC dialogue consists of largely separate, interleaved conversations. Between the exchanges of this dialogue, the controller is talking with several other pilots. Pilots can hear the controller's conversations with other pilots, so they are forming and updating beliefs about the actions and intentions of other pilots. Often these beliefs do not affect the pilot's own actions, and when that is the case the conversations may be extracted and analyzed separately with minimal distortion.

5.3 Characteristic Patterns in ATC Communications

The Sundance 512 dialogue includes several exchanges that are characteristic of ATC communications: confirmation exchanges (utterances 108 and 109), directions (utterances 110 through 128), explicit acknowledgments (utterances 109 and 111, for example), and complex transmissions (utterances 108, 131, and 132). In this section, the confirmation exchange will be examined in detail to illustrate the reasoning behind the analysis in Appendix E. The other patterns will be discussed in more general terms.

5.3.1 Confirmation Exchanges

ATC dialogue is unusual in that the conversants often hold strong expectations about what the other knows and what the other will say. The ILS approach procedure is particularly predictable, even for an ATC task. Pilot and controller are both trained in ILS approach procedures; each knows at least approximately what the other should do under normal circumstances. The system is not infallible, however, and circumstances are not always normal, so a certain amount of conversation is directed toward confirming information that each thinks the other probably already knows or expects. In tracking the conversants' changing beliefs during such exchanges, the following pattern will frequently be seen:

- Pilot and controller separately form similar beliefs.
- One conversant attempts to confirm the belief.
- The conversants each update their beliefs to reflect the newly-established mutuality of the belief.

This pattern is discussed in more detail in the following paragraphs.

Initially, pilot and controller will separately form similar beliefs, generally derived from different evidence. Each expects the other to hold similar beliefs, but they have not yet confirmed that this is so. For instance, the pilot may believe that the aircraft is flying at an airspeed of 190 knots and may think that the controller should know this from the information displayed on the radar console. Because this has not been confirmed — that is, it is not yet mutually known — the belief representation will show the pilot and controller in different belief groups, like this:

```
blf (airspeed ([sun512], 190),
     true,
     [[sun512], [approach]])
```

If this belief appeared in the pilot's belief space, it might be read "the pilot of Sundance 512 believes that pilot and controller each believe that Sundance 512's airspeed is 190 knots. Pilot and controller have not yet confirmed this with each other, however."

One conversant will then attempt to confirm the belief. This attempt is motivated primarily by the fact that FAA procedures recommend or require this confirmation. There is also a safety consideration, however, in that the evidence from which the beliefs were derived may be in error: flight progress slips may be incorrectly entered, flight conditions may have changed, instruments may be miscalibrated. Thus the conversants are highly motivated to ensure that safety-critical beliefs are mutually understood. A common method of confirming a belief is to report the belief and give the other an opportunity to confirm or correct it. One conversant therefore forms the intention to report the belief to the other, resulting in a belief structure that follows this pattern:

```
blf (intend ([sun512],
             report ([approach],
                     blf (airspeed ([sun512], 190),
                          true,
                          [[sun512], [approach]]))),
     true,
     [[sun512], [approach]])
```

This structure represents an agent's belief that the pilot of Sundance 512 intends to confirm with the approach controller the belief that Sundance 512's airspeed is 190 knots. The agent believes that the pilot intends to make this report and that the controller expects it. Until knowledge of the airspeed is confirmed to be mutual, however, pilot and controller are represented in separate belief groups. This belief is an example of the nested belief structure discussed in the previous chapter. The inner belief clause represents the belief that will be reported, while the outer belief clause represents the agent's belief that pilot and controller expect the report to occur. Note that this belief could appear in any agent's belief space, representing that agent's expectation that the pilot will make the report.

The agent's intention is expressed through some communicative act, often by speaking a concise phrase identifying the belief. In this model, the intentions of the speaker are represented in terms of acts that motivate the phrases that the speaker selects for the communication. For example, the pilot's intention to report the airspeed is seen as motivating the phrase **one ninety on the speed** that concludes utterance 108.

Upon hearing the utterance, the hearer attempts to infer the motivating acts that underlay the speaker's choice of words. The acts intended by the speaker may not always match the acts inferred by the hearer, resulting in misunderstandings. This conversation contains no evidence of misunderstanding, however, so in depicting the speech acts of this dialogue only the intended acts of the speaker will be presented; the hearer's interpretation will be assumed to be similar.

After an utterance is made, the speaker's beliefs are updated. These updates typically reflect:

- the speaker's belief that the act occurred,
- expectations that the hearer will give some evidence of understanding (acknowledgment),
- expectations that the hearer will give some evidence of agreement (confirmation).

Upon successful completion of an utterance, the speaker's belief space is updated to reflect the speaker's perception that the intended acts did indeed take place, e.g.,

```
blf  (occurred ([sun512],
                report ([approach],
                        blf (airspeed ([sun512], 190),
                            true,
                            [[sun512], [approach]]))),
     true,
     [[sun512], [approach]])
```

As discussed in the previous chapter, "occurred" represents the agent's memory of past events. This structure, then, would represent an agent's belief that the pilot had attempted to confirm Sundance 512's airspeed; the controller has not yet confirmed this, though.

The speaker also forms expectations that the hearer will respond with evidence of understanding, some acknowledgment of the attempted act:

```
blf  (intend ([approach],
            acknowl ([sun512],
                     report ([approach],
                             blf (airspeed ([sun512], 190),
                                 true,
                                 [[sun512], [approach]])))),
     true,
     [[sun512], [approach]])
```

That is, the agent expects the approach controller to acknowledge the pilot's report. Note that this does not represent an expectation of agreement, only expectation that the utterance will be acknowledged. The distinction is perhaps subtle, but it is necessary to express the distinction between agreeing with a statement and merely understanding it.

In this dialogue, the conversants are attempting to confirm information that the other should already have. The speaker will therefore also form strong expectations that the hearer will in fact confirm the speaker's statement:

```
blf (intend ([approach],
            report ([approach],
                   blf (airspeed ([sun512], 190),
                       true,
                       [[sun512, approach]]))),
     true,
     [[sun512], [approach]])
```

This structure represents the agent's expectation that the approach controller will confirm the report. Notice that only now does the representation begin to reflect mutuality of knowledge; the agent expects that the controller will confirm that Sundance 512 is expected at PDX (the inner, mutual belief) but has not yet received the confirmation (the outer, non-mutual belief).

The speaker also believes that the intended speech act has been attempted, although it is not yet known whether the act was successful, that is, whether the intended recipient actually heard and correctly interpreted the utterance. The belief that originally represented the speaker's intention to perform the speech act, thereby serving as its motivation, has served its purpose and may be removed. If the act was in fact unsuccessful, then the speaker will become aware of that fact when the hearer fails to acknowledge the utterance appropriately. At that time, an intention to re-attempt the communication can be reinstated from the unfulfilled expectation of acknowledgment.

Meanwhile, the beliefs of any agents who hear the utterance must also be updated. If the transmission was understood and if the agents hold the expected beliefs — or at least do not hold conflicting beliefs — the effect on the hearers will parallel the changes made in the speaker's beliefs:

- The hearers interpret the utterance as reflecting the speaker's intention to commit certain speech acts.
- The hearers form expectations that the agent being addressed will give some evidence of understanding.
- The hearers form expectations that the agent being addressed will give some evidence of agreement.

For simplicity, only the effects on the beliefs of the agent being addressed will be tracked in this analysis (Appendix E), but in general it will be necessary to model the beliefs of all agents able to hear the transmission.

5.3.2 Directions

Another common form of interaction occurs when a controller directs the pilot to perform some action. In utterances 110 through 128 (Figure 3), for example, the approach controller guides the aircraft through a series of speed, heading, and altitude changes until it reaches a point from which it can begin flying the published approach procedure.

Recall that in the confirmation exchange, pilot and controller drew on different sources of information to form similar beliefs. Directions, however, are marked by imprecise expectations on the part of the pilot. The pilot does not usually know in advance exactly what directions to expect, only that directions will be given. In the dialogue analyzed here, for instance, the controller is guiding the pilot through the ARSA to the approach gate. The pilot probably does not know the precise route that the controller will use in directing the aircraft to the approach gate. The pilot can only expect an indefinite number of directions that culminate in a final turn near the approach gate. This indefinite expectation is represented by the notation <change> appearing in place of specific values for airspeed, direction, and altitude. For example:

```
blf (intend ([approach],
            direct ([sun512],
                    direction ([sun512], <change>))),
      true,
      [[sun512], [approach]])
```

represents an agent's belief that the approach controller will direct Sundance 512 to change heading, although the agent does not know exactly what the new heading will be. These generalized expectations differ from the more specific expectations seen in the confirmation exchange in that they are not considered fulfilled and withdrawn when the pilot has heard a single instruction of that type. Instead, the expectations remain in force until the pilot is cleared onto the approach course, the point at which the pilot is authorized to begin flying the approach procedure without further instructions from the controller.

5.3.3 Acknowledgments

In the dialogue in Figure 3, utterances occur in pairs; each initial transmission is verbally acknowledged by the other. This acknowledgment often consists of a *readback*, a verbatim repetition of a portion of the first speaker's presentation:

(114) Approach: Sundance five twelve turn left heading zero five zero.

(115) Sun 512: Left zero five zero Sundance five twelve.

Notice that the controller does not reply to the pilot's acknowledgment, nor does the pilot expect a reply. This is in accordance both with recommended ATC procedure and with Clark's rules for evidence of understanding [Clark 89]. The initial acknowledgment was made in the form of the fairly strong *display* of understanding. It should be followed by some weaker evidence of understanding, either *initiation of the next relevant contribution* (in this case, there was none) or with *continued attention* (silence). Strictly speaking, this would be represented as an expectation of acknowledgment that is fulfilled after the pilot has passed up an opportunity to initiate a correction or clarification to the controller's contribution. For simplicity, this step will not be represented explicitly in the analysis in Appendix E.

5.3.4 Complex Transmissions

As was discussed in Section 3.2.2, the radio channel is a constrained resource, especially in the crowded airspace around busy airports. Conversants respond to the crowded channels by packing multiple speech acts into a single complex transmission which may be hard to understand and which potentially requires an equally complex response. The half-duplex modality makes the overhead of providing conversational feedback high, yet safety considerations make confirmation vital. The result is a tension between the need for communicative efficiency and the need for a high degree of assurance that mutuality has been achieved.

This tension can be seen in two complex transmissions found in the Sundance 512 dialogue. In the first exchange, the conversants' strong expectations and the mutuality of their belief permit an extremely terse, efficient response to a complex transmission. In the second example, less-precise expectations prompt a more verbose response, although mutuality of belief still explains several omissions in the readback. These contrasting examples will be examined more closely in Sections 5.3.4.1 and 5.3.4.2. In Section 5.3.4.3, the representation issues raised by these complex transmissions will be discussed.

5.3.4.1 Communicative Efficiency

The constraints of the radio modality encourage the use of complex transmissions; instead of contributing a single item and waiting for evidence of understanding before continuing, conversants tend to package multiple speech acts together and transmit them in a single rapid-fire utterance. For example, the dialogue opens with this exchange:

(108)	**Sun 512:**	**Approach Sundance five twelve's with you at five thousand, one hundred on the heading one ninety on the speed.**
(109)	**Approach:**	**Sundance five twelve Portland Approach roger.**

In utterance 108, the pilot intends four speech acts. The pilot's primary purpose is to establish contact with the controller; this is signalled by the use of the phrase **with you**. As part of the same transmission, however, the pilot also attempts to confirm that the altitude, direction, and airspeed information shown on the cockpit instruments matches the information being displayed on the controller's console.

In this case, the controller is able to respond quite efficiently to the pilot's complex transmission. The phrase **Sundance five twelve Portland Approach** formally acknowledge and confirm that contact has been made, and the single word **roger** suffices to acknowledge and confirm the pilot's other three speech acts. This communicative efficiency is possible because of the strong expectations and beliefs that underlie this exchange. Each conversant believes that the other should know the aircraft status, although each derives the information from a different source. Both therefore strongly expect the controller to acknowledge the acts in utterance 108 and to confirm the altitude, heading, and airspeed information. These strong expectations will encourage both parties to accept even moderately weak evidence of understanding as acknowledgment; anything less than an explicit statement of disagreement is likely to be taken as both confirmation and agreement. This is reflected in the terseness of the controller's reply. In the absence of evidence to the contrary, the strong expectations set up by procedures and by the multiple sources of information within the ATC system permit the conversants to interpret an extremely efficient response correctly.

What might have constituted "evidence to the contrary" had the controller not intended to acknowledge and confirm the utterance? The problem-filled conversation shown in Figure 1 offers an example. In utterances 82 and 83, the controller understood enough of the transmission to realize that someone had attempted to contact Portland Approach. The controller was unable to understand the aircraft identification, however:

(82) Ford 645: **Portland Approach, Ford Trimotor niner six four five.**

(83) Approach: **((garbled)) Portland Approach say again?**

The controller's intentions in this exchange would be represented as:

```
blf (intend ([approach],
            acknowl( [<unknown>],
                    contact ([approach])),
    true,
    [[approach]])

blf (intend ([approach],
            report ([<unknown>],
                  blf (occurred ( [<unknown>],
                                  contact ([approach])),
                      false,
                      [[approach]])),
    true,
    [[approach]])
```

The controller acknowledged the attempt to contact by replying **Portland Approach**, but indicated that the attempt was unsuccessful with the request for clarification **say again?** The formal contact is not successfully completed until the controller finally confirms the aircraft identification in utterance 91 (Figure 1).

5.3.4.2 Assurance of Mutuality

The opening exchange of the Sundance 512 dialogue demonstrates the communicative efficiency that can be achieved when conversants are confirming information that each believes the other already knows. In contrast, utterances 131 and 132 are typical of the exchanges that occur when mutuality is less certain. In this exchange, the approach controller authorizes the pilot of Sundance 512 to begin the approach procedure:

(131)	**Approach:**	**Sundance five twelve eight miles from Laker turn left heading three one zero. Maintain four thousand five hundred till established localizer cleared ILS two eight right approach, maintain one seven zero knots until Laker.**
(132)	**Sun 512:**	**Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.**

The controller's transmission is the longest in the entire dialogue, using 34 words to express five speech acts:

- A report of the aircraft position (**eight miles from Laker**).
- An instruction to the pilot to turn to intercept the final approach course (**turn left heading three one zero**).
- An authorization for the pilot to execute a particular approach procedure (**cleared ILS two eight right approach**).
- Altitude and speed assignments to be observed until the aircraft is established on a segment of the approach procedure (**Maintain four thousand five hundred till established localizer ... maintain one seven zero knots until Laker**).

Despite the amount of information contained, this transmission is delivered at the same rapid-fire pace as the rest of the dialogue. One would expect this to be difficult or impossible to follow, and in fact non-pilots do find this utterance to be nearly unintelligible [Novick 92]. How does the pilot understand it? One explanation is that although the pilot probably cannot anticipate exactly what instructions the controller will give, the pilot can strongly expect to receive the instructions in the order given. The arrival instructions are largely boilerplate; FAA procedure requires the controller to include all but the speed assignment. The standardized format helps the pilot to focus on the most important words: **eight miles, heading three one zero, two eight right, four thousand five hundred,** and **one seven zero knots**.

The pilot's response in utterance 109 supports the boilerplate hypothesis and points up the importance of mutuality of belief in explaining the acceptability of a non-standard response. Notice that the pilot's response consists mostly of the "important" words identified above:

(132)	**Sun 512:**	**Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.**

There are notable omissions in the pilot's readback of this long and complex transmission: the identifying words **heading** and **knots** have been dropped; the pilot says the word **speed** but fails to repeat the actual speed assignment; the position report has disappeared altogether. Despite these omissions, the controller is apparently satisfied with the response; there is no request for clarification.

This non-standard response — and its acceptability to the controller — can be explained in terms of the conversants' ability to make use of the growing mutuality of knowledge to increase the efficiency of their exchange. The pilot gives the strong *display* evidence of understanding to show that the new information has been successfully understood, but omits the redundant words **heading** and **knots** for the sake of communicative efficiency; the controller is expecting this information, so it is easily interpretable in the current context of the conversation. The pilot acknowledges the speed restriction even more generally with **we'll maintain uh, speed**, indicating the pilot's understanding that the current aircraft speed is not to change. Note that neither the actual speed assignment nor the position report is read back. Both of these are items of information that each conversant believes the other already knows; mutuality of knowledge permits them to be understandably omitted.

Another example of general conversational principles explaining an acceptable deviation from standard phraseology is seen in the final exchange in the dialogue:

(157) **Approach:** **Sundance five twelve contact tower one one eight point seven.**

(158) **Sun 512:** **Five twelve good-bye.**

The pilot's response **good-bye** is not approved phraseology. It is, however, a *relevant next contribution* [Clark 89], and as such strongly signals understanding and acceptance of the controller's contribution.

5.3.4.3 Representation of Complex Transmissions

How should complex transmissions such as utterance 131 be represented? This analysis presents them in terms of multiple speech acts (e.g., utterance 131 in Table 46). If the conversants do in fact consider the first part of utterance 131 to be a single long standard phrase[1] then perhaps it might be better expressed as a composite instruction consisting of four sub-instructions that may, in other conditions, occur independently. This would argue for a hierarchy of speech acts that is not captured by this model.

The main contribution of the notion of a composite instruction would be the expectation that all components will be issued together. This same expectation can be captured without introducing the composite act by appealing to the task structure and the need for mutually confirming critical information: this is the last vectoring instruction that the controller will issue before the pilot begins executing the published approach procedure. While executing the approach procedure, the pilot will maneuver without explicit instructions from a controller, so the controller takes this last opportunity to confirm certain critical information — the aircraft position — before stating the altitude, speed, and direction assignments that are to be observed until the pilot can begin the standard approach procedure. Thus, the complex instruction can be motivated in terms of task within the simpler, non-hierarchical representation.

5.4 Shortcomings

The belief representation presented in this thesis fails to capture certain nuances of the Sundance 512 dialogue, particularly in representing the imprecise expectations that mark the pilot's acceptance of the controller's instructions. It does not capture "reasonable-

1. That is, the utterance up to the speed assignment. The speed assignment is not part of the standard arrival instructions, but may have been added to ensure that minimum aircraft separation requirements were met. Notice that it occurs after the standard parts of the arrival instructions and is preceded by a slight pause, which is consistent with the view that the controller did not think of it as part of the arrival instructions but as a separate instruction.

ness checks" in the conversants' expectations; for instance, the pilot does not expect to be directed to slow to 40 knots and would certainly protest if asked to do so.

Strength or immediacy of expectation is not captured either. For instance, a pilot flying under instrument flight rules may receive instructions to change course, speed, or altitude at any time. At some level, then, a pilot must always hold "background" expectations that capture this expectation while reflecting the pilot's lack of foreknowledge as to when directions will be given and precisely what they will be. At this stage of the flight, these expectations are intensified and made more immediate; instead of merely recognizing that the controller *may* issue vectors, the pilot now positively expects that the controller *will* issue vectors to guide the aircraft to the approach gate. In fact, the pilot would be quite surprised if the controller failed to do so. This representation does not differentiate between the expectation that something may reasonably occur and the expectation that something should certainly occur. This distinction has implications for motivating and explaining dialogue. One would not expect someone to question the lack of an instruction that was recognized as possible but not strongly expected. A pilot would almost certainly say something if the controller failed to begin issuing vectors to the approach within a reasonable length of time, though.

Because this model does not capture time or more general conditional instructions, the "until" clauses seen in utterance 131 are not represented. Although this is not adequate for ATC instructions in general, it is acceptable in this case; the "until" clauses refer to conditions that would normally supercede the controller's instructions anyway. Once established on the localizer, the pilot is expected to descend along the glideslope defined by the localizer. After passing the Laker beacon, the pilot should be adjusting speed in accordance with approach procedure and the performance characteristics of the aircraft. A more general representation might include a "conditions" clause on all instructions, however.

5.5 Summary

This chapter discussed the issues that arose in modelling a complete ATC dialogue in terms of the beliefs and goals of the conversants. From this analysis, typical patterns emerged. Conversants separately form similar beliefs based on separate inputs, then attempt to confirm the mutuality of the belief. As utterances are made and acknowledged, speaker and hearer form beliefs about what speech acts are intended by the utterance and form expectations that the act will be acknowledged and confirmed. These expectations form a strong context that permit the conversants to easily recognize the speech acts motivating terse responses that might otherwise be ambiguous. Finally, the fulfilled expectations are deleted and beliefs about the state of the world are updated to reflect the newly-confirmed information.

The model succeeds in accounting for several aspects of the dialogue that a strict reading of the approved FAA phraseology does not. The abbreviated readbacks seen un utterances 109 (discussed in appendix E.2), utterance 132 (appendix E.4), and utterance 158 (appendix E.5) are non-standard, but the strong expectations of the conversants coupled with general conversational principles allowed the hearer to correctly deduce the acts intended by the speaker. In utterance 110, a potentially confusing instruction to "maintain" a new speed was correctly interpreted in a similar fashion (discussed in appendix E.3).

The model does not, however, capture several nuances of the exchange, notably reasonableness limits in the conversants expectations, immediacy or intensity of belief, time, and conditional instructions.

