

The Basic Block Model of Extended Explanations

David J. Mooney and Sandra Carberry and Kathleen F. McCoy

Department of Computer and Information Sciences

103 Smith Hall

University of Delaware

Newark, Delaware 19716

Abstract

In this paper we argue that current generation methodologies are inadequate for determining the high-level structure characteristic of naturally-occurring extended explanations. Our analysis of such explanations indicates that high-level structure - composed of a *unifying framework* and its associated *basic blocks* - must be determined by bottom-up processes that attempt to satisfy speaker, listener, and *compositional* goals, after which top-down strategies can be used to organize the material about the selected framework. In addition to a description of this structure, this paper describes three types of repetition whose use is dependent on this high-level structure; their use not only contributes to the cohesiveness of extended explanations, but supports our thesis of the non-recursive nature of the high-level structure. We conclude with an outline of our computational strategy for generating this structure.

INTRODUCTION

In this paper, we describe the high-level structure of naturally-occurring, extended explanations and contend that the organizational strategies currently employed for structuring short explanations are inadequate for generating this characteristic high-level structure. Our analysis suggests that text structure is not completely recursive as others have claimed ([Grosz and Sidner, 1986], [Reichman, 1978], [Polanyi, 1986], [Mann and Thompson, 1988]), but rather that the high-level structure of extended explanations is determined by processes separate from those which organize text at lower levels. Evidence for the non-recursive nature of this structure is provided by repetition, which is frequently used in extended explanations for a wide variety of purposes. This device, whose use has been neglected by current explanation research, exhibits a distinct relationship to the structure of extended explanations, implying that it, too, cannot be generated by recursive strategies.

Our analysis is based on several sources, the primary of which is the testimony before congressional committees regarding the Three Mile Island (TMI) acci-

dent ([Representatives, 1979], [Representatives, 1980]). These explanations were punctuated with numerous interruptions and characteristically involved the operation of complex physical devices and extended causal chains of events. The majority of these explanations were planned beforehand and required many pages of text - typically five, ten, or more, exclusive of interruptions. Furthermore, the testimony comes from a variety of sources: operators who were on the scene as the accident developed; engineers who gave post-accident evaluations; executives from the company that designed the plant. This variety allowed examination of explanations of the same set of events from many perspectives.

A secondary source was transcripts of a medical class involved in patient diagnosis. While these transcripts differed greatly from those of TMI - in particular, the explanations provided were relatively short; they were unplanned; and they were concerned with diagnostic *strategy* as opposed to causal events - they exhibited many of the characteristics found in the TMI explanations.

The next section provides a brief overview of current models for structuring text, followed by a description of the *basic block*, the unit of discourse on which our model of discourse is based. We then identify three types of repetition whose use supports our contentions of the non-recursive nature of basic block structure. We conclude by describing the characteristics of high-level structure of extended explanations, followed by a brief description of our strategy for computationally generating this structure. A more detailed description of our strategy and its role in a complete generational system is contained in [Mooney, *et al.* 1989].

THE BASIC BLOCK MODEL OF EXTENDED DISCOURSE

Given any large body of information to convey, there are many possible combinations of that material, some more cohesive than others. Frequently, deciding *how* to organize a large body of material is more difficult than determining *what* to include. Our research is concerned with the identification of a coherent unifying framework

about which an extended explanation can be organized and the criteria for selecting from among several frameworks when more than one viable alternative exists.

CURRENT MODELS OF TEXT STRUCTURE

COMPUTATIONAL MODELS Two general methodologies have been applied to the structuring of explanations: schemas ([McKeown, 1985], [McCoy, 1985], [Paris, 1987]) and plans based on rhetorical structure theory (RST) ([Hovy and McCoy, 1989], [Hovy, 1988a], [Moore and Paris, 1988a], [Moore and Swartout, 1988b]). Each methodology assumes that text is hierarchically structured and that its operators can be recursively expanded at any level within the discourse hierarchy. While these methods have proven to be effective for organizing *short* pieces of text, we maintain that they are inadequate for generating the characteristic high-level structure of *extended* explanations and certain forms of repetition. Rather, we maintain that this high-level structure can only be captured by a separate, bottom-up process. Once this has been completed, the integration of appropriate repetition can proceed in conjunction with, but independently of, the application of recursive strategies to the organization of text at lower levels.

