DANDELION

Discourse Functions and Discourse Representation:
An Empirically and Linguistically Motivated,
Interdisciplinarily-Oriented Approach to Natural Language Texts

ESPRIT Basic Research Project 6665

Deliverable R1.1.2: a,b

Coherence relations: analysis and specification

John Bateman and Klaas Jan Rondhuis

October 5, 1994

—— CLS Tilburg/Nijmegen ——
— University of Edinburgh GMD-IPSI Darmstadt Universität Saarbrücken —
—— Universidad Complutense Madrid ——

Document ID	:	Deliverable R1.1.2: a,b: *Coherence relations: analysis and specification*
Issue date	:	October 5, 1994
Responsible	:	GMD-IPSI Darmstadt
Version	:	final
Status	:	Public
WP number	:	WP 1: Empirical Studies
Distribution	:	DANDELION Consortium, CEC
Keywords	:	Coherence Relations, Rhetorical Structure Theory, Conjunctive Relations, Discourse Structure, Causation

Coherence relations: analysis and specification

John Bateman and Klaas Jan Rondhuis

GMD-IPSI Darmstadt

Abstract:

This deliverable considers the problem of providing 'coherence relation' analyses for natural texts. A critical review of the state of the art reveals a number of problems which require attention before any particular mode of analysis can be adopted. These problems are addressed on the basis of empirical analyses of texts in order to arrive at some improved guidelines for what is to be interpreted as a 'coherence relation' and for establishing an improved specification of coherence relations in general. We illustrate both some of the benefits of the specification for linguistic analysis and some possibly problematic areas for computational specification. This latter stands as direct input for deliverable R2.3.3 whose task it is to complete such a specification for importation in the Dandelion discourse analyst's workbench.

Contents

1 Function and Scope of the present Deliverable

The original intention of this deliverable as set out in the Dandelion technical annex was to examine the motivations for the use of coherence relations between units of texts and to provide a specification of the same. Particular attention was to be given to the possible existence of more than one relation between text parts as this is a major theme in current discussions of rhetorical and discouse organization. The assumption was made that the kinds of coherence relations extant and their theoretical status were sufficiently clear as to discuss their multiple occurence in texts. However, after a wider review of the available approaches and an attempted application of those approaches to real texts during empirial analysis, we came to the conclusion that this initial assumption was not valid. The kinds and statuses of coherence relations are in fact unclear. Our initial focus in this deliverable therefore of necessity became that of improving our theoretical understanding of possible coherence relations driven by empirical analysis of texts; this is reported in Section 3. The result of this work is a new proposal for the types and interrelationships of a wider set of entities previously loosely grouped under the general rubric 'coherence relations'. This proposal, although in our opinion empirically and linguistically more adequate, directly raises a number of difficult issues concerning its possible computational specification. These topics are central to the computational specification to be given in the relevant discourse grammar of Deliverable 2.3.3.

The framework specification we provide at the conclusion of this part of the deliverable is not, of course, complete in itself—what it does set out to achieve, and the basis on which it should be evaluated, is to support:

- a breadth, depth, and density of coverage fully compatible with the original Dandelion goals (i.e., involving all levels of description from text types to concrete surface phenomena),
- a level of specificational detail that supports immediate transfer to the Dandelion discourse analysis workbench, also placing constraints on the necessary form of representations used within that workbench,
- a sufficiently complex and realistic treatment of coherence relations to evaluate the general applicability and value of the Dandelion discourse workbench for discourse analysis.

Further framework content and revision is to be expected in the concrete application of the Dandelion workbench as undertaken as part of its initial development (Task and deliverable 2.2.1) and subsequent evaluation by use (during tasks and deliverables: 2.3.1, 2.3.2 and 2.3.3), as well as further considerations of implementational and formal requirements drawn from ongoing related work in other projects (e.g., [Brew *et al.*, 1993, Gardent, 1994, Bloembergen and Gardent, 1994]). It is an important aspect of the current work to show the relation and, where appropriate, compatibility, of our results to such work.

2 First words

A general investigation of coherence relations is essential for an understanding of how texts work as texts; such an investigation is therefore equally essential for any project whose main focus is text and its organization. A considerable difficulty facing any such investigation is the lack of detailed consensus concerning what may or may not be admitted as a coherence relation. There is a common assumption of a single class of coherence relations for which an account is required. [Hobbs *et al.*, 1992], for example, suggest that their account applies equally for any sufficiently defined theory of coherence relations, such as 'rhetorical structure theory' (see Section 3.1.1) or Hobbs' own (see Section 3.1.2). [Asher, 1993] also adopts relations such as those drawn from rhetorical structure theory, [Polanyi, 1985] and Hobbs for linking discourse segments (see Section 3.1.4.1). Similarly, attempts at bringing order to the proposals for coherence relations such as those of [Hovy, 1990, Maier and Hovy, 1993] explicit classify coherence relations from very diverse theoretical positions into a single hierarchy. However, the assumption of uniformity notwithstanding, almost all approaches differ in the concrete relations and definitions proposed for their analyses.

A detailed overview of the assumptions and definitions made in the different approaches suggests that one possible reason for the lack of consensus is that, in fact, a rather heterogeneous set of relations are being investigated with diverse and complex interrelationships. A further possible contributing factor is that most approaches implicitly restrict the range of evidence upon which their definitions are based—in particular, the *linguistic evidence* appealed to is generally narrow. Our first phase of empirical analysis addresses these issues directly. We seek to widen the range of linguistic evidence appealed to in order to clarify the distinct types of coherence relations and their likely interrelations. The space of possible coherence relations is thus to be provided with more internal structure. This is the main content of Section 3.

It is also a major concern of this deliverable to propose a firmer methodology for motivating coherence relations and providing testable definitions for such relations. We propose more modularity in the account in order to be able to factor out those aspects not directly concerned with coherence relations as we will define them. It is generally recognized that understanding coherence in texts relies heavily on general knowledge—it is not our aim here, however, to explicate general reasoning capabilities. A tighter modularity of coherence relations is then essential in stating where and how general reasoning capabilities concering commonsense knowledge play a role. This is clarified in our initial specification for coherence relations given in Section 5. Finally, here, we set out the initial uncompleted tasks necessary for the Discourse Grammar deliverables of the second set of work packages (Deliverables R2.3.1, R2.3.2, R2.3.3).

The account we propose is in certain respects significantly more complex than previous views of coherence relations. But it is also underspecified in a number of important areas. This underspecification raises difficulties for a computational specification. We seek to further constrain that underspecification by means of several direct empirical investigations of possible interrelationships between distinct kinds of coherence relations and other linguistic constructs. This is reported in Section 4.

Following the pre-specification, we illustrate the use of our account by applying it to some problems currently being discussed in the literature of coherence relations and rhetorical structure (Section 6).

3 Empirical Analysis: first phase

In this section we will introduce the most developed proposals for handling coherence relations available to date, show a representative subset of those proposals in action analysing actual texts, and use these analyses to specify more precisely how the proposals differ, how they are similar, and what their relative strengths and weaknesses are. Our own specification attempts to synthesize these current research results and to take that synthesis further on the basis of the empirical analysis. The approach therefore deliberately differs from those which attempt to stay within a single existing view of coherence relations, extending this ever further in the face of counter-examples[1] As our overview of approaches should make clear, there is a sufficiently wide range of information being handled in the differing perspectives that to accept any one view as a basis at this stage would be premature and wasteful of progress in other approaches.

The organization of the section will be as follows. First, the approaches will be introduced, their basic tenets explained and illustrated with simple examples. Second, we show the analyses using these

[1] An example of such work is given for the case of Rhetorical Structure Theory in Section 3.2.3.3 below.

approaches for two selected texts. We present all of the analyses for a single text together, rather than presenting each mode of analysis for each text, in order to give a better comparative sense of the differences/similarities between approaches. Third, we discuss those differences/similarities in order to refine our view of what kinds of coherence relations are in fact necessary.

3.1 The range of approaches

This section[2] gives an overview of the current approaches to coherence relations in order to clarify the possible starting points for analysis. The approaches presented are, in fact, a subset of those we have investigated; we restrict our presentation here to those that directly influence the content of our proposed account of Section 5.

We begin our discussion with two approaches that have been developed and tested in computational linguistics. Probably the most extensively used theory of coherence relations in the field of text generation is that of Rhetorical Structure Theory (RST) [Mann and Thompson, 1987]; we introduce RST in Section 3.1.1. A further account, particularly well oriented to a speaker's knowledge and inferential capabilities is that of [Hobbs, 1990]; we introduce Hobbs' extensive view of discourse theory in Section 3.1.2. Text organization has also long been the concern of functional approaches to linguistics. Here one of the most developed accounts of text coherence is that of [Martin, 1992]; we introduce this in Section 3.1.3. There is also considerable work on mechanisms for supporting coherence relation-related processing. Particularly relevant here are approaches based on Discourse Representation Theory (DRT) [Kamp, 1981] and the formal discourse work of [Polanyi, 1986, Scha and Polanyi, 1988]. Among these, the 'Structured DRT' framework developed in, e.g., [Lascarides and Asher, 1991, Asher, 1993] is the most developed; we discuss this approach in Section 3.1.4 below.

A further important general orientation to investigating the relational coherence of discourse anchors itself in cognitive psychology; here discourse relations or structures are sought that are psychologically 'plausible' in a sense supportable by empirical study. Well-developed approaches here include [Sanders, 1992] and [Knott and Dale, 1992]. These are not directly discussed in the deliverable, although aspects of their work have influenced our general approach. It should also be noted that the approaches selected for attention are those that are more crucially concerned with the relational nature of coherence. There are also approaches that concentrate more specifically on the structural, constituency, or hierarchical nature of discourse organization. Such approaches include, e.g., [Kamp, 1981, Kamp and Reyle, 1993, Reichman, 1978, Reichman, 1985, Polanyi, 1986, Scha and Polanyi, 1988, Grosz and Sidner, 1986]. Also highly relevant here, although not yet available for inclusion, is the work in the LRE Project 'Towards a declarative theory of discourse'. Although in this deliverable these approaches will receive less attention, it will become clear below (Section 5.4) in our listing of possible problems for a computational specification of our account that we also need to include detailed structural information in any complete specification of discourse organization. This issue will, therefore, be returned to when the discourse grammar implications are considered.

For each approach selected for attention[3], we will give examples where relevant and ask a number of orienting questions before giving a preliminary appraisal. Each approach is rated along several dimensions of evaluation; in particular, we consider what aspects of these approaches can we re-use, and where problems surface when they are applied. Where an approach is descriptively broad, we will focus more on the general area of 'causation' in order to limit the discussion somewhat; this will be

[2] Including in part preliminary work by Liesbeth Degand and John Bateman [Degand and Bateman, 1993].

[3] One possible candidate for a useful treatment of coherence relations that, due to time-constraints, has not been considered in the present deliverable is the *dynamic exchange* approach of [Ventola, 1984, Ventola, 1987]. This has primarily been applied in the study of dialogues but there are indications that it might also be of use in monologue texts of the type studied within Dandelion; the approach shares with dynamic semantics (e.g., [Groenendijk and Stokhof, 1991, Heim, 1983]) a central concern with the ongoing dynamic development of discourse. An appraisal and possible incorporation of Ventola's approach should be considered for future work.

one of the central areas taken up in the corresponding discourse grammar specification.

3.1.1 Mann and Thompson's Rhetorical Structure Theory

Rhetorical Structure Theory [Mann and Thompson, 1987] is a theory of discourse relations which seeks to provide a general way to describe the relations among segments of a text, whether or not they are grammatically or lexically signalled. A set of relations was posited on the basis of detailed descriptive linguistic analyses, originally of some 400 texts of varying content and genres. In contrast to work such as that of [Sanders *et al.*, 1992], the account is not a psychologically motivated account. The approach was also not itself constructed as a 'computational' theory; the analyses called extensively on the intuition of the discourse analyst. For this reason, the theory as a whole has been criticised on the basis that it is 'merely' a descriptive account (e.g., [Knott and Dale, 1992]), although we will suggest below that this is somewhat of an oversimplification.

RST defines approximately 25 relations which are hypothesized adequate for the majority of texts. These relations, although generally presented as a list in earlier papers, contained from the outset suggestions of possible interrelationships and subclassifications. Thus, for example, the definition of the relation ELABORATION explicitly contains the alternative of including 6 subtypes (set-member, abstract-instance, whole-part, process-step, object-attribute, and generalization-specific [Mann and Thompson, 1987, p52]) and ANTITHESIS is stated to be a subtype of CONTRAST [Mann and Thompson, 1987, p11]. Subsequent presentations of more extensive sets of RST relations have, therefore, generally moved to proposals for a classification hierarchy of such relations. This is clearly also more suitable for their computational representation and use in planning. Two kinds of extensions are typically found, extensions in kind and extension by subtype.

A central defining feature of rhetorical relations is the recognition that the text segments related by a relation are often not of equal textual 'status' or importance. The two segments related are therefore assigned differing status: one is typically the *nucleus*, the other the *satellite*. It is claimed that ommitting a nucleus from a text would destroy a text's coherence whereas omitting a satellite is relatively harmless. As we shall see below, however, there are rhetorical relations (particularly 'interpersonal') where this is not so clear. There are also a small set of relations which do not assign differing textual status; these are more analogous to paratactic relations.