Chapter 6

Conclusions

The goal of this work is to develop a computational model for dialogue at the speech act level in a real-world domain. Such models are expected to prove useful in future spoken language understanding systems for two reasons. First, dialogue models provide a means for tracking the state of the dialogue at a high level, which may help limit the number of possibilities that the speech recognition component must consider in interpreting the acoustic signal. Second, speech act models provide a basis for linking an utterance to the speaker's goals and beliefs, thus providing a basis for reasoning about the speaker's intended meaning.

The domain chosen for this study was that of air traffic control dialogue. ATC dialogue exhibits several characteristics that make it particularly suitable for dialogue modeling: Context, roles, and vocabulary are unusually structured; tasks are short and well-documented; the reliance on radio communications minimizes such difficult-to-capture phenomena as non-verbal communications and overlapping speech. Despite these characteristics, ATC dialogue is unconstrained conversation between people engaged in real-world tasks.

This thesis presents a computational representation in which the evolving context of the dialogue is represented in terms of the mutuality of belief established by the conversants. An agent's perception of the dialogue is represented as a set of beliefs about:

- Acts that agents have performed,
- Acts that agents intend to perform,

- The state of objects in the domain.

Utterances that the agent hears or produces are interpreted or intended as speech acts in light of that agent's beliefs. Each agent's belief space is separate; thus, agents may hold conflicting beliefs and may place conflicting interpretations upon the same utterance (misunderstanding).

A key contribution of this model is the representation of mutuality of belief in terms of belief sets. Belief sets capture both an agent's understanding of who believes a given piece of information and the mutuality that the agents holding that belief have established among themselves. This representation allows an agent to hold beliefs about other agents' possibly conflicting beliefs ("I believe A but he believes B.") as well as allowing agents to hold beliefs about the mutuality of their knowledge ("She and I have established that we both believe A; he should also believe A, but we have not yet confirmed that."). The belief set representation is flexible enough to represent and reason about mutuality combinations involving any number of agents, thus supporting the modelling of multi-agent conversations.

To test and refine the initial model, radio exchanges between air traffic controllers and pilots were taped and transcribed. A complete dialogue, consisting of all exchanges between the controller and the pilot of a commercial flight approaching the airport to land, was selected for detailed analysis. A small set of speech acts and associated domain facts was developed to represent the dialogue events seen in that conversation. The dialogue was then explicated at the speech act level in terms of the beliefs and intentions of the conversants. From this analysis, it can be seen that the model captures sufficient conversational context to motivate and explain actual dialogue.

This study is part of a larger effort directed toward improving the performance of spoken language understanding systems. Future work, then, will include expanding this model to account for a larger subset of ATC dialogue and incorporating the resulting model into a working spoken language understanding system. The first steps in this long-range goal will be to implement a working example of this model and to define the rules needed to effect the transition from one conversational state to the next. These will be tested against a broader sample of ATC dialogue using saso, a rule-based shell for modeling multi-agent interactions [Novick 90].

Bibliography

[Allen 89] James Allen, Stephane Guez, Louis Hoebel, Elizabeth Hinkelman, Keri Jackson, Alice Kyburg, and David Truam, "The Discourse System Project," Technical Report 317, University of Rochester (1989).

[Allen 91a] James F. Allen, "Discourse Structure in the TRAINS Project," *Proceedings of the Fourth DARPA Workshop on Speech and Natural Language* (Feb. 1991).

[Allen 91b] James Allen, "Time and Time Again: The Many Ways to Represent Time," *International Journal of Intelligent Systems*, Vol. 6, pp. 341-359 (1991).

[Austin 62] J. L. Austin, *How to Do Things with Words*, Oxford University Press, London (1962).

[Clark 81] Herbert H. Clark and Catherine R. Marshall, "Definite Reference and Mutual Knowledge," pp. 10-63 in *Elements of Discourse Understanding*, ed. Ivan A. Sag, Cambridge University Press, Cambridge (1981).

[Clark 86] Herbert H. Clark and Deanna Wilkes-Gibbs, "Referring as a Collaborative Process," *Cognition*, Vol. 22, pp. 1-39 (1986).

[Clark 89] Herbert H. Clark and Edward F. Schaefer, "Contributing to Discourse," *Cognitive Science*, Vol. 13, pp. 259-294 (1989).

[Cohen 79] Philip R. Cohen and C. Raymond Perrault, "Elements of a Plan-Based Theory of Speech Acts," *Cognitive Science*, Vol. 3(3), pp. 177-212 (1979).

[Cohen 84] Philip R. Cohen, "The Pragmatics of Referring and the Modality of Communication," *Computational Linguistics*, Vol. 10(2), pp. 97-146 (Apr 1984).

[Cook 90] Guy Cook, "Transcribing Infinity: Problems of Context Presentation," *Journal of Pragmatics*, Vol. 14, pp. 1-24 (1990).

[Doddington 85] George R. Doddington, "Speaker Recognition - Identifying People by their Voices," *Proceedings of the IEEE*, Vol. 73(11), pp. 1651-1664 (Nov. 1985).

[Doring 83] B. Döring and A. Knauper, "A Simulation Study with a Combined Network and Production Systems Model of Pilot Behavior on an ILS-approach," *Automatica*, Vol. 19(6) pp. 741-747 (1983).

[Dunham 91] Scott Dunham, personal communication, June 10, 1991.

[FAA 89] FAA, *Air Traffic Control*, (7110.65F) (1989).

[FAA 91] FAA, *Airman's Information Manual* (1991).

[Fanty 92] Mark Fanty, Ronald A. Cole, and Krist Roginski, "English Alphabet Recognition with Telephone Speech," *Advances in Neural Information Processing Systems 4*, ed. John E. Moody, Steven J. Hanson, and Richard P. Lippmann, Morgan Kaufmann Publishers, San Mateo, CA (1992).

[Fink 86] Pamela K. Fink and Alan W. Biermann, "The Correction of Ill-Formed Input using History-Based Expectation with Applications to Speech Understanding," *Computational Linguistics*, Vol. 12(1) (Jan. 1986).

[Fischetti 86] Mark A. Fischetti, Tekla S. Perry, Paul Wallich, Glenn Zorpette, John Voelcker, and John Horgan, "Our Burdened Skies," *IEEE Spectrum* (Nov. 1986).

[Goguen 83] Joseph A. Goguen and Charlotte Linde, "Linguistic Methodology for the Analysis of Aviation Accidents," NASA contractor report 3741 (7110.65F), National Aeronautics and Space Administration (1983).

[Goodine 91] David Goodine, Stephanie Seneff, Lynette Hirschman, and Michael Phillips, "Full Integration of Speech and Language Understanding in the MIT Spoken Language System," *Second European Conference on Speech Communication and Technology* (Sept. 1991).

[Grimshaw 80] Allen D. Grimshaw, "Mishearings, Misunderstandings, and Other Nonsuccesses in Talk: A Plea for Redress of Speaker-Oriented Bias," *Sociological Inquiry*, Vol. 50(3/4), pp. 31-74 (1980).

[Grosz 86] Barbara J. Grosz and Candace L. Sidner, "Attention, Intentions, and the Structure of Discourse," *Computational Linguistics* Vol. 12(3) pp. 175-204 (1986).

[Hall 88] Robert Francis Hall, "Voice Recognition and Artificial Intelligence in an Air Traffic Control Environment," AFIT/CI/NR 88-171 (1988). Master's Thesis.

[Harold 91] Lloyd Harold, FAA Academy, personal communication, Jan. 29, 1991.

[Hauptmann 88] Alexander G. Hauptmann, Sheryl R. Young, and Wayne H. Ward, "Using Dialog-Level Knowledge Sources to Improve Speech Recognition," *AAAI '88*, pp. 729-733 (1988).

[Heppenheim 90] T. A. Heppenheimer, "Air Traffic Control," *Aviation Week & Space Technology*, pp. S1-S11 (Dec. 3, 1990).

[Hovy 91] Eduard H. Hovy, "Pragmatics and Natural Language Generation," *Artificial Intelligence* Vol. 43, pp. 153-197 (1990).

[Jelinek 90] Federick Jelinek, "Self-Organized Language Modeling for Speech Recognition," pp. 450-506 in *Readings in Speech Recognition*, ed. Alex Waibel and Kai-Fu Lee, Morgan Kaufmann Publishers, San Mateo, CA (1990).

[Lee 90] Kai-Fu Lee, Hsiao-Wuen Hon, and Raj Reddy, "An Overview of the SPHINX Speech Recognition System," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 38(1), pp. 35-45 (Jan. 1990).

[Litman 87] Diane J. Litman and James F. Allen, "A Plan Recognition Model for Subdialogues in Conversations," *Cognitive Science*, Vol. 11, pp. 163-200 (1987).

[Malcolm 90] Janet Malcolm, "The Morality of Journalism," *The New York Review*, p. 20 (Mar. 1, 1990).

[Novick 88] David G. Novick, *Control of Mixed-Initiative Discourse through Meta-locutionary Acts: A Computational Model* (Technical Report No. CIS-TR-88-18). Department of Computer and Information Science, University of Oregon (1988).

[Novick 90] David G. Novick, "Modeling Belief and Action in a Multi-Agent System," pp. 34-41 in *AI, Simulation and Planning in High Autonomy Systems*, ed. Bernard Zeigler and Jerzy Rozenblit, IEEE Computer Society Press, Los Alamitos, CA (1990).

[Novick 92] David G. Novick, Ron Cole, Benjamin Corliss, *The Effect of Context on the Intelligibility of Air Traffic Control Communication*, in preparation (1992).

[Oviatt 88] Sharon L. Oviatt and Philip R. Cohen, "Discourse Structure and Performance Efficiency in Interactive and Noninteractive Spoken Modalities," Technical Note 454, SRI International (Nov. 1988).

[Perrault 80] C. R. Perrault and J. F. Allen, "A Plan-Based Analysis of Indirect Speech Acts," *American Journal of Computational Linguistics*, Vol. 6 (3-4), pp. 167-182 (July-December 1980).

[Power 79] Richard Power, "The Organization of Purposeful Dialogues," *Linguistics*, Vol. 17, pp. 107-152 (1979).

[Price 91] Patti Price (Ed.), *Proceedings of the Fourth DARPA Workshop on Speech and Natural Language* (1991).

[Rosenbaum 88] Rosenbaum, S. L. (1988). *A User's View of the Air Traffic Control (ATC) System*, Internal Memorandum 46321-881130-01.IM, AT&T Bell Laboratories (Nov. 1988).

[Reddy 76] D. Raj Reddy, "Speech Recognition by Machine: A Review," *IEEE Proceedings* Vol. 64(4), pp. 502-531 (April 1976).

[Saville-Troike 85] Muriel Saville-Troike, "The Place of Silence in an Integrated Theory of Communication," pp.3-18 in *Perspectives on Silence*, ed. Deborah Tannen and Muriel Saville-Troike, Ablex Publishing Co., Norwood, New Jersey (1985).

[Schober 89] Michael F. Schober and Herbert H. Clark, "Understanding by Addressees and Overhearers," *Cognitive Psychology* Vol. 21, pp. 211-232 (1989).

[Schwartz 92] Randal L. Schwartz, personal communication, July 2, 1992.

[Schwartz 90] Richard Schwartz and Yen-Lu Chow, "The N-Best Algorithm: An Efficient and Exact Procedure for Finding the N Most Likely Sentence Hypotheses," *ICASSP '90* (1990).

[Searle 69] John R. Searle, *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press, Cambridge (1969).

[Searle 75] John R. Searle, "Indirect Speech Acts," pp. 59-82 in *Syntax and Semantics, Volume 3: Speech Acts*, ed. J. L. Morgan, Academic Press, New York (1975).

[Searle 85] John R. Searle and Daniel Vanderveken, *Foundations of Illocutionary Logic*, Cambridge University Press, Cambridge (1985).

[Soong 90] Frank K. Soong and Eng-Fong Huang, "A Tree-Trellis Based Fast Search for Finding the N Best Sentence Hypotheses in Continuous Speech Recognition," *Proceedings of the 1990 International COnference on Spoken Language Processing* (1990).

[Steedman 90] Mark Steedman, "Structure and Intonation in Spoken Language Understanding," Technical Report MS-CIS-90-23, University of Pennsylvania (April 1990).

[Stubbs 83] Michael Stubbs, *Discourse Analysis*, Chicago Press, Chicago, IL (1983).

[Summers 88] W. Van Summers, David B. Pisoni, Robert H. Bernacki, Robert I. Pedlow, and Michael A. Stokes, "Effects of Noise on Speech Production: Acoustic and Perceptual Analyses," *J. Acoust. Soc. Am.*, Vol. 84(3), pp. 917-928 (Sept. 1988).

[Ward 90a] Karen Ward, David Novick, and Carolyn Sousa, "Air Traffic Control Communications at Detroit Metro - Wayne County Airport," Technical Report No. CS/E 90-024, Oregon Graduate Institute (Dec. 1990).

[Ward 90b] Karen Ward, David Novick, and Carolyn Sousa, "Air Traffic Control Communications at Portland International Airport," Technical Report No. CS/E 90-025, Oregon Graduate Institute (Dec. 1990).

[Ward 91] Wayne Ward, "Current Status of the CMU ATIS System," in *Fourth DARPA Workshop on Speech and Natural Language*, ed. Patti Price, (1991).

[Wesson 77] Robert B. Wesson, *Problem-Solving with Simulation in the World of an Air Traffic Controller*, University of Texas at Austin (1977). Dissertation.

[Young 89a] Sheryl R. Young, Alexander G. Hauptmann, Wayne H. Ward, Edward T. Smith, and Philip Werner, "High Level Knowledge Sources in Usable Speech Recognition Systems," *Communications of the ACM* Vol. 32(2) (1989).

[Young 89b] Sheryl R. Young, Wayne H. Ward, and Alexander G. Hauptmann, "Layering Predictions: Flexible Use of Dialog Expectation in Speech Recognition," *IJCAI 89*, (1989).

[Young 91] Sheryl Young, "Using Semantics to Correct Parser Output for ATIS Utterances," in *Fourth DARPA Workshop on Speech and Natural Language*, ed. Patti Price (1991).

Appendix A

Glossary of Air Traffic Control Terms

This glossary includes ATC terms that are used in this document plus local terms that occur in the transcript in Appendix C. The definitions given are deliberately simplified; the intention is to provide sufficient information to permit the non-pilot to understand their use in this document.

28 Right Approach	See ILS RWY 28R Approach.
Active Runway	Many airports have several runways to accommodate various weather conditions. The one that is currently being used is called the active runway.
Alfa	FAA-approved pronunciation of the letter "A." FAA documentation uses the spelling "alfa" instead of "alpha."
Approach Control	The controller working the Approach position in a TRACON. This controller is responsible for coordinating aircraft as they approach a moderately-busy airport, e.g., Portland International. Also called Approach, Departure.
ARSA	Approach Radar Service Area, the airspace around a moderately-busy airport (e.g., Portland International).
ARTCC	Air Route Traffic Control Center, the radar facility responsible for coordinating traffic between airports.
ATC	Air Traffic Control.
ATIS	Automatic Terminal Information Service. A continuous radio broadcast of recorded information about airport conditions, e.g., wind, visibility, altimeter setting, active runway. The ATIS broadcast frequency is identified on the pilot's charts.
City, the	The downtown area of Portland, Oregon; about 5 NM from PDX.

Cleared	Authorized. Controllers use this term in authorizing pilots to perform certain actions, e.g., to begin flying a standard approach procedure.
Contact	To establish two-way communications with.
Control Tower	The controller position that has responsibility for coordinating traffic on the runways.
Departure	See Approach Control.
Enroute Controller	The controller position that has responsibility for coordinating traffic between airports.
Established on the Localizer	The aircraft is flying along the path defined by the localizer.
Flight Progress Slip	A summary of the flight plan filed by the pilot. The flight strip is automatically printed at each facility along a pilot's route shortly before the aircraft is scheduled to arrive and is updated by controllers as the flight progresses.
Flight Plan	Information filed by an IFR pilot describing the intended flight. It includes aircraft information, route, and expected time of arrival.
Foxtrot	FAA-approved pronunciation of the letter "F," used in utterance 151 (Appendix C) to refer to the current ATIS broadcast.
Ground Controller	The controller position that has responsibility for coordinating traffic on the taxiways within an airport. Also called Ground.
Handoff	The act of formally transferring responsibility for an aircraft from one controller to the next.
Heading	Compass direction, given in degrees.
Ident	A request (instruction) for a pilot to activate the aircraft transponder identification feature. This has the effect of highlighting the radar image of the aircraft on the controller's radar display so that the controller can more easily associate the aircraft with the correct radar image.
IFR	Instrument Flight Rules. The regulations that govern aircraft operation when flying by reference to instruments only, e.g., in poor weather.
ILS	Instrument Landing System. A IFR navigation system that is used to guide aircraft to the airport and to the active runway.

ILS 28 Right Approach	ILS RWY 28R Approach, an instrument approach to PDX that results in aircraft landing on the Runway 28. See Appendix F.
Laker	A Non-Directional Radiobeacon southeast of PDX. Laker is one of the marker beacons that define the ILS 28R approach into PDX. See Appendix F.
Localizer	A directional radio beam, part of an ILS. The localizer defines the path that the aircraft should follow in descending to the runway.
Looking	The pilot is looking for traffic (nearby aircraft) that the controller has just pointed out.
McMinnville	A smaller, uncontrolled airport about 30 NM southwest of PDX.
Niner	FAA-approved pronunciation of the number nine.
NM	Nautical mile(s).
PDX	Portland International Airport, located in Portland, Oregon. This is a moderately-busy airport with an ARSA. See Appendix F.
Readback	The recipient of a transmission repeats back all or part of the transmission to verify that it was understood correctly.
River, the	The Columbia River. Portland International (PDX) is located on the south bank of the Columbia. See Appendix F.
Roger	"I have received your last transmission." In dialogue terms, the FAA defines this to be an acknowledgment and not an agreement. However, it is often used to convey agreement.
Seattle Center	The ARTCC facility that controls the airspace over most of Washington and northern Oregon.
Squawk	A request (instruction) for a pilot to set the aircraft transponder to emit a particular code. This code is displayed on the controller's radar console so that the controller can more easily associate the radar image with the correct aircraft.
Tower	The Tower Controller. See Control Tower.
TRACON	Terminal Radar Approach Control, the facility from which an Approach Controller operates. This facility coordinates traffic around moderately-busy airports.
Traffic	Other aircraft nearby. The controller uses this term to notify a pilot of nearby aircraft.

Traffic	Other aircraft nearby. The controller uses this term to notify a pilot of nearby aircraft.
Traffic no longer a factor	Other aircraft are no longer nearby. This phrase is used to cancel a traffic warning. See Traffic.
Transponder	A radar transmitter/receiver installed in an aircraft. It responds to certain radio signals with information about the aircraft, e.g., a squawk code and altitude information. This information is displayed on the controller's radar console.
Troutdale	A smaller, controlled airport about 8 NM east of PDX. See Appendix F.
Vectors	Instructions from controller to pilot to change speed, heading, or altitude.
VFR	Visual Flight Rules. The set of regulations that govern aircraft operations when navigating by visual reference to the ground.
Visual 28 Left	An approach procedure into Portland International Airport. See Appendix F.

Appendix B

Transcription Conventions

In this transcript, an utterance is defined to be a single transmission directed toward a single recipient. That is, an utterance is ended when either:

- the speaker begins addressing a different recipient, or
- the speaker ends the transmission (by closing the microphone).

Each utterance is numbered and labeled with the speaker, if known, or with "Unknown" if the speaker's identity could not be determined.

The end of a transmission is typically accompanied by a pause of a second or so while other potential speakers realize that the channel is available. These brief pauses are not noted in the transcript. Pauses longer than two or three seconds are marked as untimed pauses (see below), and indicate that no conversant who could transmit wished to do so. The determination of whether a pause had occurred was made subjectively by the transcriber.

Except for the notation indicated below, words are spelled using standard spellings. No attempt was made to depict alternate pronunciations that may have been used (e.g., "goin" for "going").

: (colon)	Denotes a lengthened sound within a word.
- (dash)	Indicates a stop, a cutting off of sound. In mid-utterance, it denotes a verbal hesitation. It also occurs at

	the end of an utterance when the microphone is closed before the speaker finishes speaking.
(words in parenthesis)	Transcriber was uncertain of the words.
((words in double parenthesis))	Transcriber notes.
. (period)	A full stop, with falling intonation.
? (question mark)	A rising intonation.
, (comma)	A short pause with slight falling intonation, less than that indicated by a period.
(.)	Untimed pause between utterances.

Appendix C

Full Transcript Containing Sundance 512 Dialogue

This appendix is extracted from a longer corpus of ATC dialogue [Ward 90b]. It begins with the first utterance of the Sundance 512 dialogue and includes all transmissions made on the channel through the end of the Sundance 512 dialogue. The intent is to provide additional context for this conversation. Notation and transcription conventions are described in Appendix B. Local terms are defined in Appendix A.

		(.)
(108)	Sun 512:	Approach Sundance five twelve's with you at five thousand, one hundred on the heading one ninety on the speed.
(109)	Approach:	Sundance five twelve Portland Approach roger.
		(.)
(110)	Approach:	Sundance five twelve maintain speed one seven zero.
(111)	Sun 512:	One seven zero, Sundance five twelve.
		(.)
(112)	Approach:	Delta fourteen forty three turn left heading three one zero join the localizer.
(113)	Delta 1443:	Delta fourteen forty three, three one zero join the local-.
(114)	Approach:	Sundance five twelve turn left heading zero five zero.

