


AD-A216 463

DTIC FILE COPY

ISI Reprint Series (4)
ISI/RS-89-245
November 1989

University
of Southern
California



Eduard H. Hovy

Approaches to the Planning of Coherent Text

Presented at the 4th International
Workshop on Text Generation,
Catalina Island, CA, July 1988.

To appear in *Natural Language
in Artificial Intelligence and
Computational Linguistics*, 1990.

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DTIC
ELECTE
JAN 3 1990
S B D

INFORMATION
SCIENCES
INSTITUTE



213/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

90 01 03 057

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT This document is approved for public release; distribution is unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ISI/RS-89-245			5. MONITORING ORGANIZATION REPORT NUMBER(S) -----		
6a. NAME OF PERFORMING ORGANIZATION USC/Information Sciences Institute		6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION Office of Naval Research (DARPA)		
6c. ADDRESS (City, State, and ZIP Code) 4676 Admiralty Way Marina del Rey, CA 90292-6695			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, VA 22217		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA AFOSR		8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DARPA (ONR): F49620-87-C-0005 AFOSR: N00014-82-K-0149		
8c. ADDRESS (City, State, and ZIP Code) --over--			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. -----	PROJECT NO. -----	TASK NO. -----
			WORK UNIT ACCESSION NO. -----		
11. TITLE (Include Security Classification) Approaches to the Planning of Coherent Text (Unclassified)					
12. PERSONAL AUTHOR(S) Hovy, Eduard H.					
13a. TYPE OF REPORT Research Report		13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) 1989, November	
15. PAGE COUNT 24					
16. SUPPLEMENTARY NOTATION --over--					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
09	02		coherence, computational linguistics, discourse, natural language generation, natural language processing, Penman, Rhetorical Structure Theory, text planning		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
This paper discusses the planning of multisentential text by computer. In order to construct coherent paragraphs, we have been using relations from Rhetorical Structure Theory (RST) operationalized, as plans. The paper first describes, in some detail, the current method of planning a paragraph using operationalized RST relation/plans. It then makes two points that illustrate why RST relation/plans are the ideal tool for planning paragraphs. First, these relation/plans can be shown to combine the best features of paragraph-sized schemas and clause-sized planning rules under a top-down planning regime in a way which affords much flexibility to the user. Second, RST relation/plans can support both standard top-down planning and open-ended conversation-like behavior; a small difference in treatment gives rise to either paradigm.					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Victor Brown Sheila Coyazo			22b. TELEPHONE (Include Area Code) 213/822-1511		22c. OFFICE SYMBOL

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

(8c continued)

Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Air Force Office of Scientific Research
Bolling Air Force Base, Building 410
Washington, DC 20332



Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

*University
of Southern
California*



Eduard H. Hovy

Approaches to the Planning of Coherent Text

Presented at the 4th International
Workshop on Text Generation,
Catalina Island, CA, July 1988.

To appear in *Natural Language
in Artificial Intelligence and
Computational Linguistics*, 1990.

INFORMATION
SCIENCES
INSTITUTE



213/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

APPROACHES TO THE PLANNING OF COHERENT TEXT

Eduard H. Hovy

Information Sciences Institute
of the University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292-6695
U.S.A.

Telephone: 213-822-1511
ARPA address: HOVY@ISI.EDU

June 1988, revised July 1989

Abstract

This paper discusses the planning of multisentential text by computer. In order to construct coherent paragraphs, we have been using relations from Rhetorical Structure Theory (RST) operationalized as plans. The paper first describes, in some detail, the current method of planning a paragraph using operationalized RST relation/plans. It then makes two points that illustrate why RST relation/plans are an ideal tool for planning paragraphs. First, these relation/plans can be shown to combine the best features of paragraph-sized schemas and clause-sized planning rules under a top-down planning regime in a way which affords much flexibility to the user. Second, RST relation/plans can support both standard top-down planning and open-ended conversation-like behavior; a small difference in treatment gives rise to either paradigm.

Presented at the 4th International Workshop on Text Generation, Catalina Island, CA, July 1988.

This work was supported in part by the Advanced Research Projects Agency monitored by the Office of Naval Research under contract N00014-82-K-0149. It was also supported by AFOSR contract F49620-87-C-0005.

1 Introduction

In the past two years, the Penman project at ISI/USC has been investigating the planning of coherent multisentential paragraphs of text by computer. As our basic planning operators, we use RST relations from Rhetorical Structure Theory (see [Mann & Thompson 83, 87]), which posits that approximately 20 relations suffice to relate adjacent blocks of text (sentences and groups of sentences) in the ways that are coherent in English. We are operationalizing these relations and using them as plans in a top-down hierarchical expansion planning system (see [Hovy 88a]) to develop the dependency structure which underlies each coherent paragraph of text. This work, and an example, are described in the first section of this paper.

In the remainder of the paper, we describe the possibilities that operationalized RST relation/plans afford. The first point is prompted by the observation that two major approaches to planning paragraphs exist at present, but that the relationship between them has not yet been made clear. In one approach, structures called schemas (developed by [McKeown 82]) are used as recipes for building paragraphs. In the other, rules about forming illocutionary acts, surface speech acts, and parts of sentences, are used to plan one- or two-sentence texts that achieve communicative goals (see [Appelt 81, 85]); this theory assumes that longer texts will also be coherent as a side-effect of proper planning. When these two theories are compared along the dimensions of data structure and planning process, it becomes apparent that each approach has a strength and a weakness, and can be made complementary when RST relation/plans are used instead of schemas in an appropriate way. With respect to the data structures involved:

- schemas straightforwardly provide the structure of paragraph-sized spans of text (a strength);
- planning rules provide only far more limited (i.e., clause-sized) structure, which must then be assembled somehow into paragraphs.