Among its 25 relations, RST defines five coherence relations within the causal domain: Volitional Cause, Non-Volitional Cause, Volitional Result, Non-Volitional Result, and Purpose. Cause and Result relations are distinguished by their nuclearity; Cause relations have the causing situation as their satellite, whereas in Result relations the causing situation is nuclear and the caused situation is less central. Within each of the Cause and Result pairs, Volitional and Non-Volitional variants are distinguished. Purpose is neutral in this respect, it includes both volitional and non-volitional cases.[4]

Volitional Cause

> *Writing has almost become impossible* (causing situation - satellite) **so** *we had the typewriter serviced* (caused situation - nucleus) *(and I may learn to type decently after all these years).*

Non-Volitional Cause

> *(we've been able to mine our own iron ore, coal, manganese, dolomite, all the materials we need to make our own steel.) And* **because** *we can mine more than we need,* (causing-

[4] In each of the subsections here, unless otherwise stated, examples given are drawn from the original sources describing the individual approaches.

situation - satellite) *we've had plenty of manganese and iron ore for export.* (caused situation - nucleus)

Volitional Result

Using thumbs is not the problem but heredity is, and the end result is no use of thumbs if I don't do something now. (causing situation - nucleus) *Writing has almost become impossible so we had the typewriter serviced and I may learn to type decently after all these years* (caused situation - satellite).

Non-Volitional Result

The blast, the worst industrial accident in Mexico's history, destroyed the plant and most of the surrounding suburbs. (causing situation - nucleus) *Several thousand people were injured, and about 300 are still in hospital.* (caused situation - satellite)

Purpose

To *see which Syncom diskette will replace the ones you're using now,* (satellite) *send for our free "Flexi-Finder" selection guide and the name of the supplier nearest to you.* (nucleus)

The distinctions between rhetorical relations rests on the relations that have been found necessary to propose convincing descriptions of the writer's intentions in constructing a text. The motivations for the relations is, therefore, not to be found in linguistic realization. Relations do not, for example, necessarily distinguish between different linguistic realizations, although it is possible that more subtle patternings of linguistic realizations are being covered (for example, [Matthiessen and Thompson, 1989]'s link of nuclearity with grammatical taxis), or that, e.g., nuclearity is important in its own right for discourse/coherence accounts.[5] Nuclearity is in some cases clearly used for differentiating between relations; for example the distinction between (Non-)Volitional-Cause and (Non-)Volitional-Result relations is solely whether nuclear status is attributed to the Cause or the Effect respectively.

A typical example of an RST definition is that shown in Figure 1 for CONDITION drawn from [Mann and Thompson, 1987, p65]. Here we can see that definitions are expressed informally, but nevertheless place constraints on the kinds of meanings that the nucleus and satellite may express and the consequences (the 'effect') of using/recognizing the relation. The constraints and conditions of the relations overall are in fact often sufficient to produce good inter-coding reliability—a result which is perhaps surprising given the apparent informality of the definitions.

RST received an initial computational operationalization by [Hovy, 1988] and has since been incorporated in a wide range of natural language generation systems (see, e.g., [Moore and Paris, 1988, Scott and de Souza, 1990, Bateman *et al.*, 1991, Dobeš and Novak, 1991, Cross, 1992, Fawcett and Davies, 1992, Hovy *et al.*, 1992, Rösner and Stede, 1992, Vander Linden, 1993, André and Rist, 1993]).

In current computational contexts, rhetorical relations are modelled as one type of communicative goal. There is now a distinction drawn between higher level communicative goals and lower level goals—only the latter are directly related to rhetorical relations. Rhetorical relations are therefore formalized in terms of 'plan operators'. An example of such a plan operator from the system of Moore and Paris [Moore and Paris, 1993] is the following:

[5] For example, Gisela Redeker (personal communication) suggests that this might be important even for the way that we model the generation process: we may wish to proceed from 'nuclear' information, giving it focused thematic status, or we may wish to be able to proceed forward or backward in the causal chain (in the knowledge base) in searching for a suitable satellite. Then it would be useful to maintain the distinction, since a cause-satellite is going to do other kinds of work than a result-satellite, and a cause-nucleus is going to differ from a cause-satellite.

relation name:	CONDITION
constraints on N:	none
constraints on S:	S presents a hypothetical, future, or otherwise unrealized situation (relative to the situational context of S)
constraints on the N+S combination:	Realization of the situation presented in N depends on realization of that presented in S.
the effect:	R recognizes how the realization of the situation presented in N depends on the realization of the situation presented in S.
locus of the effect:	N and S

Figure 1: RST definition of relation CONDITION

```
EFFECT: (PERSUADED ?hearer (DO ?hearer ?act))
CONSTRAINTS: (AND (STEP ?act ?goal)
                  (GOAL ?hearer ?goal)
                  (MOST-SPECIFIC ?goal)
                  (CURRENT-FOCUS ?act)
                  (SATELLITE))
NUCLEUS: (FORALL ?goal
                  (MOTIVATION ?act ?goal))
SATELLITES: nil
```

The successful application of this operator has the effect that a state of the hearer being 'persuaded' (a discourse intention) to do some act is achieved. The operator may be applied when the specified constraints hold. When this is the case, a rhetorical structuring involving 'motivation' is constructed. Information selection is thus achieved as a side-effect of binding variables in the operator's constraints. Further such plan operators then decompose the rhetorical relation 'motivation' until sequences of surface speech acts are reached. The system of Moore and Paris contains approximately 150 such plan operators and is considered sufficiently stable for use in various application systems.

There is no doubt that RST forms the most influential approach to coherence relations found in text generation to date.

Preliminary conclusion

RST appears to be a theory which is useful for Text Generation on a higher level of discourse structure. It's aim is indeed to describe the organization of texts, characterizing their structure primarily in terms of relations that hold between parts of the text. An interesting claim of the theory is that the same rhetorical relations that hold between spans realized by individual clauses also account for the relationships between larger segments of text, irrespective of the size of the text span involved. This makes it possible to characterize the structure of a text in terms of a hierarchy of rhetorical relations that connect together the parts. On the other hand, this general text structure approach keeps it at considerable distance from the text to be produced. The essential claim of being 'pre-realizational' makes its consequences for linguistic realization at best indirect. If we want RST to be useful in this respect, the next step is to determine how the established rhetorical relations of a given text can be mapped onto the lexicogrammatical realizations, whether these are intrasentential or intersentential. This mapping is certainly not straightforward, which means RST has to be completed by other theories, the combination of which should constrain the number of possible realizations. Note that there are regular correlations between aspects of rhetorical structure and lexicogrammatical structure (cf., e.g., the detailed work on the relation between RST structures—particularly nuclearity status—and grammatical taxis (hypotaxis and embedding) by [Matthiessen and Thompson, 1989] or the relation between RST structures and thematicity [Bateman and Matthiessen, 1990]), but these are usually more subtle than the identification of single cue phrases.

3.1.2 Hobbs' Coherence relations

Hobbs' [Hobbs, 1990] approach to coherence relations places the primary source of evidence for relations in the 'world' that texts describe: that is, a text is coherent because it describes a coherent state of affairs or a coherent sequence of events. His definitions of coherence relations are therefore rooted in conditions that must hold of the world for coherence of the designated kinds to be recognised.

Coherence relations are then seen as strategies by which speakers and writers can construct texts. Four classes of such 'text-building strategies' are distinguished:

- **Occasion**: where a message is coherence because it tells about coherent events in the world [Hobbs, 1990, p86]. Subtypes of this relation include **cause** and **enablement**.[6]
- **Evaluation**: where what is said is related to some goal of the discourse [Hobbs, 1990, p89].
- **Ground-figure** and **Explanation**: where relations are drawn to the listener's or reader's prior knowledge [Hobbs, 1990, p90], providing 'background', or a general 'geography'.
- **Expansion**: which expands the discourse "in place rather than carrying it forward" [Hobbs, 1990, p92]. Such relations are seen as aiding the reader's or listener's inferential processes[7]

Most of these in fact have subtypes; an overview is given in Figure 2.

As with RST, each such coherence relation is accompanied by a specification of the conditions that must be satisfied for it to hold. However the form of these constraints and their theoretical commitments differ considerably from those of RST. For example, the coherence relation 'elaboration' is defined to hold between two discourse segments S_0 and S_1 when we can:

> Infer the same proposition P from the assertions of S_0 and S_1. [Hobbs, 1990, p95]

Similarly, the relation 'occasion' requires that:[8]

> 1. A change of state can be inferred from the assertion of S_0, whose final state can be inferred from S_1.
> 2. A change of state can be inferred from the assertion of S_1, whose initial state can be inferred from S_0. [Hobbs, 1990, p87]

As well as providing recognition criteria, this form of definition allows relationships between the coherence relations to be defined perhaps more easily than with RST. Thus, 'elaboration' can be seen as a subtype of the relation 'parallel', since the latter relation is defined as obtaining when:

> Infer $p(a_1, a_2, \ldots)$ from the assertion of S_0 and $p(b_1, b_2, \ldots)$ from the assertion of S_1, where a_i and b_i are similar, for all i. [Hobbs, 1990, p93]

[6] Later Hobbs also mentions subtypes such as 'solutionhood' and 'set of events plus outcome'; the list of relations is thus clearly not intended to be closed.

[7] This metaphor can, of course, also be applied to rhetoric, in its more traditional sense, since the listener's inferential processes are helped in certain very specific directions. Whether the metaphor really helps in this case is not, however, so clear.

[8] Note that it remains unclear to us at this time, what precisely is meant by 'state' or 'event' in these definitions. It does not seem to play a decisive role in the applicability of relations, unlike the case in the account of Lascarides and Asher described below.

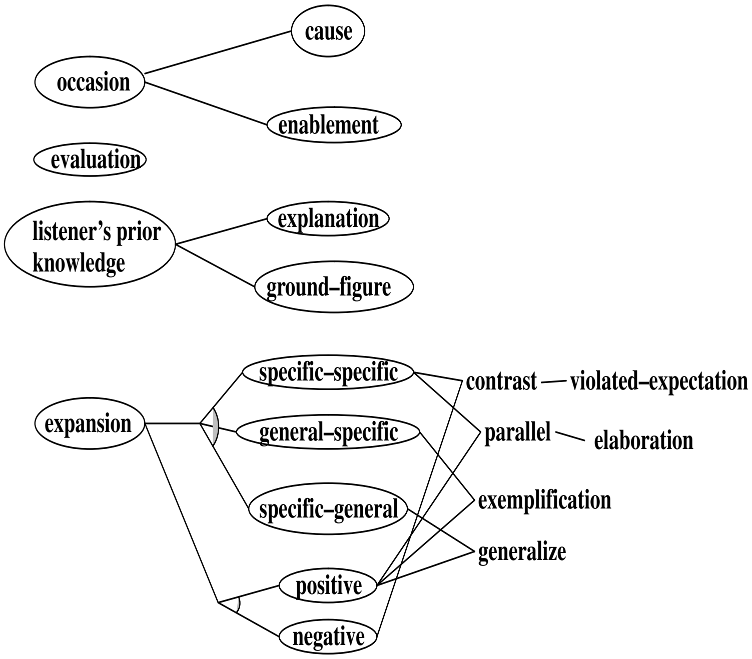

Figure 2: Overview of Hobbs' coherence relations

If a logical rendering of the RST-relation definitions were possible, this difference would, however, disappear. A more significant distinction is that Hobbs' definitions appear to refer to states of affairs or conditions that hold in the world, whereas RST definition always refer to the *change* that recognition of a relation brings about (see again, for example, the effects field of the definition of CONDITION shown in Figure 1 above). We will make more of this distinction below.

An interesting case of the use of logical inference over the relation definitions is given by Hobbs in the analysis of the text fragment:

(1) John can open Bill's safe.
He knows the combination.

Since both the propositions of being able to open the safe and knowing the combination of the safe entail that John is, in fact, able to open the safe, the conditions for 'elaboration' hold. However, as Hobbs notes, there is also a tendency to interpret a kind of explanation to hold between then two, that is the reason for John being able to open the safe is that he knows the combination. Hobbs' definition for 'explanation' is:

> Infer that the state or event asserted by S_1 causes or could cause the state or event asserted by S_0. [Hobbs, 1990, p91]

Given this,

> "Elaboration and explanation can blend into each other for the following reason. To recognize explanation, we need to infer S_1 *cause* S_0. To recognize elaboration, we must infer a P such that S_0 *imply* P and S_1 *imply* P. Very often the P is just the assertion

of S_0 itself, that is, from each of S_0 and S_1 we infer S_0. If in addition to inferring S_0 from S_1, we recognize that we have drawn that inference, we have thereby inferred S_1 *imply* S_0. Implication can be viewed as a kind of bloodless causality; it plays the role in information systems that causality plays in physical systems, and it seems to me likely that we understand implication by analogy with causality. That is, the inference S_1 *imply* S_0 is a variety of S_1 *cause* S_0." [Hobbs, 1990, p96/7]

This link between implication and causality supports the kind of blend noted.

Hobbs in fact places 'causality' in a central position in his account of coherence relations. Of the relations described above, those of 'occasion', 'evaluation' and 'explanation' are explicitly related directly to causality.[9] As he suggests, one possible way of grounding the definition of coherence relations is by assuming that

> "these coherence relations are instantiations in discourse comprehension of more general principles of coherence that we apply in attempting to make sense out of the world we find ourselves in, principles that rest ultimately on some notion of cognitive economy. ...Recognizing coherence relations may thus be just one way of using certain very general principles for simplifying our view of the world." [Hobbs, 1990, p101]

And causation, and its recognition, appears a good candidate for one such very general principle.