(115)	Sun 512:	Left zero five zero Sundance five twelve.
(116)	Approach:	Ford Trimotor six four five contact approach one one eight point one.
(117)	Ford 645:	Eighteen one, roger.
(118)	Horizon 182:	Portland Approach Horizon, uh one eighty two visual (two eight left).
(119)	Approach:	Horizon one eighty two roger contact tower one one eight point seven.
(120)	Horizon 182:	((Garbled)) (one eighty two).
(121)	Approach:	Delta fourteen forty three you've got niner miles to Laker maintain three thousand till established on the localizer cleared ILS two eight right approach maintain one seven zero knots until Laker.
(122)	Delta 1443:	Delta fourteen forty three, cleared for the ILS to two eight right and uh, maintain three thousand, till on the localizer.
(123)	Approach:	Delta fourteen forty three that's correct and maintain one seven zero knots until Laker.
(124)	Delta 1443:	Copy one seven zero knots till Laker Delta fourteen forty three, thank you.
		(.)
(125)	Approach:	Sundance five twelve turn left heading zero one zero.
(126)	Sun 512:	Left zero one zero Sundance five twelve.
		(.)
(127)	Approach:	Sundance five twelve descend and maintain four thousand five hundred.
(128)	Sun 512:	Four thousand five hundred, Sundance five twelve.

		(.)
(129)	Expo 92:	Portland Approach Expo nine two out of sixty five.
(130)	Approach:	Expo nine two Portland Approach roger.
(131)	Approach:	Sundance five twelve eight miles from Laker turn left heading three one zero. Maintain four thousand five hundred till established localizer cleared ILS two eight right approach, maintain one seven zero knots until Laker.
(132)	Sun 512:	Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.
(133)	Approach:	Expo nine two reduce speed to one seven zero descend and maintain four thousand five hundred.
(134)	Expo 92:	One seven zero on the speed, and down to forty five.
		(.)
(135)	Approach:	Delta fourteen forty three traffic one-o-clock two miles maneuvering vicinity Troutdale altitude unknown.
(136)	Delta 1443:	Fourteen forty three: looking.
		(.)
(137)	Bonanza 64A:	Portland Approach Bonanza three three six four alpha, three point five.
(138)	Approach:	Expo nine two turn left heading zero two zero.
(139)	Expo 92:	Zero two zero, Expo nine two.
(140)	Approach:	Bonanza six four alpha Portland Approach roger.

(.)

(141) Approach: Expo nine two turn left heading three six zero.

(142) Expo 92: Three six zero Expo ninety two.

(143) Approach: Delta fourteen forty three traffic no longer a factor contact tower one one eight point seven.

(144) Delta 1443: Delta fourteen forty three.

(.)

(145) Bonanza 64A: Portland Approach Bonanza three three six four alpha, request immediate lower traffic ILS maintain VFR over.

(146) Approach: Bonanza six four alpha maintain at or above two thousand five hundred will that be uh low enough correction high enough (for you)?

(147) Bonanza 64A: Think so, thank you very much.

(.)

(148) Approach: Expo nine two you're niner miles from Laker turn left heading three one zero, maintain three thousand till established on the localizer, cleared ILS runway two eight right approach, maintain one seven zero knots until Laker.

(149) Expo 92: That's two one, uh zero on the heading and, uh- one seven till Laker.

(.)

(150) United 169: (United) one sixty nine Portland.

(.)

(151) United 169: Portland Approach United one uh sixty nine with you out of nine seven with foxtrot.

(152)	Approach:	United one sixty nine Portland Approach descend and maintain five thousand five hundred.
(153)	United 169:	OK down to five thousand five hundred and uh what was (about the) speed?
(154)	Approach:	Bonanza or correction United one sixty nine normal speed.
(155)	United 169:	Normal speed down fifty five hundred, is that affirmed?
(156)	Approach:	That's affirmative United one sixty nine.
(157)	Approach:	Sundance five twelve contact tower one one eight point seven.
(158)	Sun 512:	Five twelve good-bye.
		(.)

Appendix D

Act and State Definitions

The acts and states used in the analysis are defined here. Note that this is by no means a complete list of the acts needed to explain all ATC dialogue; it represents only the acts and states needed to represent this particular dialogue.

D.1. Acts

Often acts are divided into speech acts, which accomplish their goals through the use of language, and domain acts, which are everything else. In this definition I have made no distinction between the two. I justify this with the observation that many so-called "speech" acts can be accomplished equally well without words and so the distinction in many cases between speech acts and domain acts is arbitrary. In this case, the purpose is to model the dialogue that actually occurred, so acts which pertain more to the physical world are either absent or represented only at a very high level.

Acts are being modelled at a high level. Preconditions and effects are represented only to the degree necessary to understand the communication, and are necessarily incomplete. For instance, the preconditions for a departure would technically include — among other things — a requirement that the pilot be legally and physically able to fly the aircraft in question, that the plane be fueled and ready to fly, that the pilot had filed a flight plan (if this is to be an IFR flight), and so on. These considerations do not generally enter in to the dialogue between pilot and approach controller, though, and so may be omitted without compromising understanding.

Acknowledge

```
acknowl ([target-agent], act)
```

Preconditions: To perform or recognize an acknowledgment, an agent must believe that:

- some action was performed by one "acting" agent,
- the action was directed toward another "target" agent,
- the "target" agent has not yet acknowledged the "acting" agent's action,
- the "acting" agent expects an acknowledgment.

Effects: Upon successful completion of an acknowledge act, the agent will believe that the "acting" agent and the "target" agent mutually believe that:

- the action was acknowledged by the "target" agent.

Example: The approach controller acknowledges Sundance 152 's attempt to contact approach:

```
approach:
    act(acknowl([sun512],contact(approach))
```

Authorize

authorize ([target-agent], act)

Preconditions: To perform or recognize an authorization, an agent must believe that:

- the "target" agent wishes to perform the action,
- the "target" agent has not yet been authorized to perform the action,
- the "target" agent is able to perform the action,
- the "target" agent should not perform the action unless authorized,
- the "authorizing" agent is able to authorize the other agent to perform the action (in this domain, this generally means that the "authorizing" agent is a controller),

Effects: Upon successful completion of an authorize act, the agent will believe that the "authorizing" agent and the "target" agent mutually believe that:

- the "target" agent is authorized to perform the action.

Example: The approach controller authorizes the pilot of Sundance 152 to contact the tower controller:

```
approach:
    act(authorize([sun512],contact(tower))
```

Contact

```
contact ([target-agent])
```

Preconditions: To perform or recognize a contact act, an agent must believe that:

- the "contacting" agent and the "target" agent are not in two-way communication,
- the "contacting" agent intends to establish two-way communication with the "target" agent.

Effects: Upon successful completion of a contact act, the agent will believe that the "contacting" agent and the "target" agent mutually believe that:

- they have established two-way communication with each other.

Example: The approach controller believes that the pilot of Sundance 512 intends to establish communications with Approach.

```
approach:
    blf(intend( [sun512],
                contact([approach])
        true,
        [[approach],[sun512]])
```

Notes: **contact** in Air Traffic Control is the formal domain task of establishing two-way communication between pilot and controller. Its successful completion requires that both pilot and controller have addressed each other by identification, that is, the contact act is not completed if the controller has failed to understand and repeat the pilot's call sign.

Direct

```
direct ([target-agent], act)
```

Preconditions: To perform or recognize a directing act, an agent must believe that:

- the "directing" agent intends for the "target" agent to perform the action,
- it is not obvious that the "target" agent would otherwise perform the action,
- the "target" agent is able to safely perform the action,
- the "directing" agent is in a position of authority over the "target" agent. In this domain, the "directing" agent will often be a controller and the "target" agent will be a pilot.

Effects: Upon successful completion of a directing act, the agent will believe that the "directing" agent and the "target" agent mutually believe that:

- the "target" agent has been directed to perform the action,
- the "target" agent will either perform the action or will object to the instruction.

Example: Seattle Center directs the pilot of Sundance 512 to contact the approach controller:

```
seattle:
    act(direct([sun512],contact([approach]))
```

Report

report ([target-agent], belief)

Preconditions: To perform or recognize a report, an agent must believe that:

- the "reporting" agent and the "target" agent have not yet established the belief as being mutually known between them,
- the" reporting" agent intends that the "target" agent know that the "reporting" agent holds the belief.

Effects: Upon successful completion of a report act, the agent will believe that the "reporting" agent and the "target" agent mutually believe that:

- the "reporting" agent holds the belief.

Example: The pilot of Sundance 512 reports to the approach controller that the aircraft is at an altitude of 5000 feet.

```
sun512:
    act(report([approach],
               blf(altitude([sun512],5000),
                   true,
                   [[sun512]])))
```

Request

```
request ([target-agent], act)
```

Preconditions: To perform or recognize a request act, an agent must believe that:

- the "requesting" agent intends for the "target" agent to perform the action,
- it is not obvious to both that the "target" agent would otherwise perform the action,
- the "target" agent is able to safely perform the action.

Effects: Upon successful completion of a request act, the agent will believe that the "requesting" agent and the "target" agent mutually believe that:

- the" requesting" agent has requested that the other agent perform the action,
- the "target" agent will respond with a report indicating whether the "target" agent intends to perform the action or not.

Example: Cessna N69016 requests the approach controller to authorize the pilot to contact the tower controller:

```
N69016:
    act(request( [approach],
                 authorize([N69016],
                           contact([tower])))
```

Notes: Note that, unlike a direct act, a request carries no expectation that the agent will comply. Note also that the preconditions for a request differ from the preconditions of a direction primarily in the relative authority of the agents.

D.2. States

A small number of domain states are mentioned in the example dialogue. These are defined briefly here.

approach-in-use (facility, approach-procedure)

Definition: An airport may have several published landing procedures, designed to handle various weather conditions, etc. At any given time, some subset of these procedures will be in use.

Example: Seattle Center informs the pilot of Sundance 512 that the ILS-28R approach is in use at PDX:

```
seattle:
    act(report( [sun512],
                    approach-in-use(pdx,ILS-28R)))
```

airspeed ([object], number)

Definition: An object's speed with respect to the air, also called Indicated Airspeed. This is the speed shown on the aircraft airspeed indicator, and is the speed normally used in pilot/controller communications. Note that this is not generally the same as groundspeed (the speed with respect to the ground) or the True Airspeed (the airspeed relative to undisturbed air) [FAA 89, p. A-7].

Example: Sundance 512 informs Portland Approach that the aircraft's airspeed is 190 knots:

```
sun512:
    act(report( [approach],
                  airspeed([sun512],190)))
```

altitude ([object], number)

Definition: The height of a level, point, or object measured in feet Above Ground Level (AGL) or from Mean Sea Level. [FAA 89, p. A-9]. In these dialogues, often refers to indicated altitude, that is, the altitude as shown by an altimeter (e.g., in the aircraft).

Example: Sundance 512 informs Portland Approach that the aircraft's indicated altitude is 5000 feet:

```
sun512:
    act(report(approach],
                altitude([sun512],5000)))
```

destination ([aircraft], facility)

Definition: The agent's destination is the facility.

Example: Approach believes that Sundance 512's destination is Portland International:

```
approach:
    blf(destination([sun512],pdx),
        true
        [[approach,sun512]])
```

direction ([aircraft], degrees)

Definition: The compass heading that the aircraft is flying, in degrees.

Example: Approach believes that Sundance 512 is flying on a heading of 050 degrees:

```
approach:
    blf(direction([sun512],050),
        true
        [[approach,sun512]])
```

position ([aircraft], marker, distance)

Definition: The position of an aircraft relative to some reference point.

Example: Approach believes that Sundance 512 is eight miles from the Laker marker beacon:

```
approach:
    blf(position([sun512],laker,8),
        true
        [[approach,sun512]])
```

Appendix E

Speech Act Representation of a Sample Dialogue

In this appendix, a representative dialogue (Figure 4) is represented in terms of speech acts. The analysis is presented in sections corresponding to the steps of the task. First, the initial beliefs of the conversants are represented, then the initial callup (Section E.2), vectors (Section E.3), arrival instructions, (Section E.4), and the handoff (Section E.5). Within each section, the task and the issues that arose in representing it are discussed first, then the representation itself is presented with explanatory notes.

(108)	Sun 512:	Approach Sundance five twelve's with you at five thousand, one hundred on the heading one ninety on the speed.
(109)	Approach:	Sundance five twelve Portland Approach roger.
(110)	Approach:	Sundance five twelve maintain speed one seven zero.
(111)	Sun 512:	One seven zero, Sundance five twelve.
(114)	Approach:	Sundance five twelve turn left heading zero five zero.
(115)	Sun 512:	Left zero five zero Sundance five twelve.
(125)	Approach:	Sundance five twelve turn left heading zero one zero.
(126)	Sun 512:	Left zero one zero Sundance five twelve.
(127)	Approach:	Sundance five twelve descend and maintain four thousand five hundred.
(128)	Sun 512:	Four thousand five hundred, Sundance five twelve.
(131)	Approach:	Sundance five twelve eight miles from Laker turn left heading three one zero. Maintain four thousand five hundred till established localizer cleared ILS two eight right approach, maintain one seven zero knots until Laker.
(132)	Sun 512:	Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.
(157)	Approach:	Sundance five twelve contact tower one one eight point seven.
(158)	Sun 512:	Five twelve good-bye.

Utterance numbers are from the original transcript [Ward 90b]. See Appendix C for context.

Figure 4. Dialogue between Sundance 512 and Portland Approach

E.1 Initial Beliefs

As the exchange begins, Sundance 512, a scheduled commercial flight, has just been handed off to Portland Approach. This means that the approach controller has accepted responsibility for the aircraft from the previous controller, probably Seattle Center (represented in controller belief CB1 in Table 2 as "artcc"), and the pilot has been instructed to contact Portland Approach.

Pilot and controller have not yet established radio contact. Each begins the conversation with strong expectations about the knowledge held by the other, however. The approach controller's initial beliefs, shown in Table 2, are derived from several sources. The controller has a flight strip for Sundance 512 which supplies — among other things — the flight's destination, route, and estimated time of arrival. Thus, the controller is expecting the flight before the handoff takes place. In accepting the handoff from Seattle Center, the controller acknowledges (to the Seattle Center controller) that Sundance 512 appears on the approach controller's radar display. From this display, the controller can see position, direction, approximate airspeed, and altitude information for approaching aircraft.

The approach controller believes that, as part of the handoff procedure, the ARTCC controller has instructed the pilot to contact Portland Approach. The controller expects that the pilot will comply with this direction. Recall that "contact" is a formal act in this domain; the pilot must formally establish two-way radio contact with the approach controller before entering the ARSA. The pilot is expected to initiate communications, and the controller's beliefs reflect that (see controller belief CB2).

As part of the contact procedure, the controller expects the pilot to confirm the aircraft altitude, heading, and airspeed (represented in Table 2 as controller beliefs CB4, CB6, and CB8). Although the controller has already formed expectations based on the information displayed on the workstation console (CB3, CB5, and CB7), those beliefs must be checked against the pilot's report. There are many reasons why the pilot's beliefs may not match the controller's beliefs exactly. Some possible causes might include: faulty transponder on the aircraft; wind causing an unexpected discrepancy between the groundspeed (reported on the radar) and the airspeed (reported in the cockpit); or instrument error on either side. In the crowded airspace around a busy airport, it is particularly important that these beliefs be confirmed as mutual or that some idea of the magnitude of the discrepancy be established.

Finally, the controller believes (CB10) that the ILS-28R Approach Procedure is currently in use at PDX. This belief is based on the controller's awareness of conditions at the airport and on information displayed on the controller's console. The controller believes that the pilot should have been informed which approach procedure was in use as part of the handoff procedure (the pilot may have been told directly by the Seattle Center controller, or may have listened to the ATIS, a continuously-broadcast recording of airport conditions).

The pilot has also formed a set of initial beliefs, shown in Table 3. Although similar in form to those of the approach controller — there was no evidence of misunderstanding in this dialogue — most of the beliefs come from different sources. The pilot was indeed instructed by Seattle Center to contact Portland Approach and had formed the intention to do so (PB1, PB2). The pilot's beliefs about the altitude, heading (direction), and airspeed of the aircraft are based on the information displayed on the cockpit instruments. These beliefs are represented in pilot beliefs PB3, PB5, and PB7. The Seattle Center controller probably informed the pilot that the ILS-28R Approach Procedure is in use (PB10), although the pilot may also have learned this from the ATIS broadcast. The pilot forms the intention to report and thereby attempt to confirm this information, as called for in the handoff procedure (PB4, PB6, and PB8).

Table 2. Initial Beliefs of Controller

Representation	Notes
CB1: blf (occurred ([artcc], direct ([sun512], contact ([approach])), true, [[approach, artcc], [artcc, sun512]])	(CB1) The controller believes that Seattle Center instructed the pilot of Sundance 512 to contact Portland approach (from handoff procedure).
CB2: blf (intend ([sun512], contact ([approach])), true, [[approach], [sun512]])	(CB2) The controller believes that the pilot intends to establish communications with approach (derived from CB1).
CB3: blf (altitude ([sun512], 5000), true, [[approach], [sun512]])	(CB3) The controller believes that the aircraft is maintaining an altitude of 5000 feet (displayed on the controller's console).
CB4: blf (intend ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])), true, [[approach], [sun512]])	(CB4) The controller expects the pilot to confirm the aircraft altitude. This expectation is derived from the handoff procedure and from controller belief CB3.
CB5: blf (direction ([sun512],100), true, [[approach], [sun512]])	(CB5) The controller believes that the pilot is flying heading 100 (displayed on console).
CB6: blf (intend ([sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB6) The controller expects the pilot to confirm the aircraft direction. This expectation is derived from the handoff procedure and from controller belief CB5.

Table 2. Initial Beliefs of Controller (Continued)

Representation	**Notes**
CB7: blf (airspeed ([sun512],190), true, [[approach], [sun512]])	(CB7) The controller believes that the aircraft is flying at an airspeed of 190 knots (derived from information displayed on console).
CB8: blf (intend ([sun512], report ([approach], blf (airspeed ([sun512], 190), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB8) The controller expects the pilot to confirm the aircraft speed. This expectation is derived from the hand-off procedure and from controller belief CB7.
CB9: blf (destination ([sun512], pdx)], true, [[approach], [sun512]])	(CB9) The controller believes that the pilot plans to land at PDX (from flight progress strip).
CB10: blf (approach-in-use (pdx, ils-28r), true, [[approach], [sun512]])	(CB10) The controller believes that ILS-28R approach is currently in use, and that the pilot has been informed of this by Seattle Center (from procedure).

Table 3. Initial Beliefs of Pilot

Representation	Notes
PB1: blf (occurred ([artcc], direct ([sun512], contact ([approach])), true [[sun512, artcc], [artcc, approach]])	(PB1) The pilot believes that Seattle Center instructed the pilot to contact Portland Approach, and the pilot believes that Approach knows this from the handoff procedure.
PB2: blf (intend ([sun512], contact ([approach])), true, [[sun512], [approach]])	(PB2) The pilot intends to establish communications with Portland Approach, and Approach expects this (derived from handoff procedure and PB1).
PB3: blf (altitude ([sun512], 5000)), true, [[sun512], [approach]])	(PB3) The pilot believes that the aircraft is maintaining an altitude of 5000 feet (from aircraft instruments) and that Approach can see this on the console.
PB4: blf (intend ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB4) The pilot intends to confirm the belief that the aircraft is maintaining an altitude of 5000 feet. The pilot believes that the controller expects this (derived from handoff procedure and pilot belief PB3).
PB5: blf (direction ([sun512], 100), true, [[sun512], [approach]])	(PB5) The pilot believes that the aircraft is flying on a heading of 100 (from aircraft instruments) and that Approach can see this on the console.
PB6: blf (intend ([sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB6) The pilot intends to confirm the belief that the aircraft is flying a heading of 100 with the controller, and the pilot believes that the controller expects this (derived from handoff procedure and Sun512's belief PB5).

Table 3. Initial Beliefs of Pilot (Continued)

Representation	**Notes**
PB7: blf (airspeed ([sun512], 190), true, [[sun512], [approach]])	(PB7) The pilot believes that the aircraft is maintaining an airspeed of 190 knots (from aircraft instruments), and the pilot believes that Approach can see this on the console
PB8: blf (intend ([sun512], report ([approach], blf (airspeed ([sun512], 190), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB8) The pilot intends to confirm the belief that the aircraft is maintaining an indicated airspeed of 190, and the pilot believes that the controller expects this (derived from handoff procedure and the pilot's belief PB7).
PB9: blf (destination ([sun512], pdx), true, [[sun512], [approach]])	(PB9) The pilot plans to land at PDX, and the pilot believes that Approach knows this from the flight plan.
PB10: blf (approach-in-use (pdx, ils-28r]), true, [[sun512], [approach]])	(PB10) The pilot believes that the ILS-28R approach is in use (from ATIS), and the pilot believes that the controller knows this.
PB11: blf (intend ([approach], authorize ([sun512], approach (pdx, ils-28r))), true, [[sun512], [approach]])	(PB11) The pilot expects to be authorized to execute the ILS-28R instrument approach procedure. This belief is derived from PB9 and PB10.