RHETORIC Rhetoric, the formal study of the art of good writing, provides general strategies for organizing text at a high level that are absent from the computational models. Analysis - "the method of explanation whereby a subject is divided into its separate component parts"¹ - is possibly the most instrumental of these strategies. There are no hard-and-fast rules for determining what constitutes an appropriate analysis of a subject. As [Wicker and Albrecht, 1960] observes, a subject may be classified in as many ways as it has characteristics/parts/stages/etc. However, there are three criteria which experts ([Wicker and Albrecht, 1960], [Arena, 1975], [Thompson, 1957], [Daniel, 1967], [Kane and Peters, 1966]) mutually consider essential for a satisfactory organizational strategy:

1. The scheme should be logical; a single, consistent criterion should be used for the analysis (e.g., time, steps in a process).
2. The scheme should exhaust all of the possibilities; everything to be conveyed should be encompassed by the scheme.
3. The resultant categories should be mutually exclusive; nothing should belong to more than one.

While the type of explanation with which this paper is concerned exhibits a high-level organization reflective of these criteria, the criteria by themselves do not

¹[Arena, 1975] page 107

Now, what happened at Three Mile Island was that a feedwater transient was experienced, and by that I mean, simply, we lost feedwater to the plant momentarily.

Now, with loss of feedwater to the steam generator, the plant will experience a momentary pressurization above its normal pressure. This pressurization is released by a relief valve located at the top of the pressurizer. It simply opens and lets out a little bit of steam to take care of the excess pressure that is < *interruption occurs here* >

Then at 15 seconds into the event-keeping in mind that the valves opened maybe 5 seconds into the event-at 15 seconds the pressure started coming down because the valve had opened and cut off the pressure.

The valve should have reclosed when it got back down to about 2,250 pounds; it did not reclose. The pressure proceeded to come on down. At about 30 seconds into the event, this water here started disappearing, of course, because now you are continuing to remove a very large amount of heat here, which is then coming off as secondary side steam generation, and this water will proceed to disappear if you do not start replacing it.

And the auxiliary feedwater which normally comes on to make sure this does not go dry came on at about 30 seconds into the event. And at least the pumps were running. So this picture here is just the first 30 seconds and-or thereabouts.

And this figure here is indicative of the situation from 1 to 6 minutes into the event.

Figure 1: A Typical Primary Segment

provide the specificity necessary for computational generation. In addition, these guidelines include no suggestions for dealing with situations in which no logical, all-inclusive framework can be identified, nor do they offer suggestions for selecting among several organizational schemes which meet the prescribed criteria equally well. Furthermore, the guidelines are not sufficient in-and-of themselves to account for all of the observed phenomena discussed in the following sections.

BASIC BLOCKS

Our model is based on a discourse unit which we have termed a *basic block*. A basic block consists of two elements:

1. an organizational focus, such as a person or location, and
2. a set of propositions to be conveyed that are related to that focus.²

The focus is what makes a cohesive unit of the material in the block; it is the thread common to all of this material, whether directly or indirectly.

²The term *basic block* is also used to refer to the segment of text generated from these propositions.

A basic block will be realized as a primary segment of text which occupies the first level of the discourse hierarchy. In a coherent discourse, the foci on which the basic blocks are based are themselves related, each representing a different aspect of some *unifying framework*. This framework is analogous to the organizational scheme of rhetorical analysis.