Hobbs also, very relevantly for Dandelion, discusses the concept of text 'genres' and their relation to coherence structures. In one example text analysis that he offers, two structures described as subtypes of the *occasion* relation are introduced—*problem-solution* and 'a set of events and its outcome'—and placed in a larger structure. Of this larger structure Hobbs says:

> "It is likely that this structure is a very common pattern for stories in our culture. It is a coherence structure, but not just any coherence structure. In this convention, the *occasion* relations are constrained to be a problem-solution relation and an event-outcome relation, and the *contrast* has to be between two possible solutions." [Hobbs, 1990, p107].

This is an example of what Hobbs terms a 'genre': i.e., while the coherence relations allow arbitrary patterns of coherence structures to be constructed, some subset of these is picked out by a culture and made conventional.

Preliminary conclusion

The relations proposed by Hobbs are very general and are given straightforward definitions appealing to a logical representation of a presumed 'knowledge base' available to speaker and listener, writer and reader. Reasoning about the relationships between relations is not possible without these kinds of definitions. There are no direct links drawn with linguistic realization. Linguistic realization provides evidence for an interpretation, but is in no cases absolutely required. This is then equivalent to the 'ideal' version of RST in which the relations are purely 'pre-realizational' [Mann and Thompson, 1987, p8]. Also, just as with RST, particular genres are presumed to be conventionalized patterns of coherence relations.

It is worth noting that Hobbs does not propose many constraints on the application of his coherence relations and that it is accepted that there may for any text be very many possible analyses. [Hobbs *et al.*, 1992] describe a cost-weighted abductive mechanism which provides a way of ordering constraint application, but little is said about how such weights might be determined.

[9] Theories of these relations may therefore be taken as importing from a microtheory of causation as suggested in Hobbs's work on commonsense ontology (e.g. [Hobbs, 1986]).

3.1.3 The Conjunction System of Martin's discourse semantics

Martin's [Martin, 1992] central concern is to describe the relationship between discourse structures and grammatical ones and to focus on the complementarity between discourse and grammar. This makes his approach very valuable to our concern, since in text generation this relationship has to be made explicit in order to choose the specific lexicogrammatical realizations of the given relations on the higher discourse level. About this realizational link, he notes that *the semantic system* CONJUNCTION *tends to be realised through paratactic [coordinating] and hypotactic [subordinating] relationships within the clause complex, and through 'cohesive' conjunctions relating clause complexes to each other.* But these logico-semantic relations are also very general ones which are manifested on various places in the lexicogrammar. This poses the problem of granularity, which we will develop further below.

Martin recognises four main types of logico-semantic relations: additive, comparative, temporal and consequential; the latter one, which is of our concern here since it covers causation, being further subdivided into purpose, condition, consequence and manner. All of them can be combined with the feature concession (covering [Sanders *et al.*, 1992]'s polarity operation). Following [Halliday and Hasan, 1976], he also distinguishes between *internal* and *external* relations. Internal relations *obtain in the organisation of the text itself rather than in the organisation of the world the text describes.* External relations are *oriented to what is going on outside the text, rather than within.* ([van Dijk, 1977], [Redeker, 1990] and [Sanders *et al.*, 1992] distinguish between semantic and pragmatic connectives/relations along similar lines). This distinction is central in Martin's account since some conjunctions can only function internally, while others are exclusively external in use. We first turn to external consequential relations and then to internal ones.

3.1.3.1 External consequential relations

Consequential relations are constituted by (at least) two propositions (Martin uses the term 'message' or 'event'), so that one proposition is seen as enabling or determining the other. They all have the experiential structure [Cause-Effect]. The first opposition to be drawn is that between **Manner** and the other consequential relations grouped under the label **Cause**. With Manner the relationship is one of enablement/potentiality, i.e. the Cause enables the Effect, as in:

- *How did you win?*
- *By training hard.*

The Cause relationship between the propositions/events is one of determination/obligation, i.e. the Cause determines the Effect, as in:

- *Why did you win?*
- *Because we trained hard.*

Successive modularisations of the cause relations, which obtain through diverging traversals of the network given in Figure 3 — and which are thus, according to systemic linguistics, motivated by different semantic choices — result in the following realizations (for each example a 'cohesive' and a 'hypotactic' realization is given).

cause:consequence

a. We arrived late. **So** *we didn't have much time to prepare.*
b. **Because** *we arrived late, we didn't have much time to prepare.*

cause:consequence/concession

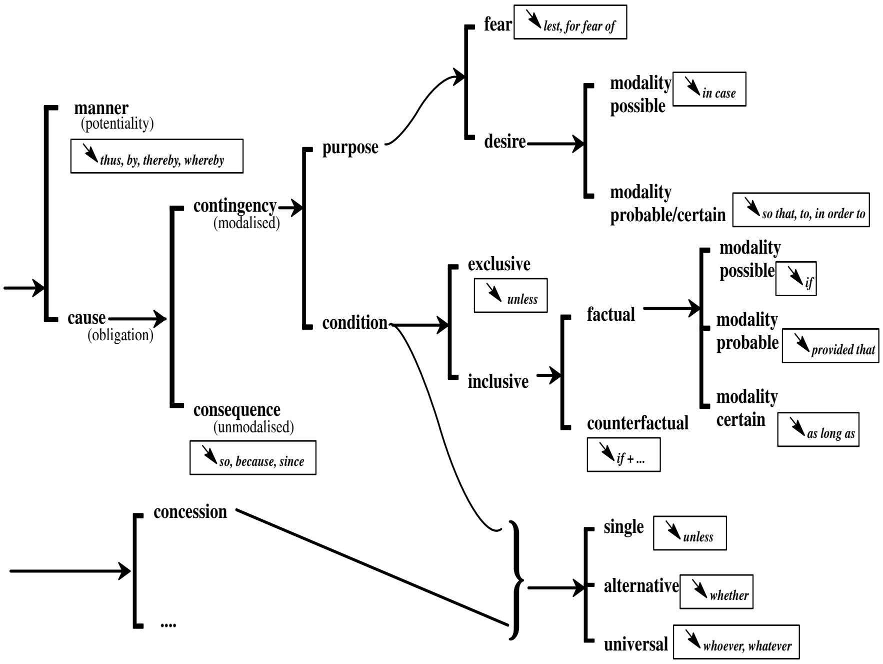

Figure 3: The external consequential relations

a. *We didn't arrive late.* **However** *we didn't have much time to prepare.*
b. **Although** *we didn't arrive late, we didn't have much time to prepare.*

cause:contingency:condition:exclusive

a. *Our dog had to stand still.* **Otherwise** *we wouldn't have won.*
b. **Unless** *our dog had stood still, we wouldn't have won.*

cause:contingency:condition:inclusive:counterfactual

a. *He* **might** *have arrived on time.* **Then** *he would have won.*
b. **If** *he'd arrived on time he would have won.*

cause:contingency:condition:inclusive:factual:modality possible

a. *We* **might** *go.* **Then** *we'll win.*
b. **If** *we go, we'll win.*

cause:contingency:condition:inclusive:factual:modality probable

a. *We'll* **probably** *turn up* **Then** *we'll win.*
b. **Provided that** *we show up, we'll win.*

cause:contingency:condition:inclusive:factual:modality certain

a. We'll **certainly** *show up* **Then** *we'll win.*
b. **As long as** *we show up, we'll win.*

cause:contingency:condition/concession:single

Even if *we get that judge again we'll win.*

cause:contingency:condition/concession:alternative

Whether *we get that judge again or not we'll win.*

cause:contingency:condition/concession:single

Whichever *judge we get we'll win.*

cause:contingency:purpose:fear

a. We **didn't want** *to lose.* **So** *we didn't enter.*
b. We didn't enter **for fear of** *losing.*

cause:contingency:purpose:desire:modality possible

a. We **wanted** *to win and had a* **chance**. **So** *we entered.*
b. We entered **in case** *we won.*

cause:contingency:purpose:desire:modality probable/certain

a. We **wanted** *to win and had a* **good chance**. **So** *we entered.*
b. We entered **so that** *we would win.*

3.1.3.2 Internal consequential relations The internal resources to code non-concessive consequential relations are fairly limited, so that for the most part the same conjunctions are used to express internal and external relations of this type. As for the concessive relations, they are realized through a number of distinctively internal conjunctions. The network displaying the internal consequential relations is given in Figure 4, and some of the corresponding realizations are given in the examples below.

concessive:objection:concede:modality possible and probable

(...) **Admittedly** *there are individuals who clearly have inherently imperfect grammars, (...).* **Of course**, *there is no shortage of differences between grammars, (...).*

concessive:objection:concede:modality certain

(...) **Needless to say**, *no psychologist would dream of claiming that this was possible, even in principle.*

concessive:objection:dismiss

Dinner's ready.
- But I'm not hungry.
In any case, *I'm going out.*
Sorry I forgot to tell you.

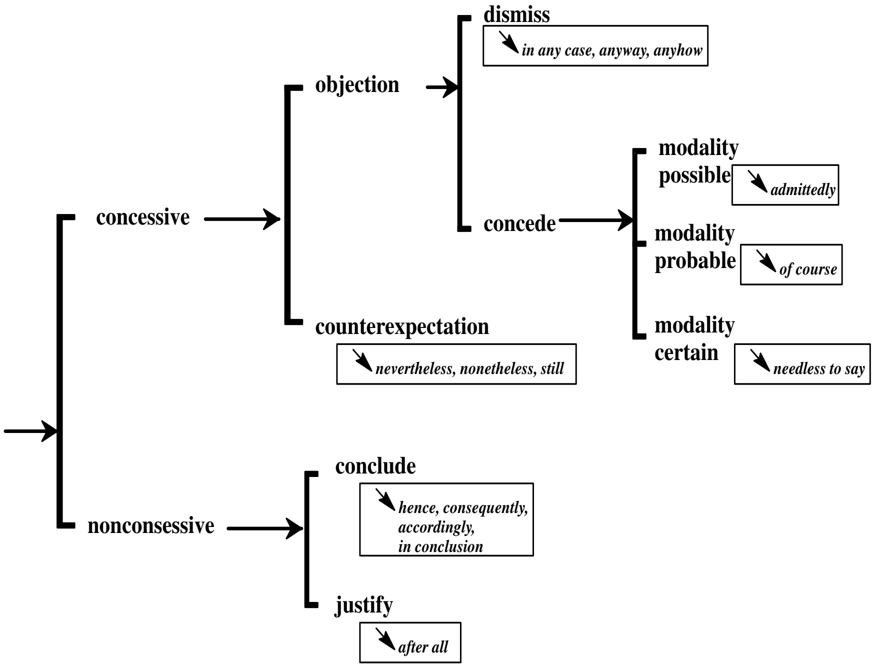

Figure 4: The internal consequential relations

3.1.3.3 'Incongruent' realizations of consequential relations The realizations of the consequential relations considered so far were all within the clause complex, expressed by means of paratactic and hypotactic conjunctions, or at the intersentential level where the relations are expressed through 'cohesive' conjunctions relating simplex and/or complex clauses. Martin expands considerably, however, on the possible domain of application of conjunctive relations by incorporating Halliday's proposal of possible *incongruencies* in lexicogrammatical realizations [Halliday, 1985a, Chapter 10]. The realization of semantic configurations in terms of clause complexes and independent sentences is treated as the 'congruent' realization (i.e., the most 'natural' and 'direct' realization given the current historical state of the language). For any such relation it is also possible to have a systematically related range of incongruent realizations. Thus, instead of realizing conjunctive relations as logico-semantic relations between processes, the relations may also be realized 'within', rather than between processes—as, for example, Circumstances, within the structure of a relational clause, or within the structure of a nominal group.

Expressing logical relations within a process depends on grammatical metaphor; (see [Halliday, 1985a, p321]) *one or both of the processes involved must be 'thing-ized' to some extent - via embedding or nominalisation. These realisations of logical relations can be taken as incongruent and are more characteristic of written than spoken text* [Martin, 1992]. Some divergent patterns of realization are summarized in Figure 5 (from [Martin, 1992]), and they are illustrated on the basis of [Halliday, 1985a, p378-384]'s systematic account of this diversification for causal relations.

All of the patterns shown in Figure 5 vary both in their interpretation of the event and in their distribution of the information, some may seem more plausible than others, *but each one could be derived as the form that would be predicted by a particular configuration of the features of the context of situation* [Halliday, 1985a, p384]. If we want to generate the right text for the given situation the task is precisely to determine these features.

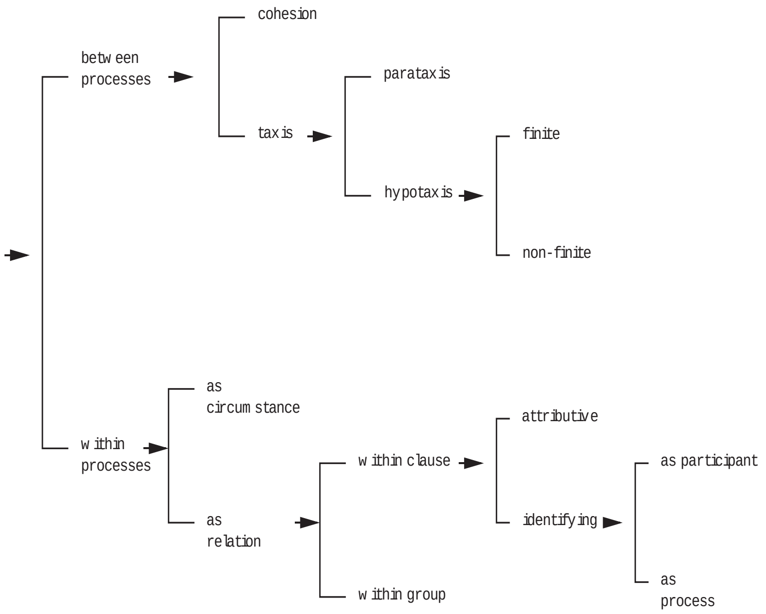

Examples.
Cohesion:
She didn't know the rules. Consequently, she died.
Parataxis:
She didn't know the rules; so she died.
Hypotaxis-finite:
She died because she didn't know the rules.
Hypotaxis-non-finite:
She died through not knowing the rules.
Process as Circumstance:
Her death/That she died resulted from her ignorance of the rules.
Process as Relation within clause-attributive:
Her death was due to ignorance of the rules.
Process as Relation within clause-identifying as participant:
Her ignorance of the rules was the cause of her death.
Process as Relation within clause-identifying as process:
Her ignorance of the rules caused her death.
Process as Relation within group:
Her death resulting from her ignorance of the rules...