E.2 Initial Callup

The initial callup begins with utterance 108 from the pilot of Sundance 512. In this utterance the pilot intends four speech acts, shown in Table 4. The pilot's primary intention is to establish contact with the controller; this is signalled by the use of the phrase "with you." During this transmission, the pilot also attempts to confirm that the altitude shown on the cockpit instruments matches the altitude being displayed on the controller's console (which is broadcast by the aircraft's Mode C transponder). The pilot also reports the direction and speed assignments that were received from Seattle Center before handoff.

Immediately after this first utterance, the pilot's beliefs are updated to reflect the pilot's attempt to contact the approach controller and confirm certain information (represented in Table 5). These beliefs represent the pilot's view of what just happened; the pilot intended to perform those acts (Table 3), the pilot attempted to perform them in making the initial callup (Table 4), the pilot then believes that those acts have been performed (Table 5). The belief structure still reflects a certain amount of uncertainty as to the mutuality of these beliefs, however. The pilot does not yet know for certain whether the transmission has been received, or whether the pilot's words have been understood as intended, or whether the controller's information in fact matches the pilot's. The pilot assumes that all is routine, though, and forms expectations that the controller has understood the utterance as intended (pilot beliefs PB2 through PB14 in Table 5) and that the controller will both acknowledge and confirm the reports (pilot beliefs PB15 through PB22, Table 5).

In this case, the approach controller has in fact heard and correctly interpreted the pilot's transmission. The controller's beliefs are updated to reflect that fact, as shown in Table 6. The information is still not confirmed mutual, however; the parties have not yet succeeded in establishing two-way communication. The controller's belief space is therefor updated to show that the controller believes that both pilot and controller understand the same information to be true, but that the controller has not yet communicated that fact to the pilot. The controller intends to acknowledge and confirm the pilot's report, and believes that the pilot expects this.

At this point, both pilot and controller have independently formed strong expectations that the controller will acknowledge the acts in utterance 108 and will furthermore confirm the altitude, heading, and airspeed information. These strong expectations will encourage both parties to accept even moderately weak evidence of understanding as acknowledgment; anything less than an explicit statement of disagreement is likely to be taken as confirmation. This is reflected in the terseness of the controller's reply (utterance 109, Table 7). The controller formally completes the contact as specified by FAA procedures [FAA 89] by addressing the pilot ("Sundance 512") and repeating the facility identification ("Portland Approach") to confirm that the pilot is in communication with the correct facility. The single word "roger" suffices to acknowledge and confirm the remaining three speech acts; in the absence of evidence to the contrary, the strong expectations set up by procedures and by the multiple sources of information within the ATC system permit the conversants to interpret the response correctly.

After the controller's acknowledgment in utterance 108, the controller and pilot believe that they are in communication and have confirmed altitude, heading, and airspeed information to be mutually known true. The changes to the controller's and pilot's beliefs are shown in Table 8 and Table 9. Notice that the pilot does not reply to Approach's acknowledgment in utterance 109, nor does the controller expect a reply. This is in accordance both with recommended ATC procedure and with Clark's [Clark 89] rules for evidence of understanding; one acknowledges an acknowledgment with either the initiation of the next relevant contribution (in this case, there was none) or with continued attention (silence). Strictly speaking, this would be represented as an expectation of acknowledgment that is fulfilled after the pilot has passed up an opportunity to initiate a correction or clarification to the controller's contribution. For simplicity in this example, this step will not be represented explicitly.

Table 4. Intentions of Pilot in Utterance 108

```
(108) Sun 512:   Approach Sundance five twelve's with you
                 at five thousand, one hundred on the
                 heading one ninety on the speed.
```

Representation	Notes
PA1: act (contact ([approach]))	(PA1) The pilot attempts to establish formal communications with the controller. This act is motivated by the pilot's belief PB2 and is expressed by the phrase `Sundance five twelve's with you`.
PA2: act (report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])))	(PA2) The pilot attempts to confirm with the controller the belief that the aircraft is maintaining an altitude of 5000 feet. The act is motivated by the pilot's belief PB4 and is expressed by the phrase `at five thousand`.
PA3: act (report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]])))	(PA3) The pilot attempts to confirm that the aircraft is flying a compass heading of 100. This act is motivated by the pilot's belief PB6 and is expressed by the phrase `one hundred on the heading`.
PA4: act (report ([approach], blf (airspeed ([sun512],190), true, [[sun512], [approach]])))	(PA4) The pilot attempts to confirm that the aircraft is maintaining an airspeed of 190 knots. This act is motivated by the pilot's belief PB8 and is expressed by the phrase `one ninety on the speed`.

Table 5. Changes in Pilot's Beliefs After Utterance 108

Representation	Notes
PB12: blf (occurred ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB12) The pilot has attempted to confirm that the controller believes that the aircraft is maintaining an altitude of 5000 feet. This belief is derived from the pilot's attempt to perform act PA2, and it replaces belief PB4.
PB13: blf (occurred ([sun512], report ([approach], blf (direction ([sun512],100), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB13) The pilot has attempted to confirm that the aircraft is maintaining a heading of 100. This belief is derived from the pilot's attempt to perform act PA3, and it replaces PB6.
PB14: blf (occurred ([sun512], report ([approach], blf (airspeed ([sun512],190), true, [[sun512], [approach]]))), true, [[sun512], [approach]])	(PB14) The pilot has attempted to confirm that the controller believes that the aircraft is maintaining an airspeed of 190 knots. This belief is derived from the pilot's attempt to perform act PA4, and it replaces PB8.
PB15: blf (intend ([approach], acknowl ([sun512], contact ([approach]))), true, [[sun512], [approach]])	(PB15) The pilot expects the controller to acknowledge the pilot's attempt PA1 to contact Approach. This expectation is derived from procedures and from the pilot's intention PB2 to perform act PA1.

Table 5. Changes in Pilot's Beliefs After Utterance 108 (Continued)

Representation	Notes
PB16: blf (intend ([approach], report ([sun512], blf (occurred ([sun512], contact ([approach]), true, [[sun512, approach]]))), true, [[sun512], [approach]])	(PB16) The pilot expects the controller to confirm that the pilot's attempt to contact Portland Approach in utterance 108 was successful.
PB17: blf (intend ([approach], acknowl ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])))), true, [[sun512], [approach]])	(PB17) The pilot expects the controller to acknowledge the pilot's attempt to perform PA2. This expectation is derived from procedures and from the pilot's intention PB12 to perform act PA2.
PB18: blf (intend ([approach], report ([sun512], blf (altitude ([sun512],5000), true, [[approach, sun512]]))), true, [[sun512], [approach]])	(PB18) The pilot expects the controller to confirm the pilot's report PA2 that the aircraft altitude is 5000 feet.

Table 5. Changes in Pilot's Beliefs After Utterance 108 (Continued)

Representation	Notes
PB19: blf (intend ([approach], acknowl ([sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]])))), true, [[sun512], [approach]])	(PB19) The pilot expects the controller to acknowledge the pilot's act PA3. This expectation is derived from procedures and from the pilot's intention PB13 to perform act PA3.
PB20: blf (intend ([approach], report ([sun512], blf (direction ([sun512], 100), true, [[sun512, approach]]))), true, [[sun512], [approach]])	(PB20) The pilot expects the controller to confirm the pilot's report PA3 that the aircraft is flying heading 100.
PB21: blf (intend ([approach], acknowl ([sun512], report ([approach], blf (airspeed ([sun512], 190), true, [[sun512],[approach]])))), true, [[sun512], [approach]])	(PB21) The pilot expects the controller to acknowledge the pilot's act PA4. This expectation is derived from procedures and from the pilot's intention PB14 to perform act PA4.
PB22: blf (intend ([approach], report ([sun512], blf (airspeed ([sun512], 190), true, [[sun512, approach]]))), true, [[sun512], [approach]])	(PB22) The pilot expects the controller to confirm the pilot's report in PA4 that the aircraft is flying at 190 knots.

Table 6. Changes in Controller's Beliefs After Utterance 108

Representation	Notes
CB11: blf (occurred ([sun512], contact ([approach])), true, [[approach], [sun512]])	(CB11) The controller recognizes act PA1 from the use of the phrase **`Sun-dance five twelve's with you`** and from controller expectation CB2. CB11 replaces controller belief CB2.
CB12: blf (occurred ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB12) The controller recognizes act PA2 from the use of the phrase **`at five thousand`** and from the expectation CB4 that the pilot will report that the aircraft is flying at 5000 feet. CB12 replaces CB4.
CB13: blf (occurred ([sun512], report ([approach], blf (direction ([sun512],100), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB13) The controller recognizes act PA3 from the use of the phrase **`one hundred on the heading`** and from the expectation CB6 that the pilot will report that the aircraft is flying heading 100. CB13 replaces CB6.
CB14: blf (occurred ([sun512], report ([approach], blf (airspeed ([sun512],190), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB14) The controller recognizes act PA4 from the phrase **`one ninety on the speed`** and from the expectation CB8 that the pilot will report that the aircraft is maintaining an airspeed of 190 knots. CB14 replaces CB8.
CB15: blf (intend ([approach], acknowl ([sun512], contact ([approach]))), true, [[approach], [sun512]])	(CB15) The controller intends to acknowledge the pilot's attempt to formally contact the approach facility (from recommended procedure and the controller's belief CB11).

Table 6. Changes in Controller's Beliefs After Utterance 108 (Continued)

Representation	Notes
CB16: blf (intend ([approach], report ([sun512], blf (occurred ([sun512], contact ([approach])), true, [[sun512, approach]]))), true, [[approach], [sun512]])	(CB16) The controller intends to confirm that the pilot has succeeded in contacting the approach facility (from recommended procedure and the controller's belief CB11).
CB17: blf (intend ([approach], acknowl ([[sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]]))), true, [[approach], [sun512]])	(CB17) The controller intends to acknowledge the pilot's report that the aircraft is at 5000 feet (from controller belief CB12 and recommended procedure).
CB18: blf (intend ([approach], report ([sun512], blf (altitude ([sun512], 5000), true, [[approach, sun512]]))), true, [[approach], [sun512]])	(CB18) The controller further intends to confirm the report that the aircraft is at 5000 feet (from controller beliefs CB3 and CB12).
CB19: blf (intend ([approach], acknowl (sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512],[approach]]))), true, [[approach], [sun512]])	(CB19) The controller intends to acknowledge the pilot's report that the aircraft is maintaining a heading of 100 (from controller belief CB13 and from recommended procedure).

Table 6. Changes in Controller's Beliefs After Utterance 108 (Continued)

Representation	Notes
CB20: blf (intend ([approach], report ([sun512], blf (direction ([sun512],100), true, [[approach,sun512]]))), true, [[approach], [sun512]])	(CB20) The controller further intends to confirm the report that the aircraft is maintaining a heading of 100 (from controller beliefs CB5 and CB13 and from recommended procedure).
CB21: blf (intend ([approach], acknowl ([sun512], report ([approach], blf (airspeed ([sun512], 190), true, [[sun512],[approach]]))), true, [[approach], [sun512]])	(CB21) The controller intends to acknowledge the pilot's report that the aircraft is maintaining an airspeed of 190 knots (from controller belief CB14 and recommended procedure).
CB22: blf (intend ([approach], report ([sun512], blf (airspeed ([sun512], 190), true, [[approach, sun512]]))), true, [[approach], [sun512]])	(CB22) The controller further intends to confirm as mutual the belief that the aircraft is maintaining an airspeed of 190 knots (from controller beliefs CB7 and CB14 and from recommended procedure).

Table 7. Intentions of Controller in Utterance 109

(109) Approach: Sundance five twelve Portland Approach roger.

Representation	Notes
CA1: act (acknowl ([sun512], contact ([approach])))	(CA1) The controller attempts to acknowledge the pilot's attempt to contact the approach facility. This act is motivated by the controller's belief CB15. This act is expressed by the controller's confirming self-identification **Portland Approach.**
CA2: act (report ([sun512], blf (occurred ([sun512], contact ([approach]), true, [[sun512, approach]])))	(CA2) The controller attempts to confirm that the pilot has succeeded in contacting the approach facility. This act is motivated by controller intention CB16 and is expressed by the use of the confirming term **roger**.
CA3: act (acknowl ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])))))	(CA3) The controller attempts to acknowledge the report of the aircraft altitude. This act is motivated by controller intention CB17 and is expressed by the use of the confirming term **roger** and by the absence of a request for clarification.
CA4: act (report ([sun512], blf (altitude ([sun512], 5000), true, [[approach, sun512]])))	(CA4) The controller attempts to confirm the pilot's report of the aircraft altitude. This act is motivated by controller intention CB18. The act is expressed by the use of the confirming term **roger** and by the absence of a correction to utterance 108.

Table 7. Intentions of Controller in Utterance 109 (Continued)

Representation	**Notes**
CA5: act (acknowl ([sun512], report ([approach], blf (direction ([sun512],100), true, [[sun512], [approach]]))))	(CA5) The controller attempts to acknowledge the report of the aircraft heading. This act is motivated by controller intention CB19. The act is expressed by the use of the confirming term `roger` and by the absence of a request for clarification.
CA6: act (report ([sun512], blf (direction ([sun512],100), true, [[approach, sun512]])))	(CA6) The controller attempts to confirm the report of the aircraft heading. This act is motivated by controller intention CB20. The act is expressed by the use of the confirming term `roger` and by the absence of a correction to utterance 108.
CA7: act (acknowl ([sun512], report ([approach], blf (airspeed ([sun512],190), true, [[sun512], [approach]])))	(CA7) The controller attempts to acknowledge the pilot's report of the airspeed. This act is motivated by controller intention CB21. The act is expressed by the use of the confirming term `roger` and by the absence of a request for clarification.
CA8: act (report ([sun512], blf (airspeed ([sun512],190), true, [[approach, sun512]])))	(CA8) The controller attempts to confirm the report of the aircraft airspeed. This act is motivated by controller intention CB22. The act is expressed by the use of the confirming term `roger` and by the absence of a correction to utterance 108.

Table 8. Changes in Controller's Beliefs After Utterance 109

Representation	Notes
CB23: blf (occurred ([approach], acknowl ([sun512], contact ([approach]))), true, [[approach, sun512]])	(CB23) The controller believes that the pilot's attempt to contact the approach facility has been acknowledged. This belief is derived from the controller's attempt to perform act CA1 and replaces belief CB15.
CB24: blf (occurred ([sun512], contact ([approach])), true, [[approach, sun512]]))	(CB24) The controller believes that Sundance 512 has successfully contacted Portland Approach. This belief is derived from the controller's attempt to perform act CA2 and replaces belief CB16.
CB25: blf (occurred ([approach], acknowl ([sun512], report ([approach], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])))), true, [[approach, sun512]])	(CB25) The controller believes that the pilot's altitude report has been acknowledged. This belief is derived from the controller's attempt to perform act CA3 and replaces controller belief CB17.
CB26: blf (occurred ([approach], report ([sun512], blf (altitude ([sun512], 5000), true, [[approach, sun512]]))), true, [[approach, sun512]])	(CB26) The controller believes that the pilot's altitude report has been confirmed. This belief is derived from the controller's attempt to perform act CA4 and replaces controller belief CB18.
CB27: blf (altitude ([sun512], 5000), true, [[approach, sun512]])	(CB27) The controller believes that they now mutually believe that the aircraft is at 5000 feet. This belief replaces controller belief CB3.

Table 8. Changes in Controller's Beliefs After Utterance 109 (Continued)

Representation	Notes
CB28: blf (occurred ([approach], acknowl ([sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]])))), true, [[approach, sun512]])	(CB28) The controller believes that the pilot's direction report has been acknowledged. This belief is derived from the controller's attempt to perform act CA5 and replaces controller belief CB19.
CB29: blf (occurred ([approach], report ([sun512], blf (direction ([sun512],100), true, [[approach, sun512]])))), true, [[approach, sun512]])	(CB29) The controller believes that the report of the aircraft heading has been confirmed. This belief is derived from the controller's attempt to perform act CA6 and replaces controller belief CB20.
CB30: blf (direction ([sun512], 100), true, [[approach, sun512]])	(CB30) The controller believes that it is mutually known that Sundance 512 is maintaining heading 100. This belief replaces controller belief CB5.
CB31: blf (occurred ([approach], acknowl ([sun512], report ([approach], blf(airspeed ([sun512], 190), true, [[sun512], [approach]])))), true, [[approach, sun512]])	(CB31) The controller believes that the pilot's airspeed report has been acknowledged. This belief is derived from the controller's attempt to perform act CA7 and replaces controller belief CB21

Table 8. Changes in Controller's Beliefs After Utterance 109 (Continued)

Representation	Notes
CB32: blf (occurred ([approach], report ([sun512], blf (airspeed ([sun512],190), true, [[approach, sun512]]))), true, [[approach, sun512]])	(CB32) The controller believes that the pilot's report of the airspeed has been acknowledged and confirmed. This belief is derived from the controller's attempt to perform act CA8 and replaces controller belief CB22.
CB33: blf (airspeed ([sun512], 190), true, [[approach, sun512]])	(CB33) The controller believes that controller and pilot mutually know that Sundance 512 is maintaining an airspeed of 190 knots. This belief replaces controller belief CB7.

Table 9. Changes in Pilot's Beliefs After Utterance 109

Representation	Notes
PB23: blf (occurred ([approach], acknowl ([sun512], contact ([approach]))), true, [[sun512, approach]])	(PB23) The pilot believes that the controller has acknowledged PA1. This belief is derived from pilot expectation PB15 and from the controller's use of the phrase `Portland Approach`. This belief replaces PB15.
PB24: blf (occurred ([approach], report ([sun512], blf (occurred ([sun512], contact ([approach]), true, [[sun512, approach]]))), true, [[sun512, approach]])	(PB24) The pilot believes that the controller has confirmed PA1, the pilot's attempt to establish contact. This belief is derived from PB15 and from the controller's phrase `Sundance five twelve`. This belief replaces PB15.
PB25: blf (occurred ([approach], acknowl ([sun512], report ([sun512], blf (altitude ([sun512], 5000), true, [[sun512], [approach]])))), true, [[sun512, approach]])	(PB25) The pilot believes that the controller has acknowledged PA1, the pilot's report of the aircraft altitude. This belief is derived from pilot expectation PA17 and from the lack of a request for clarification from the controller. This belief replaces PA17.
PB26: blf (occurred ([approach], report ([sun512], blf (altitude ([sun512], 5000), true, [[approach, sun512]]))), true, [[sun512, approach]])	(PB26) The pilot believes that the controller has confirmed the report of the aircraft altitude. This belief is derived from the controller's reply `roger` and from the pilot's expectation PB18. This belief replaces PB18.

Table 9. Changes in Pilot's Beliefs After Utterance 109 (Continued)

Representation	Notes
PB27: blf (altitude ([sun512], 5000), true, [[sun512, approach]])	(PB27) The pilot believes that they now mutually believe that Sundance 512 is maintaining an altitude of 5000 feet. This belief replaces pilot belief PB3.
PB28: blf (occurred ([approach], acknowl ([sun512], report ([approach], blf (direction ([sun512], 100), true, [[sun512], [approach]])))), true, [[sun512, approach]])	(PB28) The pilot believes that the controller has acknowledged the report of the aircraft direction. This belief is derived from the pilot's expectation PB19 and from the absence of a request for clarification from the controller. This belief replaces pilot belief PB19.
PB29: blf (occurred ([approach], report ([sun512], blf (direction ([sun512],100), true, [[approach, sun512]]))), true, [[sun512, approach]])	(PB29) The pilot believes that the controller has confirmed the report of the aircraft heading. This belief is derived from the pilot's expectation PB20, and from the controller's reply `roger`. This belief replaces PB20.
PB30: blf (direction ([sun512], 100), true, [[sun512, approach]])	(PB30) The pilot believes that they mutually know that the aircraft is flying heading 100. This belief replaces pilot belief PB5.

Table 9. Changes in Pilot's Beliefs After Utterance 109 (Continued)

Representation	**Notes**
PB31: blf (occurred ([approach], acknowl ([sun512], report ([sun512], blf (airspeed([sun512], 190), true, [[sun512], [approach]])))), true, [[sun512, approach]])	(PB31) The pilot believes that the controller has acknowledged PA4, the report of the airspeed. This belief is derived from the pilot's expectation PB21 and from the absence of a request for clarification from the controller. This belief replaces PB21.
PB32: blf (occurred ([approach], report ([sun512], blf (airspeed ([sun512],190), true, [[approach,sun512]]))), true, [[sun512, approach]])	(PB32) The pilot believes that the controller has confirmed the report of the airspeed. This belief is derived from the controller's reply `roger` and from the pilot's expectation PB22. This belief replaces PB22.
PB33: blf (airspeed ([sun512], 190), true, [[sun512, approach]])	(PB33) The pilot believes that it is mutually known that the aircraft is maintaining an airspeed of 190. This belief replaces pilot belief PB7.

E.3 Vectors

There was a pause between utterance 108 and 110. During this time, the controller may have been engaged in some other task or may have been waiting for Sundance 512 to reach a particular position. The pilot was proceeding along the assigned heading.

With the next instruction to Sundance 512 (utterance 110, Table 12), the controller begins the process of directing the aircraft to a point where it can turn onto the approach path and intercept the localizer (a navigation beacon). The approach controller must keep traffic flowing smoothly while maintaining required separations between aircraft in a relatively congested airspace. During these exchanges (utterances 110 through 128), the controller's utterances are motivated primarily by the progress of the aircraft along the route to the approach gate and by the controller's knowledge of the positions and performance characteristics of other aircraft in the vicinity.