The testimony from which the basic block in Figure 1 was extracted³ illustrates these points. This block references a particular time frame: zero to thirty seconds of the accident at TMI. The remaining blocks of that testimony (not shown) are similarly constructed around time frames, e.g., one to six minutes, six minutes to one hour, etc. Observed frameworks demonstrate a gamut of types: properties of the concepts (location, time), planning strategies in which events are involved (medical diagnosis), and characteristics that are not only inherent in the material but also due in part to the speaker's perception of them (significant factors). There appears to be no limit to what can constitute an acceptable framework, *only that it is derived from the material itself* and not from an independent device solely concerned with text structure. What may be a potential framework for one set of material may be totally inadequate for another. Note that these features are reflective of the guidelines suggested by analysis.

In addition to forming a cohesive unit, basic block structure is explicitly distinguished in the following two ways. First, it is often explicitly marked. In Figure 1, the speaker closes the block with explicit reference to its focus: "*So this picture here is just the first 30 seconds and-or thereabouts.*" The subsequent block in this testimony is also well marked, this time by its initial sentence (which is the last sentence of the figure): "*And this figure here is indicative of the situation from 1 to 6 minutes.*" While the explicit marking of basic blocks in discourse provides persuasive evidence for their distinctive role in the discourse hierarchy, even stronger testimony is supplied by the manner in which certain forms of repetition are used.

REPETITION

Our analysis of naturally-occurring extended explanations has identified three types of repetition whose use reflects a close dependence on the basic block structure of extended explanations: *Final Summaries*, *Transition Summaries*, and *Back-on-Track* repetitions.

FINAL SUMMARIES

The term *Final Summary* is self-descriptive: it appears at the end of a discourse⁴, emphasizing the major points

³Space limitations prevent inclusion of the complete text.

⁴We loosely interpret the term *discourse* to refer to a piece of text that can *stand as an independent discourse* even when extracted from a larger text of which it may be a component, as in the case of a chapter from a technical book.

-
1. They are getting a little out of place here.
 2. So at the end of one hour things look pretty good.
 3. The fact is around this period of time he also got the water up in the steam generator, but he had plenty of heat removal capability.
 4. But I am sure he was getting a little concerned about the reactor coolant pumps because they were - their flow was dropping off, and if you read some of the material produced by the operators and the utility, it indicates they were getting vibration, and they were getting ammeter surging, the surging effect on the ammeter itself.
 5. So there was something strange about pumping this fluid, whatever it was; he did not and has no way to know that this is here, and he has no way to know that this condition has developed over here, no direct way. [Indicating]
 6. He does know that his pressurizer is in good condition, and he is trained, I believe, to think that when the pressurizer is full, the system is in good shape.⁵

⁵From [Representatives, 1979]

Figure 2: Transition Summary: Events

of the text. The final summary serves several purposes:

- It is a means of emphasizing material that an author finds most important.
- It is a vehicle for including information that does not fit within the framework of the main body of text.
- It is a "traditional" closing device.

In the final summaries which we have examined, there is a strong correlation between the points included in the summary and the basic block structure of the text. Although a final summary is not constrained to address each basic block, there is usually a one-to-one relation between the foci of the blocks and the items stressed in a summary. These features are demonstrated by the rather extensive final summary contained in the appendix⁵.

TRANSITION SUMMARIES

Unlike the situation for final summaries, both the conditions calling for the use of a transition summary and the determination of its contents are relatively clear-cut. Transition summaries are similar to final summaries in that they summarize important information, but the similarities end there. While final summaries are used to conclude a discourse, transition summaries are used to signal a shift in circumstances that

⁵The summary has been condensed from the original, which was punctuated with numerous questions. Although the speaker indicates six major points, only five are discussed, the extended questioning preventing him from concluding his summary.

are not implicated by the preceding material and for which the listener is probably unprepared; where a final summary stresses the major points of a complete discourse, a transition summary emphasizes that material instrumental to a listener's understanding of this change in direction and consequently prevents potential confusion. Transition summaries consist of a summary of the material which implicates an expected, but unrealized, continuance of the discourse. Additionally, reference to the material which is responsible for preventing this expected outcome and an indication that everything is not as it seems may also be included. Note that this information, unlike that in a final summary, may play a very minor role in the overall text. Transition summaries also exhibit a relation to basic blocks. While final summaries are found only at the *end* of a text, transition summaries have only been observed at the end of blocks within the *body* of the text.

