Perceptual control and layered protocols in interface design: II. The general protocol grammar

ARTICLE *in* INTERNATIONAL JOURNAL OF HUMAN-COMPUTER STUDIES · JUNE 1999

Impact Factor: 1.29 · DOI: 10.1006/ijhc.1998.0258 · Source: DBLP

3 AUTHORS, INCLUDING:

Justin G Hollands

Defence Research and Dev…

57 PUBLICATIONS **1,717** CITATIONS

SEE PROFILE

Int. J. Human-Computer Studies (1999) **50,** 521–555
Article No. ijhc.1998.0258
Available online at http://www.idealibrary.com on IDEAL®

Perceptual control and layered protocols in interface design: II. The general protocol grammar

M. M. Taylor
Martin Taylor Consulting, 369 Castlefield Avenue, Toronto, Ont., Canada M3M 3B9. email: mmt@interlog.com

P. S. E. Farrell and J. G. Hollands
Defence and Civil Institute of Environmental Medicine, Box 2000, Toronto, Ont., Canada M3M 3B9. email: philip.farrell@dciem.dnd.ca, justin.holland@dciem.dnd.ca

Perceptual control theory (PCT) is a framework theory for psychology, based on the tenet "All behaviour is the control of perception." Layered protocol theory (LPT) is PCT applied to the special case of communication between cooperating partners, each controlling their own perceptions and many levels of abstraction within a dialogue. This paper discusses some perceptual control processes that occur within a single dialogue level, in the form of a General Protocol Grammar that is asserted to be valid for every level of every dialogue. A companion paper is concerned with LPT applied to the design and analysis of human–machine interfaces.

1. Introduction

1.1. PERCEPTUAL CONTROL AND LAYERED PROTOCOLS

In a companion paper (Farrell *et al*., 1999, this issue) we discuss the use of perceptual control theory (PCT, Powers, 1973) and layered protocol theory (LPT, Taylor 1988*a*, *b*, 1989—a specialization of PCT for communication) to analyse and design human–machine interfaces. In the form of LPT, Farrell *et al.* used PCT to analyse the interface to helicopter flight management system (the control display unit, CDU) and to sketch the design of an improved version (Farrell *et al*., 1999, this issue; Farrell & Semprie, 1997). The present paper discusses in more detail a central feature of LPT, the "General Protocol Grammar" (GPG), which describes the ways in which possible communication errors are avoided or corrected at any one level of the Layered Protocol hierarchy.

Both PCT and LPT are based on the statement that "All behaviour is the control of perception" (Powers, 1973). "Perception," in this statement, refers to an internal state or "signal". The value of a perceptual signal corresponds to some state of the environment that can be sensed through sensor systems. Being only the value of a signal, which might be a rate of neural firing, a voltage, or a chemical concentration, perception need not be conscious. It occurs at many different levels of abstraction, corresponding to a perception of the intensity of a sound, through the perception of a word, to the perception of a complex concept such as "democracy". To "control" anything is to bring it to

1071-5819/99/060521 + 35 $30.00/0

a reference condition and to maintain it there in the face of external influences that would alter it; to control perception is to act on the outer world so that the corresponding internal state is maintained near some reference value.

Perceptions, whether controlled or uncontrolled, are simply signal values. Although the nature of the signals may be different in biological organisms and constructed systems, LPT suggests that dialogue between humans and machines can be analysed functionally in the same way as dialogue between humans. The present paper and its companion treat human–human dialogue as the more general case, and refers specifically to human–machine dialogue only where it seems useful to point out important restrictions of the machine. We do not treat human–computer interaction as *analogous* to human–human interaction; we treat each as an expression of the same underlying universe of necessity.

1.2. COMMUNICATION: HUMAN–HUMAN AND HUMAN–MACHINE

1.2.1. Communication and message

According to PCT, all actions occur because some perceived state of the world differs from a reference (goal) value for that state. Actions are performed to change the perceived state of the world, to bring it closer to the goal state. If the changed state is that of another person (or, we claim, of a sufficiently complex machine), we call the actions "communicative actions", or "messages". Messages are "sent" by an originator to a recipient, and are successfully communicated if the originator perceives the recipient to have come to the desired state. To ease the discussion, we give the originator and recipient names beginning with "O" and "R", and of different genders. We call the originator Oliver and the recipient Rachel. Either name may represent a machine or a human.

Communication need not be cooperative. Oliver may, for example, be concerned only that Rachel perform some action, and may not even want her to perceive that he is acting to influence her. Rachel, on the other hand, may want to determine what Oliver wants of her without letting Oliver know that she is observing him. Military or commercial deception may take either of these forms. In either case, there would be communication, but no dialogue.

For a dialogue to occur, not only must Oliver have a goal whose accomplishment involves Rachel in some way, but Rachel also must have a goal that involves Oliver—more specifically Oliver's actions in respect of her. For the dialogue to be "cooperative", one of Rachel's goals must be to discover what state Oliver wishes to bring about in her—perhaps a belief Oliver wants her to accept, an action Oliver wants her to perform, or some information he wants her to provide.

If there is a dialogue, Oliver originates a "message", which Rachel receives. To "originate a message" is to control a perception that the dialogue is bringing the other partner (the recipient) to some desired condition—a goal. To "receive a message" is to understand the condition desired by the originator, though not necessarily to come to the desired condition. It is not even necessary that Rachel be aware that a communication is taking place, or even know that Oliver exists, provided Oliver can perceive Rachel's state and can influence it toward the state he wants for her. Of course, if Rachel is unaware of the existence of a communication, she is not being cooperative, and the derived grammar will differ from the one we develop for cooperative communication.

In this paper we deal only with cooperative dialogue, in which both parties want to be satisfied with the success of the message transmission. In human–machine interaction, it is reasonable to assume that the machine has been designed to help the user do what the user wants, and that the user, in using the machine, is trying to get it to do what is wanted. Both are necessarily being cooperative—which does not mean that the interaction will always be easy or successful!

1.2.2. The General Protocol Grammar as description of perceptual control

In a cooperative dialogue Rachel controls a perception that Oliver is satisfied with the effect of the dialogue, but does not control (i.e. stabilize against external influences) the perception of the message content. Oliver likewise controls a perception that Rachel is satisfied that Oliver is satisfied. The control of such perceptions and their further recursions is the topic of this paper. Together, they induce what we call the "General Protocol Grammar" (GPG). The GPG is a central element of what we call a "protocol", the knowledge and procedures the partners use to communicate a message at a single level of abstraction.

The GPG is no compendium of rules and structures that must be followed if the dialogue is to be "grammatical"; rather, it is a categorization of the kinds of things participants ordinarily do at various stages of a dialogue to bring their various controlled perceptions to the corresponding reference values. Any rules and structures that may be discovered by a non-participant analyst are by-products of Oliver's intention that he come to perceive Rachel as achieving some state, and (in a cooperative dialogue) of Rachel's intention that she come to perceive Oliver as satisfied that she has come to the state he desires.

"Rules" in dialogue are no more than patterns that have been seen by analysts to recur in various dialogues within a culture. They recur because the partners have found or been taught that particular patterns of action are often effective in serving their communicative purposes. Although in some dialogues formal, explicit, patterns may constrain the participants' actions, in most dialogues the participants use no overt rules, and may not know explicitly the regularities that occur in their dialogues.

The GPG is a compendium of the things that are reasonably likely to happen in the control of the set of perceptions involved in the communication of a single message. The same kinds of controlled perceptions may be common to every abstraction level of an interaction. The reference values of most of these perceptions are predetermined by the assumption that the dialogue is cooperative. It is this commonality of reference values that allows the derivation of a protocol grammar that is general across all the levels of abstraction in a dialogue.

We derive the GPG by considering the relations between what the participants in a dialogue want to perceive (reference values for controlled perceptions) and what they may actually be perceiving at different stages of the dialogue. Each participant in a dialogue deals only with his, her, or its individual goals and perceptions. Oliver has some task goal external to the dialogue that might be aided by the recipient, Rachel. To achieve his goal, Oliver acts so as to influence Rachel. He "originates a communication" that Rachel receives; he "sends a message" to Rachel.

From the viewpoint of an omniscient external analyst, a message has been successfully passed by means of a cooperative dialogue when Rachel correctly understands or does

what Oliver wants her to understand or do, and both Oliver and Rachel are satisfied that this is the case. "To understand" means for Rachel either to come to the state Oliver desires, or at least to perceive correctly what that state is (though she may not want to come to that state).

The participants are not omniscient, and neither can ever know for sure that Rachel understands; nor can either Rachel or Oliver ever know for sure that the other is satisfied that she understands. From Oliver's viewpoint, communication is successful when he *perceives* that Rachel understands—when he perceives her to have come to the state he desires, or at least to know what that state is—which is very different from correctly knowing that she understands. From Rachel's point of view (if she intends to cooperate) the communication is successful when she believes she understands and she perceives Oliver to be satisfied that she does. Each of them acts throughout the dialogue to bring these and other related beliefs to their reference values.

Oliver cannot know for sure that Rachel has understood correctly, though he may see that her actions satisfy his goal; and Rachel cannot know for sure that what she has understood is what Oliver wanted her to understand, though she may see that he seems to be satisfied with the result of the communication. The best they can do is to gather evidence that can lead to beliefs about various aspects of the state of the other and of the other's beliefs. For example, even though Oliver can never know that Rachel understands, he can strongly believe it; to come to this strong belief, one piece of evidence he might use is a perception that Rachel believes she understands, which he can link with an existing belief (trust) that Rachel actually understands when she thinks she does. "Trust" is an important component of the GPG, as we shall see.

LPT uses no concept analogous to "mutual belief" (e.g. Perrault, 1989), since only an omniscient external observer could determine that a mutual belief exists at any moment. The participants cannot, and therefore the state of a "mutual belief" can never affect the course of a dialogue.

1.3. CONTROLLED PERCEPTIONS AND BELIEFS

1.3.1. Analogies between Layered Protocols and Perceptual Control

We start our analysis of dialogue by taking the viewpoint of the originator (Oliver), who has a goal that Rachel comes to a particular state. Oliver perceives Rachel to be in some state and perhaps to be changing state in some direction. If he perceives her to be in or to be moving toward the state he desires, then he need not act to influence her. Therefore, we must assume that Oliver now perceives Rachel's state both not to be as he desires, and also not likely to be as he desires in the near future. If she is to come to the state he desires, he must act. For example, Oliver may perceive Rachel as not knowing something he wants her to know, or perhaps she may not be doing something he wants her to do.

We can put these notions into the framework of either LPT or PCT, since LPT is a special case of PCT. In PCT, the basic unit is called an "Elementary Control Unit", consisting of a Perceptual Input Function that constructs a perceptual signal from its many inputs, a Comparator that produces an error signal from the difference between the reference signal and the perceptual signal, and an Output Function that produces an output signal from the error signal. The output signal eventually affects the inputs to the Perceptual Input Function, thus completing the loop.

In LPT, the elementary unit corresponding to the Elementary Control Unit of PCT is the "Protocol Node". The correspondences between the Elementary Control Unit and the Protocol Node are explored in the companion paper. To summarize, the Decoder corresponds to the Perceptual Input Function of the Elementary Control Unit. The Decoder produces a perception of Rachel's current state (in respect of several variables which are discussed below). The state in which Oliver wants to perceive Rachel is called the "Primal Message" (corresponding to the "reference perception" of PCT). The difference between the Primal Message and output of the Decoder is constructed in a component of the Protocol Node called the "Model". This difference corresponds to the error signal that drives the output function in the PCT Elementary Control Unit. A "Coder" that transform the error signal into a "virtual message" perceptible to Rachel corresponds to the Output Function of the Elementary Control Unit.

Rachel has a corresponding Protocol Node. It has a "Decoder" that interprets the virtual message and also allows her to perceive aspects of Oliver's state. She does not control the content of the virtual message, but perceives it passively. She does, however, control aspects of Oliver's state relating to this perception of her success in interpreting the message, and her perceptions of those aspects are compared in her Model with her reference values for them. The resulting error is transformed by her Coder into virtual messages that are the inputs to Oliver's Decoder, completing what we call a "protocol loop", or "channel" (Figure 1). We call Oliver's Protocol Node a "transmitting node" and Rachel's a "receiving node".

Several individual perceptual control loops exist within a single protocol loop. In our analysis we identify 19 of them, as discussed below. Ten of these loops concern control by Oliver of a state in Rachel, displayed by Rachel through virtual messages to Oliver, and nine concern control by Rachel of a state in Oliver. The "state in which Oliver wants to perceive Rachel" includes not only the task-related goal (Primal Message) that led him to initiate the dialogue, but also components relating to whether he wants the dialogue to be cooperative and to the progress of the communication. If he does want the dialogue to be cooperative, the state in which he wants to perceive her includes that she believes him to be satisfied with the state she has achieved. Rachel's reference state for Oliver includes similar elements. Most particularly, if she wants to be cooperative, it includes a reference perception that she should perceive Oliver to be satisfied that she has adequately interpreted his Primal Message.