And with respect to the planning process,

- schema-based planning resembles script instantiation, affording relatively little dynamic variability;
- planning with rules is amenable to standard planning techniques and affords much flexibility (a strength).

While the two approaches make use of different planning methods and data structures, they are similar enough to suggest the existence of a generalized set of structures and planning method that subsume both. In our investigations, we have found that RST relation/plans can be used as such structures (being easily adaptable to act either as macro-level schemas or as micro-level planning rules), and our method of planning text structure with them can be used as the subsuming planning method.

The second point is the following: When being used to plan out larger spans of structure, RST relation/plans require that their growth points (explained later) be treated as mandatory planning rules. However, when the interpretation of growth points is relaxed to treat them as suggestions, more open-ended planning behavior ensues. Concomitantly, opportunities to apply additional criteria of control (such as focus and text balance) present themselves. The relation/plans' ability to simultaneously support differentially flexible growth points makes them a very useful vehicle with which to conduct text planning experiments.

2 Structuring Paragraphs using RST

2.1 Background to the planning

The text structuring planner we have developed operates antecedent to Penman, the systemic generator being built at ISI [Mann 83], [Mann & Matthiessen 83], [Matthiessen 84]. The structurer¹ accepts inputs from the domain of discourse and rewrites the inputs into a common form (called here input units) which consist of collections of input characteristics. Using RST relation/plans to satisfy the communicative goals specified by the user, it assembles the input units into a tree that expresses the paragraph structure, in which branch nodes are instantiated RST relation/plans and leaves are input units. It then traverses the tree and dispatches its contents to be generated by Penman. An early version of the structurer is described in [Hovy 88a].

The structurer plans out coherent paragraphs to achieve communicative goals posted by the host system to affect the hearer's knowledge in some way. These goals are rhetorical; they pertain to the organization of the paragraph. Each goal is associated with one or more RST relation/plans.

The central data structure is the paragraph structure tree, which is composed primarily of instantiated RST relation/plans. All adjacent independent clauses in the text are linked by some RST relation/plan; adjacent pairs of RST relation/plans are themselves linked by another relation/plan, and so forth, until one relation/plan spans all (though not all relations need be binary, those implemented were made so for convenience). In the planning process, it is assumed that each relation/plan achieves some rhetorical goal, and hence acts as a plan or subplan in the plan tree. The understanding that a relation can be used as a plan (forming an operator of dual nature) is the principal insight that made this method of planning texts possible (see [Hovy 88a]; see also [Moore & Paris 89], Moore and Swartout, and Paris, both in this volume, for work which developed out of that realization).

The structurer's plans are operationalizations of the rhetorical relations of RST. Each relation/plan has two parts, a *nucleus* and a *satellite*, and relates some unit(s) of the input (cast as nucleus) to other unit(s) of the input (cast as satellite) in a unique way; the underlying theory is that when two or more units are successfully related, the resulting juxtaposition will help bring about coherent text. In order to admit only properly formed relations, nuclei and satellites contain requirements that must be matched by characteristics of the input units. (Thus, for example, the PURPOSE relation cannot be used to relate some input state or condition to some input action unless it can be proved (to the structurer's satisfaction, using the PURPOSE requirements) that the state was in fact the purpose of the action.) Thus the nucleus and satellite requirements are semantic in nature.

When studying relations between clauses, it can be seen that additional material is often included, its type depending on the type of relation. For example, when describing a problem and its solution, people often add background material to the problem to explain why it is a problem, and, to the solution, material explaining further results of the solution. For another example, in their study of salience in generation, [Conklin & McDonald 82] found that people always included an elaborating clause when

¹Since the task of text planning includes diverse subtasks, we will use the word *structurer* instead of *planner* for the task of organizing paragraphs of text from the given input.

introducing the main subject in a descriptive paragraph. In order to capture this capability, we built so-called *growth points* into RST relation/plans. Each growth point contains one or more goals to search for and include additional material of a type suitable to the relation. Thus growth points contain rhetorical information.

For this work, we assumed that the information to be generated had already been selected by some process. This assumption conforms to the typical relationship between text generators and host systems such as data base retrieval systems or report generation systems. However, as shown by the work of Moore and Swartout and of Paris (this volume), this assumption is unnecessarily strong, since a planning method of this type can also be used simultaneously to retrieve the information to be said from a collection of well-structured data, given powerful enough relation/plans.

2.2 An example of the planning process

We initially tested the structurer in three domains of discourse (before Moore and Paris wrote their text planner, described in this book). In this paper, we present an example from the third domain, a Naval application in which the structurer and Penman are part of a larger system that presents data base information about U.S. Navy vessels to a user using maps, tables, and text [Arens et al. 88]. Given its simplicity, this domain is ideal for preliminary experimentation. In this section, we describe the planning of the following text (where "C4" denotes a level of operational readiness):

Knox, which is C4, is en route to Sasebo. It is at 18N 79E, heading SSW. It will arrive on 4/24. It will load for 4 days.