Figure 5: Divergent grammaticalisations of the system CONJUNCTION

Questions:

- Is the question of congruent vs. incongruent realizations one which should be treated together with the coherence relations?
- If it is not, where should this be accounted for?
- To what level of granularity do we want to go?
- Are all the distinctions drawn by Martin on the basis of modality a matter for the type of relations themselves? Could they not be considered as a general constraint coming from interpersonal aspects?

Preliminary conclusion

The lexicogrammatical distinctions made by Martin are the finest of the accounts available. Furthermore, in contrast to similarly fine-grained work such as [Knott and Dale, 1992], he can draw on a detailed account spanning most of the lexicogrammatical resources that are available in English (described at length syntagmatically in [Halliday, 1985a] and paradigmatically in [Matthiessen, 1992a]), rather than limiting his evidence to cue phrases. He is also the only one to point to the incongruent realizations of the coherence relations, leaving the question of how to decide between a congruent and an incongruent realization open. Since these realizations are actually possible realizations of the coherence relations, it seems reasonable to take them into account. An immediate consequence of this would be to enforce the proposition as the minimal discourse segment instead of the clause. His approach shows how rich the lexicogrammatical potential is, but this richness may be a handicap if it is not well-organized and most of all if it is not integrated into the overall structure of the text to be generated.

But there are some additional features of Martin's proposals that make them stand out for our present purposes. First, there are the relationships drawn to a wider range of surface linguistic forms. Second, there are the relationships drawn to other kinds of semantics (e.g., to 'modulation', 'obligation', 'identification'—some of these we will illustrate briefly below). And third, there are the relationships drawn to context (with external conjunctive relations) and to genre (internal conjunctive relations) [Martin, 1992, p180]. These combined make Martin's account of considerable interest both for an understanding of coherence relations and for the particular goals of the Dandelion project.

3.1.4 Asher, Lascarides and Asher

3.1.4.1 Asher In order to account for discourse structure and abstract entity anaphora within the DRT framework, [Asher, 1993] proposes a new view of discourse updating. The notion of updating in the original DRT-fragment of [Kamp, 1981] is a very simple one. The content of the discourse as a whole is represented by a large DRS which is constructed by the DR-theoretic union of the DRSs for each constituent sentence. Discourse structure is disregarded by this method. To overcome this shortcoming Asher postulates an additional level of structure which is imposed on top of the logical structure. This new structure is called a *segmented DRS* or SDRS. A SDRS in its most simple form consists of a set of DRSs and a set of conditions on those DRSs. The conditions indicate the discourse relations that hold between the DRSs. Asher postulates a set of discourse relations which can be divided into two classes: rhetorical relations and coherence relations. The former are relations over propositions whereas the latter are based on world knowledge. The rhetorical relations used by Asher are: Elaboration, Explanation, Instance, Purpose, Continuation, Consequence, Generalisation and Contrast. An instance of a coherence relation is Cause. Neither the coherence relations nor the rhetorical relations need to be signalled explicitly by a cue phrase. Furthermore, the number of possible discourse relations is infinite.

Some relations, such as Elaboration and Continuation, are based on the notion of *topic*. The principal function of a topic is "*to contribute an important structural component of SDRSs. [] A topic is a proposition that summarizes the content of a constituent in a SDRS*"[Asher, 1993]. Asher distinguishes explicitly and implicitly represented topics. Topic-based discourse relations give a hierarchical character to discourse structure. Rather than introduce simple n place predicates in the SDRS, they create complex conditions in which one constituent is dominated by another. This phenomenon is known as *discourse domination* or *d-domination*. It is comparable to domination in a taxonomy of some domain. The distinction between topic-based (structural) relations and non-structural ones gives rise to two kinds of SDRS update. Before we discuss the update process, we look at the constraints on possible attachment sites for new constituents. The first constraint is the Rule of Relevance: Given an SDRS α for an initial segment of a discourse D, any new constituent derived from further processing of D must be attached to some constituent of α, if D is coherent. Asher defines the notions of *d-freedom* and *openness* which not only determine at which site a new constituent may be attached, but also distinguish between topic-based and non-topic-based relations. A topic-based relation requires an open attachment point, whereas a non-topic-based relation can only be used at an attachment site which is d-free as well as open.

The topic-based update is a complex process in that it it not only affects the SDRS as a whole, but also the topic of the constituent in which the new information is entered. So updating is not a purely additive operation, it may even involve the replacement of a topic in the old SDRS with a new topic. The update operation for non-structural relations is rather straightforward compared to the other one. If one has a discourse relation R, one of whose terms is some element γ of α, and a new constituent β then β is simply added to the sets of constituents in α and the condition $R(\gamma, \beta)$ to the condition list of α.

In addition to topic the notion of *theme* is introduced. Theme plays an important role in the Parallel/Contrast relation. The two relations can only hold if the constituents they relate have a particular semantic structure and content. The required semantic structure is defined in terms of *embedding trees*. An embedding tree of a constituent not only depicts the hierarchical structure of the constituent SDRs but of the DRSs in the SDRs as well. A Parallel relation requires a pairing of nodes in the embedding trees of its constituents such that each pair contains structurally similar and semantically similar objects. A Contrast relation requires a pairing of nodes in the embedding trees of its constituents such that at least some pair contains structurally similar and semantically dissimilar objects. Similarity/dissimilarity is represented by means of a polarity value (+,-) assigned to nodes. Now theme comes into play. Each constituent has a theme. If two constituents have a common theme, then parallelism is maximized (assignment of positive polarity); if the themes are complementary or contraries then contrast is maximized (negative polarity assignment). Furthermore there are explicit linguistic markers indicating a Parallel/ Contrast relation, namely *too* and *but*. If no common theme or two complementary/ contrary theme can be distinguished, then no polarity value is assigned.

Preliminary conclusion

Asher defines his discourse relations in a very tight and formal way. The algoritm for building the discourse structure is equally well formalised. His theory is heavily based on Discourse Representation Theory. The texts which he uses to illustrate how the relations are constructed are all small in size, usually two till four sentences. The sentence structure is fairly simple as well. These small texts already produce a sizeable analysis. It is impracticable to do an analysis of the 'real' texts in our Gulf War article/fundraising letter corpus, since these not only have a more complex sentence structure, but also are many times larger in size.

3.1.4.2 Lascarides and Asher Lascarides and Asher [Lascarides and Asher, 1991] develop a dynamic theory of discourse structure in which a discourse structure is built up through the processing of successive clauses in the text. They concentrate on the temporal discourse relations. Five relations are distinguished: Narration, Background, Result, Explanation and Elaboration. The relations are expressed in the form of defeasible rules. The Result relation rule has the following form

$$(\alpha, \beta) \wedge cause(me(\alpha), me(\beta)) > Result(\alpha, \beta)$$

where the operator $>$ in $\alpha > \beta$ is to be read as "α then β, unless there is information which says otherwise". (α, β) means that the clause α is part of the already existing discourse structure and β is the clause to be attached with some discourse relation. The abbreviation $me(\alpha)$ stands for the main eventuality that is described by α. To give an example, $me(John\ yawned)$ represents the event of John yawning. Each relation is accompanied by an axiom which is *not* defeasible and which contains temporal information. This comes of use in particular when the relation itself provides no clue to the temporal ordering. The axiom on Result has the following form:

$$\Box(Result(\alpha, \beta) \rightarrow me(\alpha) \prec me(\beta))$$

where $\prec$ means "occurs wholly before"

Apart from the discourse relations and their axioms, there are several defeasible laws which express world knowledge (in combination with linguistic knowledge). They can be created if the need arises. For instance, if one wants to say that a boring film can be the cause of yawning then the following law is created:

Yawn Law I:

$(\alpha, \beta) \wedge Info(\alpha, \beta) > cause(me(\beta), me(\alpha))$ where $Info(\alpha, \beta)$ is a gloss for "$me(\alpha)$ is John yawning and $me(\beta)$ is the film being boring" The static classification of *boring*, $Info(\alpha, \beta)$ entails $state(me(\beta))$.

What is needed now is an inference mechanism for inferring the underlying discourse structure of a text. Lascarides and Asher distinguish three patterns of non-monotonic inference:

- **Defeasible Modus Ponens**: if one default rule has its antecedent verified, then its consequent is defeasibly inferred.
- **Penguin Principle**: if there are conflicting default rules that apply and the antecedent of one entails that of the other, then the consequent of the more specific rule (the former one) is defeasibly inferred.
- **Nixon Diamond**: no conclusion is drawn if there are conflicting default rules that apply whose antecedents aren't logically related.

The following example will illustrate how the inference process works. First we'll give the definitions of all relations and axioms.

- **Narration**
  $(\alpha, \beta) > Narration(\alpha, \beta)$
- **Axiom on Narration**
  $\Box(Narration(\alpha, \beta) \rightarrow me(\alpha) \prec me(\beta))$
- **Explanation**
  $(\alpha, \beta) \wedge cause(me(\beta), me(\alpha)) > Explanation(me(\alpha), me(\beta))$
- **Axiom on Explanation**
  $\Box(Explanation(\alpha, \beta) \rightarrow \neg me(\alpha) \prec me(\beta))$

- **Elaboration**
  $(\alpha, \beta) \wedge prep(me(\beta), me(\alpha)) > Elaboration(me(\alpha), me(\beta))$
- **Axiom on Elaboration**
  $\Box(Elaboration(\alpha, \beta) \rightarrow \neg me(\alpha) \prec me(\beta))$
- **States Overlap**[10]
  $(\alpha, \beta) \wedge state(me(\alpha)) > overlap(me(\alpha), me(\beta))$
- **Background**
  $(\alpha, \beta) \wedge overlap(me(\alpha), me(\beta)) > Background(\alpha, \beta)$
- **Axiom on Background**
  $\Box(Narration(\alpha, \beta) \rightarrow overlap(me(\alpha), me(\beta)))$
- **Result**
  $(\alpha, \beta) \wedge cause(me(\alpha), me(\beta)) > Result(me(\alpha), me(\beta))$
- **Axiom on Result**
  $\Box(Result(\alpha, \beta) \rightarrow me(\alpha) \prec me(\beta))$

Now, for the discourse segment:

John yawned. The film was boring.

the antecedents of Narration, States Overlap and Yawn Law I are satisfied. Yawn Law I conflicts with States Overlap and the latter conflicts at its turn with Narration. The antecedent of Yawn Law I entails that of States Overlap which entails that of Narration. The Penguin Principle comes into force and Yawn Law I (having the most specific antecedent) is inferred, yielding $cause(me(\beta), me(\alpha))$. The antecedent of Explanation is now satisfied, but its axiom still has to be satisfied for the relation to apply. The main eventuality of John not yawning (i.e. $\neg me(\alpha)$) has to occur wholly before the film being boring. For the moment we assume that this is the case and consequently the Explanation relation holds. However, if $\neg me(\alpha)$ overlaps with $me(\beta)$ (John only starts yawning halfway through the film) then the axiom prevents Explanation from applying. The next specific rule after Yawn Law I to have its antecedents satisfied is States Overlap and $overlap(me(\alpha), me(\beta))$ is inferred. This conclusion satisfies the antecedents of Background. The corresponding axiom doesn't cause any conflicts and the Background relation holds. The remaining candidate (Narration) can't hold (although its antecedents are satified), because of its axiom which says that John's yawning should occur wholly before the film being boring.

The above example makes clear that the inference mechanism does very well in selecting an appropriate coherence relation between two clauses, but that the temporal restrictions which Lascarides and Asher assume on some of the relations are so strong that they may prevent an otherwise perfect coherence relation from holding. This problem seems to occur in particular when one of the clauses in the discourse segment is a state. To further illustrate this point we will look at a new discourse segment which is the inverse of the previous one.

The film was boring. John yawned.

In order to account for this segment we will have to slightly adapt Yawn Law I.

Yawn Law II:

[10] [Lascarides and Oberlander, 1993] propose two versions of this rule. The first one (presented here) has a stative first clause, whereas in the second version the second clause is stative. We will use both versions.

$(\alpha, \beta) \wedge Info(\alpha, \beta) > cause(me(\alpha), me(\beta))$ where $Info(\alpha, \beta)$ is a gloss for "$me(\alpha)$ is the film being boring and $me(\beta)$ is John yawning" The static classification of *boring*, $Info(\alpha, \beta)$ entails $state(me(\alpha))$.