When we talk about the "Primal Message" we usually mean only the task-related components of the state Oliver wants to perceive Rachel as achieving, because the perceptions that relate to the progress of the communication are the same for all dialogues and at every level of a dialogue. Each partner's reference perceptions include perceptions relating to the other's satisfaction; these reference perceptions, being always assumed, are not ordinarily included in the graphical representations of LPT. The GPG, on the other hand, is specifically concerned with the control of these beliefs, and assumes without further analysis that there exist content-specific mechanisms to deal with task-related perceptions. Interface design must, of course, go beyond the GPG, to specify those mechanisms.

1.3.2. Interpretation and cooperation

In what follows, we concentrate on the beliefs that relate to the progress of the dialogue. When we talk about interpreting the Primal Message—Rachel coming to, or coming to

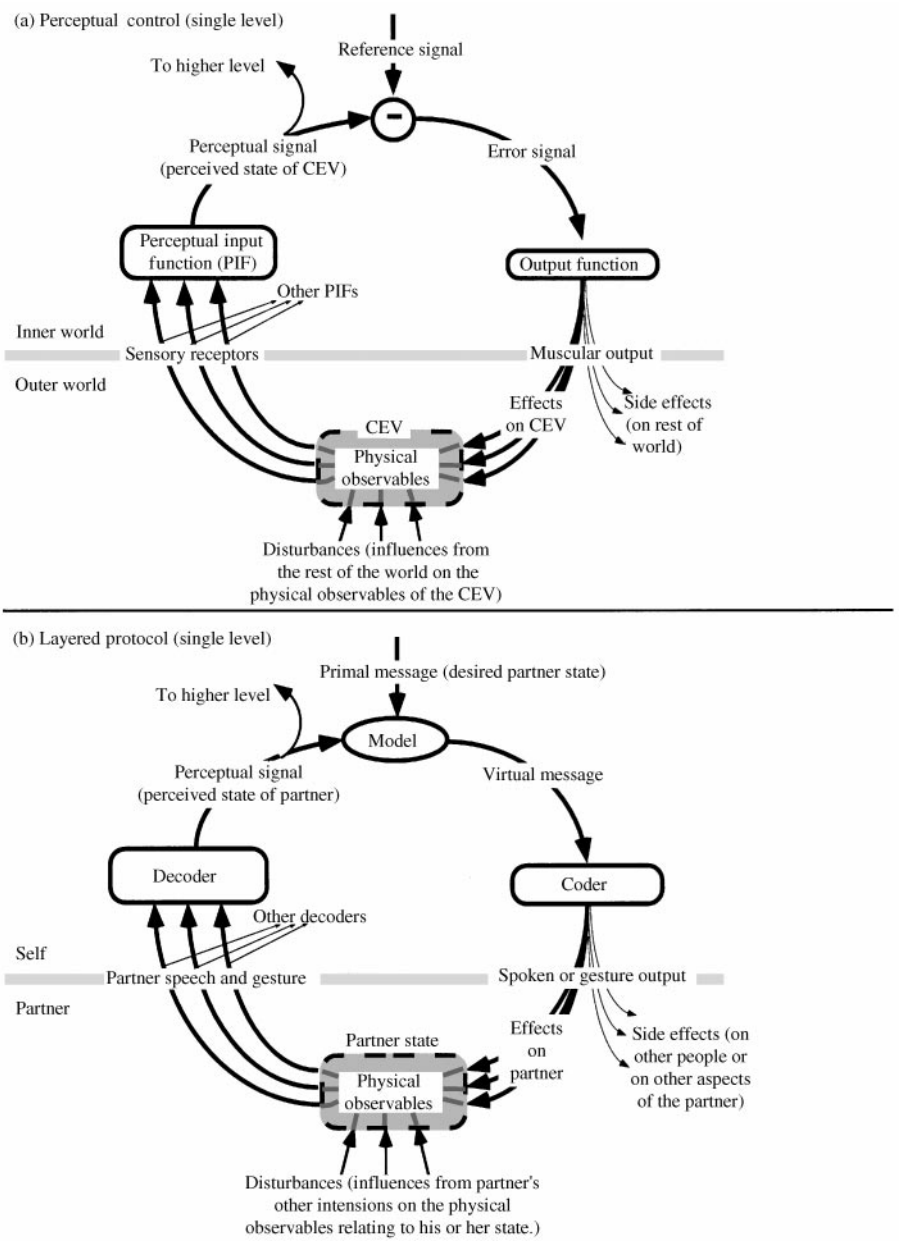

FIGURE 1. Analogy between the signal types in PCT and the message types in LPT. Just as the PCT input and output signals seem to come from and to act directly on the CEV although they actually pass through lower level controls units, so the LPT virtual messages seem to come from and go to the partner directly although they are actually implemented by lower-level protocol nodes.

know, the state Oliver intends for her—we will ignore the mechanisms Rachel uses to form her interpretation, simply taking it for granted that such mechanisms exist. When she has made an adequate interpretation, either Rachel will be in the state Oliver desires, or she will at least perceive what Oliver desires. But the question still exists as to whether either of them believes her to have arrived at the desired state, or believes the other to believe it. This is the question that the GPG addresses.

We can identify two distinct conditions for Oliver to initiate a dialogue: (1) that Oliver perceives Rachel's state to differ from his goal (reference) state of her and (2) that Oliver cares about the difference strongly enough to act so as to bring her state closer to his reference. Putting it semi-formally, we can define, as a reference belief for Oliver (and, in a cooperative dialogue, for Rachel as well) a statement *S*.

S: The recipient has adequately interpreted the Primal Message.

"Adequately" here is the important word. It is in Oliver's eyes that the interpretation must be adequate. Oliver's reference value (that *S* be true) for his perception of Rachel's state constitutes a general goal for the originator of any communication. For Rachel, if she is being cooperative, it is not enough that she believes that she has made an adequate interpretation; she must also come to perceive that Oliver believes she has, which defines for her another reference belief that we discuss below.

S is a statement about the world. Neither Oliver nor Rachel ever know how nearly Rachel has come to the state Oliver want. Oliver necessarily has a goal to perceive *S* to be true. Rachel may or may not have a belief in the truth of *S* as a goal, depending on whether she is being cooperative. At any moment, Oliver's observation of Rachel may lead him to believe or disbelieve *S* weakly or strongly, or have no opinion on it. Rachel likewise may have some level of belief in the truth of *S*, whether or not *S* is true in fact, and her belief about *S* may differ from Oliver's perception of her belief about *S*. All of these beliefs, and more, about *S* and about each other's beliefs are important in analysing the dialogue.

If Oliver strongly believes *S* to be true, he will believe that he has done all that this particular message can do to bring Rachel to the state in which he wants her. If Rachel believes *S* to be true, she believes she has come to the state Oliver wants for her, since this is the definition of "adequately interpreting the Primal Message".

If either believes *S* to be true, then what remains (in a cooperative dialogue) is for that partner to ensure both that the other believes *S* to be true and that the other believes they believe *S* to be true. In this statement lie the seeds of a potentially infinite recursion of "I believe that you believe that I believe that…" In practice, however, we find that only three levels of recursion are required.

In order for Rachel to have made an adequate interpretation of the Primal Message—for *S* to be true—she must have made some interpretation. Furthermore, the communication mechanism must have been of sufficiently high quality to allow her interpretation to be adequate. We can therefore divide the statement *S* into two parts, which we call *P*1 and *P*2.

P1: The recipient has made some interpretation of the Primal Message.

*P*1 will be true if Oliver is currently acting to change Rachel's state and her state is actually changing as a consequence. In other words, *P*1 says that Oliver is performing

a communicative act—sending a Primal Message for Rachel to interpret—whether or not Rachel knows it, and whether or not Oliver can perceive her consequent state change.

P2: The quality of the communication mechanism is sufficient for an adequate interpretation.

*P*2 is independent of whether Oliver is now trying to change Rachel's state. It says that to the degree that Oliver has sent (or will send) a Primal Message, any interpretation Rachel has made (or will make) is (or will be) adequate for Oliver's purposes. In *P*2, the "quality of the communication mechanism" includes Rachel's interpretive ability, and by implication her related background knowledge.

If Oliver is not at the moment trying to send a message, his level of belief in *P*2 may be based on previous interactions with Rachel that will allow him to design a message suited to her; it may be based on having observed Rachel, or others he perceives as being "like" Rachel, receiving messages from other originators; or it may be based on Oliver's experience with the particular communication medium (a noisy radio link, for example). But if Oliver is currently sending a message, his real-time observations of Rachel allow him to vary his estimate how well the message is being interpreted at that moment, augmenting any considerations based on past experience. His degree of belief about *P*2 can go up and down as he observes Rachel's behaviour. No matter how Oliver's belief in *P*2 changes over time, he must eventually come to believe it strongly if the messages is to be completed satisfactorily. Rachel, of course, can have her own beliefs about *P*2, based on similar considerations, and for her to believe the message to have been satisfactorily communicated, she also must come to have a strong belief in *P*2.

The nature of *P*2 may be exemplified by an extreme example (Figure 2). If Oliver and Rachel are in different houses on different sides of town, Oliver probably does not believe that Rachel would understand him if he spoke. In other words, he does not believe *P*2. Now Oliver picks up the telephone and dials Rachel. Before he starts to speak he believes that Rachel will understand what he is about to say—he now has a strong belief in *P*2, because he has phoned Rachel often in the past, with good results. If, however, Rachel complains about a buzzing noise that makes him hard to understand, he begins to reduce his level of belief in *P*2. He talks louder and more simply, and Rachel says she understands—she believes *P*2. If Oliver trusts Rachel when she says she understands, he also believes *P*2. But then Rachel may say something that leads Oliver to question her understanding, decreasing his belief in *P*2. He tries to improve her understanding, increasing his belief in *P*2 again.

The example shows how Oliver's belief in *P*2 can be high or low before he starts to send a message, and it can rise and fall during the communication. Oliver's belief in *S*, however, is more likely to show general increase, because his belief in *P*1 increases as he sends more and more of his message. Much of the GPG is concerned with variations in the beliefs of the partners about *P*2 and about each other's beliefs relating to *P*2.

*P*1 and *P*2 together form a statement about the adequacy of the communication mechanism, rather more precise than the original statement *S*. *P*1 says that Oliver is trying to bring Rachel to some state, and *P*2 says that the circumstances are such that his message will have its desired effect. If both participants believe both *P*1 and *P*2, then the communication has been successful, though they may not yet know it. For them to

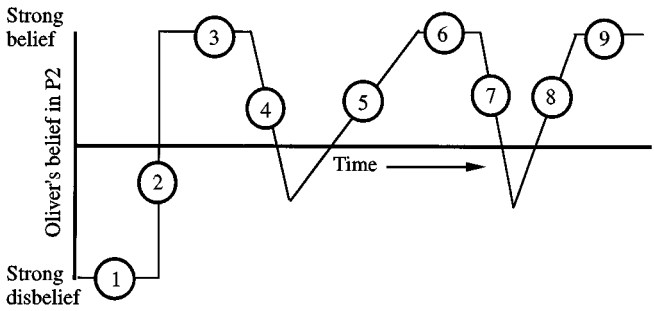

FIGURE 2. Changes over time on Oliver's belief in *P*2, as a consequence of prior experience and current data before and during a telephone call with a noisy line: (1) Oliver at home, Rachel across town; (2) Oliver picks up phone and dials Rachel; (3) Oliver hears Rachel answer the phone; (4) Rachel complains of buzzing noise; (5) Oliver talks loudly and simply; (6) Rachel says she understands; (7) Rachel says something that suggests to Oliver that she misunderstands; (8) Oliver corrects the misunderstanding; (9) message completed.

believe that the communication has been successful, each has to believe that the other believes *P*1 and *P*2. We will argue further that there must be one, and only one, further level of recursion, each having to consider the degree to which they believe the other's beliefs about themselves.

It may not be clear why *S* should be split into *P*1 and *P*2, since for Rachel to interpret the message (i.e. to come to know or to arrive at the state that is Oliver's goal state for her), there must surely *be* a message. Another example may suggest how *P*2 can be true when *P*1 is not:

Oliver: I'm going now.
Rachel: OK.

Rachel's "OK" has not told Oliver what interpretation she has made of Oliver's message. She has told him only "Yes, I believe you have sent a message" (she believes *P*1; in the General Protocol Grammar this is called "Neutral Normal Feedback"). Is her "OK" an adequate indication to Oliver that she has properly interpreted his message? That depends on whether Oliver believed *P*2 before *P*1 became true (i.e. before he sent the message). If Rachel appeared to be listening, Oliver may well have assumed that anything he would say would be understood if Rachel did not question it. Before he spoke, he believed *P*2, and because Rachel does not raise a question, he still believes *P*2; since he how believes *P*1 as well, he now believes *S*.