The data base consists of lists of assertions about entities and actions that are defined in a property-inheritance network written in the language NIKL [Kaczmarek et al. 86]. The network, which is accessible by Penman and the structurer, contains a full taxonomy of the entities found in the data base Navy world. A typical set of input elements, describing the ship Knox, appears in Figure 1. From this input the structurer builds six distinct units, each of which eventually becomes a clause, using a number of domain-specific rules; for example, the unit **ARRIVE11400**, representing the arrival of the Knox at Sasebo, is built according to a domain rule that an arrival occurs when a moving ship employment is followed by a stationary employment. Unit **ARRIVE11400** is then linked to unit **E105** by the addition of the term (**NEXT-ACTION.R E105 ARRIVE11400**) to the aspects of **E105**, this being the domain's way to represent temporal succession between events. The input units are shown in Figure 2.

Next follows the structure planning task. The structurer seeks to satisfy the following goal, posted by the host system:

(BMB SPEAKER HEARER (POSITION-OF E105 ?NEXT))

which can be glossed as: from the input units, tell the hearer the sequence of events of which **E105** is a principal part (i.e., its temporal **POSITION**). More precisely, the goal is to be read as: achieve the state in

Figure 1: Unstructured data from the Navy domain.

((SHIP.EMPLOYMENT A105)	
(SHIP.R A105 KNOX)	(ENROUTE E105)
(SHIP.COURSE.R A105 195)	(EBEG.R E105 870420)
(CURRENT.POSITION.R A105 P102)	(EEND.R E105 870424)
(POSITION P102)	(DESTINATION.R E105 SASEBO)
(LONGITUDE.R P102 79)	(LOAD E107)
(LATITUDE.R P102 18)	(EBEG.R E107 870425)
(READINESS.LEVEL.R A105 C4)	(EEND.R E107 870428))
(NEXT.MAJOR.EMPLOYMENT.R A105 E107)	
(CURRENT.MAJOR.EMPLOYMENT.R A105 E105)	

Figure 2: Inputs to the structurizer.

((ENROUTE E105)	((ARRIVE ARRIVE11400)
(SHIP.R E105 KNOX)	(SHIP.R ARRIVE11400 KNOX)
(DESTINATION.R E105 SASEBO)	(TIME.R ARRIVE11400 870424)
(HEADING.R E105 HEADING11416)	(NEXT-ACTION.R ARRIVE11400 E107))
(READINESS.R E105 READINESS11408)	
(NEXT-ACTION.R E105 ARRIVE11400))	((POSITION POSITION11410)
	(SHIP.R POSITION11410 KNOX)
((HEADING HEADING11416)	(LONGITUDE.R POSITION11410 79)
(SHIP.R HEADING11416 KNOX)	(LATITUDE.R POSITION11410 18))
(SHIP.COURSE.R HEADING11416 195)	
(POSITION.R HEADING11416 POSITION11410))	((LOAD E107)
	(SHIP.R E107 KNOX)
((READINESS READINESS11408)	(EBEG.R E107 870425)
(SHIP.R READINESS11408 KNOX)	(EEND.R E107 870428))
(READINESS.LEVEL.R READINESS11408 C4))	

which the hearer believes that it is the intention of the speaker that they mutually believe that the event E105 is followed by some other event².

The structurer starts with this goal, which matches the *results* field of only one relation/plan, namely SEQUENCE (see Figure 3). In the match, ?PART is bound to E105. With this binding, the structurer begins searching for an appropriate nucleus. First it checks for an input unit whose contents match the combined nucleus and satellite requirements

((BMB SPEAKER HEARER (NEXT-ACTION.R E105 ?NEXT)))

The unit E105 can satisfy this requirement since its characteristics match when ?NEXT is bound to ARRIVE11400. With this binding for ?NEXT throughout the relation, the independent nucleus and satellite requirements become, respectively,

((BMB SPEAKER HEARER (TOPIC E105)))

and

((BMB SPEAKER HEARER (TOPIC ARRIVE11400)))

which can directly be fulfilled by simply saying E105 as the nucleus and ARRIVE11400 as the satellite of the SEQUENCE relation/plan (the TOPIC terms are simply a programming convenience for binding variables). In this way the structurer relates E105 to ARRIVE11400, ensuring that the hearer will understand their sequentiality.

Notice that the growth points for both nucleus and satellite remain unfulfilled (had any of the input units' features matched them, they would have been satisfied and removed). Therefore, the growth points are posted on the structurer's agenda as extant goals to be achieved.

We continue this example one step further. The next goal on the agenda is the growth-point-turned-goal

((BMB SPEAKER HEARER (CIRCUMSTANCE-OF E105 ?CIRC)))

which matches the results field of only one relation/plan, CIRCUMSTANCE, given in Figure 4. (As can be seen from the nucleus requirements, this plan is tailored to the Navy domain. In a more sophisticated scheme, the terms for HEADING and TIME could be replaced by a single generalization such as SPATIO-TEMPORAL-LOCATION, and the matching process adapted accordingly.) After binding ?X to E105, the first requirement within the OR term matches the unit E105, since it contains the proposition

((BMB SPEAKER HEARER (HEADING.R E105 HEADING11416)))

This match permits the structurer to form a new instantiation of the CIRCUMSTANCE relation to relate E105 to HEADING11416. The new relation fulfills one of the growth-points-turned-goals of the original SEQUENCE nucleus, and is therefore added into the paragraph tree as follows:

²This syntax, and the term BMB, are from the modal operator language developed by Cohen and Levesque, with which they tried to derive complex speech acts from a small set of primitive operators; see [Cohen & Levesque 85]. Though they have subsequently retracted the conclusions of that paper, and hold instead that one requires also a notion of commitment when performing such derivations, the utility of this notation remains unaffected for our descriptive domains.