The discourse segment satisfies the antecedents of Narration, States Overlap and Yawn Law II. The revised Yawn Law conflicts with States Overlap which, at its turn, conflicts with Narration. The antecedent of Yawn Law II entails the antecedent of States Overlap which entails that of Narration. The Penguin Principle comes into force and Yawn Law II (having the most specific antecedent) is inferred, yielding $cause(me(\alpha), me(\beta))$. The antecedent of Result is now satisfied, but the axiom gives rise to a conflict (the film being boring doesn't occur wholly before the yawning: there is overlap). Therefore the Result relation cannot hold. The next specific rule after Yawn Law II to have its antecedents satisfied is States Overlap and $overlap(me(\alpha), me(\beta))$ is inferred. This conclusion satisfies the antecedents of Background. The corresponding axiom doesn't cause any conflicts and the Background relation holds. The remaining candidate (Narration) can't hold (although its antecedents are satified), because of its axiom which says that John's yawning should occur wholly before the film being boring.

One of the shortcomings of the earlier versions of this account were a lack of contact with communicative intentions. This has now been addressed in [Asher and Lascarides, 1994] where axioms and inference rules concerning the intentions of speakers and hearers are added into the general model. There is it argued interestingly that the rules/axioms for reasoning about discourse structure differ in form from those for intentions. The account remains in a very early stage however.

Preliminary conclusion

Lascarides and Asher develop a formal account of temporal discourse relations. They use defeasible rules and laws to encode world knowledge and linguistic knowledge together with an inference mechanism to build up the discourse structure. The definitions of the discourse relations are based entirely on semantics. Although Lascarides and Asher do say that (temporal) information can be gained from syntactic markers and, in [Lascarides and Asher, 1993], show how tense has an effect on discourse attachment, other surface syntactic evidence is sporadically treated. Furthermore, nothing is said about how discourse relations can be linguistically realised (if they are). The examples that are given have all implicit discourse relations. We'll have to wait and see how good the analysis works on larger 'real' texts. The main benefit of the Lascarides and Asher approach seems to lie in the tight formal specification of the discourse relations, which should be of much help in our specification of coherence relations.

3.2 Applying the approaches in analysis

The above review of approaches suggests that there is much partial overlap of information within the various accounts used for coherence relations but all adopt different perspectives which foreground different kinds of information. We now need to assess these differences more finely and on the basis of text analysis. In particular, we ask:

- what points of overlap are there?
- do some accounts have particular strengths or appeal to particular kinds of information that should be maintained in a more comprehensive view?
- do some accounts have particular weaknesses that are more adequately handled by others?

It is common for approaches to coherence relations to work within a single perspective. A further motivation for considering a range of perspectives in unison—as we do here—is to establish whether a broadening of perspectives improves our ability to investigate text. There have been several attempts to combine the disparate approaches to coherence relations in order to produce a single kind of description. This has been attempted in a very direct way by, for example, [Jordan, 1988] and [Hovy, 1990, Maier and Hovy, 1993], where coherence relations from a variety of sources are combined into a single set. It is also been assumed to be possible implicitly in work where the notion of 'coherence relation' is taken as a given. We seek to establish whether such a simple combination method is applicable, or whether more complex relations between the information foregrounded in distinct approaches are required; this will be of particular concern in Section 4 below. In particular, our analyses attempt to establish if there are areas of descriptive work which are in fact duplicated across accounts or which serve to stretch an account beyond its natural borders of application.

Our method during the first phase was to take texts primarily from the Dandelion Gulf War corpus and the fund-seeking letter corpus used in [Abelen *et al.*, 1993] and analyze these according to the principal perspectives outlined in Section 3.1 and as justified by the critical evaluations given. We focused on problems arising in the performance of the analyses as well as on the analysis results themselves. Indeed, sometimes we were not able to complete a satisfactory analysis according to some mode of investigation—which is itself an interesting result which we evaluate. The problems encountered were investigated to reveal their most likely sources and a specific evaluation of each mode of analysis and its likely combination with other modes was prepared.

We illustrate the kind of analysis in detail[11] by showing what happens when the modes of analysis adopted for investigation are applied in the analyses of two selected texts. These texts were picked since

- they are short enough for the analyses to be shown in full,
- they show interesting, although in no way atypical, phenomena,
- and they belong to two distinct genres or text types.

Both texts are from the Dandelion Gulf War corpus. Text I is taken from the *Glasgow Herald* (city edition; 17.1.91); Text II is taken from the *Daily Record* (16.1.91).

3.2.1 Analyses of Text I

Text I is shown below, divided into numbered units for purposes of analysis presentation.

[11] Note that since our initial aim in this deliverable has become to improve our understanding of coherence relations as such, we were not able to pick out particular conjunctions or linguistic phenomena, such as tense or a particular (set of) cue phrase(s), for a focused analysis. We needed to construct analyses of texts as a whole relying on as much linguistic detail as possible. This is naturally time-consuming and reduces the quantity of text analysis that it is possible to carry out. However, the degree of detail of the individual analyses is hoped to counterbalance the reduced range of texts analysed.

"Round the clock air bombardment"

(1)	First strikes launched by allied fighter bombers and high level American B52 strategic bombers appear to be directed against the 16 major air bases in Iraq and two in occupied Kuwait, chemical weapons plants, Scud B missile sites, and command centres in and around Baghdad.
(2)	The waves of attackers were preceded early today by special "Wild Weasel" war planes designed to suppress anti-aircraft fire by launching missiles which home on emissions from ground control radars.
(3a)	The US has 1800 aircraft in the area
(3b)	including 450 strike bombers on six aircraft carriers.
(4)	Britain has three squadrons of Tornado ground attack aircraft at Dhahran and Tabuk in Saudi Arabia and Bahrain and a further 12 Jaguar low level bombers at Thumraid in Oman.
(5)	Some, or all of these, are believed to be involved in the assault.
(6)	The bombing offensive, codenamed Operation Desert Storm, will go on around the clock for the next few days in the most concentrated air action since World War II.
(7a)	The aim is to cripple Iraq's own air power on the ground
(7b)	before it can launch counter attacks
(7c)	and to try to eliminate the threat of chemical attacks from long range missiles.
(8)	There are unconfirmed reports of "hostile action" around the main allied press centre at Dhahran.
(9a)	A few minutes ago the centre went off the air
(9b)	and cannot be contacted.

3.2.1.1 Lascarides/Asher The analysis in terms of Lascarides and Asher's account of discourse relations is shown in graphical form in Figure 6. The motivation for this analysis is as follows.

1-2 Elaboration

The Narration relation is ruled out because of the axiom $(me(\alpha) \prec me(\beta))$. Unit 2 is not the cause of unit 1 and vice versa, so Explanation and Result do not apply. There are no states overlapping and therefore the Background relation doesn't hold. However it is possible to consider the event in unit 2 as being in the preparatory phase of unit 1, so the Elaboration relation is possible. The axiom on Elaboration is satisfied as well $(\neg me(\alpha) \prec me(\beta))$.

1-3 Background

Unit 3 and unit 2 have no coherence relation. There is a topic shift which closes off unit 2 as an available attachment point for unit 3. The only remaining candidate for attachment is now unit 1. Since the state expressed by unit 3 overlaps with the event in unit 1, States Overlap applies and the following relations are ruled out: Narration, Explanation, Elaboration and Result. The Background relation is triggered by States Overlap and its axiom $(\Box(Narration(\alpha,\beta) \rightarrow overlap(me(\alpha), me(\beta))))$ doesn't give rise to any conflict, so this relation applies.

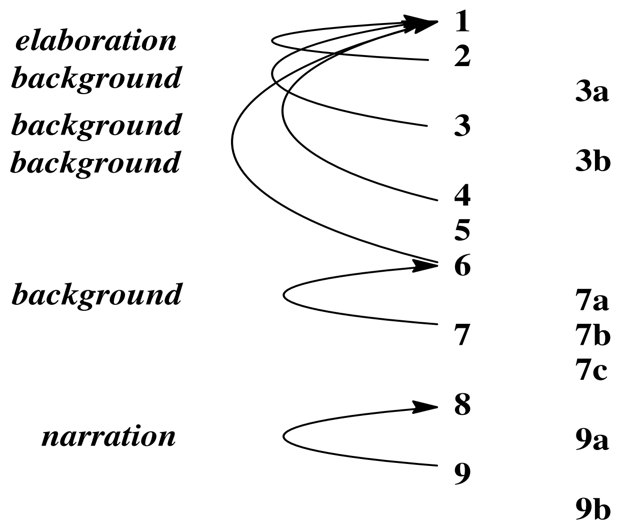

Figure 6: Lascarides/Asher-style analysis of Text I

1-4 Background

> Unit 3 and 4 have the same topic (fighter aircraft). They are related to one another by Continuation [Lascarides and Asher, 1993, p252]. Consequently, unit 4 is not only connected to the same unit as unit 3 (i.e. unit 1), but also the same coherence relation holds (Continuing Discourse Patterns). For more information on topic continuity see Section 3.1.4.1.

There is here a nice ambiguity in where to attach unit 5; possibilities from the discourse structure include units 4, 3 and 1. This would presumably be a good example of the need for interaction between discourse structure reasoning and informational reasoning as argued and demonstrated for some very simple cases in [Asher and Lascarides, 1994]. It is unlikely that all of the 1800 US aircraft plus the three squadrons of British aircraft were involved in the first strikes mentioned. To this we can add the fact that the text was produced in a British newspaper, which raises the likelihood that British involvement might receive more salience. Thus the most plausible attachment point would have to be 4, although the mechanisms to achieve this are certainly complex.

1-6 Background

> It is not possible to attach unit 6 to unit 5 because of the difference in topic. An available attachment point is unit 1. Unit 6 can be seen as the expression of a state which overlaps with the event in unit 1. The antecedent of States Overlap is satisfied and consequently also Background. Narration is ruled out because of its axiom. There is no causal relation between units 1 and 6, so Explanation and result do not apply. Unit 6 is not in the preparatory phase of unit 1 either and therefore Elaboration is ruled out as well.

6-7 Background

> The state expressed by unit 7 overlaps with the event in unit 6. States Overlap and consequently Background is satisfied. There is no conflict with the axiom on Background.

8-9 Narration

> Unit 8 and 9 are not in any causal relationship, which excludes Explanation and Result. Unit 9 is not in the preparatory phase of unit 8, so Elaboration can not be triggered. There is no overlap between 8 and 9, which rules out Background. The only remaining relation is Narration. The axiom causes no problems and the relation holds.

The Lascarides and Asher analysis of text I does not provide a wide range of linguistic information. The prevailing coherence relation is Background. The cause for this prevalence may very well lie in a too rigid formulation of the axioms, as already has been pointed out in the Lascarides and Asher overview. We will return to this when analysing text II. The merit of the analysis doesn't lie in the coherence relations it yields, but in the tight and formal way the relations are defined.

3.2.1.2 Martin The analysis of Text I in terms of Conjunctive relations is shown graphically in Figure 7, following the form of presentation proposed by [Martin, 1992]. Relations to the left of the figure are internal; those to the right are external.

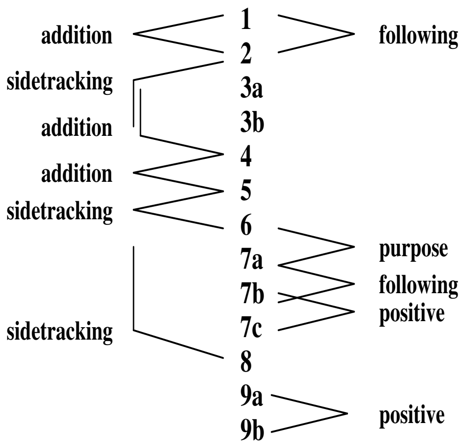

Figure 7: Conjunctive relations analysis of Text I

The motivations for this analysis are as follows.

Within unit 1 several nominal groups are tied together by an additive coherence relation. There are four of them and they all can be classified as follows:

[conjunction:additive:external:addition:positive]
with realisation: *and*, parataxis,explicit.

Unit 2 is connected to unit 1 by a temporal coherence relation which is realised metaphorically by *preceded*. Its hypotactic counterpart would be *after*.

[conjunction:temporal:external:successive:following:anterior dependent:non deictic]
with realisation: *preceded*, within process,explicit.

In unit 2 we also find a description of how (the manner in which) the "Wild Weasel" air planes strike. this is realised by non-finite hypotaxis.

[conjunction:consequential:external:manner]
with realisation: *by*, hypotaxis non-finite,explicit.

Within unit 4 four nominal groups are paired by two additive coherence relations, which can be classified as follows:

[conjunction:additive:external:addition:positive]
with realisation: *and*, parataxis,explicit.

Two pronominal groups in unit 5 are combined by an additive coherence relation.

[conjunction:additive:external:alternation]
with realisation: *or*, parataxis,explicit.

Unit 6 is connected to unit 7 by a metaphorically realised coherence relation. Although not explicitly listed in Martin's networks, its function corresponds to that of *in order to*. In fact, the replacement of the metaphor by its hypotactic counterpart results in the finite hypotactic realisation of the coherence relation.

[conjunction:consequential:external:cause:contingency:purpose:desire:modality certain]
with realisation: *The aim is to*, within process,explicit.

The three subparts of unit 7 are held together by explicit conjunctive relations. the first one is realised hypotactically and connects 7a and 7b. At its turn 7ab is connected to 7c by explicit parataxis (*and*).

[conjunction:temporal:external:successive:following:posterior dependent]
with realisation: *before*, hypotaxis finite,explicit.

[conjunction:additive:external:addition:positive]
with realisation: *and*, parataxis,explicit.