The route through the ARSA to the approach gate is probably a standard one from the controller's standpoint but may not be known to the pilot. These routes are not part of the published procedures available to pilots, although a pilot who flies into PDX frequently may learn the routes and come to anticipate the directions. Because the pilot generally has no way to anticipate the exact vectors, this analysis assumes only that the pilot expects to hear instructions to change direction, speed, and altitude. These expectations are represented in Table 10.

The representation in Table 10 fails to capture certain nuances of the pilot's expectations: "reasonableness checks" in the conversants' expectations; strength or immediacy of expectation, that is, the difference between the expectation that something may reasonably occur and the expectation that something should certainly occur; imprecise expectations. These are discussed in greater detail in Chapter 5.

In utterance 110, the approach controller directs the pilot to "maintain speed one seven zero." In isolation, this might appear to be an instruction for the pilot to continue flying at a speed of 170 knots. In utterance 108, however, the pilot reported the aircraft's speed as 190 knots, and Approach acknowledged this in utterance 109. Furthermore, an aircraft being vectored by a controller, as this one is, should not deliberately change speed unless directed to do so. Thus, the aircraft should still be travelling at an airspeed of 190 knots and both the pilot and the controller should know this. Furthermore, they believe that this knowledge is mutual, that is, that they have explicitly established in utterances 108 and 109 that both know that Sundance 512 is flying at 190 knots. In light of this, it is possible for the pilot to interpret Approach's utterance as an instruction to slow to and then maintain 170 knots.

Sometime after completing the exchange in utterances 110 and 111, the pilot completes the process of slowing the aircraft to 170 knots. This is represented for both pilot and controller as a mutually-known belief (PB41, CB42) even though the conversants do not confirm this verbally. This assumption is justified, however, from the earlier exchanges that have established that their instruments agree well enough and that the pilot intends to make the speed change. The controller expects the pilot to slow to approximately 170 knots, and only if the radar display shows something substantially different would the controller be expected to say something. Likewise, the pilot believes that the controller expects the speed change and can see it on the radar display. Only if the pilot encounters some difficulty in maintaining the requested speed would there be another transmission about the speed change.

As was seen in the initial callup, the strong expectations engendered by the task and the authority relationship between the conversants — a pilot is expected to comply with ATC directions if able to do so safely — allow for an extremely terse and efficient reply. The pilot is able to signal an intention to comply merely by acknowledging the direction without objection. In the notes for utterances in which the pilot is responding to directions, I have described the pilot's implied assent to the instructions as being signalled by the lack of a statement that the pilot is unable to comply (see, for example, PA8 in Table 23). FAA procedure suggests that the pilots use the phrase **unable** to inform the controller that they cannot comply with an instruction; the asymmetrical authority relationship may lead to more indirect formulations, however.

Table 10. Changes in Pilot's Beliefs Between Utterances 108 and 110

Representation	Notes
PB1: blf (intend ([approach], direct ([sun512], airspeed ([sun512], <change>))), true, [[sun512], [approach]])	(PB1) The pilot expects to be directed to change speed to accommodate other traffic in the area (derived from recommended procedure).
PB2: blf (intend ([approach], direct ([sun512], direction ([sun512], <change>))), true, [[sun512], [approach]])	(PB2) The pilot of Sundance 512 expects to be given vectors to the approach gate. The exact vectors are unknown to the pilot, however. This belief is derived from PB9 and PB10.
PB3: blf (intend ([approach}, direct ([sun512], altitude ([sun512], <change>))), true, {[sun512], [approach}})	(PB3) The pilot of Sundance 512 expects to be given instructions to change altitude. The exact altitude assignment is unknown to the pilot, however (derived from recommended procedures).

Table 11. Changes in Controller's Beliefs Between Utterances 108 and 110

Representation	Notes
CB34: blf (intend ([approach], direct ([sun512], direction ([sun512], <change>))), true, [[approach], [sun512]])	(CB34) The controller intends to issue a series of vectors to direct the pilot to the ILS-28R approach gate. The controller has a plan for doing this, probably utilizing a standard routing; the transcriber does not know what that route is, however. This belief is derived from controller belief CB9 of the pilot's destination and CB10 of the approach in use at PDX.
CB35: blf (intend ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[approach]])	(CB35) The controller intends to direct the pilot to slow the aircraft to 170 knots. This intention is formed from knowledge of recommended procedures, from CB34 above, and from controller beliefs CB26, CB30,and CB33 of the aircraft speed, altitude, and course.

Table 12. Intended Act of Controller in Utterance 110

```
(110) Approach:   Sundance five twelve maintain speed one
                  seven zero.
```

Representation	Notes
CA9: act (direct ([sun512], airspeed ([sun512], 170)))	(CA9) The controller attempts to direct the pilot to fly at an airspeed of 170 knots. This act is motivated by the controller's belief CB35 and is expressed by the phrase `maintain speed one seven zero.`

Table 13. Changes in Controller's Beliefs After Utterance 110

Representation	Notes
CB36: blf (occurred ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[approach], [sun512]])	(CB36) The controller believes that the pilot has been directed to slow the aircraft to 170 knots. This belief is derived from the controller's attempt to perform act CA9 and it replaces controller belief CB35.
CB37: blf (intend ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170)))), true, [[approach], [sun512]])	(CB37) The controller expects the pilot of Sundance 512 to acknowledge the controller's direction to slow to 170 knots. This belief is derived from CB36.
CB38: blf (intend ([sun512], report ([approach], intend ([sun512], airspeed ([sun512], 170)))), true, [[approach], [sun512]])	(CB38) The controller expects the pilot to inform the controller of the pilot's intention to comply with the direction. This belief is derived from CB36.

Table 14. Changes in Pilot's Beliefs After Utterance 110

Representation	Notes
PB34: blf (occurred ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[sun512], [approach]])	(PB34) The pilot of Sundance 512 believes that the controller has directed the pilot to slow the aircraft to 170 knots. This belief is derived from the pilot's belief PB33 that the current airspeed is mutually known to be 190 and from the pilot's expectation PB1 that the controller will direct the pilot to change the aircraft speed.
PB35: blf (intend ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170)))), true, [[sun512], [approach]])	(PB35) The pilot intends to acknowledge the controller's direction to slow to 170 knots. This belief is derived from PB34.
PB36: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512], [approach]])	(PB36) The pilot intends to comply with the direction. This belief is derived from PB34, recommended procedures, knowledge of the immediate situation, and knowledge of the flight characteristics of the aircraft.
PB37: blf (intend ([sun512], report ([approach], intend ([sun512], airspeed ([sun512], 170)))), true, [[sun512], [approach]])	(PB37) The pilot intends to inform the controller of the intention PB36 to comply with the direction.

Table 15. Intended Acts of Pilot in Utterance 111

(111) Sun 512: One seven zero, Sundance five twelve

Representation	Notes
PA5: act (acknowl ([approach], direct ([sun512], airspeed ([sun512], 170))))	(PA5) The pilot of Sundance 512 attempts to acknowledge the controller's direction to slow to 170 knots (derived from pilot belief PB35 and expressed with the phrase one seven zero).
PA6: act (report ([approach], intend ([sun512], airspeed ([sun512], 170))))	(PA6) The pilot attempts to inform the controller of the pilot's intention to comply with the direction (from pilot belief PB37).

Table 16. Changes in Pilot's Beliefs After Utterance 111

Representation	Notes
PB38: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170)))), true, [[sun512, approach]])	(PB38) The pilot of Sundance 512 believes that the controller's direction in utterance 110 has been acknowledged. This belief is derived from the pilot's intention PB35 to perform act PA5. It replaces pilot belief PB35.
PB39: blf (occurred ([sun512], report ([approach], intend ([sun512], airspeed ([sun512],170)))), true, [[sun512, approach]])	(PB39) The pilot believes that the controller has been informed of the pilot's intention to comply with the request to slow to 170 knots. The belief is derived from the pilot's intention PB37 to perform act PA6. It replaces PB37.
PB40: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512, approach]])	(PB40) The pilot believes that pilot and controller mutually know that the pilot intends to comply with the request to slow to 170 knots. The belief is derived from PB39. It replaces PB36.

Table 17. Changes in Controller's Beliefs After Utterance 111

Representation	Notes
CB39: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170)))), true, [[approach, sun512]])	(CB39) The controller believes that the pilot of Sundance 512 has acknowledged the controller's direction to slow to 170 knots. This belief is derived from CB37 and from the pilot's phrase **`one seven zero`**. CB39 replaces controller belief CB37.
CB40: blf (occurred ([sun512], report ([approach], intend ([sun512], airspeed ([sun512], 170)))), true, [[approach, sun512]])	(CB40) The controller believes that the pilot has informed the controller of the pilot's intention to slow to 170 knots. This belief is derived from CB38 and from the absence of a statement from the pilot indicating inability to comply. This belief replaces controller belief CB38.
CB41: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[approach, sun512]])	(CB41) The controller believes that the pilot intends to comply with the request to slow to 170 knots. This belief is derived from CB40.

Table 18. Changes in Controller's Beliefs Between Utterances 111 and 114

Representation	Notes
CB42: blf (airspeed ([sun512], 170)), true, [[approach, sun512]])	(CB42) The controller believes that the aircraft has slowed to 170 knots, from information displayed on the console. This belief replaces controller beliefs CB33 and CB41.
CB43: blf (intend ([approach]; direct ([sun512], direction ([sun512], 050)))), true, [[approach]])	(CB43) The controller intends to direct the pilot to change heading to 050. This intention is formed from knowledge of recommended procedures and from controller belief CB30 about the aircraft direction.

Table 19. Changes in Pilot's Beliefs Between Utterances 111 and 114

Representation	Notes
PB41: blf (airspeed ([sun512], 170)), true, [[sun512, approach]])	(PB41) The pilot believes that the aircraft has slowed to 170 knots, from information displayed on the aircraft instruments. This belief replaces pilot beliefs PB33 and PB36.

Table 20. Intended Act of Controller in Utterance 114

(114) Approach: Sundance five twelve turn left heading zero five zero.

Representation	Notes
CA10: act (direct ([sun512], direction ([sun512], 050)))	(CA10) The controller attempts to direct the pilot to fly heading 050. This act is motivated by controller belief CB43 and is expressed with the phrase `turn left heading zero five zero`.

Table 21. Changes in Controller's Beliefs After Utterance 114

Representation	Notes
CB44: blf (occurred ([approach], direct ([sun512], direction ([sun512], 050))), true, [[approach], [sun512]])	(CB44) The controller believes that the pilot of Sundance 512 has been directed to change to heading 050. This belief is derived from the controller's intention CB43 to perform act CA10; it replaces CB43.
CB45: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 050)))), true, [[approach], [sun512]])	(CB45) The controller expects the pilot to acknowledge the controller's direction to change direction. This belief is derived from CB44.
CB46: blf (intend ([sun512], report ([approach], intend ([sun512], direction ([sun512], 050))), true, [[approach], [sun512]])	(CB46) The controller expects the pilot to inform the controller of the pilot's intention to comply with the direction (from CB44).

Table 22. Changes in Pilot's Beliefs After Utterance 114

Representation	Notes
PB42: blf (occurred ([approach], direct ([sun512], direction ([sun512], 050))), true, [[sun512], [approach]])	(PB42) The pilot believes that the controller has directed the pilot to change heading to 050. This belief is derived from the pilot's belief PB30 that the current heading is mutually known to be 100, from the pilot's expectation PB2 that the controller will instruct the pilot to change direction, and from the controller's use of the phrase `turn left heading zero five zero`.
PB43: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 050)))), true, [[sun512], [approach]])	(PB43) The pilot intends to acknowledge the controller's direction to change heading. The belief is derived from PB42.
PB44: blf (intend ([sun512], direction ([sun512], 050)), true, [[sun512], [approach]])	(PB44) The pilot intends to comply with the direction. This belief is derived from PB42 and from the pilot's knowledge of current conditions.
PB45: blf (intend ([sun512,] report ([approach], intend ([sun512], direction ([sun512],050)))), true, [[sun512], [approach]])	(PB45) The pilot intends to inform the controller of PB44.

Table 23. Intended Acts of Pilot in Utterance 115

`(115) Sun 512:    Left zero five zero Sundance five twelve.`

Representation	Notes
PA7: act (acknowl ([approach], direct ([sun512], direction ([sun512], 050)))	(PA7) The pilot of Sundance 512 attempts to acknowledge the controller's direction to change heading to 050. This act is motivated by pilot belief PB43 and is expressed by the phrase **`left zero five zero`**.
PA8: act (report ([approach], intend ([sun512], direction ([sun512], 050))))	(PA8) The pilot attempts to inform the controller of the intention to comply. This belief is motivated by pilot belief PB45 and is expressed by the absence of a statement that the pilot is unable to comply.

Table 24. Changes in Pilot's Beliefs After Utterance 115

Representation	Notes
PB46: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction (sun512, 050)))), true, [[sun512, approach]])	(PB46) The pilot of Sundance 512 believes that the controller's direction in utterance 114 has been acknowledged. This belief is derived from the pilot's intention PB43 to perform PA7. It replaces PB43.
PB47: blf (occurred ([sun512], report ([approach], intend ([sun512], direction ([sun512],050)))), true, [[sun512, approach]])	(PB47) The pilot believes that the controller has been informed of the pilot's intention to comply with the request to change to heading 050. This belief is derived from the pilot's intention PB45 to perform PA8. It replaces and PB45.
PB48: blf (intend ([sun512], direction ([sun512], 050))), true, [[sun512, approach]])	(PB48) The pilot believes that pilot and controller mutually know that the pilot intends to comply with the request to change to heading 050. This belief is derived from PB47 and replaces pilot belief PB44.

Table 25. Changes in Controller's Beliefs After Utterance 115

Representation	Notes
CB47: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 050)))), true, [[approach, sun512]])	(CB47) The controller believes that the pilot of Sundance 512 has acknowledged the controller's direction to change heading to 050. This belief is derived from CB45 and from the pilot's use of the phrase `left zero five zero`. This belief replaces controller belief CB45.
CB48: blf (occurred ([sun512], report ([approach], intend ([sun512], direction ([sun512], 050)))), true, [[approach, sun512]])	(CB48) The controller believes that the pilot has informed the controller of the pilot's intention to comply with the request to change heading to 050. This belief is derived from CB46 and from the absence of a statement that the pilot is unable to comply. This belief replaces controller belief CB46.
CB49: blf (intend ([sun512], direction ([sun512], 050)))), true, [[approach, sun512]])	(CB49) The controller believes that the pilot intends to comply with the request to change heading to 050. This belief is derived from CB48.

Table 26. Changes in Pilot's Beliefs Between Utterances 115 and 125

Representation	Notes
PB49: blf (direction ([sun512], 050)), true, [[sun512,approach]])	(PB49) The pilot believes that the aircraft is now flying heading 050, from information displayed on the aircraft instrument panel. This belief replaces pilot beliefs PB30 and PB44.

Table 27. Changes in Controller's Beliefs Between Utterances 115 and 125

Representation	Notes
CB50: blf (direction ([sun512], 050)), true, [[approach, sun512]])	(CB50) The controller believes that the aircraft is now flying heading 050, from information displayed on the controller's console. This belief replaces CB30 and CB49.
CB51: blf (intend ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[approach]])	(CB51) The controller intends to direct the pilot to change heading to 010. This intention is derived from knowledge of recommended procedures and from controller belief CB50.

Table 28. Intended Act of Controller in Utterance 125

`(125) Approach:   Sundance five twelve turn left heading zero one zero.`

Representation	Notes
CA11: act (direct ([sun512], direction ([sun512], 010)))	(CA11) The controller attempts to direct the pilot to fly heading 010. This act is motivated from controller belief CB51 and is expressed with the phrase `turn left heading zero one zero`.

Table 29. Changes in Controller's Beliefs After Utterance 125

Representation	Notes
CB52: blf (occurred ([approach], direct ([sun512], direction ([sun512], 010))), true, [[approach], [sun512]])	(CB52) The controller believes that the pilot has been directed to change direction to heading 010. This belief is derived from the controller's attempt to perform act CA11 and it replaces controller belief CB51.
CB53: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[approach], [sun512]])	(CB53) The controller expects the pilot to acknowledge the controller's direction to change heading. This belief is derived from CB52 and recommended procedure.
CB54: blf (intend ([sun51]2, report ([approach], intend ([sun512], direction ([sun512], 010))), true, [[approach], [sun512]])	(CB54) The controller expects the pilot to inform the controller of the pilot's intention to comply with the direction (from CB52).

Table 30. Changes in Pilot's Beliefs After Utterance 125

Representation	Notes
PB50: blf (occurred ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[sun512], [approach]])	(PB50) The pilot believes that the controller has directed the pilot to change heading to 010. This belief is derived from the pilot's expectation PB2 that the controller will direct the pilot to change direction, from pilot belief CB50 about the current heading, and from the controller's use of the phrase `turn left heading zero one zero`.
PB51: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[sun512], [approach]])	(PB51) The pilot intends to acknowledge the controller's direction to change heading to 010. This belief is derived from PB50.
PB52: blf (intend ([sun512], direction ([sun512], 010)), true, [[sun512], [approach]])	(PB52) The pilot intends to comply with the direction. This belief is derived from the pilot's knowledge of current conditions.
PB53: blf (intend ([sun512], report ([approach], intend ([sun512], direction ([sun512], 010)))), true, [[sun512], [approach]])	(PB53) The pilot intends to inform the controller of the intention to comply with the direction (from PB52 and recommended procedure).

Table 31. Intended Acts of Pilot in Utterance 126

`(115) Sun 512:    Left zero one zero Sundance five twelve.`

Representation	Notes
PA9: act (acknowl ([approach], direct ([sun512], direction ([sun512], 010)))	(PA9) The pilot of Sundance 512 attempts to acknowledge the controller's direction to change heading to 010 (motivated by PB51 and expressed with the phrase `left zero one zero`).
PA10: act (report ([approach], intend ([sun512], direction ([sun512], 010))))	(PA10) The pilot attempts to inform the controller of the pilot's intention to comply with the direction. This act is motivated by PB53 and expressed by the absence of a statement that the pilot is unable to comply.

Table 32. Changes in Pilot's Beliefs After Utterance 126

Representation	Notes
PB54: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[sun512, approach]])	(PB54) The pilot of Sundance 512 believes that the controller's direction in utterance 125 has been acknowledged. This belief is derived from the pilot's intention PB51 to perform act PA9. It replaces pilot belief PB51.
PB55: blf (occurred ([sun512], report ([approach], intend ([sun512], direction ([sun512], 010)))), true, [[sun512, approach]])	(PB55) The pilot believes that the controller has been informed of the pilot's intention to comply with the request to change to heading 010. This belief is derived from the pilot's intention PB53 to perform act PA10. It replaces pilot belief PB53.
PB56: blf (intend ([sun512], direction ([sun512], 010)), true, [[sun512, approach]])	(PB56) The pilot believes that pilot and controller mutually believe that the pilot intends to comply with the request to change to heading 010. This belief is derived from PB55 and replaces pilot belief PB52.

Table 33. Changes in Controller's Beliefs After Utterance 126

Representation	Notes
CB55: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 010)))), true, [[approach, sun512]])	(CB55) The controller believes that the pilot has acknowledged the direction to change heading to 010. This belief is derived from CB53 and from the pilot's use of the phrase `left zero one zero`. This belief replaces CB53.
CB56: blf (occurred ([sun512], report ([approach], intend ([sun512], direction ([sun512], 010)))), true, [[approach, sun512]])	(CB56) The controller believes that the pilot has informed the controller of the pilot's intention to comply with the request to change heading to 010. This belief is derived from CB54 and from the absence of a statement from the pilot indicating inability to comply. This belief replaces CB54.
CB57: blf (intend ([sun512], direction ([sun512],010)), true, [[approach, sun512]])	(CB57) The controller believes that the pilot intends to comply with the request to change heading to 010. This belief is derived from CB56.

Table 34. Changes in Pilot's Beliefs Between Utterances 126 and 127

Representation	Notes
PB57: blf (direction ([sun512], 010)), true, [[approach, sun512]])	(PB57) The pilot believes that the aircraft has changed direction to heading 010. This belief is derived from information displayed on the pilot's instrument panel. It replaces pilot belief PB52.

Table 35. Changes in Controller's Beliefs Between Utterances 126 and 127

Representation	Notes
CB58: blf (direction ([sun512], 010)), true, [[approach, sun512]])	(CB58) The controller believes that the pilot has changed direction to heading 010. This belief is derived from information displayed on the controller's console. It replaces controller belief CB57.
CB59: blf (intend ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[approach]])	(CB59) The controller intends to direct the pilot to descend to 4500 feet. This intention is formed from knowledge of recommended procedures and of current conditions, and from controller belief CB27 about the aircraft altitude.

Table 36. Intended Act of Controller in Utterance 127

(127) Approach: Sundance five twelve descend and maintain four thousand five hundred.

Representation	Notes
CA12: act (direct ([sun512], altitude ([sun512], 4500)))	(CA12) The controller attempts to direct the pilot to descend to an altitude of 4500 feet. This act is motivated by controller belief CB59 and expressed with the phrase **descend and maintain four thousand five hundred**.