Figure 3: The RST relation/plan SEQUENCE

Name: SEQUENCE

Results:

((BMB SPEAKER HEARER (POSITION-OF ?PART ?NEXT)))

Nucleus+Satellite requirements/subgoals:

((BMB SPEAKER HEARER (NEXT-ACTION.R ?PART ?NEXT)))

Nucleus requirements/subgoals:

((BMB SPEAKER HEARER (TOPIC ?PART)))

Nucleus growth points:

((BMB SPEAKER HEARER (CIRCUMSTANCE-OF ?PART ?CIR))

(BMB SPEAKER HEARER (ATTRIBUTE-OF ?PART ?VAL))

(BMB SPEAKER HEARER (PURPOSE-OF ?PART ?PURP)))

Satellite requirements/subgoals:

((BMB SPEAKER HEARER (TOPIC ?NEXT)))

Satellite growth points:

((BMB SPEAKER HEARER (ATTRIBUTE-OF ?NEXT ?VAL))

(BMB SPEAKER HEARER (DETAILS-OF ?NEXT ?DETS))

(BMB SPEAKER HEARER (POSITION-OF ?NEXT ?FOLL)))

Order: (NUCLEUS SATELLITE)

Relation-phrases: (" "then" "next")

Activation-question:

"Could 'A be presented as start-point, mid-point, or
end-point of some succession of items along some
dimension? -- that is, should the hearer know that
'A is part of a sequence?"

Figure 4: The RST relation/plan CIRCUMSTANCE

Name: CIRCUMSTANCE

Results:

((BMB SPEAKER HEARER (CIRCUMSTANCE-OF ?X ?CC)))

Nucleus+Satellite requirements/subgoals:

((OR (BMB SPEAKER HEARER (HEADING.R ?X ?CIRC))
(BMB SPEAKER HEARER (TIME.R ?X ?CIRC))))

Nucleus requirements/subgoals:

((BMB SPEAKER HEARER (TOPIC ?X)))

Nucleus growth points:

((BMB SPEAKER HEARER (ATTRIBUTE-OF ?X ?ATT)))

Satellite requirements/subgoals:

((BMB SPEAKER HEARER (TOPIC ?CIRC)))

Satellite growth points:


((BMB SPEAKER HEARER (ATTRIBUTE-OF ?CIRC ?VAL)))

Order: (NUCLEUS SATELLITE)

Relation-phrases: ("")

Activation-question:

"Could 'A be presented within some spatial, temporal,
or situational framework? -- that is, should the hearer
know to interpret 'A in some appropriate context?"




The nucleus E105 is moved to become the nucleus of the newly-formed relation CIRCUMSTANCE, which replaces it in the paragraph tree. The remaining two unfulfilled nucleus growth points of the SEQUENCE relation are propagated along with it, and are joined by the unfulfilled nucleus growth point of the CIRCUMSTANCE relation. Together, these points potentially give rise to tree growth at the new position of E105. Similarly, the remaining growth point of the satellite HEADING11416 is added to the agenda.

Planning proceeds as follows: one of the SEQUENCE nucleus growth points is fulfilled by an ELABORATION-ATTRIBUTE relation between E105 and READINESS11400, the unit that represents the operational status of the Knox. The remaining satellite growth point of the CIRCUMSTANCE gives rise to an ELABORATION-ATTRIBUTE relation between HEADING11416 and POSITION11410, which represents the ship's latitude and longitude. Finally, the growth point term POSITION-OF in the ARRIVE11400 satellite is fulfilled by a SEQUENCE relating that node to E107. Though at the end some growth points remain unsatisfied, all the input units have been related, and the structuring process stops with the paragraph structure in Figure 5.

As is clear, the planning cycle is the following: a new growth-point-turned-goal is taken from the agenda; zero or more relations are found to fulfill it; these relations (if any) furnish requirements for nucleus and satellite fillers; if any input units match, the relations are instantiated with the units and added to the tree. Unfulfilled growth points are added to the agenda of goals. When more than one relation/plan can be added, new trees are formed, identical except for the new relations, and the structurer proceeds to plan out all alternatives. (Not committing to any particular tree growth during planning is an experimentation strategy; since we have no hearer model or stylistic criteria on which to base a preference for one relation over another, and since the issue is clearly related to the hearer and the desired style of the text, we have chosen to plan out all options rather than commit to an option arbitrarily.)

When all input units have been used, or when no goals are left on the agenda, planning ceases. The structurer applies a simple evaluation metric to select the most comprehensive tree, if more than one have been built, as follows: the trees containing the most input units are collected, and of those the tree(s) with the fewest remaining unsatisfied growth point goals; in case of a resulting tie, one tree is selected at random.

Figure 5: Paragraph structure and corresponding Navy text



Knox, which is C4, is en route to Sasebo. It is at 18N 79E,
heading SSW. It will arrive on 4/24. It will load for 4 days.

3 Relation/plans, Schemas, and Clause-Level Rules

3.1 Relation/plans and Schemas

Paragraphs of text can be planned and generated in various ways. The primary method of planning paragraphs by computer (other than the method described in this paper) makes use of data structures called schemas. Schemas capture the stereotypical arrangement of clauses in common expository texts such as descriptions of objects. First developed by [McKeown 82] and then extended by the inclusion of the sophistication of the hearer's knowledge as decision criterion [Paris 87], schemas are an extremely useful method of controlling the collection and structuring of material from collections of data into lists of clause-sized units that each support the production of a sentence.