Figure 2 provides a typical example of a transition summary. Lines 3 and 6 describe the facts, presented in the testimony preceding this summary, that imply to the operator that everything is under control; lines 2 and 6 summarize this condition. However, the facts contained in lines 4 and 5 (also presented in the earlier testimony), indicate that all is not as well as the operator may have thought. In fact, the conditions they describe lead to the next sequence of events, in which the serious problems become evident. Thus, this summary provides a transition from a situation in which things seem to be under control to one in which they are not.

BACK-ON-TRACK REPETITION

Back-on-track repetition is a recovery device, usually brief, which occurs after an interruption. Our transcript analysis suggests that its primary use is to refocus the dialogue on the interrupted segment. Its form is the same regardless of the cause of the interruption, consisting of one of the following:

- a repetition of the last concept presented (or partially presented) before the interruption. This often consists of an (almost) exact reproduction of the original phrase;
- one or two lines summarizing the main point(s) immediately preceding the interruption;
- a sentence that refers generally to the interrupted context.

Figure 3 provides an example of the first type.

Research (e.g., [Reichman, 1978], [Grosz and Sidner, 1986]) indicates that an interrupting segment is explicitly closed before returning to the one interrupted. Although repetition cannot serve as a closing device itself, it frequently accompanies those that do. While the conditions for use of back-on-track repetition are vague, we have observed that back-on-track repetitions very frequently accompany interruptions which occur

At this point the operator was obliged to shut off the reactor coolant pumps, the reason being as follows: we know the fact that there were significant vibrations as determined by vibrational logs, and there were ammeter problems. As the operator observed the ammeter, it was oscillating.

< *interruption* >

OK. *He was obliged to trip the reactor coolant pumps.*^a

^aFrom [Representatives, 1979]. Italics mine.

Figure 3: Back-On-Track: Literal Repeat

within a basic block, but not *between* them. This can be explained by hypothesizing that if a speaker has completed a block, his mental stack is essentially empty (except for the highest level goal); consequently there is nothing to which to return. However, when the block has not been completed, the speaker needs to return attention to the open focus space on top of the stack. This serves not only as a clue to the listener but may also allow the speaker to collect his thoughts.

The relationship between basic blocks and repetition is strong evidence for the distinctive role played by basic blocks in the discourse hierarchy; it also argues the inadequacy of recursive strategies to generate these rhetorical devices. Completely recursive organizational strategies assume that their operators can be expanded at any level within the discourse hierarchy. Consequently, we could expect them to generate transition summaries at the end of *any* segment of text, not just at the *end* of a basic block; we could expect them to generate final summaries at the end of *any* segment which contained several constituent subsegments, not just at the *end* of an entire discourse; and we could expect them to generate back-on-track repetition after *any* interruption, not only those occurring *within* a basic block. However, the above types of repetition exhibit a distinct relationship to the basic block structure and *to no other level of segmentation smaller than the basic block within the hierarchy*. Since the application of a recursive operator is independent of the level of recursion, such organizational strategies cannot guarantee that these types of repetition will be generated only in relation to the basic block structure.

CHARACTERISTIC HIGH-LEVEL STRUCTURE

Given the existence of a high-level structure as evidenced by explicit markings and repetition, we must consider how such a framework is chosen. If the only consideration were the identification of a set of related

foci which can partition the material to be conveyed, then *any* set of related concepts about which the material can be cohesively organized would suffice. Consider the motivation behind the block of Figure 1. On the surface, the events in this block all occurred within the first thirty seconds of the accident. However, it is doubtful whether the driving force behind the construction of this block was to communicate the time frame in which these events occurred; rather, what is of importance is their relative sequence in the total series of events, their cause-effect relations, and their impact on the resulting accident. One could argue that the individual events in this block represent a cause-effect chain, and hence their mutual grouping; but this chain is continued in the next block of the testimony. Apparently, other factors beyond the ability to cohesively juxtapose clauses contribute to the segmentation and the high-level framework about which it is constructed.