Table 37. Changes in Controller's Beliefs After Utterance 127

Representation	Notes
CB60: blf (occurred ([approach], direct ([sun512], altitude ([sun512], 4500))), true, [[approach], [sun512]])	(CB60) The controller believes that the pilot has been directed to descend to 4500 feet. This belief is derived from the controller's intention CB59 to perform act CA12 and it replaces controller belief CB59.
CB61: blf (intend ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[approach], [sun512]])	(CB61) The controller expects the pilot to acknowledge the controller's direction to change altitude (from CB60 and recommended procedure).
CB62: blf (intend ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500))), true, [[approach], [sun512]])	(CB62) The controller expects the pilot to inform the controller of the pilot's intention to comply with the direction (from CB60).

Table 38. Changes in Pilot's Beliefs After Utterance 127

Representation	Notes
PB58: blf (occurred ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB58) The pilot believes that the controller has directed the pilot to change altitude to 4500 feet. This belief is derived from the pilot's belief PB27 that the current altitude is mutually known to be 5000 feet, from the pilot's expectation PB3 that the controller will direct the pilot to change altitude, and from the controller's use of the phrase `descend and maintain four thousand five hundred.`
PB59: blf (intend ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB59) The pilot intends to acknowledge the controller's direction to change altitude (from PB58 and recommended procedure).
PB60: blf (intend ([sun512], altitude ([sun512], 4500)), true, [[sun512], [approach]])	(PB60) The pilot intends to comply with the direction (from PB58 and from knowledge of current conditions).
PB61: blf (intend ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB61) The pilot intends to inform the controller of the intention to comply with the direction (from PB64 and recommended procedure).

Table 39. Intended Acts of Pilot in Utterance 128

(128) Sun 512: Four thousand five hundred, Sundance five twelve.

Representation	Notes
PA11: act (acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))	(PA11) The pilot attempts to acknowledge the controller's direction to descend to 4500 feet. This act is motivated by pilot belief PB59 and expressed with the phrase **four thousand five hundred.**
PA12: act (report ([approach], intend ([sun512], altitude ([sun512], 4500))))	(PA12) The pilot attempts to inform the controller of the pilot's intention to comply with the direction. This act is motivated by pilot belief PB61 and is expressed by the absence of a statement indicating that the pilot is unable to comply.

Table 40. Changes in Pilot's Beliefs After Utterance 128

Representation	Notes
PB62: blf (occurred ([sun512], acknowl ([approach, direct ([sun512], altitude ([sun512], 4500)))), true, [[sun512, approach]])	(PB62) The pilot believes that the controller's direction in utterance 128 has been acknowledged. This belief is derived from the pilot's intention PB59 to perform act PA11. It replaces pilot belief PB59.
PB63: blf (occurred ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500)))), true, [[sun512, approach]])	(PB63) The pilot believes that the controller has been informed of the pilot's intention to comply with the request to descend to 4500 feet. This belief is derived from the pilot's intention PB61 to perform act PA12; it replaces pilot belief PB61.
PB64: blf (intend ([sun512], altitude ([sun512], 4500)), true, [[sun512, approach]])	(PB64) The pilot believes that pilot and controller mutually know that the pilot intends to comply with the direction. This belief is derived from PB63 and replaces PB64.

Table 41. Changes in Controller's Beliefs After Utterance 128

Representation	Notes
CB63: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[approach, sun512]])	(CB63) The controller believes that the pilot of Sundance 512 has acknowledged the controller's direction to change altitude to 4500 feet. This belief is derived from controller expectation CB61 and from the pilot's use of the phrase `four thousand five hundred`; it replaces controller belief CB61.
CB64: blf (occurred ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500)))), true, [[approach, sun512]])	(CB64) The controller believes that the pilot has informed the controller of the pilot's intention to comply with the request to change altitude to 4500. This belief is derived from controller expectation CB62 and from the absence of a statement indicating that the pilot is unable to comply. This belief replaces controller belief CB62.
CB65: blf (intend ([sun512], altitude ([sun512], 4500)), true, [[approach, sun512]])	(CB65) The controller believes that the pilot intends to comply with the request to change altitude to 4500 (derived from CB64).

Table 42. Changes in Controller's Beliefs Between Utterances 128 and 131

Representation	Notes
CB66: blf (altitude ([sun512], 4500), true, [[approach, sun512]])	(CB66) The controller believes the aircraft has descended to an altitude of 4500 feet. This belief is derived from information displayed on the controllers console; it replaces controller belief CB27 and CB65.

Table 43. Changes in Pilot's Beliefs Between Utterances 128 and 131

Representation	Notes
PB65: blf (altitude ([sun512], 4500), true, [[sun512, approach]])	(PB65) The pilot believes the aircraft has descended to an altitude of 4500 feet. This belief is derived from information displayed on cockpit instruments; it replaces pilot beliefs CB27 and CB65.

E.4 Arrival Instructions

The aircraft is now two to three miles from the approach gate for the ILS-28R Approach, and the controller issues arrival instructions (utterance 131, table Table 46). This is the last vectoring instruction that the controller will issue before the pilot begins executing the published approach procedure. While executing the approach procedure, the pilot will maneuver without explicit instructions from a controller. In this transmission, the controller confirms certain critical information — the aircraft position — and states the altitude, speed, and direction assignments that are to be observed until the pilot can begin the standard approach procedure.

This utterance is the longest transmission in the entire dialogue, using 34 words to express five speech acts:

- a position report (**`eight miles from Laker`**),
- a turn to intercept the final approach course (**`turn left heading three one zero`**),
- authorization to execute a particular approach procedure (**`cleared ILS two eight right approach`**),
- altitude and speed assignments to be observed until the aircraft is established on a segment of the approach procedure (**`Maintain four thousand five hundred till established localizer`** ... **`maintain one seven zero knots until Laker`**).

Despite the amount of information contained, this transmission is delivered at the same rapid-fire pace as the rest of the dialogue. One would expect this to be difficult or impossible to follow, and in fact non-pilots do find this utterance to be nearly unintelligible [Novick 92]. How does the pilot understand it? One explanation is that the pilot strongly expects to receive this set of instructions in the order given. The arrival instructions are largely boilerplate; FAA procedure requires the controller to include all but the speed assignment. Furthermore, much of the content is confirming information that the pilot should already know. The pilot is therefore able to quickly recognize and focus on the most important words: **`eight miles`**, **`heading three one zero`**, **`two eight right`**, **`four thousand five hundred`**, and **`one seven zero knots`**.

Because this model does not capture time or more general conditional instructions, the "until" clauses of this instruction are not represented. Although this is not adequate for ATC instructions in general, it is acceptable in this case; as mentioned above, the "until" clauses refer to conditions that would normally supercede the controller's instructions anyway. Once established on the localizer, the pilot is to descend along the glideslope defined by the localizer. After passing the Laker beacon, the pilot should be adjusting speed in accordance with approach procedure and the performance characteristics of the aircraft.

Table 44. Changes in Controller's Beliefs Before Utterance 131

Representation	Notes
CB67: blf (position ([sun512], laker, 8), true, [[approach], [sun512]])	(CB67) The controller believes that the aircraft is approximately eight miles from the Laker beacon. This belief is derived from information displayed on the controller's console.
CB68: blf (intend ([approach], authorize ([sun512], approach (pdx, ils-28R))), true, [[approach], [sun512]])	(CB68) The controller intends to authorize the pilot to execute the ILS-28R approach procedure into PDX. This belief is derived from controller beliefs CB9 (the pilot's destination) and CB10 (the approach in use) and from information displayed on the console.
CB69: blf (intend ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]]))), true, [[approach], [sun512]])	(CB69) The controller intends to inform the pilot of the aircraft position. This belief is derived from controller belief CB67 and from recommended procedure.
CB70: blf (intend ([approach], direct ([sun512], altitude ([sun512], 4500))), true, [[approach]])	(CB70) The controller intends to direct the pilot to maintain an altitude of 4500 feet until the aircraft begins executing the ILS-28R approach procedure. This belief is derived from CB66 and from information displayed on the controller's console.
CB71: blf (intend ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[approach]])	(CB71) The controller intends to direct the pilot to maintain an airspeed of 170 knots until the aircraft flies over the Laker beacon. This belief is derived from CB44 about the aircraft altitude and from information displayed on the console.

Table 44. Changes in Controller's Beliefs Before Utterance 131 (Continued)

Representation	Notes
CB72: blf (intend ([approach], direct ([sun512], direction ([sun512], 310))), true, [[approach]])	(CB72) The controller intends to direct the pilot to change heading to 310. This belief is derived from the controller's beliefs about the aircraft position (CB67) and heading (CB58) and from information displayed on the controller's console.

Table 45. Changes in Pilot's Beliefs Before Utterance 131

Representation	Notes
PB66: blf (position ([sun512], laker, 8), true, [[sun512], [approach]])	(PB66) The pilot believes that the aircraft is approximately eight miles from the Laker beacon. This belief is derived from information displayed on the cockpit instruments.

Table 46. Intended Acts of Controller in Utterance 131

(131) Approach: Sundance five twelve eight miles from Laker turn left heading three one zero. Maintain four thousand five hundred till established localizer cleared ILS two eight right approach, maintain one seven zero knots until Laker.

Representation	Notes
CA13: act (report ([sun512], blf (position ([sun512], laker, 8) true, [[approach], [sun512]])))	(CA13) The controller attempts to inform the pilot of the controller's belief that the aircraft is approximately eight miles from the Laker beacon. This act is motivated by controller belief CB69 and is expressed by the use of the phrase **eight miles from Laker**.
CA14: act (direct ([sun512], direction ([sun512], 310)))	(CA14) The controller attempts to direct the pilot to fly heading 310. This act is motivated by controller belief CB72 and expressed by the use of the phrase **turn left heading three one zero**.
CA15: act (direct ([sun512], altitude ([sun512], 4500)))	(CA15) The controller attempts to direct the pilot to maintain an altitude of 4500 feet until the aircraft begins executing the ILS-28R approach procedure. This act is motivated by controller belief CB70 and expressed by the use of the phrase **maintain four thousand five hundred till established localizer**.

Table 46. Intended Acts of Controller in Utterance 131 (Continued)

Representation	**Notes**
CA16: act (authorize ([sun512], approach (pdx, ils-28R)))	(CA16) The controller intends to authorize the pilot to execute the ILS-28R instrument approach procedure into PDX. This act is motivated by CB68 and is expressed by the use of the phrase **`cleared ILS two eight right approach.`**
CA17: act (direct ([sun512], airspeed ([sun512], 170)))	(CA17) The controller intends to direct the pilot to maintain an airspeed of 170 knots until the aircraft flies over the Laker beacon. This act is motivated by controller belief CB71 and is expressed by the use of the phrase **`maintain one seven zero knots until Laker`**.

Table 47. Changes in Controller's Beliefs After Utterance 131	
Representation	**Notes**
CB73: blf (occurred ([approach], authorize ([sun512], approach (pdx, ils-28R))), true, [[approach], [sun512]])	(CB73) The controller believes that the pilot has been authorized to execute the ILS-28R instrument approach procedure into PDX. This belief is derived from the controller's intention CB68 to perform CA16. This belief replaces CB68.
CB74: blf (intend ([sun512], acknowl ([approach], authorize ([sun512], approach (pdx, ils-28R)))), true, [[approach], [sun512]])	(CB74) The controller expects the pilot acknowledge the authorization. This belief is derived from CB73 and recommended procedure.
CB75: blf (occurred ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]]))), true, [[approach], [sun512]])	(CB75) The controller believes that the pilot has been informed of the controller's belief that the aircraft is approximately eight miles from the Laker beacon. This belief is derived from the controller's intention CB69 to perform act CA13. This belief replaces controller belief CA13.
CB76: blf (intend ([sun512], acknowl ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]])))), true, [[approach], [sun512]])	(CB76) The controller expects the pilot to acknowledge the controller's report of the aircraft position. This belief is derived from controller belief CB75.

Table 47. Changes in Controller's Beliefs After Utterance 131 (Continued)

Representation	Notes
CB77: blf (intend ([sun512], report ([approach], blf (position ([sun512], laker, 8), true, [[sun512, approach]]))), true, [[approach], [sun512]])	(CB77) The controller further expects the pilot to confirm the controller's report of the aircraft position. This belief is derived from controller belief CB67 about the aircraft position and from controller belief CB75 that the position report occurred.
CB78: blf (occurred ([approach], direct ([sun512], altitude ([sun512], 4500))), true, [[approach], [sun512]])	(CB78) The controller believes that the pilot has been directed to maintain an altitude of 4500 feet until the aircraft begins executing the ILS-28R approach procedure. This belief is derived from the controller's intention CB70 to perform act CA15. This belief replaces CB70.
CB79: blf (intend ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[approach], [sun512]])	(CB79) The controller expects the pilot to acknowledge the altitude assignment. This belief is derived from the controller's belief CB78 that the assignment occurred.
CB80: blf (intend ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500)))), true, [[approach], [sun512]])	(CB80) The controller expects the pilot to inform the controller of the pilot's intention to comply with the altitude assignment. This belief is derived from controller act CB78.
CB81: blf (occurred ([approach], direct ([sun512], direction ([sun512], 310))), true, [[approach], [sun512]])	(CB81) The controller believes that the pilot has been directed to fly heading 310. This belief is derived from the controller's intention CB72 to perform act CA14. This belief replaces controller belief CB72.

Table 47. Changes in Controller's Beliefs After Utterance 131 (Continued)

Representation	**Notes**
CB82: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 310)))), true, [[approach], [sun512]])	(CB82) The controller expects the pilot to acknowledge the heading assignment. This belief is derived from CB81.
CB83: blf (intend (sun512, report (approach, intend (sun512, direction ([sun512],310)))), true, [[approach], [sun512]])	(CB83) The controller expects the pilot to inform the controller of the pilot's intention to comply with the heading assignment. This belief is derived from CB81.
CB84: blf (occurred ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[approach], [sun512]])	(CB84) The controller believes that the pilot has been directed to maintain an airspeed of 170 knots until the aircraft flies over the Laker beacon. This belief is derived from the controller's intention CB71 to perform CA17. This belief replaces controller belief CB71.
CB85: blf (intend ([sun512], acknowl ([approach], direct ([sun512, airspeed ([sun512], 170)))), true, [[approach], [sun512]])	(CB85) The controller expects the pilot to acknowledge the speed assignment. This belief is derived from CB84.
CB86: blf (intend ([sun512], report ([approach], intend ([sun512], airspeed ([sun512], 170)))), true, [[approach], [approach]])	(CB86) The controller expects the pilot to inform the controller of the pilot's intention to comply with the speed assignment. This belief is derived from CB84.

Table 48. Changes in Pilot's Beliefs After Utterance 131

Representation	Notes
PB67: blf (occurred ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]]))), true, [[sun512], [approach]])	(PB67) The pilot believes that the controller has reported that the aircraft is approximately eight miles from the Laker beacon. This belief is derived from the controller's use of the phrase `eight miles from Laker`.
PB68: blf (intend ([sun512], acknowl ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]])))), true, [[sun512], [approach]])	(PB68) The pilot intends to acknowledge the controller's position report. This belief is derived from pilot's belief PB67 and from recommended procedure.
PB69: blf (intend ([sun512], report ([approach], blf (position ([sun512], laker, 8), true, [[sun512, approach]])))), true, [[sun512], [approach]])	(PB69) The pilot further intends to confirm the controller's position report. This belief is derived from pilot's belief PB67 and from the pilot's belief PB66 about the aircraft position.
PB70: blf (occurred ([approach], direct ([sun512], direction ([sun512], 310))), true, [[sun512], [approach]])	(PB70) The pilot believes that the controller has directed the pilot to fly heading 310. This belief is derived from the controller's use of the phrase `turn left heading three one zero` and from the pilot's expectation PB2.

Table 48. Changes in Pilot's Beliefs After Utterance 131 (Continued)

Representation	Notes
PB71: blf (intend ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 310)))), true, [[sun512], [approach]])	(PB71) The pilot intends to acknowledge the heading assignment. This belief is derived from pilot belief PB70 and from recommended procedure.
PB72: blf (intend ([sun512], report ([approach], intend ([sun512], direction ([sun512], 310)))), true, [[sun512], [approach]])	(PB72) The pilot intends to comply with the heading assignment. This belief is derived from pilot belief PB55 and from knowledge of the aircraft capabilities and of the current conditions.
PB73: blf (intend ([sun512], report ([approach], intend ([sun512], direction ([sun512], 310)))), true, [[sun512], [approach]])	(PB73) The pilot intends to inform the controller of the pilot's intention to comply with the heading assignment. This belief is derived from pilot belief PB72 and from recommended procedure.
PB74: blf (occurred ([approach], direct ([sun512], altitude ([sun512], 4500))), true, [[sun512], [approach]])	(PB74) The pilot believes that the controller has directed the pilot to maintain an altitude of 4500 feet until the aircraft begins executing the ILS-28R approach procedure. This belief is derived from the controller's use of the phrase **maintain four thousand five hundred till established localizer**.

Table 48. Changes in Pilot's Beliefs After Utterance 131 (Continued)

Representation	**Notes**
PB75: blf (intend ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB75) The pilot intends to acknowledge the altitude assignment. This belief is derived from pilot belief PB74 and from recommended procedure.
PB76: blf (intend ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB76) The pilot intends to comply with the altitude assignment. This belief is derived from pilot belief PB65 that the aircraft is at 4500 feet and from knowledge of the flight characteristics of the aircraft and of the current conditions.
PB77: blf (intend ([sun512], report ([approach], intend ([sun512], altitude ([sun512], 4500)))), true, [[sun512], [approach]])	(PB77) The pilot intends to inform the controller of the pilot's intention to comply with the altitude assignment. This belief is derived from pilot belief PB76 and from recommended procedure.
PB78: blf (occurred ([approach], authorize ([sun512], approach (pdx, ils-28R))), true, [[sun512], [approach]])	(PB78) The pilot believes that the controller has authorized the pilot to execute the ILS-28R instrument approach procedure into PDX. This belief is derived from the controller's phrase `cleared ILS two eight right approach` and from pilot expectation PB11.

Table 48. Changes in Pilot's Beliefs After Utterance 131 (Continued)

Representation	Notes
PB79: blf (intend ([sun512], acknowl ([approach], authorize ([sun512], approach (pdx, ils-28R))), true, [[sun512], [approach]])	(PB79) The pilot intends to acknowledge the authorization. This belief is derived from PB78.
PB80: blf (occurred ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[sun512], [approach]])	(PB80) The pilot believes that the controller has directed the pilot to maintain an airspeed of 170 knots until the aircraft flies over the Laker beacon. This belief is derived from the controller's use of the phrase `maintain one seven zero knots until Laker` and from pilot belief PB41 of the current aircraft speed.
PB81: blf (intend ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170))), true, [[sun512], [approach]])	(PB81) The pilot intends to acknowledge the speed assignment. This belief is derived from pilot belief PB80 and from recommended procedure.
PB82: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512], [approach]])	(PB82) The pilot intends to comply with the speed assignment. This belief is derived from knowledge of the aircraft characteristics and the current conditions.
PB83: blf (intend ([sun512], report ([approach], intend ([sun512], airspeed ([sun512], 170))), true, [[sun512], [approach]])	(PB83) The pilot intends to inform the controller of the pilot's intention to comply with the speed assignment. This belief is derived from pilot belief PB82 and from recommended procedure.

Table 49. Intended Acts of Pilot in Utterance 132

(132) Sun 512: **Three one zero uh four thousand five hundred we'll maintain uh, speed, (and cleared for the) ILS two eight right Sundance five twelve.**

Representation	Notes
PA13: act (acknowl ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]]))))	(PA13) The pilot attempts to acknowledge the position report. This act is motivated by PB68 and is expressed by the absence of a request for clarification to the position report.
PA14: act (report ([approach], blf (position ([sun512], laker, 8), true, [[sun512, approach]])))	(PA14) The pilot attempts to confirm the position report. This act is motivated by PB69 and is expressed by the absence of a correction to the position report.
PA15: act (acknowl ([approach], direct ([sun512], direction ([sun512], 310))))	(PA15) The pilot attempts to acknowledge the heading assignment. This act is motivated by pilot belief PB71 and from the use of the phrase **three one zero**.
PA16: act (report ([approach], blf (intend ([sun512], direction ([sun512], 310)), true, [[sun512], approach]])))	(PA16) The pilot attempts to inform the controller that the pilot intends to comply with the heading assignment. This act is motivated by PB73 and is expressed by the absence of a statement that the pilot is unable to comply.