Schemas have the advantage of being easy to build and use. For an application, a schema is defined for each type of paragraph to be generated. For each clause typically appearing in the paragraph, a predicate is built into the schema that represents the type of information it contains, as well as additional information such as the number times it can occur in the paragraph. To use a schema, a schema traversal mechanism similar to a script instantiator evaluates the conditions of use on the predicates (taking into account focus, etc.), finds the appropriate material in the data base, and records what is eventually to be realized by the realization component.

One limitation on their use is the fact that schemas do not contain a description of the functional/rhetorical role that each part of the paragraph plays with respect to the whole paragraph. This makes them unsuitable for systems requiring the ability to adapt dynamically (such as interactive systems that can be called upon to replan any particular piece of the paragraph at any time; see Paris, this volume). A second limitation is the inherent rigidity of schemas; however many optional and repeating predicates they contain, they still always mandate a paragraph of text.

The use of RST relation/plans instead of schemas overcomes some of these limitations: the paragraph trees contain a derivation history that shows the rhetorical relation between every two parts of the text, and they can control smaller spans of text than schemas — as small as two clauses in the limit. However, these advantages come at the cost of increased work, since it is more difficult to assemble an RST paragraph tree from a set of independent relations than it is to instantiate and traverse a schema, particularly under Mann and Thompson's original conception of RST, in which each relation has to be tested at every cycle of the growth of the paragraph for possible applicability (see [Mann & Thompson 87]). Clearly, a melding of the two techniques that preserves their best aspects — the selectivity of schemas, the dependencies of relation/plans — would be very useful.

By developing the notion of growth points, we have managed to achieve a suitable melding. Rather than allowing all possible RST relations to enter at each cycle of the planning process, we allow growth of only those relations whose goals appear in the growth points fields, and we build in the growth points fields only those goals that support the texts commonly found in the domain of discourse. This is, of course, exactly the same approach taken in the development of schemas.

By the introduction of growth points we establish a functional equivalence between relation/plans and schemas. To see this, consider the following argument: in the SEQUENCE relation/plan of Figure 3,

replace the CIRCUMSTANCE-OF growth point goal by the nucleus and satellite growth point goals of the CIRCUMSTANCE relation, in order of appearance (see Figure 4), and add to the ends of the nucleus and satellite requirements an additional set of requirements for a second set of clauses, namely those from the CIRCUMSTANCE relation/plan's nucleus and satellite requirements respectively. Perform the same operation for each of the other nucleus and satellite growth point goals recursively. Though it may never end, this rewriting will produce a (possibly infinitely recursive) plan that prescribes saying the nucleus, then the bottommost satellite that fulfills the requirements, then the next bottom-most, etc., via the CIRCUMSTANCE satellite, up to the initial SEQUENCE satellite, followed next by its associated relations. That is, one would compile a (possible infinitely recursive) single recipe for constructing the paragraph tree shown earlier. This recipe would contain, of course, the individual relations (such as CIRCUMSTANCE), as well as the rhetorical roles they play in the paragraph. Other than that, however, the recipe would resemble a (possible infinitely recursive) schema built for the purpose of producing paragraphs describing ships' employments, when given a finite collection of facts to be said.

Viewed this way, it is clear that a relation/plan without any growth point goals is simply a minimal two-place (if it has one nucleus and one satellite) schema. A relation/plan with growth point goals is a kind of generalized schema that builds a paragraph to achieve the communicative function of the relation/plan. Relation/plans can be combined into larger plans (see Moore and Swartout and Paris, both papers in this volume), and even into schemas (see [Mann 88]). Said differently, each RST relation/plan is simultaneously a basic rhetorical operator which can be incorporated into a schema as well as a generalized schema for building a specific type of paragraph.

3.2 Desirable Characteristics of Relation/plans

RST relation/plans exhibit a number of desirable characteristics. First, growth point goals for expanding the tree reside in the relation/plans declaratively. No additional inclusion mechanism is required to suggest directions of growth. The simple procedure is: if the plan's growth points call for something, and it can be found in the material to be included, then say it.

Second, being stated declaratively, the growth point goals are ordered. No additional ordering criteria are required; and the procedure is: simply say things in the order given by growth points in the plan.

A third characteristic is the simplicity of this structure planning method. By representing the growth points declaratively in plans, the behavior of the structure planner, the kinds of trees it can build, and the ways in which relation/plans can be altered are perspicacious and easy to change. Making growth points explicit and separate contrasts with the approach of Moore and Paris (both in this volume), in which growth point injunctions are all incorporated into the satellite fields of the plans, with the result that the rhetorical relationship between the nucleus and its particular satellite is obscured. They overcome this problem to some degree by making their plans rather small (limiting the number of satellite entries per plan).

Another desirable characteristic pertains to the use of relation/plans. In planning the structure of a coherent paragraph, the planner constructs a tree that automatically records the rhetorical role played

by each clause with respect to whole text. These roles can later be used to identify sections of the text that have been unsuccessful and can aid replanning or elaboration.