On the other hand, if Rachel appears engrossed in a newspaper, her "OK" might have been a polite mumble, and Oliver may come back with "Do you understand? I'm going now" He did not strongly believe *P*2 beforehand, and does not now. In asking "Do you understand?" Oliver tries to change Rachel's state of attention to one in which he can believe *P*2, and then restates the message "I'm going now" so that he can believe *P*1 as well, and therefore *S*. To previsage a later discussion, Rachel has been providing what Engel and Haakma (1994) call "E-feedback" that indicates her readiness or unreadiness to accept messages; Oliver tries to change her E-feedback state into showing that she is indeed ready to accept a message.

It is true that, for a message to be adequately interpreted, it has to exist. But it is not necessarily true that both partners believe it exists. *P*1 (A message is being transmitted) is, like *P*2, a fact of the world about which the two partners may have quite different beliefs. When Oliver first says "I'm going now", he believes *P*1. Does Rachel? Oliver does not know, and her mumbled "OK" might lead him to disbelieve it.

It can be true in human communication that the recipient believes *P*1 (and *P*2) when the originator does not. Body language and tone of voice can be the mechanisms whereby a person sends "unintended messages". The originator may not intend to communicate any message at all, but the recipient nevertheless may believe that there is a message (*P*1) and that it has been correctly interpreted (*P*2—slightly modified, since the correctness in that case differs from "adequacy in the view of the originator"). Beliefs in *P*1 and *P*2 can therefore differ between the two partners in a communication, and even in a dialogue, though in a dialogue much of the effort involves the reconciliation of the beliefs of the partners about these two propositions.

For deliberate communication to occur, Oliver must have a goal to believe both *P*1 and *P*2, but he need not have a goal that Rachel believe *P*2 or even *P*1, unless he wants the dialogue to be cooperative. Oliver may, for example, have as part of his goal state for Rachel that she remain unaware that Oliver is responsible for the things she is perceiving. Oliver then has goals both to believe *P*1 and to believe Rachel does not believe *P*1. The classic example of this occurs in Shakespeare's "Othello", when Iago leaves Desdemona's handkerchief for Othello to find, intending that Othello's state change to a belief that Desdemona is unfaithful. If Othello knew that Iago had left the handkerchief, the communication would fail in its purpose, since Othello would then come to believe that Iago wanted him to distrust Desdemona, and would not actually distrust her. Iago would then not be able to perceive Othello as having come to Iago's goal state for him (distrusting Desdemona).

What of Rachel's goals? If Rachel is unaware that Oliver is sending a message, she cannot have a goal to interpret it, let alone a goal that Oliver should be satisfied with her interpretation. If she does not have these goals, she cannot participate in a cooperative or collaborative dialogue, and Oliver can perceive her understanding of his message only by observing her collateral actions, such as smiling, opening a window, distrusting Desdemona, or whatever else may indicate she has come to his desired state.

In a collaborative dialogue, Rachel has a goal to interpret Oliver's message and another that he be satisfied with her interpretation. This latter goal implies that she has a goal to let him know her interpretation or at least her progress in interpreting the message, and Oliver has goals to let Rachel know both that there is a message to be interpreted and the degree to which he thinks she has adequately interpreted it.

Rachel therefore acts deliberately to help Oliver by acting so that she can perceive that his perceptions of the truth of *P*1 and *P*2 track her own changing beliefs about them. This she can do either by changing his beliefs or by changing her own. She may let Oliver know of any failure of understanding the Primal Message, or she may work harder so as to achieve a good understanding of it. Oliver likewise acts so that he can perceive Rachel's beliefs about *P*1 and *P*2 as matching his own, either by changing her beliefs about *P*1 and *P*2 or by changing his own beliefs about them.

1.3.3. Terminating a message transmission

For both *P*1 and *P*2, we have to consider not only the degree to which each of the participants believes it to be true, but also whether they care. If a participant already believes both statements to be true, then how much he or she cares does not matter, because no action is required. But if either *P*1 or *P*2 is disbelieved, then it does matter how strongly the participant cares about its truth. In the sense of PCT, there may be error, but only if the participant experiencing the error cares enough about it will they act to reduce the error. If they care that a statement should be true and it is not, they will act to make it more true, so that they can increase their belief in it.

These considerations lead to a third statement of fact, relating to how much a participant cares whether the recipient makes an interpretation of the message adequate to the purposes of the originator:

> *P3: It is not worth continuing to improve the recipient's interpretation of this message.*

Disbelief in *P*3 means that the participant believes that the dialogue should continue.

A participant in a dialogue may come to believe *P*3 (thus ending the dialogue at that level) for any one of three characteristically different reasons: (1) because *P*1 and *P*2 are both believed (the message has been successfully interpreted); (2) because the purpose of the Primal Message has been served by some other method (e.g. if the message was for Rachel to shut the window and, before Rachel had understood, someone else had done it); or (3) because the process of interpretation seems to have become too tangled and difficult. We may call the first two reasons "good" reasons for believing *P*3, because the purpose of the message has been served, and the final reason a "bad" reason. Whether good or bad, all three result in termination of message transmission at that level.

Reason (2) is worth momentary consideration, because it points up the fact that the Primal Message itself may change during the course of message transmission. The content of a message is not fixed, awaiting only in adequate interpretation by the recipient. Rather, it can vary continuously, depending on how the originator's perceptions of the partner and of the task environment change. The "error" difference between the Primal Message and the originator's perception of the recipient's state vanishes not only because the recipient made an adequate interpretation, but sometimes because the originator's reference for the recipient's state changes to match the recipient's (possibly unchanged) perceived state. Either way, *P*3 comes to be believed, for a "good" reason.

In the companion paper (Farrell *et al*., 1999, this issue), we use as an example a short dialogue in which at one level *P*3 comes to be believed for reason (2). A helicopter pilot has a high-level goal to reset a flight path so as to avoid a potentially dangerous leg on the path originally proposed by the helicopter. He specifies for the helicopter a flight path that would satisfy his goal, but the helicopter proposes a different one that would be a better way to satisfy the same goal. The pilot's original intermediate goals (defined by the path he specifies) are not satisfied. But because his higher level goal is satisfied, the reference values (Primal Messages) at the intermediate level now correspond with the state in which he believes the helicopter to be.

In further analysis and construction of the GPG, disbelief in *P*3 is ordinarily assumed, unless we are specifically dealing with how the dialogue ends.

1.4. LEVELS OF DIALOGUE

In defining the Primal Message and the three basic statements, *P*1, *P*2 and *P*3, we said nothing about the content of the Primal Message—the specific state Oliver wants to see Rachel arrive at *P*1 says only that Oliver wants to see some state in Rachel, and is acting to bring about; it may involve Rachel's overt actions, beliefs that Rachel may hold, or information Oliver is trying to get, among other possibilities. *P*2 says only that the communication mechanism is adequate to bring Rachel to the state Oliver wants to see, whatever that state might be, and *P*3 indicates only whether it is worthwhile for Oliver to continue trying to bring Rachel to the desired state.

The state Oliver is trying to bring about in Rachel may be only a local goal in bringing about in her some more global change of state. In trying to bring Rachel to a state of understanding that she is invited to a specific party, for example, Oliver may bring her to a state of believing that there is to be a party, then to a state of understanding its location, its date, appropriate dress, and so forth. The complete "party" message can be seen as analogous to a function, whose value cannot be completely determined until all its arguments (location, time, nature, purpose, etc.) have been evaluated.

All these minor changes of state—the "function arguments"—are themselves Primal Messages at a lower dialogue level. Each is analogous to a different function that itself has arguments to be evaluated. These lower-level Primal Messages form part of Oliver's major Primal Message, that Rachel should come to a state that will allow her to attend the party if she wishes. In the language of LPT, we say that the lower-level messages "support" the major message.

Oliver may, of course, believe that Rachel already knows some of these "function arguments", or that she will infer them from other sources of information. Since Oliver wants to believe that Rachel knows these items, and he does believe she knows them, Oliver does nothing. In the language of PCT, the error signal in the control system for those perceptions has a value of zero. In the language of LPT, they result in "null" virtual messages from Oliver to Rachel. "Null" virtual messages form an important aspect of the GPG, as we shall see.

The Primal Message bringing Rachel to understand that she is invited to a party may itself support a yet higher-level goal of Oliver's—such as to perceive Rachel as being happy. So one Primal Message relating to Rachel's state (to perceive Rachel as being happy) is supported by a lower-level Primal Message (to perceive Rachel as understanding that she is invited to a party), which in turn is supported by still lower-level Primal Messages (such as to perceive Rachel as knowing where the party is to be held, the appropriate dress for the occasion, etc.), and on down to a level at which the Primal Messages involve the production and recognition of words and gestures, and finally to the physical actions that affect Rachel's sensory systems.

Rachel herself may communicate Primal Messages in support of gaining an interpretation of the major Primal Message. She may raise an eyebrow. She may say things such as

"Where and when is this party?" "Who will be there?" or "What should I wear". These are messages intended to bring Oliver to a state in which he provides the desired information, but they also let Oliver know Rachel's progress toward a full interpretation of Oliver's major Primal Message. Like Oliver's messages about location, date, or dress, they too are "supporting" messages for Oliver's "party" Primal Message. Each can be analysed with its own instance of the GPG, using the same three statements $P1$, $P2$ and $P3$ in reference to its own interpretation.

Supporting messages from Rachel are often called "feedback messages". This usage is technically inaccurate, since "feedback" refers to the perception by Oliver of the consequences of his own acts (or by Rachel of hers), not to the effects on Oliver of Rachel's actions, some of which may be performed for reasons unrelated to Oliver's actions; but the inaccurate meaning is so entrenched that we will continue to use it. And many "feedback messages" do, after all, contribute to the data that allows Oliver to perceive the effect of his own messages.

At every dialogue level, what matters are the participants' beliefs about the following.

($P1$) Whether there is a non-null message at that level.
($P2$) If there is, whether the recipient has adequately interpreted it.
($P3$) Whether it is no longer worth continuing to act to improve the recipient's interpretation.

An entire dialogue is satisfactorily completed when for the top-level Primal Message, all these questions are answered "Yes", for the beliefs held by both partners, for the beliefs each partner holds about the other's beliefs about $P1$, $P2$ and $P3$, and for the beliefs each partner holds about the other's beliefs about his or her beliefs. When this is the case, $P3$ is necessarily answered "Yes" for all the supporting messages, as well. An unsatisfactory completion occurs when $P3$ is answered "Yes" while $P2$ is answered "No". If $P1$ is answered "No", $P3$ becomes moot.

2. The General Protocol Grammar

In this section of the paper, we illustrate the derivation of the GPG through the control by the two parties of their control of their beliefs about the three propositions, and of their beliefs, recursively, about each other's beliefs. We consider in detail only a small part of the GPG, to bring out its main features.

It has been argued that a grammar for dialogue is inherently impossible to construct (e.g. Good, 1989). We agree with Good's position, if by "grammar" he means rules for the entire structure of utterance interchanges between two parties who come together, converse and part. The GPG is not this kind of dialogue grammar. It is both more and less. It is more, in that it describes why the participants do certain things; it is less, because it does not specify exactly what they do to achieve those ends. It is more, in that its elements are not only utterances; they may be entire conceptual structures (at a high dialogue level) or simple phonemes, hand motions, or facial expressions (at a low level); it is less, in that it is concerned only with a well-defined part of any conversation and not with the entire conversation. Neither does the General Protocol "Grammar" have what

Good regards as an essential element—units that have prescribed and identifiable boundaries within the conversation. Instead its "lexical units" or "syntactic classes" are continuous and dynamic.

In any conversation, many GPGs are simultaneously active in each partner at different levels of abstraction. Each Primal Message being communicated at every level is associated with an active instance of the GPG. At least in part, it is the dynamic interaction among many GPGs that leads to the impossibility of constructing a global grammar for entire conservations.

2.1. THE NATURE OF THE GENERAL PROTOCOL GRAMMAR

The GPG was originally portrayed as a state transition network (Taylor 1988*a*, 1989), even though from the first it was clear that this is an inadequate way of describing what actually goes on in a dialogue. A state-transition description relates to the actual GPG much as a political cartoonist's drawing relates to the face of a politician; it emphasizes the salient points, but omits much detail and subtlety. In particular, what it omits from the GPG is the notion that the state can be, and often is, at two or more nodes simultaneously, and those nodes may not even be linked by an explicit transition arc. In this paper, we approach the GPG from the opposite side, considering to be fundamental the changes in beliefs and goals that may occur during a dialogue. From those change, we derive the "cartoon" state transition network (which turns out to be almost identical to the one presented in 1988).