In most of the dialogues examined, basic blocks from a given discourse were of approximately the same size. Furthermore, the size of basic blocks was approximately the same from dialogue to dialogue; the "ideal" size of a basic block seems to be three or four paragraphs in length for an explanation of three pages. We contend that the high-level structure of extended explanations reflects the characteristics of an "ideal" framework in which:

- The basic block foci about which the material is organized reflect various aspects of the unifying framework.
- Basic blocks are the same size.
- The length of each basic block approximates the ideal size.

Our basic block model of discourse posits that the attainment of each of these features, termed *compositional goals*, is instrumental in the selection of a high-level framework.

We postulate that a speaker attempts to identify an organizational framework that is capable of coherently expressing all of the material he wishes to convey while satisfying the compositional goals. Of course, a speaker is rarely blessed with the ideal situation. As a result, certain anomalies do occur; e.g., we have observed explanations that exhibit a well-defined high-level structure that accounts for all of the basic blocks except for the final one, which is at best loosely related to the others. In addition, some explanations contain final summaries that include material not found in the text (see **FINAL SUMMARIES**), while others are composed of blocks that are not balanced. The problem appears to be one of finding a satisfactory, rather than an optimal, unifying framework. It seems that the satisfaction of some goals will be sacrificed in favor of others so that a framework that provides the best *overall* solution will be selected.

A MODEL FOR GENERATION

We are currently developing a computational system for generating extended explanations that captures the structural characteristics observed in naturally-occurring, descriptive explanations. Our strategy, motivated by the basic block model of discourse, is based on the hypothesis that the high-level structure of a discourse can be determined by bottom-up processes that attempt to satisfy speaker, listener, and compositional goals. Once this organization has been established, top-down processes (e.g., RST or schemas. See [Hovy, 1988b] or [McKeown, 1985], for example.) can be used to organize the information within the resulting basic blocks and to supplement that information based on the choice of framework. This section describes our general strategy for identifying an organizational framework and the resulting basic blocks given an initial set of concepts to be conveyed.

IDENTIFICATION OF CANDIDATE FRAMEWORKS

We suggest that a speaker, when organizing an extended explanation, will be faced with one of the following general situations:

- He already has a well-defined organizational structure in which the material to be conveyed has already been organized.
- He has a set of goals about which he wants to structure the discourse. In this situation, the goals serve as basic block foci about which he must now attempt to partition his material.
- The speaker is starting "from scratch"; he needs to find a unifying framework about which he can structure what he wants to say. It is this situation which our strategy intends to capture.

Our strategy assumes that the generation process starts with some initial state of affairs and a goal to be achieved. Based on this, a rough set of information to be conveyed will be identified. The selection of this material is not our concern; rather, we are concerned with determining a reasonable high-level structure for this information. We assume that the material to be conveyed has been categorized into at least two levels of importance: that which must be included at all costs (the *kernel*), and that which would be nice to include, but due to time, style, or coherence may be left out. The process of identifying candidate frameworks begins with the kernel.

The implication of the discussion describing the characteristics of basic blocks suggests that a unifying framework will be some feature to which all of the kernel concepts can be related. In our model, we assume that the domain knowledge is maintained as a hierarchical network. Such a representation suggests that a node to which all of these concepts converge could serve

as the unifying framework. The basic block foci, which represent various aspects of this framework, would then correspond to children of this node.