A fifth feature supports a type of planning that is still beyond our grasp. This involves representing growth point goals in a similar way to requirements. A structure planner more powerful than ones we can build today could try to make true any unsatisfied semantic requirements in the nucleus and satellite requirements fields in exactly the way it tries to achieve the goals in the growth point fields. It may thus be able to include material that violates nucleus and satellite requirements (that is, to "lie", or at least, to present material in a way not licensed by its features!) as long as the remaining unfulfilled requirements can be satisfied by the inclusion of additional material, in exactly the same way as growth point goals are fulfilled by new material. This approach was, in fact, part of the initial attempt at developing a text structure planner; in [Hovy 88a] no explicit growth point fields occur. The nucleus and satellite requirements ensured the relational compatibility of the nucleus and satellite entities while also containing what are now the growth point goals as ancillary "requirements" to produce an adequate paragraph.

The sixth characteristic, one of the central points of this paper, is the fact that the planning paradigm using relation/plans to plan multisentence paragraphs as described is exactly the same type of planning paradigm that has been followed in some approaches to plan out the contents of single sentences. This means that a single planning process can use, at the paragraph level, appropriate RST relation/plans, and, at the sentence level, clause-sized planning rules, in a homogeneous fashion. This point is discussed in the next section.

3.3 Relation/plans and Clause-Level Rules

From the macro-sized planning of paragraphs, we turn next to the micro-sized planning of parts of individual sentences. Some of the most thorough work in this regard has been done by Appelt [Appelt 81, 85], whose planner-generator KAMP achieved communicative goals by planning, via illocutionary acts and surface speech acts, so-called utterance acts which, being straightforwardly associated with syntactic knowledge, were used to build sentences.

In developing KAMP, Appelt's primary concern was with assembling the information that had to be built into a sentence in order to provably achieve the planner's communicative goals. The goals never called for more information than could be accommodated in one or two sentences, with the result that the issue of coherence did not arise. Thus KAMP's extensibility to paragraph-sized text was never illustrated in practice.

In order to extend Appelt's work into the multisentence range, the issue of coherence must be addressed. It is not sufficient simply to prove that the information provided in the clauses of the text individually suffice to support the desired informing result; it is also necessary to ensure that the development of the material and its interclausal linking supports the derivation of the desired interpretation.

The RST relation/plans and structure planner perform part of that function. Since both the structurer and KAMP are simplified versions of NOAH [Sacerdoti 77], the RST relation/plans can be used together

with KAMP's intraclause-sized planning rules (so-called *concept activation actions*) in an appropriately general homogeneous system to plan paragraphs of text. Operating at the discourse level, relation/plans control the assembly of clause-sized units of representation into an appropriately linked tree structure, and operating at the clause level, concept activation action rules control the assembly of constituent-sized units of representation into an appropriately annotated pre-sentence structure which is then handed to a realization component.

The feasibility of such a smooth joining of discourse- and clause-level processing is enhanced by the fact that no clear boundary seems to exist between the structuring of paragraphs and the planning of clause parts such as preposition groups. That is to say, the leaves of a paragraph tree produced by the RST structurer need not always become separate clauses; very often, a number of leaves can be grouped together into a clause complex. For example, the last part of the SEQUENCE relation/plan in the tree in Figure 5 can be realized in at least two ways:

- (a). Knox, which is at 18N 79E, heads SSW, arriving on 4/24, to load for 4 days.
- (b). Knox, which is at 18N 79E, heads SSW. It arrives on 4/24. It loads for 4 days.

Going further, a leaf of the paragraph tree need not even always be realized as a clause, but can sometimes give rise to a grammatically simpler constituent such as an adjective. The ELABORATION satellite is an example, as in "the man who is big..." vs "the big man...". This issue is even more pronounced in German and Swedish than in English (for example, the italicized clause in "the suspicion *which was confirmed by this discovery* is..." need not form a separate clause in German: "die durch diese Entdeckung belegte Vermutung ist..."). To handle this issue, we have had to arbitrarily stop RST structure planning at the clause level and use a number of default rules to control realization as dependent or independent clauses (for example, ELABORATION-ATTRIBUTE satellites become dependent clauses instead of separate sentences, as in Figure 5). Making use of clause-level planning rules such as those in KAMP would help resolve these problems.

RST relation/plans can thus be seen as an extension of Appelt's clause-level planning rules to handle paragraph-level structure planning. Though their formalization has not been carried out to a degree which would support proofs of communicative correctness, this has always been their intended use, and was also one of the reasons for the selection of Cohen and Levesque's belief-based modal notation with which to represent requirements and goals [Cohen & Levesque 85].

4 Top-Down and Open-Ended Planning

4.1 A Different Treatment of Growth Points

The previous section described how RST relation/plans can be thought of as schemas which are susceptible to manipulation by a top-down planning algorithm. Though this approach has various desirable characteristics, such as explicitly representing the rhetorical role played by each clause in the paragraph

structure tree, the inflexibility that can plague schemas is also present in relation/plans' growth points. In general, in order to overcome a difficulty of this nature, one requires criteria by which to make the decisions that afford the desired variation. This section shows how a slightly generalized treatment of relation/plans supports a far more flexible planning regime.

Consider again the nucleus growth points of the SEQUENCE relation in Figure 3:

Nucleus growth points:

((BMB SPEAKER HEARER (CIRCUMSTANCE-OF ?PART ?CIR))
(BMB SPEAKER HEARER (ATTRIBUTE-OF ?PART ?VAL))
(BMB SPEAKER HEARER (PURPOSE-OF ?PART ?PURP)))

These three requirements suggest material that respectively is circumstantial to, is an attribute of, and is a purpose for, the material in the nucleus.