A protocol interchange—a "dialogue"—starts when Oliver, the originator, perceives Rachel to be in some state other than the state he desires, and he cares enough about the differences to act to reduce it. In terms of the three statements, Oliver does not believe $P1$ (that Rachel has already interpreted his message adequately—i.e. that she is in the state he wants for her). He may believe $P2$ or not (that if he were to transmit a message, Rachel would interpret it correctly). He does not believe $P3$ (that it is not worth trying to improve Rachel's interpretation).

We assume that Oliver has reference values for nine beliefs about the progress of the communication. In addition to his reference value that defines the content of the Primal Message, Oliver wants the following.

- 1, 2, 3: to come to believe $P1$, $P2$ and $P3$ respectively to be true.
- 4, 5, 6: to come to believe that Rachel believes them to be true.
- 7, 8, 9: to come to believe that Rachel believes him to believe that they are true.

(The final six beliefs are an assertion that Oliver wants the dialogue to be cooperative). We will assume also that Rachel has a corresponding set of nine beliefs, for which she has reference values that mirror Oliver's.

At this early stage, before Oliver acts on his disbelief about $P1$, what does he believe about Rachel's beliefs? And what may Rachel believe about the three statements about the world? Under what conditions might Oliver try to influence his own belief in $P2$ by improving the communication channel, before trying to modify his belief in $P1$ by sending a message? These questions, and others like them, define the GPG.

2.1.1. Notation

To address these questions, we need a notation more concise than the verbal descriptions we have been using up to this point. We need to describe Oliver's beliefs. Oliver's goals for those beliefs, and Oliver's degree of caring whether those goals are achieved (his "intentions"). And we need to do the same for Rachel. If either has a goal for a belief that differs from the actual belief, and cares enough to do something about it, he or she will act, and that action becomes a plausible element of the GPG. Either may believe a statement about the world with certainty, strongly, or weakly; he or she may be quite uncertain about it, or may disbelieve it weakly, strongly, or with certainty. In this way, degrees of belief and disbelief can be laid out on a monotonic scale from certain belief to certain disbelief.

We have found the scale from $+1$ (certain belief) to -1 (certain belief), with complete uncertainty at 0, to have some useful properties. For example (assuming Oliver trusts Rachel's beliefs completely), if Oliver is certain that Rachel is certain that X is true, Oliver's belief about X can be found by multiplying the two 1.0 values, with the result that Oliver becomes certain about X. However, if Oliver is certain that Rachel disbelieves X with certainty, multiplying the two belief values together gives Oliver a belief of -1 in X, complete disbelief, as is appropriate if he trusts Rachel's knowledge. Similarly, if Oliver is uncertain about Rachel's belief, or if Rachel is uncertain about X, Oliver's belief in X becomes zero, as it should. He neither believes nor disbelieves X, at least not insofar as his belief depends on what he knows Rachel believes about it. Finally, if Oliver only weakly believes (at a level of, say, 0.5) that Rachel weakly disbelieves X, Oliver will disbelieve X, but less strongly than does Rachel.

In the following we use a notation $O(b, P)$ or $R(b, P)$ to mean a number between $+1$ and -1 representing Oliver's (O) or Rachel's (R) degree of belief (b) about a statement about the world (P). For example, $0.3 = R(b, P)$ would mean "Rachel weakly believes P to be true". Similar notations, $O(g, P)$ and $R(i, P)$, are used for Oliver's goals and Rachel's intentions, respectively, that P be true.

The PCT concept of "gain" (crudely, the amplification of the error by the output function) has a corresponding concept in LPT, "insistence". The degree of intention is how strongly the participant "insists" on achieving the goal. To say that Oliver insists, or has a strong intention, that P should come to pass is notated $O(i, P) \approx 1$ (where "$\approx$" means "approximately equals"). If Rachel has a goal that P come to pass, but doesn't care much about it, we can write something like $0.3 = R(i, P)$ to represent the strength of her intention.

A "goal", here, is a reference value for a belief. It does not ordinarily have a numerical value other than 1.0 or -1.0 (a goal that the referenced state should not occur), because it is rarely plausible that someone wants to believe something only weakly, and will act to reduce a strongly held belief in it. In this paper, we seldom are concerned with goals, because we are always assuming cooperative dialogue, in which the goals are all predetermined (except for the content of the Primal Message). For the same reason, we assume that where there is a goal there is also a strong intention, and do not discuss intentions very much.

In practice, seldom if ever can we assign accurate numerical values to beliefs or intentions. Instead of using specific numbers, we find it convenient to label five categories: "Strong Belief" (S), "Weak Belief" (W), "Uncertainty" (U), "Weak Disbelief" (N),

and "Strong Disbelief" (D). So, to say "Oliver strongly believes that Rachel disbelieves X at least weakly" we would write "$S = O(b, N \geqslant R(b, X))$". We use the same category scale, when necessary, to reference intentions. Finally, since we are almost always talking about belief, and not about goal or intention, we drop the "b" in the notation if it can be done unambiguously, so that the preceding expression (Oliver has a strong belief that Rachel disbelieves X at least weakly) would ordinarily be written "$S = O(N \geqslant R(X))$".

2.2. STARTING THE DIALOGUE

To start a dialogue is to send a Primal Message—to begin to act so as to bring the recipient to a desired state. In our notation, $W \geqslant O(b, (P1 + P2)) + S = O(i, (P1 + P2))$ (Oliver no more than weakly believes that Rachel has adequately interpreted the message, and strongly intends that she should). We use "+" to mean that each statement is independently factual, not as an arithmetic summation operator, which would be nonsensical in this context.

It is the *difference* ($O(b, (P1 + P2)) - O(g, (P1 + P2))$) between a belief and the corresponding goal that leads to action; the numeric value of an intention indicates only how strongly the actor intends to reduce the difference. In the language of PCT rather than LPT the difference is an error value, whereas the action (a virtual message across the interface) is the result of the output function having this error value as its input. The strength of the intention corresponds roughly to the gain of the output function. We could, therefore, use the notation $O(e, (P1 + P2))$ to indicate the deviation between the goal state and the current belief Oliver has about $P1$ and $P2$.

Oliver may disbelieve ($P1 + P2$) because he disbelieves $P1$, because he disbelieves $P2$, or both. More important to how he acts initially is whether he perceives *Rachel* as believing $P1$. If he believes Rachel to be unaware that he wants to influence her state, and if he wants her to be aware of it so that she can participate in a collaborative dialogue—or, in symbolic terms, $W < O(b, U \geqslant R(b, P1)) + W \leqslant O(i, S = R(b, P1))$—then he must act to ensure that Rachel becomes aware that a message is beginning.

How does Oliver come to believe that Rachel is or is not aware of the coming message? It has to be by observation of Rachel's acts, which he may interpret, whether she is aware of them or not. Any action (or attitude) by Rachel could constitute a (possibly inadvertent) message from Rachel to Oliver about Rachel's state of belief about $P1$. Haakma and his co-workers (e.g. Engel & Haakma, 1993; Eggen, Westerink & Haakma, 1994; Haakma, 1999, this issue) have called such a message from Rachel "E-feedback". E-feedback helps the potential originator of a message to perceive what kinds of message the potential recipient is ready to receive. E-feedback is not restricted to human–human communication. It can be conveyed by the shape of a row of knobs (on a radio or in a car), by a pattern on a computer screen (such as a menu-bar), or by an explicit prompt from a computer or a person ("Tell me about Nancy"), among many possibilities.

If $S = O(b, U \geqslant R(b, P1))$, then Oliver may have to act to bring $R(b, P1)$ nearer to unity: in other words, Oliver must alert Rachel that a message is coming. This alerting constitutes a whole supporting dialogue at a lower level, for which the Primal Message is that Oliver wants to perceive Rachel to have come to a state suitable for transmission of the main Primal Message. If the recipient is a mechanical device, E-feedback might

consist of a mode indicator such as a power light being on or off, in which case the originator's alerting action becomes mode-setting, such as switching the power on.

Now Oliver has alerted Rachel and perhaps presented part of the Primal Message (begun to try to influence Rachel's state). An omniscient observer would probably be able to determine both that $S = O(b, W \geqslant R(b, P1))$ and that this belief of Oliver's is justified by the factual truth of $W \leqslant R(b, P1)$. Of course, neither the participants nor any external analyst can actually know that their beliefs are justified. They can only perceive what they perceive.

At some point, Rachel's actual beliefs and their relation to her goals and intentions become important enough for her to act. When does this happen? It may not happen at all unless Rachel is being cooperative, meaning that she has an intention to achieve at least a weak belief that she has adequately interpreted Oliver's message. If she is being cooperative, she must be in a state $W \leqslant R(i, P2) + W \geqslant R(b, P2)$. And if she is to act, she must also believe that what Oliver is doing is not leading directly enough toward her goal of $P2$ becoming true. When there is a sufficient discrepancy between her intention and her belief about her present or projected state, she will act to influence Oliver, creating a new ("feedback") Primal Message at a supporting level.

What "feedback" message will Rachel send to Oliver? This depends not only on what Rachel believes about $P1$ and $P2$, but also on what she believes Oliver believes about them and about her beliefs. Because of the E-feedback process, we can assume that both Rachel and Oliver believe $P1$ and believe that the other believes $P1$. Since the communication would be finished if they both believed $P2$ and also that the other believed $P2$, it follows that from Rachel's point of view it is not true that $S = R(P2) + S = R(S = O(P2))$. But it is Rachel's intention that these both become true, so the "feedback" message she sends to Oliver must influence whichever belief does not match that intention.

Suppose, as one extreme possibility, that Rachel strongly believes $P2$ (that she has made an adequate interpretation of Oliver's Primal Message), and that she does not (strongly) believe that Oliver thinks she has. Then she must act so as to indicate to Oliver that she believes $P2$. She tries to bring her own perception of Oliver's belief to its reference state $S = R(S = O(P2))$. We emphasize that this is a state *within herself*—a belief she has about a state within Oliver.

What Rachel does is to provide Oliver with what we call "Normal Feedback", letting him perceive that she believes she has correctly interpreted the message. We call such feedback "Normal" because at low or medium levels of abstraction, it is normal that a message is immediately interpretable. Then the only question is whether the originator realizes that the interpretation has been made, not whether an interpretation known to have been made is accurate. Normal feedback could be implemented as simply as by producing a click sound to indicate that a key press has been detected. In human communication it is often implemented by non-committal phrases such as "Yes" or "Uh-huh", or by gestures such as a slight nodding of the head, but it could be more specific, as we discuss below.

The other extreme possibility is that Rachel strongly disbelieves $P2$. If she believes Oliver also disbelieves $P2$, she may do nothing. But if she believes Oliver to disbelieve $P2$ only weakly, to be uncertain, or worse, to believe $P2$, then she will act to change Oliver's belief. In a collaborative dialogue, Rachel's intention is not only to come to the state of

having adequately interpreted Oliver's message, but also to perceive that Oliver's beliefs about her progress in interpretation continually match her own beliefs of her progress. Symbolically, she has a goal $R(b, P2) = R(b, O(b, P2))$. If $D = R(P2) + W \leqslant R(D \neq O(P2))$, Rachel has to provide what we call "Problem Feedback".

There are intermediate possibilities, in which Rachel neither strongly believes nor strongly disbelieves $P2$. Either way, she will require some further message from Oliver in order to allow her to come to her intended state $S = R(g, P2)$. Either way, she must inform Oliver somehow (if she believes he does not already believe it) that she is not in a position to terminate the interchange. In this situation, too, she must provide "Problem" feedback.

2.2.1. A digression on "turns"

Figure 3 shows a cartoon of the GPG as a state transition network, or a "node and arc" grammar. Each arc represents a lower-level message initiated by one or other partner, by Oliver if the node is a circle labelled "O", by Rachel if the node is a square labelled "R". Figure 3 therefore implies the notion of conversational "turns". However, one of the reasons that Figure 3 is only a cartoon is that the very notion of "turn" is conceptually invalid, as the earlier discussion illustrates, and as we now argue.