Table 49. Intended Acts of Pilot in Utterance 132 (Continued)

Representation	Notes
PA17: act (acknowl ([approach], direct ([sun512], altitude ([sun512], 4500))))	(PA17) The pilot attempts to acknowledge the altitude assignment. This act is motivated by PB75 and is expressed with the phrase **four thousand five hundred.**
PA18: act (report ([approach], blf (intend ([sun512], altitude ([sun512], 4500)), true, [[sun512], approach]])))	(PA18) The pilot attempts to inform the controller of the pilot's intention to comply with the altitude assignment. This act is motivated by PB77 and is expressed by the absence of a statement that the pilot is unable to comply.
PA19: act (acknowl ([approach], direct ([sun512], airspeed ([sun512], 170))))	(PA19) The pilot intends to acknowledge the speed assignment. This belief is motivated by PB81 and is expressed with the phrase **we'll maintain uh, speed.**
PA20: act (report ([approach], blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512], [approach]])))	(PA20) The pilot attempts to inform the controller of the pilot's intention to comply with the speed assignment. This belief is derived from pilot belief PB83 and is expressed by the phrase **we'll maintain uh, speed**
PA21: act (acknowl ([approach], authorize ([sun512], approach (pdx, ils-28R))))	(PA21) The pilot attempts to acknowledge the authorization. This act is motivated by PB79 and is expressed by the pilot's use of the phrase **cleared for the ILS two eight right.**

Table 50. Changes in Pilot's Beliefs After Utterance 132

Representation	Notes
PB84: blf (occurred ([sun512], acknowl ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]]))), true, [[sun512, approach]])	(PB84) The pilot believes that the position report has been acknowledged. This belief is derived from the pilot's intention PB68 to perform act PA13. This belief replaces PB68.
PB85: blf (occurred ([sun512], report ([approach], blf (position ([sun512], laker, 8), true, [[sun512, approach]]))), true, [[sun512, approach]])	(PB85) The pilot believes that the position report has been confirmed. This belief is derived from the pilot's intention PB69 to perform act PA14. This belief replaces PB69.
PB86: blf (position ([sun512], laker, 8)), true, [[sun512, approach]])	(PB86) The pilot believes that pilot and controller are mutually aware of the current position of the aircraft. This belief is derived from PB85 and replaces PB66.
PB87: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 310)))), true, [[sun512, approach]])	(PB87) The pilot believes that the heading assignment has been acknowledged. This belief is derived from the pilot's intention PB71 to perform act PA15. It replaces PB71.

Table 50. Changes in Pilot's Beliefs After Utterance 132 (Continued)

Representation	Notes
PB88: blf (occurred ([sun512], report ([approach], blf (intend ([sun512], direction ([sun512], 310)), true, [[sun512, approach]])), true, [[sun512, approach]])	(PB88) The pilot believes that the controller has been informed of the pilot's intention to comply with the heading assignment. This belief is derived from the pilot's intention PB73 to perform act PA16. It replaces PB73.
PB89: blf (intend ([sun512], direction ([sun512], 310)), true, [[sun512, approach]])	(PB89) The pilot believes that pilot and controller are mutually aware of the pilot's intention to comply with the heading assignment. This belief is derived from PB88 and replaces PB72.
PB90: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[sun512, approach]])	(PB90) The pilot believes that the altitude assignment has been acknowledged. This belief is derived from the pilot's intention PB75 to perform act PA17. It replaces PB75.
PB91: blf (occurred ([sun512], report ([approach], blf (intend ([sun512], altitude ([sun512], 4500)), true, [[sun512, approach]]))), true, [[sun512, approach]])	(PB91) The pilot believes that the controller has been informed of the pilot's intention to comply with the altitude assignment. This belief is derived from the pilot's intention PB77 to perform act PA18. It replaces PB77.

Table 50. Changes in Pilot's Beliefs After Utterance 132 (Continued)

Representation	Notes
PB92: blf (intend ([sun512], altitude ([sun512], 4500)), true, [[sun512, approach]])	(PB92) The pilot believes that pilot and controller are mutually aware of the pilot's intention to comply with the altitude assignment. This belief is derived from PB91 and replaces PB76
PB93: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], airspeed ([sun512], 170)))), true, [[sun512, approach]])	(PB93) The pilot believes that the speed assignment has been acknowledged. This belief is derived from the pilot's intention PB81 to perform act PA19. It replaces PB81.
PB94: blf (occurred ([sun512], report ([approach], blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512], [approach]]))), true, [[sun512, approach]])	(PB94) The pilot believes that the controller has been informed of the pilot's intention to comply with the speed assignment. This belief is derived from the pilot's intention PB83 to perform act PA20. It replaces PB83.
PB95: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512, approach]])	(PB95) The pilot believes that pilot and controller are mutually aware of the pilot's intention to comply with the speed assignment. This belief is derived from PB94 and replaces PB82
PB96: blf (occurred ([sun512], acknowl ([approach], authorize ([sun512], approach (pdx, ils-28r)))), true, [[sun512, approach]])	(PB96) The pilot believes that the authorization has been acknowledged. This belief is derived from the pilot's intention PB79 to perform act PA21. It replaces PB79.

Table 50. Changes in Pilot's Beliefs After Utterance 132 (Continued)

Representation	Notes
PB97: blf (authorize ([sun512], approach (pdx, ils-28r)), true, [[sun512, approach]])	(PB97) The pilot believes that both controller and pilot mutually know that the aircraft is authorized to execute the ILS-28R approach procedure. This belief is derived from PB96 and replaces PB78. At this point, the indefinite expectations that the controller will vector the aircraft to the approach gate (PB1, PB2, and PB3) are withdrawn.

Table 51. Changes in Controller's Beliefs After Utterance 132

Representation	Notes
CB87: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], direction ([sun512], 310)))), true, [[sun512, approach]])	(CB87) The controller believes that the pilot has acknowledged the heading assignment. This belief is derived from controller belief CB82 and from the pilot's phrase **three one zero**. It replaces controller belief CB82.
CB88: blf (occurred (sun512, report (approach, blf (intend (sun512, direction (sun512, 310)), true, [[sun512, approach]]))), true, [[approach, sun512]])	(CB88) The controller believes that the pilot has reported the intention of complying with the heading assignment. This belief is derived from expectation CB83 and from the absence of a statement that the pilot is unable to comply. This belief replaces CB83.
CB89: blf (intend (sun512, direction (sun512, 310)), true, [[approach, sun512]])	(CB89) The controller believes that the pilot intends to comply with the heading assignment. This belief is derived from belief CB88.
CB90: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], altitude ([sun512], 4500)))), true, [[approach, sun512]])	(CB90) The controller believes that the altitude assignment has been acknowledged. This belief is derived from controller expectation CB79 and from the pilot's phrase **four thousand five hundred**. It replaces controller belief CB79.

<table>
<tr><th colspan="2">Table 51. Changes in Controller's Beliefs After Utterance 132 (Continued)</th></tr>
<tr><th>Representation</th><th>Notes</th></tr>
<tr><td><pre>CB91:
  blf (occurred ([sun512],
                report ([approach],
                       blf (intend ([sun512],
                                   altitude ([sun512],
                                             4500)),
                           true,
                           [[sun512, approach]]))),
      true,
      [[approach, sun512]])</pre></td><td>(CB91) The controller believes that the pilot has reported the intention to comply with the altitude assignment. This belief is derived from expectation CB80 and from the absence of a statement that the pilot is unable to comply. This belief replaces CB80.
.</td></tr>
<tr><td><pre>CB92:
  blf (intend ([sun512],
              altitude ([sun512], 4500)),
      true,
      [[approach, sun512]])</pre></td><td>(CB92) The controller believes controller and pilot mutually know that the pilot intends to comply with the altitude assignment. This belief is derived from belief CB91.</td></tr>
<tr><td><pre>CB93:
  blf (occurred ([sun512],
                acknowl ([approach],
                         direct ([sun512],
                                airspeed ([sun512],
                                          170)))),
      true,
      [[approach, sun512]])</pre></td><td>(CB93) The controller believes that the speed assignment has been acknowledged. This belief is derived from controller expectation CB85 and from the pilot's phrase <code>we'll maintain uh, speed</code>. It replaces controller belief CB85.</td></tr>
<tr><td><pre>CB94:
  blf (occurred ([sun512],
                report ([approach],
                       blf (intend ([sun512],
                                   airspeed ([sun512],
                                             170)),
                true,
                [[sun512], [approach]]))),
      true,
      [[sun512, approach]])</pre></td><td>(CB94) The controller believes that the pilot has reported the intention of complying with the speed assignment. This belief is derived from controller expectation CB86 and from the pilot's phrase <code>we'll maintain uh, speed</code>. It replaces controller belief CB86.</td></tr>
</table>

Table 51. Changes in Controller's Beliefs After Utterance 132 (Continued)

Representation	Notes
CB95: blf (intend ([sun512], airspeed ([sun512], 170)), true, [[sun512, approach]])	(CB95) The controller believes that the controller and pilot mutually know that the pilot intends to comply with the speed assignment. This belief is derived from belief CB94.
CB96: blf (occurred ([sun512], acknowl ([approach], authorize ([sun512], approach (pdx, ils-28R)))), true, [[approach, sun512]])	(CB96) The controller believes that the authorization has been acknowledged. This belief is derived from controller expectation CB74 and from the pilot's phrase `cleared for the ILS two eight right`. This belief replaces CB74 .
CB97: blf (authorize ([sun512], approach (pdx, ils-28R)))), true, [[approach, sun512]])	(CB97) The controller believes that pilot and controller mutually know that Sundance 512 is cleared for the ILS-28R approach. This belief is derived from CB96 and replaces CB73.
CB98: blf (occurred ([sun512], acknowl ([approach], report ([sun512], blf (position ([sun512], laker, 8), true, [[approach], [sun512]])))), true, [[approach, sun512]])	(CB98) The controller believes that the position report has been acknowledged. This belief is derived from controller expectation CB76 and from the absence of pilot request for clarification. This belief replaces CB76.

Table 51. Changes in Controller's Beliefs After Utterance 132 (Continued)

Representation	Notes
CB99: blf (occurred ([sun512], report ([approach], blf (position ([sun512], laker, 8), true, [[sun512, approach]]))), true, [[approach, sun512]])	(CB99) The controller believes that the position report has been confirmed. This belief is derived from controller expectation CB77 and from the absence of pilot correction. This belief replaces CB77.
CB100: blf (position ([sun512], laker, 8), true, [[approach, sun512]])	(CB100) The controller believes that controller and pilot mutually know the current position of the aircraft. This belief is derived from controller belief CB99 and replaces belief CB67.

Table 52. Changes in Controller's Beliefs Between Utterance 132 and 157

Representation	Notes
CB101: blf (direction ([sun512], 310)), true, [[approach, sun512]])	(CB101) The controller believes that the Sundance 512 is flying heading 310. This belief is derived from information displayed on the console. It replaces beliefs CB89 and CB58.

Table 53. Changes in Pilot's Beliefs Between Utterance 132 and 157

Representation	Notes
PB98: blf (direction ([sun512], 310)), true, [[sun512, approach]])	(PB98) The pilot believes that the aircraft is flying heading 310. This belief is derived from information displayed on the cockpit instruments. It replaces beliefs CB89 and CB58.

E.5 Handoff

The handoff is the final exchange between Sundance 512 and the approach controller. The pilot has reached the edge of the tower controller's airspace, so the approach controller hands off the flight to the tower and then instructs the pilot to contact the tower on radio frequency 118.7. The tower frequency is published information, so the pilot expects it and in fact may have it already set on the second radio.

The pilot's response — **good-bye** — is not approved phraseology. It is, however, a relevant next contribution [Clark 89], and as such strongly signals understanding and acceptance of the controller's contribution. This, then, is an example of more general conversational principles explaining deviation from the prescribed response.

The analysis ends at this point. Note that the approach controller's belief is updated to show that the controller believes that the handoff occurred, but the pilot's beliefs are not similarly updated. This is because the controller doesn't expect to hear another transmission from Sundance 512 and so believes that all is well after a suitable interval has passed without report of a problem. The pilot's beliefs would be updated to show that the handoff occurred only after contact has been made with the tower controller — but that is part of another dialogue and so is not represented here.

Table 54. Changes in Controller's Beliefs Before Utterance 157

Representation	Notes
CB102: blf (intend [approach], direct ([sun512], contact (tower))), true, [[approach], [sun512]])	(CB102) The controller intends to direct the pilot to contact the tower controller. This belief is derived from pilot belief CB9 and from the aircraft position.

Table 55. Changes in Pilot's Beliefs Before Utterance 157

Representation	Notes
PB99: blf (intend [approach], direct ([sun512], contact (tower))), true, [[sun512], [approach]])	(PB99) The pilot expects to be handed off to the tower. This belief is derived from pilot belief PB9.

Table 56. Intended Act of Controller in Utterance 157

```
(157) Approach:  Sundance five twelve contact tower one
                 one eight point seven.
```

Representation	Notes
CA18: act (direct ([sun512], contact ([tower])))	(CA18) The approach controller attempts to direct the pilot to contact the PDX tower controller. This belief is motivated by CB102 and is expressed with the phrase `contact tower one one eight point seven`.

Table 57. Changes in Controller's Beliefs After Utterance 157

Representation	Notes
CB103: blf (occurred ([approach], direct ([sun512], contact ([tower]))), true, [[approach], [sun512]])	(CB103) The approach controller believes that the pilot has been directed to contact the tower controller. This belief is derived from the controller intention CB102 to perform act CA18.
CB104: blf (intend ([sun512], acknowl ([approach], direct ([sun512], contact ([tower])))), true, [[approach], [sun512]])	(CB104) The controller expects the pilot to acknowledge the direction. This belief is derived from CB103 and from recommended procedure.
CB105: blf (intend ([sun512], report ([approach], intend ([sun512], direct ([sun512], contact [(tower])))), true, [[approach], [sun512]])	(CB105) The controller expects the pilot to report the intention to comply with the direction to contact the tower. This belief is derived from CB103 and from recommended procedure.

Table 58. Changes in Pilot's Beliefs After Utterance 157

Representation	Notes
PB100: blf (occurred ([approach], direct ([sun512], contact ([tower]))), true, [[sun512], [approach]])	(PB100) The pilot believes that the approach controller has directed the pilot to contact the tower controller. This belief is derived from the controller's use of the phrase **`contact tower one one eight point seven`** and replaces pilot belief PB99.
PB101: blf (intend ([sun512], acknowl ([approach], direct ([sun512], contact ([tower])))), true, [[sun512], [approach]])	(PB101) The pilot intends to acknowledge the controller's handoff. This belief is derived from PB100 and from recommended procedure.
PB102: blf (intend ([sun512], contact ([tower])), true, [[sun512], [approach]])	(PB102) The pilot intends to comply with the handoff.
PB103: blf (intend ([sun512], report ([approach], intend ([sun512], direct ([sun512], contact ([tower]))))), true, [[sun512], [approach]])	(PB103) The pilot intends to inform the approach controller of the intention to comply with the handoff. This belief is derived from PB102 and from recommended procedure.

Table 59. Intended Acts of Pilot in Utterance 158

(158) Sun 512: Five twelve good-bye.

Representation	**Notes**
PA22: act (acknowl ([approach], direct ([sun512], contact ([tower]))))	(PA22) The pilot attempts to acknowledge the controller's handoff. This act is motivated by pilot intention PB101.
PA23: act (report ([approach], intend ([sun512], contact ([tower]))))	(PA23) The pilot attempts to inform the approach controller of the intention to comply with the handoff. This act is motivated by PB103 and is expressed with the phrase **good-bye**.

Table 60. Changes in Pilot's Beliefs After Utterance 158

Representation	Notes
PB104: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], contact ([tower]))), true, [[sun512, approach]])	(PB104) The pilot believes that the controller's handoff has been acknowledged. This belief is derived from the pilot's intention PB101 to perform PA22. This belief replaces PB101.
PB105: blf (occurred ([sun512], report ([approach], intend ([sun512], contact ([tower]))), true, [[sun512, approach]])	(PB105) The pilot believes that the controller has been informed of the intention to comply with the handoff. This belief is derived from the pilot's intention PB103 to perform PA23. This belief replaces PB103.
PB106: blf (intend ([sun512], contact ([tower])), true, [[sun512, approach]])	(PB106) The pilot believes that pilot an controller mutually know that the pilot intends to comply with the hand-off. This belief is derived from PB105 and replaces PB102.

Table 61. Changes in Controller's Beliefs After Utterance 158

Representation	Notes
CB106: blf (occurred ([sun512], acknowl ([approach], direct ([sun512], contact ([tower]))), true, [[approach, sun512]])	(CB106) The controller believes that the handoff has been acknowledged. This belief is derived from and replaces controller expectation CB104.
CB107: blf (occurred ([sun512], report ([approach], intend ([sun512], contact ([tower]))), true, [[approach, sun512], [approach, tower]])	(CB107) The controller believes that the pilot has informed the controller of the intention to comply with the hand-off. This belief is derived from controller expectation CB105 and the pilot's use of the phrase good-bye. This belief replaces CB105.
CB108: blf (intend ([sun512], contact ([tower])), true, [[approach, sun512], [approach, tower]])	(CB108) The controller believes that pilot and controller mutually know that the pilot intends to comply with the handoff. This belief is derived from CB107.
CB109: blf (occurred ([sun512], contact ([tower])), true, [[approach], [sun512, tower]])	(CB109) After a suitable interval, the controller believes that the pilot has contacted the tower. This belief is derived from CB108.

Appendix F

Charts

Figure F1. Portland Inset of Seattle Sectional Aeronautical Chart, 37th edition, June 29, 1989, U.S. Department of Commerce.

Figure F2. Excerpts of Entry for Portland, OR, Airport/Facility Directory, Northwest U.S., November 16, 1989, U.S. Department of Commerce, pp. 82-83.

Figure F3. ILS RWY 28R, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 213.

Figure F4. River visual RWY 28R, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 221.

Figure F5. Airport diagram, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 222.

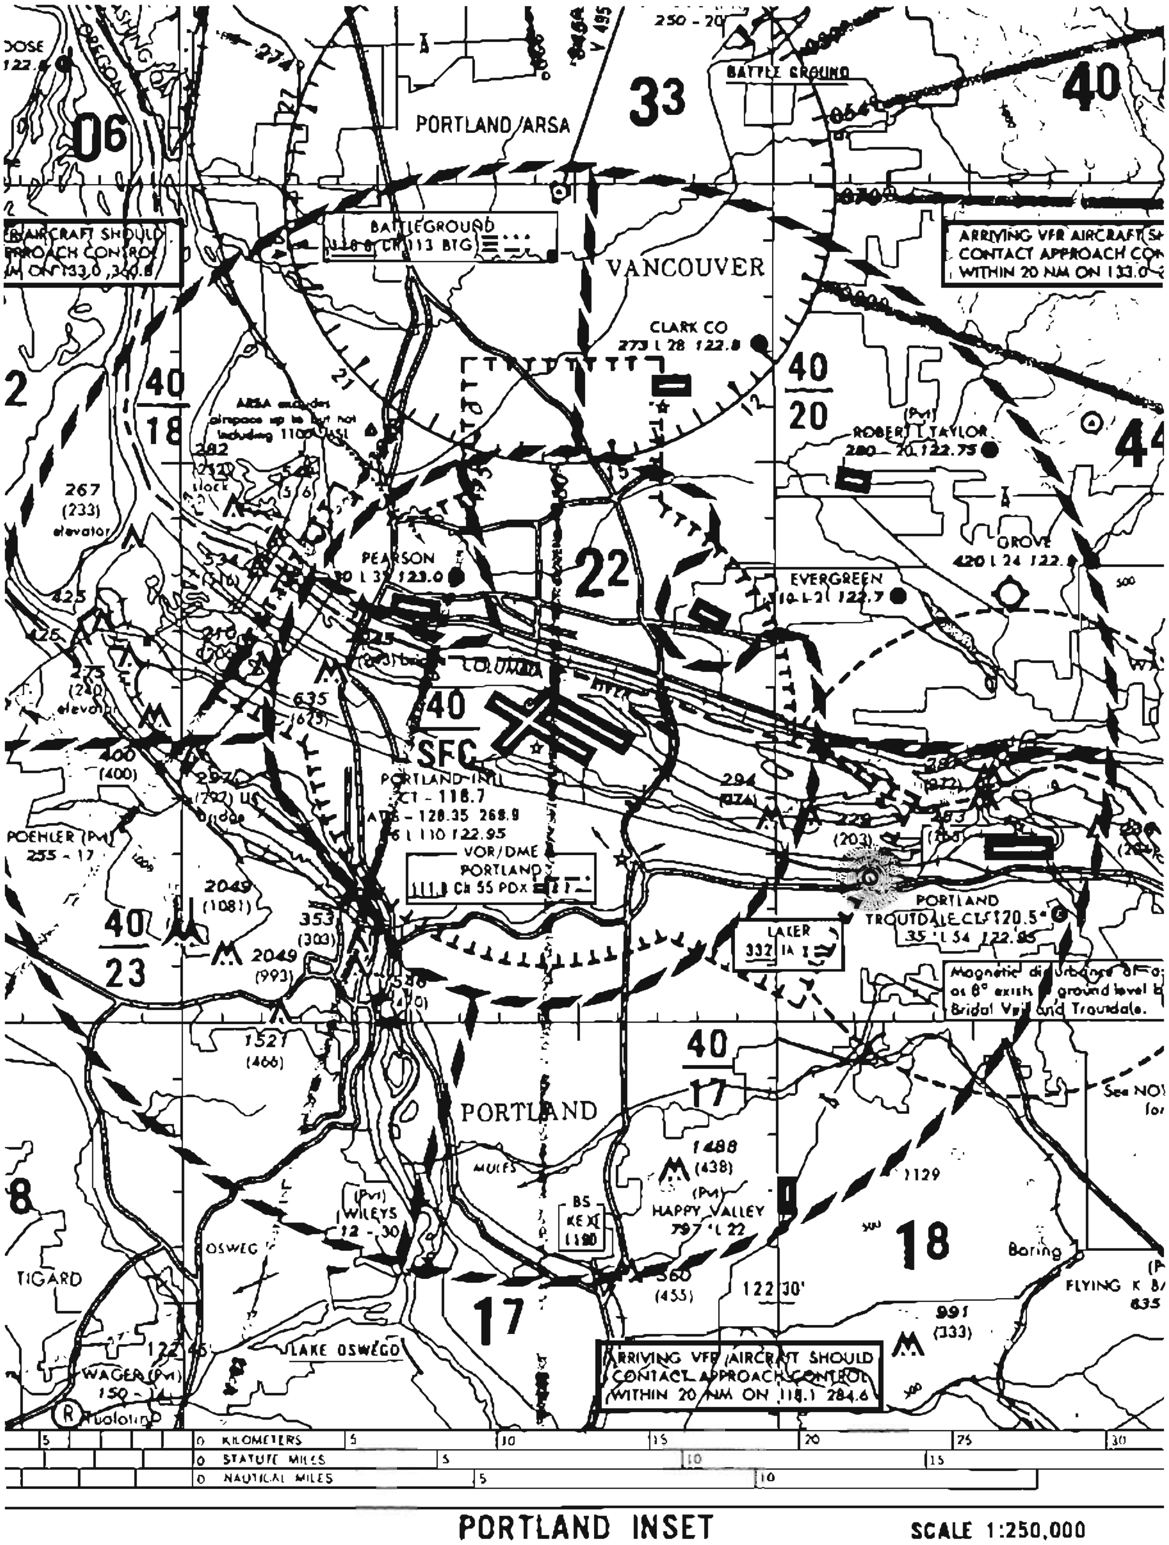

Figure F1. Portland Inset of Seattle Sectional Aeronautical Chart, 37th Edition, June 29, 1989, U.S. Department of Commerce.