As currently used, the growth points are treated as actual *injunctions* for further growth of the paragraph tree at this point, and in this order: whatever information is related to the nucleus material by any of these relations must be included, if it has been selected for inclusion by the application program. However, depending on the hearer's knowledge and expertise, the time available for the interaction, the detail and complexity of the material to be communicated, and so forth, it may very easily be the case that any elaboration of the topic as mandated by growth points is suboptimally ordered or even wholly inappropriate. What then are the ramifications of building a slightly more discriminating system? One could treat growth points merely as *suggestions* for growth, depending on additional criteria such as the amount of time or space available. In fact, in the limit, one could even altogether remove growth points from relation/plans and have the structurer attempt to grow the paragraph tree using *all* the relation/plans at its disposal in each expansion cycle (as mentioned earlier, this was Mann and Thompson's original conception of structure planning using RST), although one problem with the extreme approach is that one may lose some of the coherence afforded by limiting the growth appropriately to each relation/plan.

The essential difference between the two approaches is captured in the question "on what basis is the addition of a new block of text to the tree considered?". In the first case, the answer to this question is:

- (a) the impetus for finding a further block of text is the growth point injunction contained in the current relation/plan. No additions are considered unless they are suggested by the plan.

In the second case, the answer is:

- (b) some (or all) possible additions of further text are considered at each step in the planning cycle; the impetus for such consideration is either a suggestion in the plan or is built directly into the planning procedure itself.

Small though it may seem, this distinction gives rise to important differences in the structuring procedure. The first approach can be called the *strict top-down planning approach*. Here the RST relation/plans are plans in the traditional sense; the explicit injunctions for growth correspond to the steps, in proper order, necessary to achieve the goal. In contrast, the alternative approach just introduced can be called the *open-ended planning approach*. This approach is useful in many situations when top-down planning is inappropriate — whenever one needs less structured methods to adapt interactively to the environment

and interlocutors, such as in conversations. Here the potential to grow any matched growth point goal relation at every point affords a flexibility to make opportunistic use of the best paragraph continuation under the changing circumstances. Of course, taking this approach, one has to abandon the strict pipeline model used so far, in which the structurer completes the text plan before allowing material to be realized. A model in which planner and realizer are interleaved, as suggested in [Hovy 88b], is required. Regardless of the planning model, however, the point is that RST relation/plans support either paradigm, depending on the treatment accorded growth point goals.

However, there is no free lunch. With the enhanced opportunities to make more suitable text under the suggestion interpretation comes the responsibility of exercising choice. Criteria are required to control the inclusion and ordering of candidate growth point material. Work is in progress in developing such criteria. With respect to the inclusion of material, the following play a role:

Reasons for inclusion

1. Explicit goal to communicate the extra material
2. Rhetorical-stylistic considerations: balance/parallelism of text
3. Hearer's lack of knowledge
4. Hearer's inability to understand without extra information
5. Hearer asks for more information

Reasons for exclusion

1. Lack of time or space
2. Unpleasant/undesirable effects of extra material
3. Lack of known detail
4. Untrustworthiness of material

At present, some of these criteria seem more difficult to implement than others. Inclusion criteria 1 and 5 are simple: if the structurer is given or generates the explicit goal to include material of some kind, then such material must be searched for. Criterion 2 is more challenging, but still addressable: we can develop strategies for deciding based on an inspection of the content, balance, and parallel structure of the contents of the paragraph tree. Criteria 3 and 4 are harder, since they require a model of the hearer. The criteria for exclusion are all somewhat easier to address, since they mostly depend on the characteristics of the material itself.

A further criterion, relevant to both inclusion and ordering decisions, can be derived from the work on focus of McCoy and Cheng (this volume). At each point in the planning cycle, candidate growth point material that is neither in focus nor allowed by a legal focus move (according to McCoy and Cheng's algorithm) may not be included in the tree, and material which is initially not a legal focus move may be made so by appropriate reordering of nuclei and satellites in the RST paragraph tree. A preliminary description of the use of McCoy's focus theory to constrain RST tree structure growth appears in [Hovy & McCoy 89].

4.2 Hybridization

Though the top-down and open-ended planning paradigms seem very different, the fact that the difference hinges upon a simple change in the location and interpretation of the growth point goals means that

RST relation/plans can be used in both planning paradigms, and hence serve a useful function in the investigation of the relative merits of each.

The correspondence is so close that it is quite possible to follow a hybrid paradigm, in which the optionality of each growth point goal is explicitly contained in the relation/plans. To the extent that one can provide criteria for deciding on the inclusion and order of relatable material, to that extent one can relax the injunctive nature of the growth point to perform open-ended planning, while the remaining goals can be left as ordered injunctions. In this way the degree of hybridization is easily altered.

An allied benefit is the variable level of constraint supported by such a hybrid scheme. Consider, for example, the genre of letters, in which the address, date, salutation, and closing are almost completely fixed. In addition, usually, the first and last paragraphs are somewhat fixed with respect to tone and content, while the body proper is usually much more free-form and open-ended. A hybrid schema/plan for a letter would contain fixed injunctions for the fixed parts, some less fixed addressee-related goals for the first and last paragraphs, and a number of optional growth goals for the body (in fact, it is not clear that one could build plans for a letter-planning structurer any differently). None of the optional goals need be followed, though doing so may give rise to a rather bizarre letter (to quote an example from Bill Mann: "Dear John, I hope all is well. I hate you and never want to see you again. Love, Mary").