According to PCT, a person acts when a controlled perception fails to match a reference (goal) perception and the person cares about that difference. In LPT terms, action happens when an intention fails to match a perception of the state of the partner—or a perception of how the partner's state will evolve if left alone. Such a mismatch can

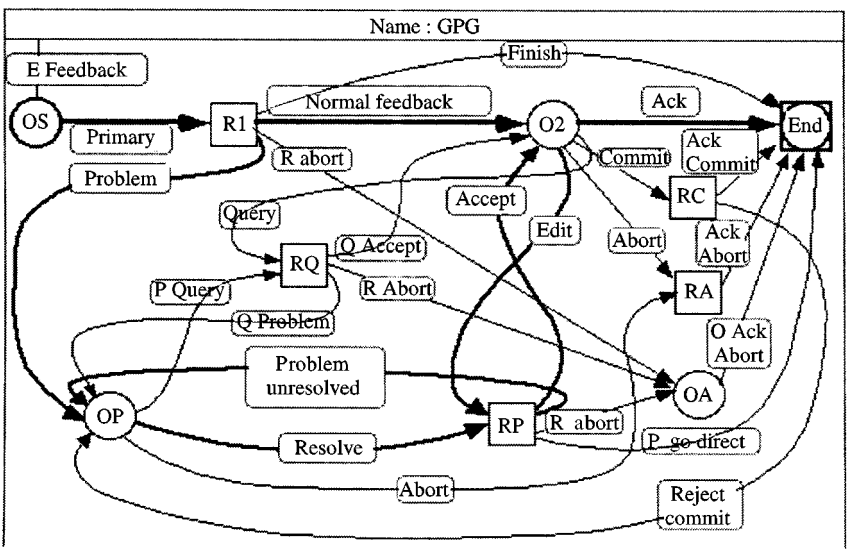

FIGURE 3. The General Protocol Grammar (GPG) shown as a state transition network, with a single arc connecting any pair of nodes. At circular nodes labelled "*Ox*" it is the originator's "turn" to act, and at square nodes labelled "*Rx*" it is the recipient's turn. This view is severely inadequate for accurate description of dialogue, but is useful for design purposes. A more detailed view would include between one and four different instantiations for each of the various arcs. A more correct view would not use a state-transition network at all.

happen at any time, whatever the partner is doing. However, if the partner is acting at any given moment, it is more probable that the perceptions of $P1$ and $P2$ (and therefore $P3$) will be moving in the direction of their goal states than they will if the partner is not acting. If one's partner is acting so as to seem to be moving $P1$ and $P2$ toward their reference states, then one is less likely to act oneself. One is more likely to act when the partner pauses before the $P1$ and $P2$ beliefs have reached their reference states. A switch of actor occurs when one partner ceases to act, while leaving some belief held by the other in a state different from its reference state. Hence, turns appear as an epiphenomenon, not as an intrinsic aspect of dialogue.

If the partner's actions suggest that the partner's understanding is departing from the reference state, the non-acting partner is very likely to interrupt in the middle of a "turn". Far from being a core concept of a dialogue grammar, the notion of "turn" seems to be an artifact of third party dialogue analysis. In most dialogues, one partner talks and then the other talks. This pattern makes "turns" look like a natural and important part of the dialogue structure, and "back-channel" or interrupting talk as an exception to be explained by the analyst. The Perceptual-Control or Layered-Protocol view suggests the opposite to be the case—turns are an epiphenomenon, and there is no intrinsic difference between turn-taking behaviour and back-channel or interrupting behaviour.

Having said that as a warning, we can return to the state-transition network ("turn-taking") view of the GPG shown in Figure 3, and cautiously use it as a guide.

2.2.2. Feedback

To determine how one's partner's state is changing, one must be able to observe it. To observe the results of one's actions on the partner's state is feedback in the engineering sense (as opposed to the "feedback messages" that Rachel sends of her own volition to Oliver). Feedback is an essential requirement of all control. Oliver wants to change some state of Rachel, and anything Rachel does that allows him to see that he has influenced the relevant state is "feedback". Likewise, anything Oliver does that affects Rachel's perception of a state in him that she is trying to change is "feedback" to her.

In cooperative communication, each partner tries to send messages that provide effective feedback for the other. These we call "protocol messages" as opposed to "content messages". "Content messages" are concerned with the nature of the Primal Message, whereas "protocol messages" are concerned with the progress of the dialogue. LPT is not concerned with how the content messages are interpreted, but it is concerned with the progress of their interpretation, and with the nature of the protocol messages that help the participants keep track of that progress.

In Figure 3, each arc represents a type of message. The two "feedback" arcs discussed above ("Normal" and "Problem") both leave node $R1$, but their messages are of different types because they are intended to correct different kinds of perceptual error. Node $R1$ represents a state in which Rachel at least weakly believes $P1$. If we are looking from the viewpoint of Oliver, the state at $R1$ is $W \leqslant O(W \leqslant R(P1))$; if we look from Rachel's viewpoint, the state is $W \leqslant R(P1)$. In order to complete the message transmission, beliefs about $P2$ are involved. If, at node $R1$, Rachel strongly believes $P2$ (the reference state), she attempts use the "Normal Feedback" arc of Figure 1. The error she is correcting is that Oliver may not believe she believes $P2$. In using the "Normal Feedback" arc, she attempts to move to node $O2$, the node at which $W \leqslant R(P1 + P2)$. Otherwise (if her

belief in $P2$ differs from its reference state) she attempts to use the "Problem" arc to move to OP, where $(W \leqslant (R(P1)) + (W \geqslant R(P2))$.

What does it mean to say "she attempts to move?" It means that she acts so as to influence Oliver's beliefs about her state of interpretation of the Primal Message—or rather, to influence her perceptions of Oliver's beliefs about her state. Likewise, when Oliver "attempts to move" within the GPG, he is acting so as to change his own perceptions of Rachel's state of interpretation of the Primal Message. Oliver's overall goal is to bring Rachel to the state implied by the Primal Message, whereas (in a cooperative dialogue) Rachel's goal is to perceive Oliver to be satisfied that she has reached the desired state (i.e. that she has adequately interpreted the Primal Message).

If Oliver believes Rachel to understand the Primal Message when she does not, he may terminate the communication prematurely. Misperceiving the dialogue situation, he may fail to achieve the reference perception (e.g. an opened window) for which he initiated the dialogue. The consequence may be that Oliver wrongly perceives Rachel to be uncooperative at the task level. If, on the other hand, he wrongly believes she misunderstands, his "corrective" actions may influence her to change her already correct understanding, causing ever more confusion. So to "attempt to move" to a node in the state-transition network view of the GPG is partly to act to align the other party's beliefs about your beliefs with the actual state of your beliefs.

If the state transition of Figure 3 is a cartoon of the actual GPG, the copy in one's own mind represents one's beliefs about the actual beliefs represented at the nodes, and one's actions in "following" any arc are attempts to influence the other party to come to a correct view of one's own beliefs. The nodes in the GPG cartoon can be seen as templates, or idealized representations of belief states that are, in practice, usually blurred. The degree to which a participant is "at" a node is the degree to which the actual belief states match the pattern characteristic of that node. Accordingly, either partner could, in principle, be more or less "at" several of the nodes simultaneously, and those nodes could be either "O" nodes or "R" nodes, depending on the extent to which the partner is prepared to act or expects the other to be acting, or both.

There is one way in which the cartoon of Figure 3 is accurate: Each arc represents a message being sent from one partner to the other. The message is a "virtual message" in that it is at the same level of abstraction as the Primal Message it is helping to transmit. Each of these virtual messages acts also as a Primal Message at the supporting level below. At that supporting level, it is transmitted using another GPG, of exactly the same form. All the same considerations apply at each level: whether the originator believes the recipient has interpreted the message, and so forth; but, of course, if the arc at the original level emerges from an "R" node, the roles of originator and recipient are reversed at the supporting level.

2.2.3. Nodes of the GPG

In Figure 3, each node represents a particular pattern of the beliefs held by one or other participant about the three basic statements we have labelled $P1$, $P2$ and $P3$, and about the partner's belief in them. In the beginning, at OS, both parties are likely to disbelieve all three statements and their recursions, if they have any opinion about them at all. Other nodes represent different waypoints on the route to a satisfactory completion of

the message transmission, when each partner has achieved the nine goal states of strong belief in $P1$, $P2$ and $P3$ and their recursions.

For example, at node $R1$, the recipient believes that a message is being transmitted. This is true whether the GPG in question is in the mind of the originator or of the recipient. If Oliver believes Rachel to believe that a message is transmitted ($S = O(W \leqslant R(P1))$), Oliver is "at" node $R1$, anticipating feedback of some kind from Rachel. If Rachel believes the same thing, then she is "at" $R1$, and will act in some way, because she has a goal that Oliver should believe not only that she believes $P1$, but also that his belief about her belief about $P2$ should match her own. If she believes $P2$, she wants Oliver to "move to" $O2$, but if not, she wants Oliver to move to OP. The template for "being at" node $O2$, for example, specifies that: both partners strongly believe $P1$ and each believes the other believes $P1$; the recipient believes $P2$ and the originator believes the recipient believes $P2$ (but nothing is specified about the originator's degree of belief or disbelief about $P2$ itself); and the template says nothing about degrees of belief about $P3$.

Nodes are labelled "O" because the originator (Oliver) is expected to act. The only reason Oliver should be expected to act is that some belief he holds about Rachel may differ from what he wants to believe. If this GPG is in mind of Rachel, her beliefs about what Oliver believes about her determine the range of possible actions she will expect from Oliver, whereas if it is in the mind of Oliver, it is what he believes about her beliefs (including her beliefs about him), as compared to his goals and intentions, that determine what he actually does. Likewise, at nodes labelled "R", what is important is something about Oliver that is believed by Rachel and that affects her actions and his expectations. The full set of correspondences between beliefs and nodes is available (Farrell & Semprie, 1997).

Progressing through the grammar, the values of the different beliefs all tend to increase, though they do not necessarily do so monotonically. The beliefs are beliefs in $P1$ and $P2$, and about the partner's beliefs in them, and about the partner's beliefs about the protagonist's own beliefs about them. For example, at $R1$ only Rachel's belief in $P1$ is defined. It is at least weakly positive, whereas her belief in $P2$ can have any value from -1 to $+1$. At $O2$, her belief in $P2$ is also asserted to be at least weak, but at OP, her belief in $P2$ is at best uncertain. This difference in Rachel's belief about $P2$ is what distinguishes node $O2$ from node OP, whether the GPG be in Rachel's mind or in Oliver's. At $O2$, it is possible for the partners to agree that the message has been successfully transmitted and (now believing $P3$) terminate the dialogue. At OP, mutually satisfactory termination is not possible, because Rachel does not believe $P2$.

At $O2$ (or OP), Oliver's belief in $P2$ is not specified. Even though at $O2$ Rachel may strongly believe $P2$, Oliver may disbelieve it. If so, he will use the "Edit-Accept" loop of the GPG. Oliver has to get Rachel to change her interpretation (temporarily reducing her strength of belief in $P2$ before restoring it in a revised interpretation) in order that he can come to believe $P2$ when she returns to node $O2$. The first time Oliver arrived at $O2$ (after Rachel provided "Normal Feedback"), his belief state might well have been $U \geqslant O(b, P2) + S = O(S = R(b, P2))$. For the message transmission to be considered successful, he wants the "$U \geqslant$" value to increase at least to "$W \geqslant$".

The discrepancy between the value Oliver attributes to $O(b, P2)$ and the value of his perception of $R(b, P2)$ presumably constitutes an error in the PCT sense, since Oliver should be expected to have a reference that Rachel and he should agree about the

adequacy of her interpretation. So Oliver has to affect not only Rachel's interpretation of the message, but also Rachel's perception of the adequacy of her interpretation—two separate problems, each requiring messages to Rachel. These messages may, of course, be multiplexed onto a single utterance or gesture (Taylor & Waugh, 1999, forthcoming), but they are indeed separate messages. What Oliver says or does to get Rachel to move down the "Edit" arc to node RP has to be enough to let her know (a) that he realizes she thinks she has understood the message, (b) that he does not believe she has and (c) in what way she is wrong.

Rachel's belief states can, and normally will, change even if Oliver does nothing, as her mind works on the available data. It is possible, therefore, for Rachel to go from node $O2$ to the "End" node without confirmatory input from Oliver. She could also go from node $O2$ to OP (the node where Rachel thinks she has a problem with her interpretation), if something she thought she understood became questionable on reflection, or from OP to $O2$ if her reflection allowed her to understand something that had been problematic.

For most message transmissions, the important nodes in Figure 3 are $R1$, $O2$, OP, RP and RQ. At RP, Rachel must get Oliver to believe either that a problem has been sorted out (moving to $O2$) or that it persists (moving to OP). At RQ, Oliver does not know what Rachel's interpretation is, and so does not have a strong belief or disbelief in $P2$. He expects Rachel to do something to allow him to move either to $O2$ or to OP, where he would be in a position to move closer to believing or disbelieving in $P2$.

2.2.4. Arcs of the GPG

At any level of a cooperative dialogue a goal of each partner is to complete the transmission of the Primal Message at that level. This occurs for Oliver when he believes Rachel's state to have become what at that moment he wants it to be. For Rachel, it occurs when she believes Oliver to be satisfied that the communication is complete. In our notation, for Oliver, the "End" node is reached when $S = O(P1 + P2)$ (which implies $S = O(P3)$). For Rachel it happens when $S = R((P1 + P2) + S = O(P1 + P2))$. But beyond these simple beliefs, the message communication is really complete only when each partner believes that the other believes it is complete, so a further recursion is also required: $S = O(S = R(S = O(P1 + P2)))$.