Units 9a and 9b are explicitly connected.

[conjunction:additive:external:addition:positive]
with realisation: *and*, parataxis,explicit.

We can also illustrate here two further very common sources of evidence for conjunctive relations—particularly for 'implicit' (as far as a surface conjunction is concerned) internal relations. Internal relations refer to the organization of the text—this organization is also carried by thematic progressions which thus often correspond with internal conjunction relation configurations (cf. [Martin, 1992, 436]). In addition, particular thematic patterns will coincide with external conjunctive relations such as contrast (cf. [Martin, 1992, p441]). Text organization can also be indicated by drawing explicit connections between the content expressed in a text, most commonly by the use of the semantic resources of *identification* [Martin, 1992, Chapter 3]: e.g., *the last chapter, for that reason*, etc. Both techniques are important for the current text since there are no 'explicitly' realized internal conjunctive relations.[12]

The internal structure of Text I can be quite simply revealed by examining the relation between themes and ideational content. The themes are as follows:

First strikes
The waves of attackers
The US
Britain
Some, or all of these,
The bombing offensive, codenamed Operation Desert Storm,
The aim
There

[12] Another good indicator is sometimes text layout: this is particularly important for the analysis of example Text II, but we will not go into this in any detail here.

A few minutes ago (the centre)

The structure of the information underlying this text is mostly centred on a 'bombing offensive' which has aims, first waves, participants, and planes of those participants. The themes do not follow any single line of informational connection, however. This can be indicated by the following indentation, where horizontal positioning shows approximate ideational constituency.

	First strikes involving aircraft		$\boxed{1}$
		Preparatory strikes involving aircraft	
	Participation		
		US are participants (and have aircraft)	
		Britain are participants (and have aircraft)	$\boxed{1}$
Bombing offensive			
	Has aim		
	Has complications		

Ideational 'discontinuities' in thematic progression thus occur whenever the horizontal indentation moves to the left (rather than staying at the same level of indentation—reminiscent of [Daneš, 1974]'s continuing thematic progression—or moving to the right—as in developing thematic progression). There are, then, indications of possible 'staging' internal conjunctive relations between 'Prepatory strikes' (2) and 'Participation' (3a) and between 'Britain are participants' (4-5) and 'Bombing offensive' (6) as shown in Figure 7. Clearly other evidence can also be brought to bear here, including the tense shifts and anaphor. Theme is often particularly effective, however.

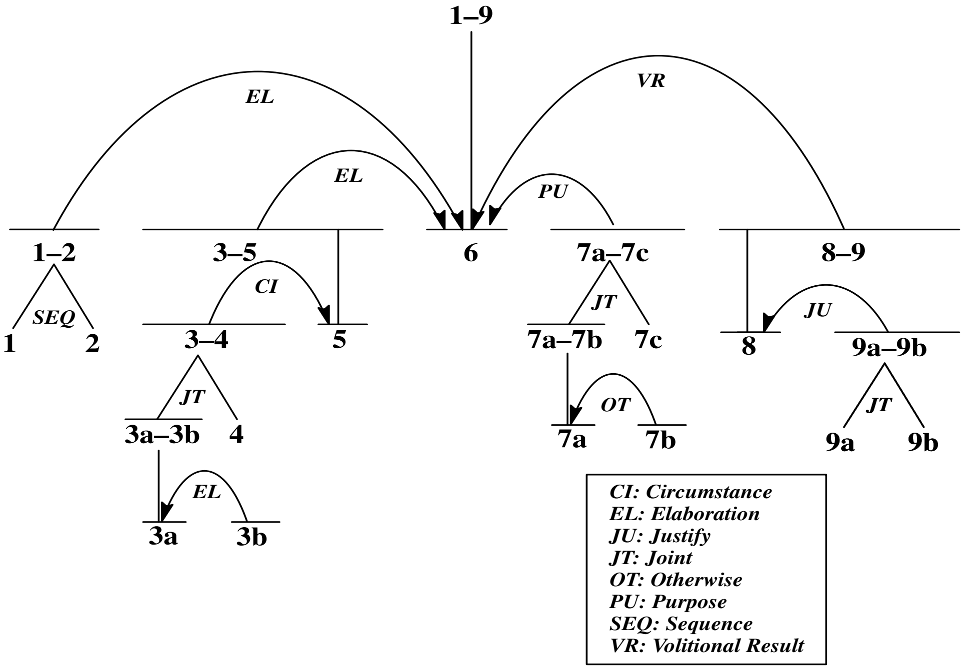

Figure 8: Rhetorical Structure Theory analysis of Text I

3.2.1.3 Mann and Thompson 1(Nucleus)-2(Nucleus) Sequence

There is a succession relationship between unit 1 and unit 2, which is linguistically realised by "preceded", i.e. the action in unit 2 occurs before the action in unit 1.

3a(Nucleus)-3b(Satellite) Elaboration

Unit 3b provides additional information for unit 3a.

7a(Nucleus)-7b(Satellite) Otherwise

In unit 7a there is an unrealised situation (Iraq's air power is crippled on the ground), as well as in unit 7b (Iraq's air power launches counter attacks). The realisation of the situation in the nucleus prevents the realisation of the situation in the satellite.

9a(Nucleus)-9b(Nucleus) Joint

Units 9a and 9b share the same subject matter (the loss of contact with the press centre) and they are combined by an explicit conjunction *and*.

3(Nucleus)-4(Nucleus) Joint

Units 3 and 4 share the same subject matter (the number of airplanes in the area). There is no conjunction present, but the two units can easily be combined by an implicit *and*.

3,4(Satellite)-5(Nucleus) Circumstance

The framework (the total number of aircraft) within which the reader is to interpret the nucleus 5 (the airplanes involved in the attack) is set by unit 3,4.

7ab(Nucleus)-7c(Nucleus) Joint

Unit 7ab and unit 7c share the same subject matter (the aim of the allied attack) and they are combined by an explicit conjunction *and*.

8(Nucleus)-9(Satellite) Justify

The relation between unit 8 and unit 9 may be considered as an Elaboration, in which the satellite 9 provides extra information for the nucleus 8. There is a more attractive alternative:[13] Justify. The writer uses unit 9 to increase the reader's readiness to accept the writer's right to present 8. In 9 he says that this is a very hot newsitem ("a few minutes ago") and that the people in the press centre cannot be contacted in order to confirm or deny the story. In this way the reader will accept his vague report.

1,2(Satellite)-6(Nucleus) Elaboration

The satellite provides additional information for the nucleus. This relation is not totally satisfactory, since the nucleus does not only mention the 24 hours bombardment of which the first attack was part, but also announces what will happen next. Unfortunately it is the best option available.

3,5(Satellite)-6(Nucleus) Elaboration

[13] personal correspondence with Gisela Redeker

Unit 3,5 gives extra information about the nucleus 6.

6(Nucleus)-7(Satellite) Purpose

The satellite presents an unrealised situation (Iraq's air power crippled on the ground and the threat of chemical attacks by long range missiles eliminated) which is to be realised through the activity in the nucleus (the round the clock air bombardment).

6(Nucleus)-8,9(Satellite) Volitional Result

The unit 8,9 presents a situation which could have arisen from a volitional action (the hostilities around the allied press centre) and the situation in the nucleus (the round the clock air bombardment) could be the cause for the situation in the satellite.

3.2.1.4 Hobbs In order to carry out a Hobbs-style analysis it is necessary to find out what is being 'asserted' by the sentences of the text. As Hobbs notes [Hobbs, 1990, p103], this can be problematic but in the present cases we assume rather straightforwardly that the sentences and their assertional content are transparent. Hobbs suggests a general method for carrying out a coherence analysis which is almost identical to that for RST [Hobbs, 1990, p111]. The steps are:

1. Recursively segment the text, top-down, in terms of its perceived major breaks.
2. Label the non-terminal nodes of the resulting tree, bottom-up, with coherence relations; here Hobbs suggests a method also suggested for conjunctive relations—i.e., inserting conjunctions in order to gauge the appropriateness of a corresponding coherence relation.
3. Make explicit the knowledge or beliefs supporting the assignment of coherence relations.
4. Validate the assumptions made in step 3 (by considering a wider corpus, etc.).

Applying this procedure to Text I runs into two problems that are probably relatable to the text type of that text. It is well known that newspaper articles, particularly 'news' articles such as the present text (and in contrast to the editorial text that we use as our second example), have a rather different kind of text structure than, for example, narratives. The text in fact has quite a normal structure for a news article—the basic news item is given, and then further details and background information concerning this item are given in several further cycles. This organization was, if anything, obscured by the Lascarides/Asher analysis and, at best, somewhat distorted by the RST analysis (which ended up placing the central weight on sentence (6)). Only the internal conjunctive analysis of martin provided some support.

First, there are very few relations that actually apply: virtually all segments are related by *elaboration*; second, the highly recursive structure of the text is perhaps not intuitively so appealing. The overall structure would be:

elab (elab (elab (elab (1, 2), ..), 6), 7)

Repeating Hobbs' definition of *elaboration* from above, the relation is defined to hold between two discourse segments S_0 and S_1 when we can:

Infer the same proposition P from the assertions of S_0 and S_1. [Hobbs, 1990, p95]

Hobbs also suggests that it is this inferred proposition P that then forms the semantics of the combined discourse unit. [Hobbs, 1990, p104]. In this case, the proposition is something rather general, such as 'there is an air bombardment in progress'. Although the links to this are in some cases rather weak (in fact, taking the inferrence allowed to be rather narrow, then the only thing inferrable would be something more like 'there are some aircraft'). In addition to these elaborations, the following relations are also suportable.

CR1: `parallel (3, 4).`
because we can

> infer $p(a_1, a_2, \ldots)$ (= 'has aircraft in the area (US)') from the assertion of S_0 (3 = 'US has 1800 aircraft in the area) and $p(b_1, b_2, \ldots)$ (= 'has aircraft in the area (Britain)') from the assertion of S_1 (4 = 'Britain has three squadrons...'), where a_i and b_i are similar, for all i. [Hobbs, 1990, p93]

CR2: `elaboration (4, 5).`
because we can

> infer the same proposition P (4 = 'Britain has aircraft in the area) from the assertions of S_0 ('Britain has aircraft in the area) and S_1 (5 = 'Some, or all of these, are believed to be involved in the assault.).

CR3: `explanation (8, 9).`
because we can

> infer that the state or event asserted by S_1 (9 = 'the centre went off the air') causes or could cause the state or event asserted by S_0 (8 = 'there are unconfirmed reports'). [Hobbs, 1990, p91]

And, also, following the RST analysis, it might be possible to argue that there is a

- `cause (1&2, CR3)`, or an
- `event-outcome (1&2, CR3)`.

relation holding if somehow access to the 'hostile action' embedded under (8) can be secured.

If we relax the suggestion that a single tree span the entire text, then a slightly more accurate coherence picture results. This is given in Figure 9. Here the highly recursive structure is broken up into a set of elaboration relationships, each elaborating further in some particular way the main news element of segment (1). The subsequence segments of the text are also marked for illustrative purposes in the figure.

3.2.2 Analyses of Text II

Text II is shown here broken into approximately into clause-size segments which are numbered to facilitate the presentation of the analysis results below. Some analyses require a finer granularity; this is introduced as necessary. In this case, in order to make the interpretation presented in our analyses clearer we also present in Figure 10 a version of the text showing more of its original formatting, which is quite revealing concerning the intended importance of various parts of the text.

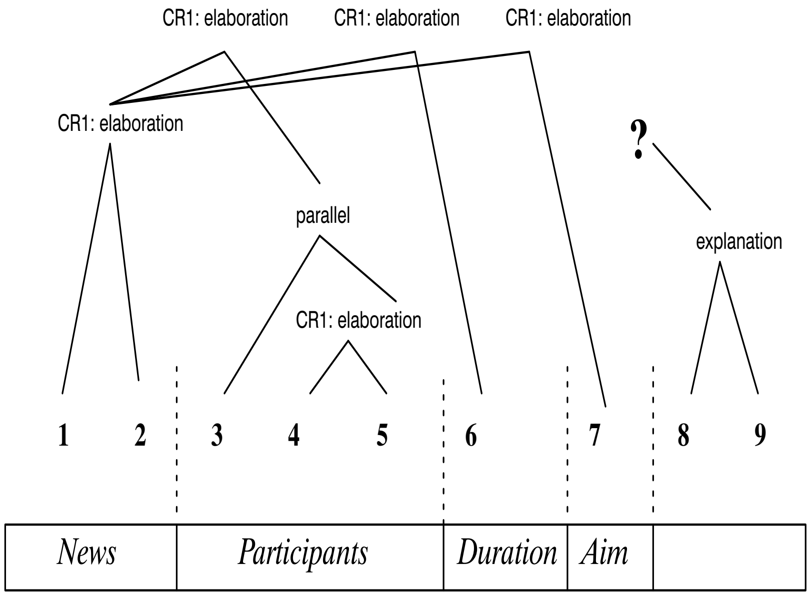

Figure 9: Hobbs-style analysis of Text I

"Stand United"

(1)	No one wants war.
(2a)	Everyone, <2b> is dreading the conflict.
(2b)	<apart from Saddam Hussein,>
(3)	And we appreciate the sincerity of those MPs who still believe sanctions instead of shells will do the job.
(4a)	But those who rebelled went too far
(4b)	by voting, in effect, against the use of force.
(5)	They failed to send the message of full-hearted and united support to our fighting forces, now facing a fearsome task.
(6)	They failed to send the message to Saddam Hussein that tyrants will be opposed – not appeased.
(7a)	And he will make use of this
(7b)	as he did with the weekend peace protests which were shown on Iraqi TV as pro-Saddam propaganda.
(8a)	As Parliament debated,
(8b)	Baghdad was rejecting even the very soft, appeasing French initiative
(8c)	and defiantly declaring that Kuwait was – and would remain – the 19th. province of Iraq.
(9a)	For five months we've sought peace
(9b)	and waited for the Iraqi leader to see sense.
(10a)	Whatever happens now
(10b)	there's only one man to blame –
(10c)	Saddam Hussein.