§ **PORTLAND INTL** (PDX) 4 NE UTC-8(-7DT) 45°35'20"N 122°35'47"W **SEATTLE**
27 B S4 **FUEL** 100, JET A OX 1, 2, 3, 4 LRA ARFF Index D **H-1A, L-1C**
RWY 10R-28L: H11011X150 (ASPH-GRVD) S-200, D-200, DT-360 HIRL CL **IAP**
RWY 10R: ALSF2 TDZ Rgt tfc. Arresting device **RWY 28L:** VASI(V4L)—GA 3.0°TCH 60'. Arresting device.
RWY 10L-28R: H8004X150 (ASPH-GRVD) S-200, D-200, DT-400 HIRL
RWY 10L: MALS. VASI(V4L)—GA 3.0°TCH 60'. **RWY 28R:** ALSF1. Levee. Rgt tfc.
RWY 02-20: H7000X150 (ASPH) S-124, D-170, DT-310 MIRL
RWY 20: REIL. VASI(V4L)—GA 4.0°TCH 33'. Thld dsplcd 489'. Road.
AIRPORT REMARKS: Attended continuously. CLOSED to non-powered acft except in Emerg. Twy inner Bravo CLOSED to acft with wing span greater than 108'. Taxiway S3B CLOSED to acft over 12,500 lbs gross weight; Exit S5 south of Rwy 10R-28L CLOSED to acft over 65,000 lbs gross weight. Exit S4 south of Rwy 10R-28L CLOSED to acft over 20,000 lbs gross weight. 180° turns by acft heavier than 12,500 lbs prohibited on all rwys and taxiways. Acft weighing more than 65,000 lbs are restricted from using taxiway G. Migratory and wintering flocks of large waterfowl on and in vicinity of arpt. Heavy seagull activity Sep.-Apr., except high number of birds year around. Rwy 28L arrivals are noise sensitive, expect apch to 28R with transition to 28L. Noise abatement procedures in effect call arpt manager, 503-231-5000, ext 407. Flight Notification Service (ADCUS) available
COMMUNICATIONS: ATIS 128.35 (503)284-1711 **UNICOM** 122.95
MC MINNVILLE FSS (MMV) TF 1-800-WX-BRIEF. NOTAM FILE PDX
Ⓡ **APP CON** 133.0 (280°-099°) 118.1 (100°-279°)
TOWER 118.7 **GND CON** 121.9 **CLNC DEL** 119.0
Ⓡ **DEP CON** 133.0 (280°-099°) 118.1 (100°-279°) 124.35
ARSA ctc **APP CON**
RADIO AIDS TO NAVIGATION: NOTAM FILE PDX.
BATTLE GROUND (H) ABVORTACW 116.6 BTG Chan 113 45°44'53"N 122°35'25"W 161° 9.6 NM to fld 250/21E. General outlook on TWEB 0600-1300Z‡. TWEB out of svc indefinitely.
(L) VORW/DME 111.8 POX Chan 55 45°35'38"N 122°36'18"W at fld. 19/20E.
VOR/DME unusable 131°-230°, 322°-039°.
LAKER NDB (MHW/LOM) 332 LA 45°32'29"N 122°27'40"W 277° 6.4 NM to fld
ILS/DME 109.9 I-PDX Chan 36 Rwy 10R Category II available
ILS 111.3 I-IAP Rwy 28R LOM LAKER NDB.
ILS/DME 108.9 I-GPO Chan 26 Rwy 20 Localizer only

§ **PORTLAND-TROUTDALE** (TTD) 10 E UTC-8(-7DT) 45°32'58"N 122°24'00"W **SEATTLE**
35 B S4 **FUEL** 100LL, JET A **H-1A, L-1C**
RWY 07-25: H5400X150 (ASPH) S-19 MIRL **IAP**
RWY 07: REIL. VASI(V4L)—GA 3.0°TCH 50'. Tree. **RWY 25:** REIL. VASI(V4L)—GA 3.0°TCH 42'. Tree.
AIRPORT REMARKS: Attended Apr-Oct 1500-0430Z‡ Nov-Mar 1500-0300Z‡. CAUTION: Migratory flocks of waterfowl on and in vicinity of arpt. ACTIVATE MIRL Rwy 07-25—CTAF. Control Zone effective 1500-0600Z‡
WEATHER DATA SOURCES: LAWRS
COMMUNICATIONS: CTAF 120.5 **UNICOM** 122.95
MC MINNVILLE FSS (MMV) TF 1-800-WX-BRIEF. NOTAM FILE TTD.
Ⓡ **PORTLAND APP CON** 133.0 (280°-099°) 118.1 (100°-279°) Ⓡ **PORTLAND DEP CON** 133.0
TROUTDALE TOWER 120.5 (1500-0600Z‡). **GND CON** 121.8
RADIO AIDS TO NAVIGATION: NOTAM FILE PDX.
BATTLE GROUND (H) ABVORTACW 116.6 BTG Chan 113 45°44'53"N 122°35'25"W 125° 14.4 NM to fld 250/21E. General outlook only 0600-1300Z‡. TWEB out of svc indefinitely
LAKER NDB (MHW/LOM) 332 LA 45°32'29"N 122°27'40"W 059° 2.6 NM to fld.
COMM/NAVAID REMARKS: Freq 121.5 not avbl at tower

Figure F2. Excerpts of entry for Portland, OR, Airport/Facility Directory, Northwest U.S., November 16, 1989, U.S. Department of Commerce, pp. 82-83.

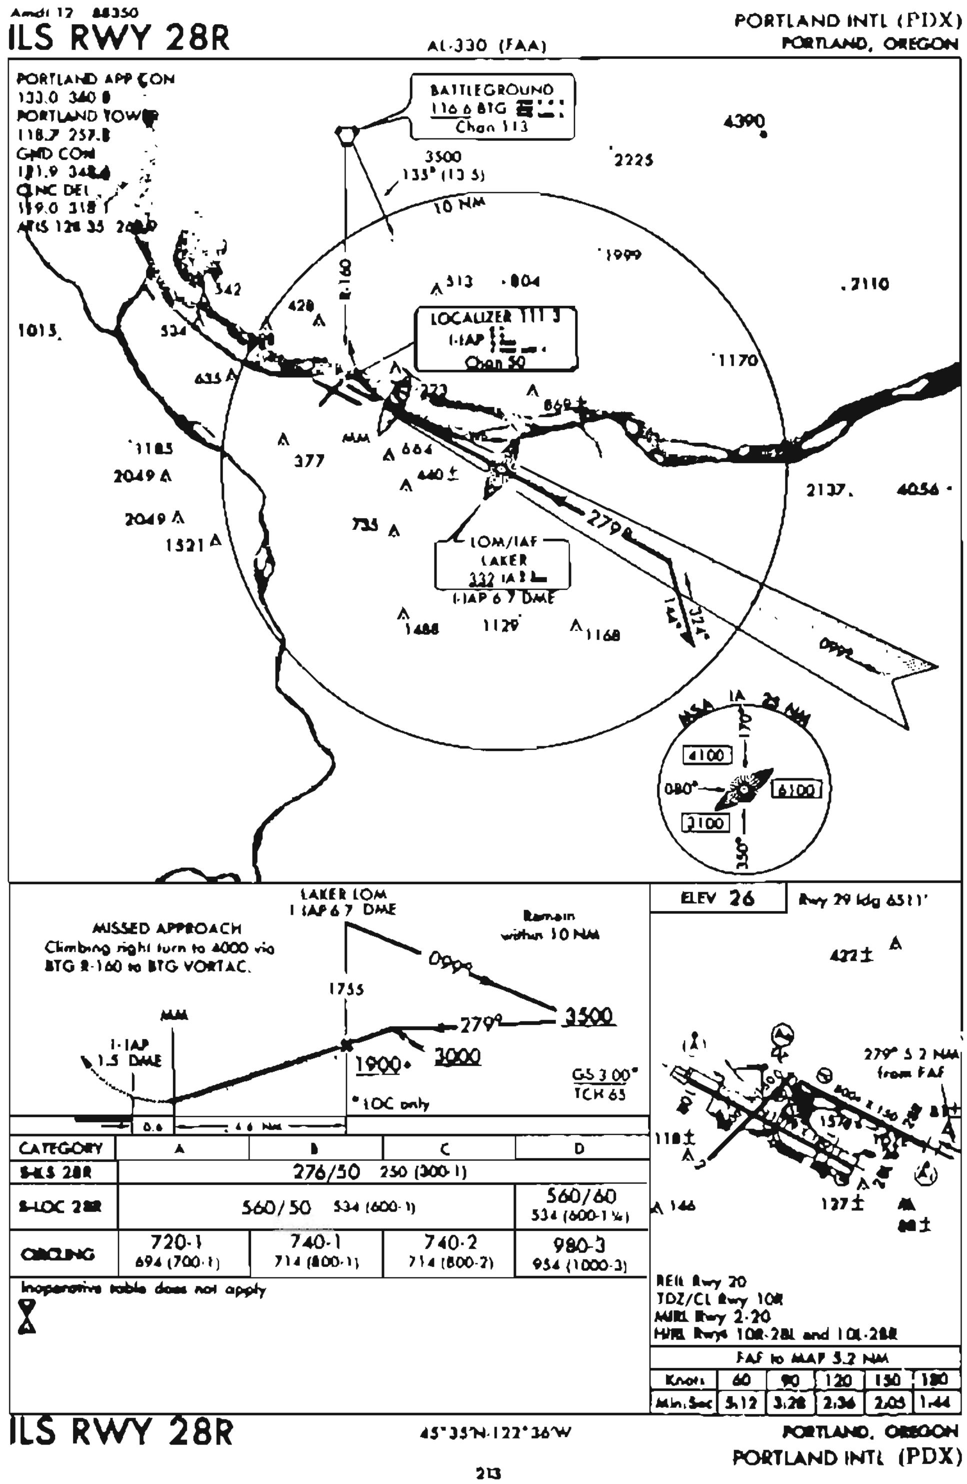

Figure F3. ILS RWY 28R, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 213.

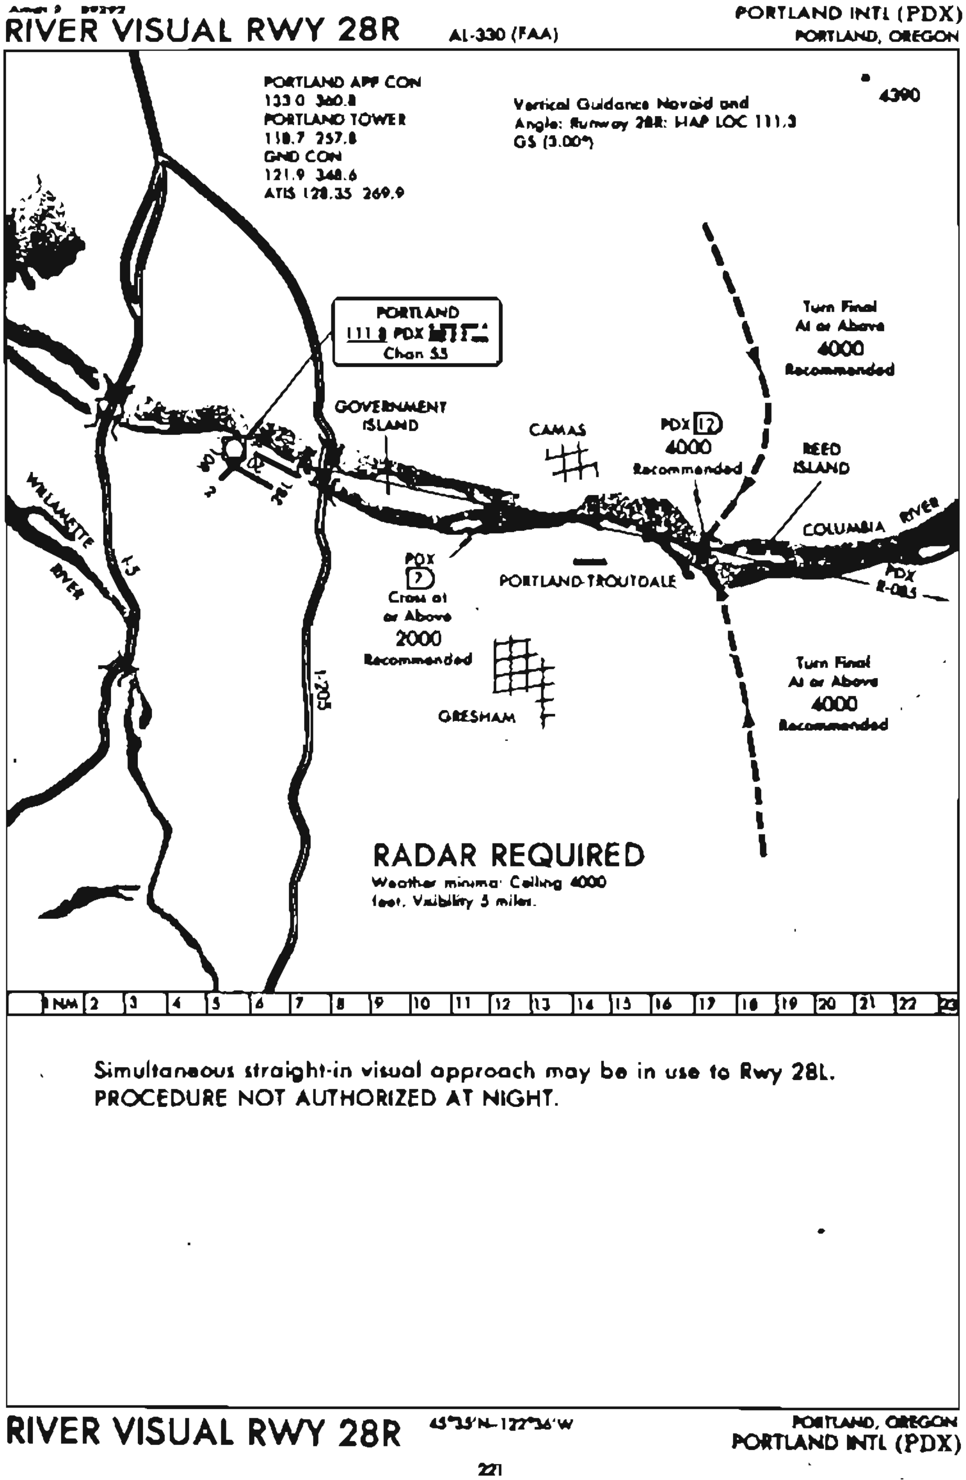

Figure F4. River Visual RWY 28R, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 221.

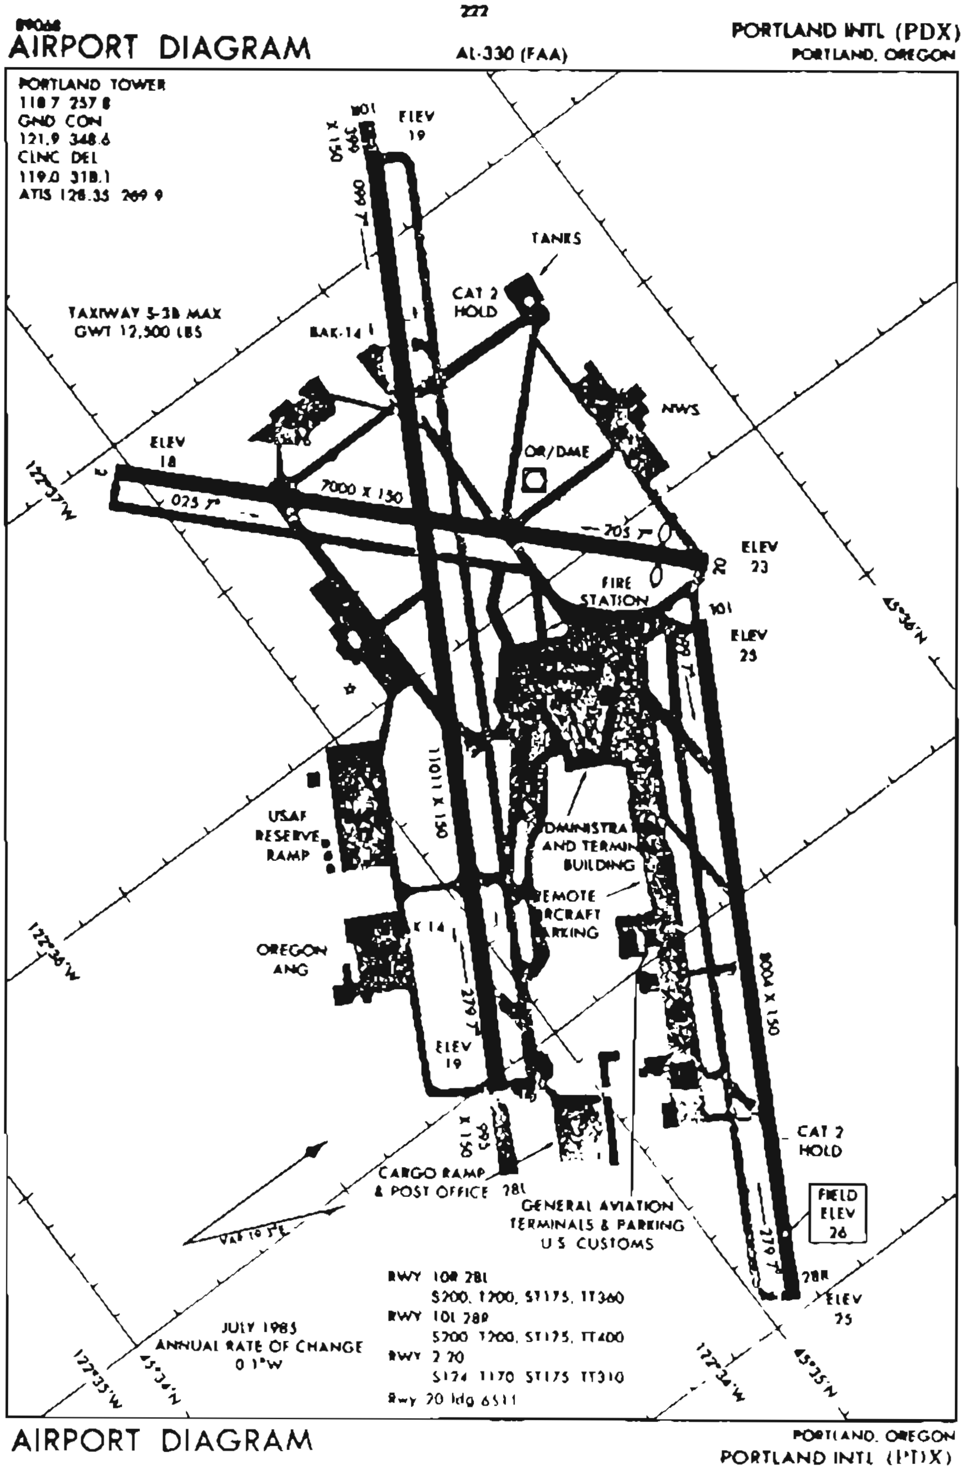

Figure F5. Airport Diagram, Portland Intl (PDX), Instrument Approach Procedures, U.S. Northwest, October 19, 1989, U.S. Department of Commerce, pg. 222.

Biographical Sketch

Despite being a native Oregonian, Karen Ward was born in Lexington, Kentucky on July 17, 1956. A few months later her parents realized the error of their ways and returned to Oregon to stay, so Karen grew up a few miles north of Portland, Oregon in the town of St. Helens.

In 1973, Karen moved to Eugene, Oregon to attend the University of Oregon. She returned to the Portland area in 1978 with a B. Sci. in Computer Science and accepted a position at Consolidated Freightways, first as a business applications programmer and later as a trainer and software standards specialist. In 1979, she moved to Portland General Electric and remained there until returning to school full time in 1991. While at PGE, Karen worked with teams responsible for developing and maintaining major software systems in support of PGE's engineering and business activities. During this time, her professional interests focused on the issues and problems encountered in designing and implementing large systems and in expanding and enhancing existing systems.

Karen entered the Ph.D. program at Oregon Graduate Institute in 1991. Her current research interests center around the use of higher-level language models in spoken language understanding systems.