In some low-level protocols, particularly in human–computer communication, it is possible to go from the start state $D = O(P1)$ to the "End" state in one move. For example, in a command-line interpreter interface, if the user correctly types the intended command line, it will almost certainly be correctly interpreted by the computer. $S = O(P2)$ is true before the line is typed, and (usually) remains true throughout. All that needs to be done by the user is to change $D = O(P1)$ into $S = O(P1)$, and that is done at the moment the line-ending "Enter" (or equivalent) key is pressed.

In most human–human protocols, matters are not so simple. It may be true that a talker (Oliver) expects that a listener will interpret correctly the words spoken, and to move for each word silently from $R1$ to $O2$ to End (at least at the "word" level of abstraction). But even at this low level, problems can arise. Rachel may question a word poorly spoken or obscured by noise, moving from $R1$ to OP instead of the expected $O2$, thereby slowing the progress through the GPG of the higher-level message to which the word contributes, while increasing the likelihood that the eventual interpretation of the arc at the higher level will be adequate.

"To be adequate" means that the person who "executes" the arc has succeeded in moving the other partner to the intended node. That, in turn, means that the person "executing" the arc has brought the other partner's belief states to what, for the moment, is appropriate.

If Oliver executes the arc, moving from an "*O*" node to an "*R*" node, Rachel's interpretation of the Primal Message is probably brought closer to what Oliver wants. If Rachel executes the arc, moving from an "*R*" node to an "*O*" node, she is attempting to bring Oliver's appreciation of her interpretation closer to what she believes it to be. At *R*1, for example, Rachel believes that Oliver does not know whether she has adequately interpreted the Primal Message. If she herself believes she has not, and she is being cooperative, she therefore has a reference to perceive Oliver as believing that she has not. Accordingly, she acts to influence him to believe that she has not, executing that arc called "Problem" from *R*1 to *OP*.

2.2.5. Types of arc instantiation

If Rachel believes Oliver does not know whether she has adequately interpreted the message what can she do? Suppose first that she believes she may have adequately interpreted the message. There are several possibilities. If, on most occasions at this protocol level, she has been correct when she thinks she has adequately interpreted a message, and she believes Oliver to concur, then all she need do is let him know that she has indeed interpreted the message. Oliver will then believe that the interpretation is adequate. Oliver, for his part, will not expect any more from her than this.

If, on the other hand, Rachel has often been right when she has believed she had an adequate interpretation, but also has often been wrong (or equally if Oliver and Rachel have just met for the first time), then Oliver will probably not find Rachel's belief in *P*2 to be enough to get him to believe *P*2. Rachel must give Oliver some evidence that *P*2 is true—that she has indeed made an interpretation that is adequate for Oliver's purposes. To some degree, she must let him know just what her interpretation is. Perhaps she repeats it to him. Perhaps she paraphrases. Perhaps she performs some action that would be appropriate only if her interpretation is correct. In some way she *informs* Oliver about the content of her interpretation.

So far, in this section we have described three qualitatively different things Rachel can do to get Oliver to move from *R*1 (where he expects Rachel to act) to *O*2 (where he is expected to act). Each of them is an "instantiation" of the "Normal Feedback" arc in Figure 3. We call these three kinds of Normal Feedback respectively "Null", "Neutral" and "Inform", instantiations of the arc.

Rachel may do nothing at all if correct interpretation is automatic on Oliver's completion of the message presentation (i.e. $S = R(S = (O(P1 + P2) + O(S = R(P1 + P2)))))$. We call this possibility "Null" feedback, by analogy with a null arc in a state-transition network. Secondly, if Rachel believes that Oliver believes that any interpretation she has made will be good (i.e. $S = R(S = O(P2)) + W \geqslant R(S = O(P1))$), Rachel may inform Oliver only that she has made an interpretation, without telling him what that interpretation is. This is Neutral feedback. Thirdly, if she is unsure of her interpretation or is unsure that Oliver is sure of it, or accepts the possibility that Oliver is unsure of it (i.e. if $W \geqslant R(P2)$ or $W \geqslant R(S = O(P2))$ or $S \geqslant R(W \geqslant O(P2))$) she can let Oliver know what her interpretation is. This is Inform feedback.

Many of the arcs in the General Protocol Grammar have these same three possibilities for instantiation. Null feedback is appropriate if the context allows the actor to be sure that the partner is at the desired "next" node by virtue of having arrived at this one. This is seldom the case except at $R1$ and $O2$. At $R1$, Rachel is the actor who may not need to do anything overt to bring Oliver to $O2$, and at $O2$ Oliver may not need to do anything to move Rachel to the "End" node. At nodes other than $R1$ and $O2$, circumstances are unlikely to admit a null move to the "next" node.

There is one node in Figure 3 from which the only plausible arc instantiation is usually an "Inform". That is the node labelled RQ. RQ is defined by the condition that Oliver is uncertain about the status of $P2$. He cannot be sure either way whether Rachel has made an adequate interpretation (i.e. $W \geqslant O(P2) \geqslant N$), and must find out from Rachel just what her interpretation is at the moment. Rachel must use an "Inform" instantiation to get to either $O2$ or OP, not because there is no other way for Oliver to determine whether she wants him to be at one or the other, but because it is the most effective way he can determine how to approach the desired state $P1 + P2$.

Clearly it is *possible* for Rachel not to know how to express her interpretation, especially if she has a problem, so a "Neutral" instantiation of an arc out of RQ is possible, especially to OP. Possible, perhaps, but not plausible as a common occurrence, so we do not normally include a "Neutral" instantiation of the arc form RQ to OP in expansions of the GPG diagram.

Figure 3 displays many arcs among its several nodes, but many other possible (but implausible) arcs, are omitted. It is, one must remember, only a cartoon showing the more salient aspects of the reality it is cartooning. Belief states can change for all sorts of reasons within and outside the dialogue itself. Transitions can occur between any two nodes in the diagram, quite apart from the fact that either partner can be "at" more than one node at a time. For example, a direct transition from OP to $O2$ can occur if, before Oliver acts, further thought on the part of Rachel resolves her problem. Oliver may not know that this shift has occurred, but Rachel can inform him using a Neutral ("Oh, I understand now") or an Inform ("Oh, you mean…") instantiation.

We have found about 47 instantiations we consider plausible among 25 arcs shown in Figure 3 connecting the 11 nodes of the GPG as documented in the appendix (from Farrell & Semprie, 1997). For example, the Normal Feedback arc has four plausible instantiations, which we label respectively Null, Neutral, Inform and Correct. Correct indicates that Rachel has reinterpreted the message while noting a communication error, and has informed Oliver about the error. A "Correct" instantiation is therefore a special case of "Inform".

There is a special problem with Neutral instantiations of the Normal Feedback and Problem arcs, since both arcs leave node $R1$. A designer of a computer interface therefore must ensure that if both Neutral instantiations are to be implemented, their implementations differ enough for the Originator to tell whether the Recipient has a Problem. How they are implemented is not an issue within the protocol being analysed or designed, but the designer of an interface must ensure that the lower-level implementations differ adequately. This is true whether the human or the machine is the Originator of Primal Messages using that protocol.

An "Inform" instantiation is defined as containing information about the content of the Primal Message. "Inform" is the most common instantiation of the "Primary" arc,

and is the only useful instantiation of the "Edit" arc (a "Neutral" instantiation of the "Edit" arc is equivalent to the originator saying "You got it wrong", providing no clue as to what might be correct, a possible but not very cooperative instantiation). "Inform" in any arc is most often used when the originator does not strongly believe the interpretation of the message to be correct, or when the recipient is either unsure of the interpretation or does not strongly believe that the originator is sure of what interpretation the recipient has made. The recipient may display an interpretation of the message and ask the originator to verify it.

"Inform" instantiations explain to the other party why the one whose "turn" it is wants to go to the chosen node, and thereby assists the other party to make an appropriate choice for the next move. "Inform" instantiations are most common in the arcs associated with problems and with editing.

Forty-seven instantiations for each protocol may seem a large number for the designer to have to contemplate, but it is of the same order as the number of letters in many alphabets, and far fewer than the number of syllables in any human language. The 47 instantiations are the lexicon for a grammar of dialogue, as required by Good (1989). In most protocols for computer–human interaction, large parts of the GPG will never be used, substantially reducing the lexicon to be implemented.

The designer uses the GPG to model, predict or analyse the interaction at each level of abstraction. The designer must think about each node, each arc, and possibly each instantiation. The designer must ask, for example, whether the device may ever need to query the user's understanding of the display (i.e. whether the RQ node of the GPG will be used when the device is the message originator). Only in a very "intelligent" interface will this be the case, but the designer should always keep it in mind as a possibility.

Each instantiation other than a "Null" involves the overt passage of a message in at least one supporting level of the protocol hierarchy. "Null" implies that the desired belief state of the recipients is the same as the current state, and no message is necessary in order to bring the two states into agreement. If the desired and current states do not agree, the difference becomes the Primal Message for one or more lower-level Protocol Nodes. A supporting lower Protocol Node transmits the message using its own GPG; each non-Null arc instantiation in that lower GPG is a Primal Message for a yet lower node, until the level of physical effects is reached.

2.3. GPG IN INTERFACE DESIGN

The GPG has the same form in every Protocol Node, at every level of abstraction, even though in any particular protocol many of the arcs may never be used. The designs of an interface may well start by considering whether there are entire regions of the GPG that can be ignored.

Does the designer of a protocol think it plausible, in a particular instance, that the recipient of the message could believe the message to be correctly interpreted while the originator believes it is not (that R believes $P1$ and $P2$ while O believes $P1$ and disbelieves $P2$)? If and only if such a situation is plausible and perceivable by O, the designer must allow for an implementation of the Edit arc, and therefore must design a supporting Protocol Node for it. If the message is a keystroke, such a conjunction of the belief states is less plausible than if the message is an architectural drawing or a parts requisition.

The designer of an interface must decide for each protocol which arcs and instantiations will be required to allow the parties to develop the appropriate beliefs, and must ensure that the interface provides for them, just at it must provide for the task-related content of the messages. This is the art of interface design, according to LPT. To specify the nodes, arcs and instantiations that should be supported in the GPG for a given protocol, the designer must imagine what combinations of the partners' belief states are plausible.

In the case of a keystroke protocol, the designer will probably not provide an Edit/Accept or a Problem/Resolve loop. Once a key is depressed, that character will inevitably be incorporated in the ongoing text construction, implying that the GPG has only an "inform" instantiation of the Primary arc, and a "Null" or "Neutral" (key-click) instantiation of the Normal Feedback arc. At the other extreme, in an architectural drawing protocol, the machine may be able to tell for some drawings that there is a problem (the symbol representing a water pipe may not be connected at one end, for example), whereas for other elements of the drawings only the architect may be able to tell whether the drawing is complete. For the first kind of drawing element, the designer should provide the arcs Problem, Resolve and Accept, whereas in a protocol for the second kind of drawing the arcs Edit and Accept, but not Problem and Resolve, should be provided.

Once the designer has determined for a particular protocol which arcs will be required, the requirement for the interface to provide the user with all the necessary instantiations drives the interface's design specifications at the next supporting protocol level. Some arcs specify things the user must enter (messages to the machine), others what the machine must display (messages to the user).

2.3.1. A brief example: communications settings in a Control Display Unit

The companion paper on the LPT uses as its primary example a Control Display Unit (CDU) in a helicopter (Farrell *et al.*, 1999, this issue). Here we consider, as a brief example of the GPG in design, only the process of establishing a communications link using the CDU. The CDU consists of a small alphanumeric display screen (3 by 5 inches with 64 pixels per inch), on each side of which is a column of soft keys whose function can be coded, and under which is an alphanumeric keyboard. It serves as the interface through which the pilot influences various functions including communication.

The pilot can use the CDU to set up radio communication with up to four stations, such as the control towers, ground stations and other aircraft. Each potential station operates on some assigned frequency, and a list of potential stations may be stored in the memory of the CDU. The CDU can set four different radios each to communicate with any one of these stations. The radios can operate in a variety of modes, including "Manual", "Marine" and "Emergency", and in each mode it can be set to operate on either a secure or an open channel. The setting of each radio is therefore determined by a location in a three-dimensional matrix with coordinates defined by station frequency, mode and security. The Primal Message (the pilot's desired state) for a single radio consists of the location of the radio's settings within this three-dimensional matrix.

If the pilot is to create an effective virtual message for the CDU, one that will allow his actual perception to match his reference perception, the changing states of the four radios must be perceptible during the setting procedure (either directly from the CDU display or from the pilot's own memory).

Initially, the pilot might have a mental model of the state of the radio settings; for example, he might earlier have performed a setup procedure that was recorded on tape or floppy disk for transfer to the helicopter. Provided his memory of the set-up is good, his perception of the current state consists of a well-defined point within the matrix of possible settings, and the virtual message intended to bring the state to its desired (Primal Message) condition need be concerned only with those components for which the remembered state differs from the desired state.