Stand united

No one WANTS war. Everyone, apart from Saddam Hussein, is dreading the conflict.

And we appreciate the sincerity of those MPs who still believe sanctions instead of shells will do the job. But those who rebelled went too far by voting, in effect, against the use of force.

- **THEY FAILED** to send the message of full-hearted and united support to our fighting forces, now facing a fearsome task.
- **THEY FAILED** to send the message to Saddam Hussein that tyrants will be opposed – not appeased.

And he will make use of this as he did with the weekend peace protests which were shown on Iraqi TV as pro-Saddam propaganda.

As Parliament debated, Baghdad was rejecting even the very soft, appeasing French initiative and defiantly declaring that Kuwait was – and would remain – the 19th. province of Iraq.

For five months we've sought peace and waited for the Iraqi leader to see sense.

Whatever happens now there's only one man to blame – Saddam Hussein.

Figure 10: Original formatting of Text II

3.2.2.1 Lascarides/Asher The analysis in terms of Lascarides and Asher's account of discourse relations is shown in graphical form in Figure 11. The motivation for this analysis is as follows.

1-2 Background

> The states in 1 and 2 overlap, therefore Narration is ruled out (Axiom: $me(\alpha) \prec me(\beta)$). The state in 2 did not cause the state in 1, therefore Explanation is also ruled out. The state 2 is not in the preparatory phase of state 1, so Elaboration is not satisfied. The state in 1 didn't cause the state in 2, thus ruling out the Result relation. Only Background has its preconditions satisfied. This example illustrates how the strong temporal restrictions which Lascarides and Asher place on their coherence relations, prevent (more appropriate) relations from holding. In this case[14], if the Elaboration relation were to be loosened to some inclusion relation between the two eventualities (and not one being in the preparatory phase of the other), then the relation between unit 1 and 2 would be an instance of this coherence relation. Hobbs' approach is more convincing here as we will see later on (Section 3.2.2.4).

2-3 Background

[14] As pointed out by Leonoor Oversteegen in personal correspondence

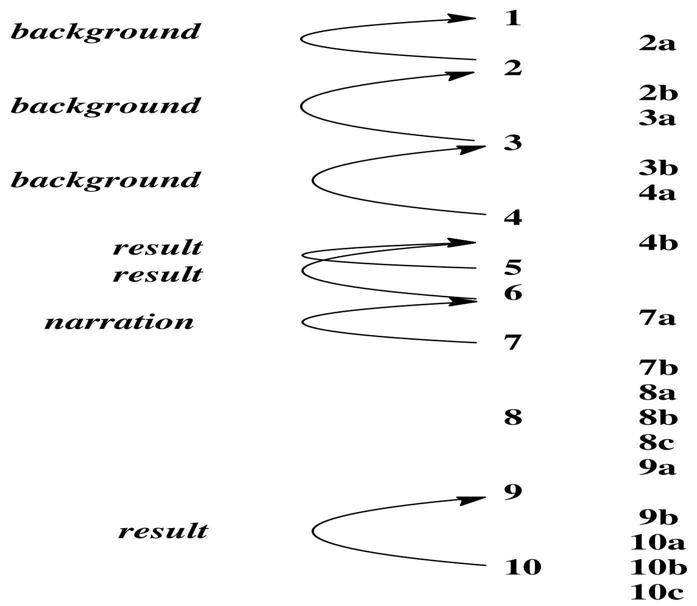

Figure 11: Lascarides/Asher-style analysis of Text II

There is also an overlap between states 2 and 3, which rules out the Narration relation (Axiom: $me(\alpha) \prec me(\beta)$). The state in 3 is not the cause of the state in 2, therefore Explanation is ruled out. The Elaboration relation can not hold because state 3 is not in the preparatory phase of state 2. The state in 2 didn't cause the state in 3, so the antecedent of Result is not satisfied. Again, Background is the only relation which has its preconditions satisfied.

3-4 Background/Contrast

The main eventualities of 3 (appreciation of the sincerity) and 4 (disapproval of rebellion) overlap and thus rule out the possibility of a Narration relation (Axiom: $me(\alpha) \prec me(\beta)$). The state in 3 is not the cause of the state in 4 and vice versa, so Result and Explanation do not apply. Elaboration doesn't hold either, since the main eventuality of 3 is not in the preparatory phase of the main eventuality of 4. The only relation to have its antecedents satisfied is Background. In [Asher, 1993] some coherence relations additional to those in [Lascarides and Asher, 1991] are presented. They are also presented in the overview (Section 3.1.4.1). One of them, Contrast, can be applied here. The themes expressed by both units are opposites (appreciation vs. approval). Furthermore, there is the linguistic marker "But" which is considered a clue to the presence of a contrast relation [Asher, 1993, p285]. The Contrast relation does not impose a temporal order on the constituents, so it is compatible with Background.

4b-5 Result

The failure of the MP's to send a message of united support to the fighting forces (main eventuality 5) is caused by their vote against the use of force (main eventuality 4b). So $cause(me(\alpha), me(\beta))$ holds and it, being the antecedent of Result, causes this relation to be triggered. Its axiom does not give rise to any conflict ($me(\alpha) \prec me(\beta)$), since 4b occurs wholly before 5.

4b-6 Result

> It is clear that there is topic continuity (the failure of the MP's) between the main eventualities of 5 and 6. The continuity is accentuated by the syntactic parallelism between 5 and 6. Therefore we can conclude that a Continuation relation holds between the two units. On the basis of the Continuing Discourse Patterns rule [Lascarides and Asher, 1993, p252] units 5 and 6 should be connected to 4b by the same coherence relation. It has already been established that a Result relation holds between 4b and 4, so 4b and 6 should have a Result relation as well.

6-7 Narration

> The Narration relation holds between 7 and 6. The axiom on Narration holds, since the failure of the MP's to send the message to Saddam ($me(6)$) occurs wholly before Saddam using their failure ($me(7)$). There is no overlap between the eventualities, so Background is ruled out. Result and Explanation do not apply since 6 is not the cause of 7 and vice versa. The main eventuality of 7 is not in the preparatory phase of 6 which excludes the Elaboration relation.

9-10 Result

> The main eventuality in 9 is the period of five months the allies have waited for the Iraqi leader to see sense. This is the cause for Saddam Hussein being the man to blame for whatever happens next (main eventuality 10). Since 9 occurs wholly before 10, the axiom on Result is satisfied and the Result relation holds.

Again, although we see more information being obtained for this text, the general result is similar as for Text I. A lot of linguistic information is not being picked up. For instance, the relation between (1) and (2) is very different to that between (2) and (3), yet both receive a 'background' relation. We have shown that the cause of this lies in too strong temporal restrictions on the coherence relations. In comparison to the other modes of analysis suggested, there is therefore rather little detail. The analysis does not cover all the text; and there is not very much differentiation in the relations found to apply (i.e., lots of backgrounds).

A further issue, which is particularly important for text generation but which also has import for analysis, is the question of the *motivation* of the use of particular relations. That is, 'background' may hold: but *why* was a background selected for the text at that particular place? Issues of this nature are crucial for a more complete specification of coherence relations. Some relevant considerations of this question within this framework have now been begun in, for example, [Asher and Lascarides, 1994]; we will return to these considerations below in our general discussion.

3.2.2.2 Martin The analysis of Text II in terms of Conjunctive relations is shown graphically in Figure 12, again following the form of presentation proposed by [Martin, 1992]. Relations to the left of the figure are internal; those to the right are external.

The motivations for this analysis are as follows.

Clause 2 is a reformulation of clause 1 with some extra information included which adjusts and corrects the statement in clause 1: nobody wants war, but there is one exception. The conjunctive relation holding is therefore one classifiable as follows according to Martin's networks:

[conjunction:comparative:internal:similarity:reformulation:adjust:correct]
with realisation: *in fact*, cohesion, implicit

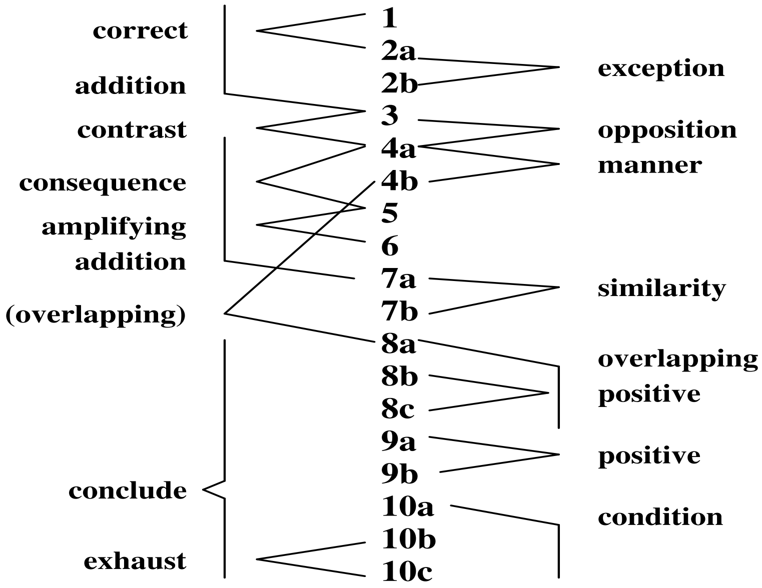

Figure 12: Conjunctive relations analysis of Text II

A congruent explicit realization of this relation is *in fact*, which could have been inserted in the text without substantially changing the textual development. However, there is no explicit cohesive tie. The actual realization of the relation in the text is implicit via cohesion.

Clause 2a and 2b are in contrast with one another: nobody wants war, but there is one person who does; he is the exception.

[conjunction:comparative:external:contrast:exception]
with realisation: *apart from*, hypotaxis non-finite, explicit

Here the center of attention shifts from the no war state of mind (1 and 2) to the members of parliament (3). The linguistic marker *and* isn't present in this place in Martin's network, but plays the same role as for instance *in addition* which does occur here.

[conjunction:additive:internal:turnbuilding:developing:addition:simple]
with realisation: *and*, cohesion, explicit

The conjunctive relation holds between two nominalisations (sanctions and shells). The use of sanctions is in contrast with the use of force.

[conjunction:comparative:external:contrast:replacement]
with realisation: *instead of*, explicit

Between (3) and (4), there is a contrast and opposition between the appreciation for the sincerity and the disapproval of the vote against use of force. This can work both internally and externally, which, as Martin notes, is distinctive of this kind of text—i.e., an editorial—where the text is being actively used to 'create' the world talked about. The realizational example offered by Martin for this latter case is: 'on the other hand'.[15]

[15] The relation [consequential: internal: concessive: contra-expectation] (with the possible realization 'nonetheless') may also be a candidate here (cf. [Martin, 1992, p227]).

[conjunction:comparative:external:contrast:opposition]
[conjunction:difference:internal:opposition:contrast]
with realisation: *but*, cohesion, explicit.

Between (4a) and (4b) the relation is straightforward:

[conjunction:consequential:external:manner]
with realisation: *by*, hypotaxis non-finite, explicit.

There is also clearly a conjunctive relation between 4 and 5. It is an internal consequential relation since it can best be paraphrased by a sentence as: 'I say X because'; if the relation were to be made explicit possible cue phrases might include 'for example', 'in particular', 'indeed', etc.: all drawn from the [conjunction:comparative:internal:similarity:reformulation] part of the network. While within (5), the following relation holds:

[conjunction:temporal:external:successive:following:anterior dependent:deictic:extension from present]
with realisation: *now*, explicit

The word *now* is not present in the above class, but *from now on* is. The troops face a fearsome task starting as of the writing/publication of the editorial or as of the decision being taken.[16]

One possible reading of the text is where of the two "failures" (5,6), the latter is the most serious. This is compatible with a central meaning in the text being that everybody wants to avoid war and that in some cases war is a necessary evil to stop a tyrant. The implicit conjunction stresses the fact that tyrants must be opposed and that war is then justified. This gives the relation:

[conjunction:additive:internal:turnbuilding:developing:addition:cumulative:amplifying]
with realisation: *moreover*, cohesion, implicit

Within (6) there is an implicit coherence relation which can be paraphrased as "tyrants will not be appeased but opposed."[17]

[conjunction:comparative:external:contrast:replacement]
with realisation: *but*, parataxis, implicit

Between (6) and (7), the following explicit relation holds:

[conjunction:additive:internal:turnbuilding:developing:addition:simple]
with realisation: *and*, cohesion, explicit.

And between (7a) and (7b), there is a similarity expressed between the use of the weekend peace protests for propaganda purposes and the possible use of the no vote.

[conjunction:comparative:external:similarity]
with realisation: *as*, hypotaxis finite, explicit

In (8a) and (8bc), there is overlapping between the debate in Parliament and the rejection of the French initiative by Baghdad. Hence, the relation that holds is:

[conjunction:temporal:external:simultaneous:overlapping:durative:anterior dependent:non stative]
with realisation: *as*, hypotaxis finite,explicit.