Given a set of propositions to be conveyed, one is randomly selected from the kernel and a trace is performed upwards from each of its arguments through the hierarchy. The traversal of the hierarchy is performed using generalization links, e.g. ISA, ISPART, and SUBGOAL, incrementing counters associated with each node that is traversed. Consider, for example, the two propositions (*agent open1 Fred1*) and (*object open1 valve1*). Using ISA links, we might expect a traversal starting from the action *open1* to pass through the nodes *opening-event*, *general-event*, etc. Furthermore, *open1* may have been a substep of a higher-level plan which would also be encountered while performing a traversal using SUBGOAL links. By the same token, we could expect a traversal starting at *Fred1* which uses ISA links to encounter *person*, *animal*, etc., while one originating from *valve1* might pass through the nodes *valve*, *mechanical-device*, etc. Since *valve1* is also part of *secondary-system1*, a traversal through *secondary-system1* is initiated in the same manner.

Once the traversal has been completed for each element of the kernel, the nodes are ranked according to their counters; this will give some idea of the number of concepts that converge on each. Only those nodes on which a majority of the kernel concepts converge will be further considered as possible candidates for the unifying framework. To continue our example, suppose that, given a kernel of sixty propositions, our traversal resulted in a convergence of fifty of those propositions on the *person* node, ten on *mechanical-device*, fifteen on *event*, and even fewer on the remaining nodes. Given that we are only interested in nodes at which the majority of propositions converge, only *person* will be given further consideration.

In all likelihood, none of the selected candidates will account for all of the kernel concepts. However, it may be possible to link a candidate to concepts not directly associated with it by finding ties between these concepts and ones for which the candidate already accounts, possibly via a property they have in common. With respect to our example, assume that (*object open1 valve1*) does not converge on *person*. However, because (*agent open1 Fred1*) does, and each of these propositions concerns the event *open1*, the first proposition can be effectively incorporated into this framework. Additionally, it may be possible to include such material in a final summary or in a "catchall" block. That speakers do introduce such material in a final summary is evidenced by the sample final summary contained in the appendix, in which all of the "significant factors" addressed in the summary had been discussed in the main body of the text except for the fourth. We claim that inclusion in a final summary is warranted if the speaker intended to include a final summary initially, and if the amount of material

is of the same order as that comprising the individual segments of the summary. The transcripts also supply evidence for speakers' use of catchall basic blocks at the end of a discourse. We claim that for such an option to be viable, there must be sufficient material to construct an adequate basic block and that the material must form a cohesive unit.⁶

SELECTION OF THE UNIFYING FRAMEWORK

Once potential frameworks have been identified, each must be evaluated according to how well it meets the criteria described above:

- How thoroughly does a candidate account for the selected material?
- How uniformly does a candidate distribute the concepts among the resulting basic blocks?
- How closely do the generated blocks conform to the ideal size?

Additionally, a candidate may be evaluated as to how well it meets the needs of a user model.

The actual blocks are constructed around a candidate's immediate descendents, not the candidate itself. So, for example, if a candidate were the node *person*, its children - actual people involved in the events - will become the basic block foci. The ideal case is one in which each child of the candidate accounts for approximately the same number of concepts. The balance each candidate achieves can be determined by comparing the counters of its children. At the same time, the candidate will be rated according to how closely its blocks conform to the ideal size. Lastly, we can rate the candidates on how well they meet the demands, if any, imposed by a user model. For example, if we know the user is familiar with the location in which the events occurred, basic blocks based on location should be given higher ratings; organizing an explanation around a framework with which the listener is familiar will facilitate his assimilation of the information.

Once this process has been accomplished, the candidate with the highest rating will be selected as the best overall framework. At this point, the material has been effectively partitioned about the resulting basic blocks; the material within each individual block can then be organized using a top-down methodology (e.g., RST or schemas), and then realized into text.