But it is more probable that the pilot has no specific information about the current state of a radio, in which case his current perception of the state fills the three-dimensional matrix, each cell having an equal probability of representing the actual state. In this case, the CDU should display the settings so that the pilot may determine which of the elements are at their reference values. The display of the current settings allows the pilot to determine a virtual message (a Primal Message in a supporting protocol) that will result in his coming to perceive the settings as having achieved their reference values. However, for the pilot to make the multiple comparisons required for an efficient virtual message may be more difficult and error-prone than for him to assert in a routine manner the desired settings for each component, ignoring the fact that some may already be at their reference value.

The CDU display of its current settings serves as E-feedback (Engel & Haakma, 1993; Haakma, 1999, this issue), perhaps indicating that the CDU is in a position to accept a message that changes the settings, and if so, what the current settings are. But the E-feedback from the CDU might tell the pilot something different: rather than the radio settings, the display might be showing the current settings of Waypoints or something else the pilot might be concerned with. Even if the E-feedback does show the current settings of the radios, the pilot may ignore their values, and take advantage only of the aspect of the message that indicates the CDU is prepared to accept new settings.

If E-feedback indicates that the CDU is not receptive to radio-setting messages, the designer has to provide the pilot with a way to send a message to the CDU itself, rather than to the radios. This message to the CDU allows it to perform a specific mode-changing action. Such a message might perhaps be implemented (at a lower dialogue level) by the pilot pressing a button labelled "COM".

The designer must consider the status of the various beliefs the partners may hold about *P*1, *P*2 and each other's beliefs. At any point in the dialogue, the designer must consider whether the CDU believes that it has a plausibly complete interpretation of the Primal Message (i.e. could the current settings be the ones the pilot wants)? If so, does it provide the pilot with Null Normal Feedback (i.e. showing no settings, assuming the pilot will be believe that the desired settings have already been asserted), Neutral Normal Feedback (i.e. showing that settings exist but not their values, assuming the pilot will believe the settings to be correct if they exist at all), or Inform Normal Feedback (assuming the pilot will be uncertain as to the values of the current settings)?

The design continues in this way, the designer considering, for each message, the likely belief states of the partners, which is to say their progress through the GPG. The designer tries to ensure that "Inform" instantiations are available when one party might be unsure of the what the other party believes the message interpretation to be, and equally importantly, that "Inform" instantiations are not used when the belief about the message

interpretation is not at issue. An excess of redundant "Inform" instantiations can cause the recipient of those messages to ignore important "Inform" instantiations.

2.3.2. Evaluation of GPG analysis and design

The full GPG analysis is part of the LP-theoretic analysis of the pilot-CDU interaction reported in Farrell and Semprie (1997). The GPG analysis clarified the proposition definitions, aspects of turn taking, nodes having membership rather than the dialogue being at a particular node, and so on. The analysis yielded the pilot–CDU interaction in a GPG and Layered Protocol framework, and some interaction deficiencies of the existing CDU design were readily apparent. An important type of deficiency in the original design of the CDU concerned the feedback at different points in the GPG. Deficiencies included missing arcs, missing forms of feedback and redundant arcs, all of which impeded the completion of the dialogue in a timely manner.

One deficiency was that the pre-existing CDU design could not distinguish at certain levels of abstraction whether a message had been acknowledged or aborted. In the GPG the messages on different arcs must be distinguishable by their recipient, because they are the way in which the recipient learns about the message originator's state of belief. Faulty assumptions about the partner's beliefs may lead to misunderstanding and ultimately impede message transmission.

The GPG framework might even predict where such a deficiency would be common. It would occur if in one protocol node many different arcs might be supported at a lower level by the same supporting protocol. Such deficiencies are likely to arise when interface real-estate is at a premium and the design tends toward multi-function buttons and displays.

3. Conclusions

LPT is a framework theory for human–human or human–machine communication that parallels, and is a special case of, PCT. PCT is based on the principle that "All behaviour is the control of perception". To "control" is to act on one's environment in such a way as to bring ones (presumably accurate) perception of some facet of the environment to a desired (reference) condition and to sustain it there. LPT deals with a specific form of behaviour—communication—and specific kinds of perception—beliefs held by one party about the other. The foundational tenet of LPT thus becomes "All communication is the control of belief". Once communication between two partners is described in terms of a control feedback loop, then knowledge of control theory can be brought to bear on the analysis of communication. Messages are seen as actions performed by one partner so as to affect that partner's beliefs about some state of the other partner. They are not seen as channels for the transmission of information from one mind to another, though such transmission may occur as a side-effect.

A GPG has been developed to express the types of belief changes that normally occur in communication. It describes the changes of belief states of each communicating partner about three basic statements or propositions relevant to the transmission of a generic message, and about the other's beliefs about the three statements. The propositions concern whether the recipient of a message has interpreted it, the reliability of the communication medium, and the value of continuing the effort to complete the

communication. The GPG structure is based on the degrees to which the partners believe the three propositions as factual statements and the degrees to which they perceive that each other believes the three propositions. The different degrees of belief determine the type and form of feedback appropriate at different stages of message transmission.

The intensity of a belief ranges from strong disbelief to strong belief. The degree of belief ranges along a continuum, but for the purpose of distinguishing different communicative actions, a five-point scale—strong or mild disbelief, uncertainty, weak or strong belief—seems to be adequate. Different states in the grammar are defined by which of these categories specifies the nine or 10 controlled perceptions of the partner being analysed.

Most dialogue analysis is based around the concept of a "turn". The GPG does not require this concept. The observation that most conversation seems to progress largely by turns is seen as an epiphenomenon predictable from the nature of control systems, which act when their perception of the variable of interest differs from its desired value and seems not to be moving rapidly enough in the desired direction.

The practical use of the GPG has been illustrated through its application to the design of a proposed interface for a CDU in a helicopter. The original design was shown to be deficient in several respects, largely because information about the CDU's "belief state" was not displayed when the pilot needed it. A revised CDU was designed which addressed the deficiencies uncovered in the analysis.

The GPG is the core of each individual layer of a Layered Protocol design for an interface, or of a Layered Protocol analysis of an existing interface, or of a human–human conversation. LPT is described in the companion paper in the context of the CDU design (Farrell *et al*., 1999, this issue).

References

ENGEL, F. L. & HAAKMA, R. (1993). Expectations and feedback in user-system communication. *International Journal of Man–Machine Studies*, **39,** 427–452.

EGGEN, J. H. HAAKMA, R. & WESTERINK, J. H. D. M. (1994). Layered protocols: hands-on experience. *International Journal of Man–Machine Studies*, **39,** 427–452.

FARRELL, P. S. E., GAMBLE, H., HOLLANDS, J. G. & TAYLOR, M. M. (1999). Perceptual control and layered protocols in interface design: I. Fundamental concepts. *International Journal of Human-Computer Studies*, **50,** 489–520.

FARRELL, P. S. E. & SEMPRIE, M. A. H. (1997). *Layered protocol analysis of a control display unit*. 97-R-70. Defence and Civil Institute of Environmental Medicine, North York, Ontario, Canada.

GOOD, D. A. (1989). The viability of conversational grammars. In M. M. TAYLOR, F. NÉEL & D. G. BOUWHUIS, Eds. *The Structure of Multimodal Dialogue*. Amsterdam: Elsevier.

HAAKMA, R. (1999). What makes novice users do what they do. *International Journal of Human-Computer Studies*, **50,** 557–570.

KNOLLE, N. T., FONG, M. W. & LANG, R. E. (1986). SITMAP: a command and control application. In L. J. PINSON & R. S. WIENER, Eds. *Applications of Object-Oriented Programming*. pp. 28–65. Reading, MA: Addison-Wesley .

NORMAN, D. A. & DRAPER, S. W., Eds. (1986). *User Centered System Design*. Hillsdale, NJ: Erlbaum.

POWERS, W. T. (1973). *Behavior*: *The Control of Perception*. Chicago: Aldine.

POWERS, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, **85,** 417–435.

POWERS, W. T., MCFARLAND, R. L. & CLARK, R. K. (1960). A general feedback theory of human behavior: Part I. Perceptual and motor skills, **11,** 71–88; Part II, Perceptual and Morot Skills, **11,** 309–323.
ROBERTSON, R. & POWERS, W. T. (1990). *Introduction to Modern Psychology*: *The Control-Theory View*. Gravel Switch, Kentucky: the Control Systems Group.
TAINSH, M. A. (1985). Job process charts and man-computer interaction within naval command systems. *Ergonomics*, **28,** 555–565.
TAYLOR, M. M. (1988*a*). Layered protocols for computer–human dialogue. I: principles. *International Journal of Man–Machine Studies*, **28,** 175–218.
TAYLOR, M. M. (1988*b*). Layered protocols for computer–human dialogue. II: some practical issues. *International Journal of Man–Machine Studies*, **28,** 219–257.
TAYLOR, M. M. (1989). Response timing is layered protocols: a cybernetic view of natural dialogue. In M. M. TAYLOR, F. NEÉL & D. G. BOUWHUIS, Eds. *The Structure of Multimodal Dialogue*. Amsterdam: Elsevier.
TAYLOR, M. M. (1993). *Principles for intelligent human-computer interaction*: *a tutorial an layered protocol theory*. DCIEM Report No. 93-32. Department of National Defence, North York, Ontario, Canada.
TAYLOR, M. M. (1994). *On helping*. Unpublished manuscript, 1995.
TAYLOR, M. M. & WAUGH, D. A. (1999*a*). Multiplexing, diviplexing, and the control of multimodal dialogue. In M. M. TAYLOR, F. NEÉL & D. G. BOUWHUIS, Eds. *The Structure of Multimodal Dialogue* II. Amsterdam: John Benjamins (to appear).
TAYLOR, M. M. & WAUGH, D. A. (1999*b*). Dialogue analysis using Layered Protocols. In H. BUNT & W. BLACK, Eds. *Pragmatics of Language Understanding Systems*.

Appendix: Belief states and transitions in the General Protocol Grammar

The reference belief state for cooperative communication is that the recipient of a Primal Message has made an adequate interpretation, and therefore it is no longer worth continuing to transmit the message, and that each party believes the other to hold this belief. "Having made an adequate interpretation" may be divided into the three propositions discussed in the paper:

*P*1: The recipient has made an interpretation of the primal message.
*P*2: The quality of the communication mechanism (or has been) sufficient for an adequate interpretation.
*P*3: It is not worth continuing to improve the recipient's interpretation of this message.

Each partner always has some level of belief for each proposition. We simplify the continuous variation of belief into five classes, which we label as follows:

Strong Disbelief = D (near -1, which represents certain disbelief)
Weak Disbelief = wd (somewhere between D and 0)
No Opinion = U (near 0)
Weak Belief = wb (somewhere between 0 and B)
Strong Belief = B (near 1, which represents certainty)

Also, we use $O(p)$ as a short-hand for the originator's belief in p, and $R(p)$ for the recipient's. For example, $D < O(P2) < 0$ means that the originator's belief that *the*

quality of the communication mechanism is (*or has been*) *sufficient for an adequate interpretation* is somewhere between Strong Disbelief and No Opinion. We also use a form analogous to $O(p) = \{D, wd, 0\}$ to indicate that the originator's belief in p belongs to one of the bracketed categories.

As discussed in the paper, some of the important belief states are about what the partner believes. That is, is it important not only what the originator believes, but also what the originator believes about the recipient's beliefs. The recipient's beliefs, of course, include what the recipient believes about the originator's beliefs. In principle, this recursion could be infinite, but we have found no instance in which a further level is required. Keeping this in mind, the idealized reference belief states for the originator are:

$O(P1) = 1, \quad O(R(P1) = 1) = 1, \quad O(R(O(P1) = 1) = 1) = 1,$

$O(P2) = 1, \quad O(R(P2) = 1) = 1, \quad O(R(O(P2) = 1) = 1) = 1,$

$O(P3) = 1, \quad O(R(P3) = 1) = 1, \quad O(R(O(P3) = 1) = 1) = 1,$

and for the recipient, they are:

$R(P1) = 1, \quad R(O(P1) = 1) = 1, \quad R(O(R(P1) = 1) = 1) = 1,$

$R(P2) = 1, \quad R(O(P2) = 1) = 1, \quad R(O(R(P2) = 1) = 1) = 1,$

$R(P3) = 1, \quad R(O(P3) = 1) = 1, \quad R(O(R(P3) = 1) = 1) = 1.$

Nodes, arcs, and instantiations of arcs in the GPG from the viewpoint of the originator and of the recipient, for collaborative dialogue

In a node-and-arc grammar (a state transition network), a node is occupied or not, and the transition is instantaneous from one node to another. This is not the case with the General Protocol Grammar (GPG) shown in Figure 3. Nodes labelled "*O*" in Figure 3 represent conditions in which the Originator may well act, whereas at those labelled *R*, the Originator is observing the Recipient and modifying his, her or its perception of the various beliefs held by the Recipient. The Originator may, and usually does, observe while acting, and it is therefore completely normal to be simultaneously at an *O* node and an *R* node, and possibly at more than one of each. Transitions from one node to another occur slowly, and it is perhaps more correct to talk about the degree to which each node is occupied, rather than which node is occupied. Nevertheless, each node is characterized by a distinct set of belief conditions, and the trajectory of the changing levels of the nine beliefs can be characterized as a passage through the graph of nodes and arcs of Figure 3.