5 Conclusion

In order to structure coherent paragraphs from given material, one requires plans. While it is possible, on the one hand, to build macro-level plans or schemas that describe whole paragraphs at a time, and, on the other, to build micro-level plans that determine the composition of individual parts of sentences, RST relation/plans (and plans of similar type) support the functionality of both. Relation/plans combine the best features of schemas (the definition of extended structures such as paragraphs; perspicuous representations) with the best features of hierarchical planning methods (power and flexibility). Thus relation/plans can be seen as generalizations of schemas that support hierarchical planning methods, and that are functional along the whole continuum from paragraphs to clause-internal constituents.

A second characteristic that makes relation/plans desirable is the ease with which they can be switched from the top-down structural planning regime required for optimally balanced, well-crafted text to the open-ended planning regime that supports the dynamic and flexible behavior required for interactive communication. By treating relation/plan growth points as injunctions, one can perform top-down planning; by treating them as suggestions and supplying criteria for their inclusion and ordering, one can perform open-ended planning.

This centrality and flexibility makes the RST relation/plans an extremely useful vehicle for investigating the planning of paragraph-length text. Opportunities for new work and unanswered questions abound; the use of plans to gather appropriate information from complex knowledge sources, and the criteria for including and excluding material are but two of the many avenues of text planning research that RST relation/plans open up.

6 Acknowledgments

Many thanks to Johanna Moore, Cécile Paris, and Kathy McCoy for comments and discussions, and to John Bateman, Christian Matthiessen, Bill Mann, HaJo Novak, and Dietmar Rösner for comments.

7 References

1. Appelt, D.E., 1981.
Planning natural language utterances to satisfy multiple goals. Ph.D. dissertation, Stanford University.
2. Appelt, D.E., 1985.
Planning English sentences. Cambridge: Cambridge University Press.
3. Arens, Y., Miller, L., Shapiro, S.C. & Sondheimer, N.K., 1988.
Automatic Construction of User-Interface Displays. *Proceedings of the 7th AAAI Conference*, St. Paul (808-813).
4. Cohen, P.R. & Levesque, H.J., 1985.
Speech Acts and Rationality. *Proceedings of the 23rd ACL Conference*, Chicago (49-59).
5. Conklin, E.J. & McDonald, D.D., 1982.
Salience: The key to the selection problem in natural language generation. *Proceedings of the 20th ACL Conference*, Toronto (129-135).
6. Hovy, E.H., 1988a.
Planning coherent multisentential text. *Proceedings of the 26th ACL Conference*, Buffalo (163-169).
7. Hovy, E.H., 1988b.
On the study of text planning and realization. *Proceedings of the AAAI Workshop on Text Planning and Realization*, St. Paul (17-29).
8. Hovy, E.H. & McCoy, K.F., 1989.
Focusing your RST: A step toward generating coherent multisentential text. *Proceedings of the 11th Cognitive Science Society Conference*, Ann Arbor (667-674).
9. Kaczmarek, T.S., Bates, R. & Robins, G., 1986.
Recent developments in NIKL. *Proceedings of the 5th AAAI Conference*, Philadelphia (978-985).
10. Mann, W.C., 1983.
An overview of the Nigel text generation grammar. USC/Information Sciences Institute Research Report RR-83-113.
11. Mann, W.C., 1988.
Text generation: The problem of text structure. In McDonald, D.D. and Bolc, L. (eds), *Natural Language Generation Systems*, New York: Springer Verlag (47-68). Also USC/Information Sciences Institute Research Report RR-87-181.
12. Mann, W.C. & Matthiessen, C.M.I.M., 1983.
Nigel: A systemic grammar for text generation. USC/Information Sciences Institute Research Report RR-83-105.
13. Mann, W.C. & Thompson, S.A., 1983.
Relational propositions in discourse. USC/Information Sciences Institute Research Report RR-83-115.

14. Mann, W.C. & Thompson, S.A., 1987.
 Rhetorical structure theory: Description and construction of text structures. In Kempen, G. (ed),
Natural Language Generation: New Results in Artificial Intelligence, Psychology, and Linguistics,
 Dordrecht: Martinus Nijhoff Publishers (85-95). Also USC/Information Sciences Institute Research
 Report RR-86-174.
15. Matthiessen, C.M.I.M., 1984.
 Systemic grammar in computation: The Nigel case. USC/Information Sciences Institute Research
 Report RR-84-121.
16. McCoy, K.F., 1985.
Correcting object-related misconceptions. Ph.D. dissertation, University of Pennsylvania.
17. McCoy, K.R. & Cheng, J., 1988.
 Focus of attention: Constraining what can be said next. Presented at 4th International Workshop
 on Natural Language Generation, Catalina Island. *This volume*.
18. McKeown, K.R., 1982.
Generating natural language text in response to questions about database queries. Ph.D. dissertation,
 University of Pennsylvania.
19. Moore, J.D. & Paris, C.L., 1989.
 Planning text for advisory dialogues. *Proceedings of the 27th ACL Conference*, Vancouver (203-211).
20. Moore, J.D. & Swartout, W.R., 1988.
 A reactive approach to explanation. Presented at 4th International Workshop on Natural Language
 Generation, Catalina Island. *This volume*.
21. Paris, C.L., 1987.
The use of explicit user models in text generation: Tailoring to a user's level of expertise. Ph.D.
 dissertation, Columbia University.
22. Paris, C.L., 1988.
 Providing tailored explanations of an expert system's behavior. Presented at 4th International
 Workshop on Natural Language Generation, Catalina Island. *This volume*.
23. Sacerdoti, E., 1977.
A structure for plans and behavior. Amsterdam: North-Holland Publishing Company.