CONCLUSIONS

This paper has identified several characteristics of naturally-occurring explanations that cannot be accounted for by recursive text organization strategies. While recursive methods such as RST and schemas are

⁶See [Mooney, *et al.* 1989] for a more complete discussion.

adequate for local organization of text, they are insufficient for the determination of high-level structure, providing neither criteria which describe what constitutes an acceptable framework for a discourse nor any clues as to how this framework should be constructed. We have presented a computational strategy for generating the high-level structure for extended explanations that addresses both issues. Our strategy is based on the *basic block*, a discourse unit consisting of an organizational focus and text structured about that focus. We contend that basic block structure cannot be adequately generated by recursive structuring methods, but must be identified by bottom-up strategies driven by the information to be conveyed. Support for our contentions is provided by repetition, a device frequently used in extended explanations for a variety of purposes.

Our analysis of high-level structure suggests a strategy for generating extended explanations which we are currently implementing as part of a larger system designed to generate extended explanations in an interactive environment. The system is intended as a testbed for identifying viable organizational strategies and recovery techniques due to user interruptions within the basic block paradigm. We believe that our analysis of basic block structure provides criteria necessary for bridging the gap between the generation of short explanations and the successful generation of extended explanations by explicating how the framework of an extended explanation can be constructed and how repetition can be cohesively incorporated into it.

References

- [Arena, 1975] L. A. Arena. *Linguistics and Composition*. Georgetown University Press, 1975.
- [Daniel, 1967] R. W. Daniel. *A Contemporary Rhetoric*. Little, Brown and Company, 1967.
- [Grosz and Sidner, 1986] B. J. Grosz and C. L. Sidner. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics*, 12, 1986.
- [Hovy, 1988a] E. Hovy. Approaches to the Planning of Coherent Text. Fourth International Workshop on Natural Language Generation, 1988.
- [Hovy, 1988b] E. H. Hovy. Planning Coherent Multisentential Text. In: *Twentysixth Annual Meeting of the Association for Computational Linguistics*, ACL, 1988.
- [Hovy and McCoy, 1989] E. Hovy and K. McCoy. Focusing Your RST: A Step Toward Generating Coherent Multisentential Text. In: *Proceedings of the Eleventh Annual Meeting of the Cognitive Science Society*, Cognitive Science Society, 1989.
- [Kane and Peters, 1966] T. S. Kane and L. J. Peters. *A Practical Rhetoric of Expository Prose*. Oxford University Press, 1966.
- [Mann and Thompson, 1988] W. C. Mann and S. A. Thompson. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8, 1988.
- [McCoy, 1985] K. F. McCoy. *Correcting Object-Related Misconceptions*. PhD thesis, University of Pennsylvania, 1985.
- [McKeown, 1985] K. McKeown. *Text Generation*. Cambridge University Press, 1985.
- [Mooney, et al. 1989] D. J. Mooney, S. Carberry, and K. F. McCoy. *The Identification of a Unifying Framework for the Organization of Extended, Interactive Explanations*. Technical Report 90-1, University of Delaware, 1989.
- [Moore and Paris, 1988a] J. Moore and C. Paris. Constructing Coherent Text using Rhetorical Relations. In: *Proceedings of the Tenth Annual Meeting of the Cognitive Science Society*, Cognitive Science Society, 1988.
- [Moore and Swartout, 1988b] J. D. Moore and W. R. Swartout. A Reactive Approach to Explanation. Fourth International Workshop on Natural Language Generation, 1988.
- [Paris, 1987] C. Paris. *The Use of Explicit User Models in Text Generation: Tailoring to a User's Level of Expertise*. PhD thesis, Columbia University, 1987.
- [Polanyi, 1986] L. Polanyi. *The Linguistic Discourse Model: Towards a Formal Theory of Discourse Structure*. Technical Report 6409, BBN, 1986.
- [Reichman, 1978] R. Reichman. Conversational Coherency. *Cognitive Science*, 2, 1978.
- [Representatives, 1979] U.S. House of Representatives, Ninety-sixth Congress: first session. Accident at the Three Mile Island Nuclear Powerplant: Oversight Hearings before a Task Force of the Subcommittee on Energy and the Environment of the Committee on Interior and Insular Affairs. U.S. Government Printing Office, 1979.
- [Representatives, 1980] U.S. House of Representatives, Ninety-sixth Congress: second session. Nuclear Powerplant Safety after TMI: Oversight Hearings before a Task Force of the Subcommittee on Energy and the Environment of the Committee on Interior and Insular Affairs. U.S. Government Printing Office, 1980.
- [Thompson, 1957] W. N. Thompson. *Fundamentals of Communication*. McGraw Hill, 1957.
- [Wicker and Albrecht, 1960] C. V. Wicker and W. P. Albrecht. *The American Technical Writer*. American Book Company, 1960.