Initial state

$O(P1) \neq B$: *O* wants to perceive *R* as having interpreted the message (reference $O(P1) = 1$) but does not, and acts to increase the level of belief in *P*1. *R*'s beliefs about *P*1, *P*2 and *P*3 are irrelevant. But if *O*'s action to affect $O(P1)$ is to have the desired effect, *O* must not believe communication to be impossible (i.e. $O(P2) \neq D$). If $O(P2) = D$ initially, then *O* must first act to increase $O(P2)$. This action might be called "Set-up", which might be accomplished in any of a variety of ways, such as by establishing a communication channel by dialing a telephone, by commanding an interface to switch modes, by awakening a partner, or by switching on a device. When the Set-up action has had its effect, *O* can be at the Start node of the GPG.

Start node

$O(P1) \neq B$, $O(P2) \neq D$, $O(R(P1) = \{wb, B\}) < B$: O does not strongly believe R has interpreted the message. Since O at least entertains the possibility that a message can be communicated, O can act to affect R's belief.

Primary Arc to node R1: O has two objectives: to increase $O(P1)$, and to perceive the value of $R(P1)$ so that O can subsequently act to bring $R(P1)$ nearer its reference level. As a by-product of acting toward these objectives, O can observe R's actions and thereby improve O's precision of belief in $P2$ and in $R(P1)$. There are therefore at least two possible instantiations:

Inform instantiation: O acts to increase both $O(P1)$ and $R(P1)$ by communicating at least some of the message content. O is likely to do this if, at the Start node, $O(R(P1) < 0) = \{wb, B\}$ (i.e. O believes R does not believe R has interpreted the message), and if $O(P2) = \{w\mathrm{b}, B\}$ (i.e. O believes the communication channel will be good enough to allow R to make some interpretation).

Neutral instantiation: If O does not strongly believe $P2$, or if O believes R may already have an interpretation (right or wrong) of the message, O may act only to enhance the precision of the belief in $P2$ and of $R(P1)$, in effect requesting R to inform O of R's belief, at least in respect of R's belief that there is a message in progress.

Node R1

$O(R(P1) > 0) = \{wb, B\}$ or $O(R(P2) < B) = \{wb, B\}$: O is observing, not acting. O is "at" node $R1$ if O believes R is likely to act in respect of a perception of the message interpretation or the poor quality of communication. O attempts to perceive whether R believes that the current interpretation is possibly adequate (going to node $O2$) or is likely to be wrong (going to node OP). O leaves node $R1$ only when $O(R(P1) = \{wb, B\}) = \{wb, B\}$ (i.e. when O believes R believes R has made some interpretation).

Normal Feedback Arc to node $O2$ $O(R(P1) = \{wb, B\} + R(P2) = \{wb, B\}) = \{wb, B\}$. Until O perceives that R believes that an interpretation has been made (to the degree that O's Primary message permits), O will be at $R1$, observing. The transition to $O2$ occurs to the degree that O perceives R both to have made an interpretation and to be satisfied that it might be an adequate interpretation. The way R acts to affect O's perceptions (instantiates the Normal Feedback arc) depends on R's perception of O's beliefs, in relation to R's reference values for those perceptions.

Null instantiation $R(P1 + P2) = B + R(O(P1 + P2) = B) = B$. R may make no overt action to confirm to O that an adequate interpretation has been made, if R believes O already believes it. This situation occurs when both parties recognize that communications of this type are almost always correctly interpreted from the primary message alone. This occurs most commonly at low levels of the dialogue, such as individual spoken words in human dialogue, or individual keystrokes on a computer keyboard. There is seldom any need for the recipient to acknowledge each individual word or keystroke (though computers often provide a Neutral Instantiation in the form of a click or a beep that acknowledges the keystroke to have been detected).

Neutral instantiation $R(P1 + P2) = B + R(O(P1) < B + O(P2) = B) = \{wb, B\}$. R believes O is unsure whether the message has been interpreted at all, and that O believes that if it has, the interpretation will have been adequate. R acts so that O can perceive that an interpretation has been made.

Inform instantiation $R(P1) = B + (R(P2) = wb$ or $R(O(P2) < B) = \{wb, B\})$. R believes that the message has been interpreted to the degree that the Primary message permits, but is unsure either as to the adequacy of the interpretation or as to whether O believes the interpretation to be adequate. R acts in such a way that O can perceive something about how R has interpreted the message.

Correction instantiation: A "correction" instantiation of Normal Feedback has the same conditions as the "Inform" instantiation, except that the Recipient believes that the Originator incorrectly formulated some element of the message in a way that the Recipient was able to detect and correct. The Recipient acts so that the Originator can perceive both the Recipient's interpretation and the potentially wrong interpretation that would have occurred had the Recipient not detected the error. The effect of a "correction" instantiation is to increase both parties' future beliefs in $P2$.

Problem Feedback Arc to node OP $O(R(P1) = \{wb, B\} + R(P2) \leqslant 0) = \{wb, B\}$. R wants O to perceive that R believes an interpretation has been made that is unlikely to be adequate.

Null instantiation $R(P2) = D + R(O(R(P2) = D) = B) = B$. R strongly disbelieves that the interpretation is adequate and believes that O believes R believes this. Such a situation is likely to occur only after a neutral instantiation of the Primary message, when both partners know that the message could not have been interpreted adequately, and when R has no basis on which to inform O as to the nature of the likely inadequacy. Except possibly from consideration of timing, a third-party analyst probably would be unable to distinguish the sequence Primary (Neutral), Problem Feedback (Null), Resolve (Inform) from an instantiation of the Primary arc that started out looking like a Neutral and continued as an Inform.

Neutral instantiation $R(P2) = \{wb, D\} + R(O(R(P2) = \{wb, D\}) = B) < B$. R disbelieves that the interpretation is adequate, but is unsure whether O recognizes this. R does not provide O with information as to the nature of the problem, but simply acts so that O can perceive that a problem exists (making it likely that O's next act will be to move to node RQ, where R is requested to inform O as to the current content of the interpretation).

Inform instantiation $R(P2) = \{wb, D\}$. Regardless of O's level of belief about whether O perceives that R has a problem, R may inform O as to the nature of the problematic interpretation. This is the usual case when a problem exists, Null and Neutral instantiations being special cases.

Node O2

$O(R(P1) = \{wb, B\} + R(P2) = \{wb, B\}) = \{wb, B\}$, O believes R believes R to have made a possibly adequate interpretation.

Ack Arc to node End $O(P1 + P2) = B$.

Null instantiation $O(R(O(P1 + P2) = B) = B) = B$: O strongly believes R strongly believes that O strongly believes that the message has been adequately interpreted. Since this state is the reference state (and all the other beliefs have reached their reference states), O need not act. Both parties accept that the message transmission is complete and that the other is also of that opinion.

Neutral instantiation $O(R(O(P1 + P2) = B) = \{wb, B\}) = \{wb, B\}$ but not $O(R(O(P1 + P2) = B) = B) = B$. Either O believes have weakly that R believes O believes the message

has been adequately interpreted, or O believes (strongly or weakly) that R only weakly believes that O believes the message to have been adequatly interpreted.

Inform instantiation: An Inform instantiation of an Ack arc is not possible, since the Ack arc is used only when O is satisfied with the interpretation and no longer will be observing R's changing interpretation.

Commit Arc to node RC $O(P1 + P2) = B$, $O(R(P1 + P2) < B) = \{wb, B\}$ or $O(R(O(P1 + P2) < B) = \{wb, B\} = \{wb, B\}$. The difference between the Commit and the Ack arcs is that the Commit arc leads to a state in which O observes R to determine whether R believes that the message communication is complete, whereas the Ack arc leads to the termination of the communication, where O no longer observes R's behaviour in respect of this message.

Inform instantiation $O(R(P2) < B) = \{wb, B\}$. O believes R to be unsure as to the adequacy of the interpretation and provides more information to assure R of O's belief in its adequacy.

Neutral instantiation $O(R(O(P2) < B) = \{wb, B\}) = \{wb, B\}$. O believes R to be unsure whether O believes the interpretation to be adequate. O settles R's belief, without further information about the content of the message.

Query Arc to node RQ $O(P2) = U$. O believes that R believes that the message may have been adequately interpreted, but is unsure whether this is true, and requests R to provide information as to the content of the interpretation.

Inform and neutral instantiations are both plausible. Which is used depends on the degree to which $O(P2)$ deviates from zero. If O has little or no idea what interpretation R has made, the instantiation will be Neutral, but if O has some notion (though not enough to justify either a weak belief or a weak disbelief in its adequacy), O may inform R about what O perceives the interpretation to be.

Edit Arc to node RP $O(P2) = \{wb, D\}$. O disbelieves that the interpretation is adequate. This could be because O believes that the information so far transmitted in necessary inadequate, or because R's actions have allowed O to perceive that the interpretation is inadequate.

Inform instantiation: This is the only commonly plausible instantiation of the Edit arc, the purpose of which is to bring R's interpretation nearer to being adequate. To do this, O would ordinarily restate or augment some part of the message. A *Neutral* instantiation may occasionally be plausible, particularly if O recognizes the statement of the message to have been ambiguous, and believes that R has arrived at the wrong one of only two plausible interpretations.

Node OP $O(R(P2) = \{wb, D\}) = \{wb, B\}$. O perceives that R believes the interpretation unlikely to be adequate.

P Query Arc to node RQ See **Query Arc** from Node $O2$ to Node RQ. The same considerations apply.

Resolve Arc to node RP O acts to bring R to a point at which R can determine whether it is plausible that a revised interpretation is adequate.

Inform instantiation: The only plausible instantiation is usually *Inform*, since O must act so that R can modify the interpretation. The exception to this is a

Neutral instantiation that says, in effect, "Even though you thought your interpretation was wrong, it is what I meant".

Abort Arc to node RA $O(P3) = B$. O believes the problem represented by being at Node OP is too difficult to be worth the effort to resolve. At Node RA, R has the opportunity to reject the Abort or to accept it.

Neutral instantiation: Ordinarily neither *Inform* nor *Null* is plausible as an instantiation of Abort, because if the problem is not resolvable, there is no point in further presentation of the message content (which would be an Inform instantiation), and the intention of O to Abort is ordinarily not detectable from a Null instantiation. A Null instantiation may sometimes be plausible in conjunction with various actions at higher levels of the dialogue, typically related to uses of the higher level "Edit-Accept" loop to correct the higher-level interpretation in a succession of different ways.

Node RP: Having acted, or being in the process of acting to modify R's interpretation, O is observing R to develop a perception of how R's interpretation is changing.

Accept Arc to node $O2$: The criteria are very like those for Normal Feedback, as arc are the instantiations. The difference is in the likelihood that R will use a Neutral, as opposed to a Null or an Inform instantiation of the arc. At RP, R is likely to believe that O has a more precise perception of R's interpretation than at $R1$, both because of O's initial observation of R at $R1$ and because of subsequent Inform instantiations in the use of either the Edit or the Resolve arc. Hence a Neutral instantiation of the "Accept" arc is relatively more likely than is a Neutral instantiation of the Normal Feedback arc.

Problem Unresolved arc to node OP: The criteria are very like those for the "Problem" arc, though the likelihood of an "Inform" instantiation is higher, since R is likely to have a better idea of wherein the problem of interpretation lies.

R Abort arc to node OA. $R(P3) = B$.

Neutral instantiation: As with the Originator's Abort that signifies an abandonment of the goal to have R act in a desired way, R can abandon the goal of cooperating with O by discovering the desired interpretation. Neutral seems the only plausible instantiation. O can accept R's withdrawal or can attempt to continue the message transmission.

Other Nodes and Arcs: The remaining nodes and arcs have to do with terminating the message transmission without completing an adequate interpretation. The considerations and criteria are the same as for the **Abort** and **R Abort** arcs described above. The parties may disagree as to whether to accept an Abort, in which case a meta-dialogue may ensue, negotiating the disagreement, but that meta-dialogue is independent of the grammar for passing the message itself. In a computer implementation of intelligent dialogue, these other nodes and arcs must be implemented or deliberately ignored by the designer, but it is not necessary to detail them here.