APPENDIX

A sample final summary

That is a very quick summary of the sequence of events. Let me turn to what we consider to be the significant factors in the accident. [Slide] We have identified six here. These are the same six factors which the NRC has developed and has discussed in some of their proceedings.

Certainly, the delayed auxiliary feedwater was a significant factor. Auxiliary feedwater block valves should have been open. They were not. They stayed closed for about 8 minutes. That produced behavior in the reactor coolant system which the operators were not familiar with from their previous experience or from training, and inhibited their ability to get the accident under control.

The second significant factor was the pilot-operated pressurizer relief valve not reseating. That was a straight mechanical failure. Probably of greater significance was the lack of recognition of that fact until 2 and one quarter hours into the event.

A third significant item was the securing of the shut-off of the high-pressure injection flow prematurely. As you recall, I indicated that the high-pressure injection came on at approximately 1,600 pounds, as designed. And some of the pumps were shut down in subsequent minutes of the transient. And our calculations indicate that had those pumps been left on and allowed to perform the high-pressure injection function, that there would not have been core damage, there would not have been release of radioactivity at the level that was observed at the Three Mile Island incident.

The fourth significant factor is related to containment isolation. We have talked a little bit about that here. The containment isolation design at Three Mile Island was one in which the isolation occurs at a 4-pounds-per-square-inch reactor building pressure. In other designs there are other events which trigger containment isolation. The significance of this isolation situation is that in the early phases of the accident GPU had an automatic transfer of fluid from the sump of the reactor building into the waste tanks. And in the early phases, when we were blowing the water, the steam, out of the relief valve into the quench tank and then alternately when the rupture disk broke into the bottom of the reactor building, fluid was being automatically transferred out of the reactor building into the waste tanks. The waste tanks overflowed and spilled onto the floor of the auxiliary building, and radiation was released from the auxiliary building. And this occurred in the period prior to the 4-pound, achieving the 4-pound pressure in the reactor building, and the isolation associated with that.

I think this certainly calls for some reassessment. What are the signals on which you isolate the reactor building, and, I think equally important, the question of what lines do you isolate and what lines do you not isolate? Because there are some lines, such as the supply to seal injection and the reactor coolant pumps, you do not want to isolate; you want to leave those open. Other lines clearly you do want to isolate. So, I think it indicates a complete reassessment of the containment isolation philosophy.

The fifth item, identified as the "inappropriate emphasis

on pressurizer level indication," early in the days following the accident there was a good deal of talk about erroneous pressurizer indication. And we have done a fair amount of investigation and satisfied ourselves that the level indication as displayed on the control panel was indicating essentially the amount of water that was in the pressurizer, that there was not an erroneous instrumentation reading, and that in fact he had an indication of how much water was in that pressurizer.

The point here is that he interpreted that, apparently - we are doing a lot of reading between the lines here - but apparently he interpreted that indication as meaning he had a full reactor system when in fact what he had was a system where he had the temperature of the water in the hot leg of the reactor coolant system at the saturation pressure, and flashing and creating steam in the reactor coolant system, even though he had indications of full pressurizer all during the time that he was continuing to blow steam out of the pilot-operated relief valves.

But I think the significant factor here is the inappropriate emphasis that was placed on the pressurizer level alone at a time when he had other indications available to him that should have triggered his concern about what conditions existed in the reactor coolant system.⁷

⁷From [Representatives, 1979]