It can also be argued that there is overlapping between (8a) which describes the debate in Parliament and (4b) which is about the vote in Parliament. As it seems that this relation is dependent on how much you want to interpret the text, it has been put in brackets in Figure 12.

Between (8b) and (8c), and also between (9a) and (9b), a straightforward external relation holds:

[16] A further interpretation is that the 'now' simply refers to an extended present.
[17] As suggested by Wilbert Spooren in personal correspondence

[conjunction:additive:external:addition:positive]
with realisation: *and*, parataxis, explicit

These lead up to the conclusion of the text, which is captured by an implicit conjunction connecting (9) and (10). (9) expresses the idea that everything possible has been done to make the Iraqi leader sees sense over a long period of time. As a consequence he is the only one to blame if war breaks out:

[conjunction:consequential:internal:non concessive:cause]
with realisation: *so*, cohesion, implicit.

Furthermore, within (10) the following relation holds between (10a) and (10b):

[conjunction:consequential:external:cause:contingency:condition:+concession:universal]
with realisation: *whatever*, parataxis, explicit

And within (10b), although *now* is not in the realizations mentioned by Martin, its function corresponds to *from now on*. The now is the time of writing/publication of the editorial. Hence:

[conjunction:temporal:external:successive:following:anterior dependent:deictic:extension from present]
with realisation: *now*, explicit

Finally, (10c) is a reformulation of (10b), and the two are connected by an implicit conjunction. It's not a reformulation in the sense of a correction, but more of a clarification. Hence:

[conjunction:comparative:internal:similarity:reformulation:rework:abstraction:exhaust]
with realisation: *that is*, cohesion, implicit.

One advantage of this approach over that of Lascarides and Asher is that there is usually something to say; there is sufficient 'density' of coverage that some relations will fit. As we also saw with respect to Text I, particularly with respect to the internal relations, this also appears to offer some useful segmentations of the text. If we again consider the thematic elements and tenses of the text as additional evidence for internal relations, we have the pattern shown in Figure 13.

Theme	*Tense*
No one	present
Everyone	present progressive
And we	present
But those	past
They	past
They	past
And he	future
As Parliament debated,	past progressive
For 5 months (we)	perfect
Whatever happens now, (only one man)	present

Figure 13: Text II: themes and tenses

This supported the assignment of internal relations and scopes of those relations such as:

- the addition relating (1-2) and (3), because of tense continuation;
- the range of the addition relating (4-6) and (7), because (4-6) are past and theme continuing and (7) switches tense and theme;

- the range of the internal conclusion leading from (8a) up to (10), due to temporal themes and tense moving towards present.

One clear disadvantage, from a computational perspective, remains the informality of the application criteria.

3.2.2.3 Mann and Thompson The RST-style analysis of Text II is shown graphically in Figure 14. The motivation for this analysis is as follows.

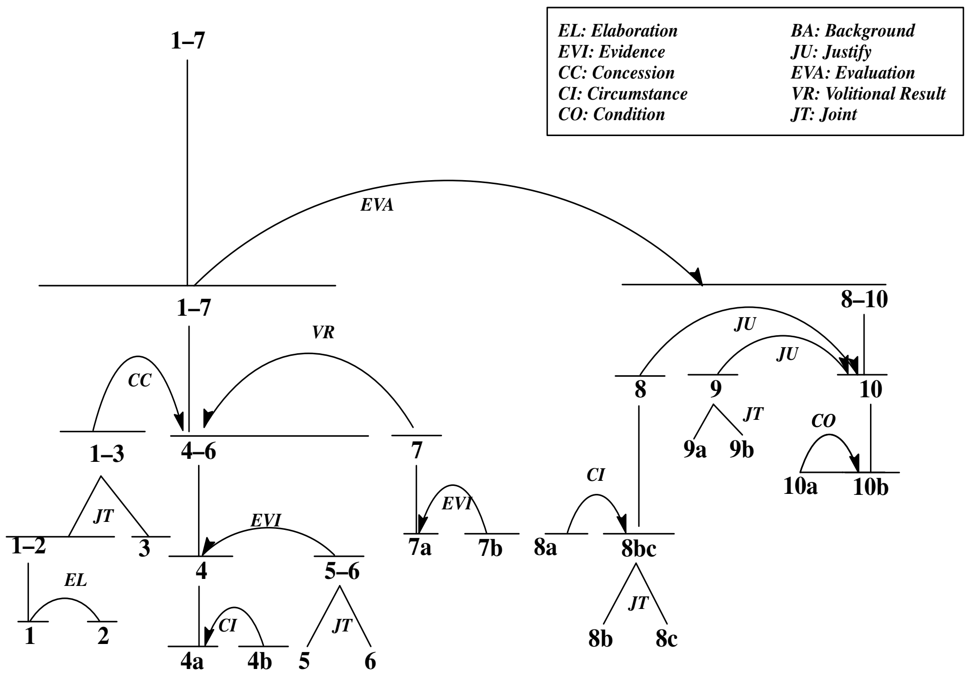

Figure 14: Rhetorical Structure Theory analysis of Text II

1(Nucleus)-2(Satellite) Elaboration

Unit 2 presents extra information (i.e. Saddam Hussein does want war) on the situation expressed in unit 1.

4a(Nucleus)-4b(Satellite) Circumstance

Unit 4b (the MP's voted against the use of force) sets the framework within which the reader is to interpret the nucleus (4a: they went too far).

5(Nucleus)-6(Nucleus) Joint

Units 5 and 6 not only have the same structure (accentuated by the typography), but also share the same subject matter (the 'failures' of the MP's who voted against the use of force).

7a(Nucleus)-7b(Satellite) Evidence

The presentation of the fact that Saddam Hussein has already used the peace protests for propaganda purposes (unit 7b) helps to increase the reader's belief in the nucleus (7a) which says that the votes against the use of force will be used as propaganda by Hussein.

8b(Nucleus)-8c(Nucleus) Joint

Units 8b and 8c share the same subject matter (the refusal of Iraq to accept a peaceful solution) and are connected by an explicit conjunction *and*.

9a(Nucleus)-9b(Nucleus) Joint

Units 9a and 9b share the same subject matter (the five months period of seeking a peaceful solution) and are connected by an explicit conjunction *and*.

10a(Satellite)-10b(Nucleus) Condition

The satellite 10a expresses an hypothetical situation (whatever happens now). A suitable paraphrase for this unit, which brings out even more the conditional aspect is: "if war breaks out". If the situation in 10a is realised, then the situation in nucleus 10b is realised as well, i.e. Saddam Hussein can be held responsible.

1,2(Nucleus)-3(Nucleus) Joint

Unit 1,2 and unit 3 share the same subject matter (the aversion towards war) and they are connected by an explicit conjunction *and*.

4(Nucleus)-5,6(Satellite) Evidence

The unit 5,6 which explains how the MP's who voted against the use of force 'failed' helps to increase the belief of the reader in the nucleus which says that those MP's went too far.

1,3(Satellite)-4,6(Nucleus) Concession

Unit 1,3 expresses the positive regard the writer has for the nucleus 4,6 (the MP's went too far by voting no, but their reasons for doing so are understandable). The concession is linguistically realised by "we appreciate ... But". The writer regards the situations in unit 1,3 and unit 4,6 as compatible and thus increases the reader's positive regard for the nucleus.

4-6(Nucleus)-7(Satellite) Volitional Result

The situation in unit 4,6 (the vote against the use of force) can be the cause for the action presented in unit 7 (Saddam Hussein uses the vote for propaganda purposes).

8a(Satellite)-8bc(Nucleus) Circumstance

Unit 8a sets the framework within which the reader is to interpret the situation in the nucleus 8bc.

9(Satellite)-10(Nucleus) Justify

> Unit 9 serves as an argument (we tried to find a peaceful solution for five months without succes) in convincing the reader of the writer's right to present unit 10 (Saddam Hussein is to blame if war breaks out).

8(Satellite)-10(Nucleus) Justify

> Unit 8 also serves as an argument (the French peace initiative, which was very soft and appeasing, was rejected) in convincing the reader of the writer's right to present unit 10 (Saddam Hussein is to blame if war breaks out). There is an alternative attachment point[18] for unit 8, namely unit 1,7. The relation that would hold then is Evaluation.

8,10(Satellite)-1,7(Nucleus) Evaluation

> The situation presented in the satellite (we have done all we could to avoid war) relates the situation in the nucleus (the MP's went too far by voting against the use of force) to the writer's positive regard for the nucleus.

3.2.2.4 Hobbs Carrying out step 1 of Hobbs' analysis method given in Section 3.2.1.4 above suggests a number of possible segmentations for Text II. One such would be to have three discourse main segments: [1–3], [4–7], [8–10]; another would be to include [4–7] within a larger segment [1–7]. The main indeterminacy appears to be the dual nature of contribution [4] as both backward-looking (by virture of the 'but') and forward-looking (by virtue of the further information concerning the same subject (the MPs who rebelled) given in contributions [5] and [6].

This question appears less important, however, when the possible coherence relations are considered independently of the informal prior segmentation. The relations that hold, with justifications drawn from Hobbs' definitions, are:[19]

CR1: `elaboration (1, 2a).`
because we can

> infer the same proposition P (1 = 'no one wants war') from the assertions of S_0 ('no one wants war') and S_1 (2a = 'everyone is dreading the conflict').

CR2: `exception (2a, 2b).`
because we can (slightly ammending Hobbs' definition of 'contrast')

> infer $p(a)$ (= 'be-against-war (everyone)') from the assertion of S_0 (2a = 'everyone is against the conflict') and $\neg p(b)$ (= 'being-for-war (Saddam Hussein)') from the assertion of S_1 (2b = 'apart from Saddam Hussein'), where b would normally be included under a.

CR3: `violated-expectation (3, 4a).`
because we can[20]

[18] Suggested by Gisela Redeker in personal correspondence

[19] We have used Hobbs' suggestion of an *exception* relation for a particular subtype of *contrast*, although Hobbs himself does not adopt this [Hobbs, 1990, p92].

[20] The following ignores the possible expectation of disapproval that a phrase such as 'we appreciate their sincerity' might evoke in the current kind of text.

infer P (= 'we approve of what the MPs do') from the assertion of S_0 (3 = 'we appreciate the MPs' sincerity...') and $\neg P$ from the assertion of S_1 (4a = 'the MPs went too far').

CR4: `evaluation (4a, 5).`
because we can

from S_1 (5 = 'asserting that the MPs failed to send a message of unity') infer that S_0 (4a = 'going too far') is a step in a plan for achieving some goal of the discourse.

The goal in the discourse is assumed to be something along the lines of 'make the speaker realize in what way the MPs went too far'; this is rather weak, however. A motivation for this would have to rely more on the influence of text type and general communicative goals, which are not emphasized in Hobbs' account. A more straightforward analysis[21] would be an **elaboration**, since the same proposition (namely, 'the MP's failed') can be inferred from both 4a and 5. We will suggest below, however, that this analysis is less resonant with the aims of the text as a whole.

CR5: `evaluation (4a, 6).`
because we can similarly

from S_1 (6 = 'asserting that the MPs failed to send a message of of opposing tyrany') infer that S_0 (4a = 'going too far') is a step in a plan for achieving some goal of the discourse.

Where the goal is as before for CR4. Here again the hypothesis of change of state rests upon the interpreted purpose of the text. We are presuming that the previous state was one where all were united in healthy debate, but now that debate has gone beyond reasonable boundaries.

CR6: `occasion (4a & 4b & 5 & 6, 7a).`
because

1. A change of state can be inferred from the assertion of S_0 ([4-6] = 'rebelling, going too far, and failing to send various messages') whose final state can be inferred from S_1 (7a = 'this being used in some way').
2. A change of state can be inferred from the assertion of S_1, whose initial state can be inferred from S_0.

CR7: `parallel (7a, 7b).`
because we can

infer $p(a_1, a_2, \ldots)$ (= 'use-for-pro-Hussein-propaganda (rebellion of MPs)') from the assertion of S_0 (7a = 'Hussein will use this') and $p(b_1, b_2, \ldots)$ (= 'use-for-pro-Hussein-propaganda (peace protests)') from the assertion of S_1 (7b = 'Hussein used weekend peace protests), where a_i and b_i are similar, for all i. [Hobbs, 1990, p93]

CR8: `contrast (8a, 8b & 8c).`
because we can

infer $p(a)$ (= 'discuss-peace (Parliament)') from the assertion of S_0 (8a = 'Parliament debating') and $\neg p(b)$ (= 'not-discussing-peace (Baghdad)') from the assertion of S_1 (8 = 'Baghad rejecting peace and making declarations'), where a and b are similar.

CR9: `occasion (9a, CR8).`
because [9a] and [CR8] are temporally related appropriately and have a clear functional connection in topic and intent.

[21] Suggested by Wilbert Spooren (personal communication).

CR10: `occasion (CR9, 10a).`
because [CR9] and [10a] are temporally related appropriately and again have a clear functional connection in topic and content.

Attempting to combine these into a single 'coherence structure' does not, however, provide a convincing picture: three fragments remain, depicted graphically in Figure 15. Where we can go further here, however, is in the principled application of non-local coherence relations. This is supported indirectly by Hobbs' allowing more than one relation to hold across propositions, although the consequences of this for a resulting 'coherence structure' are not specified.

The 'long distance' relations that can be drawn sometimes rely on a 'second order' interpretation of the text. For example, it is possible to view the text fragment CR1 (everyone disliking way) with its exception given in CR2 (except Saddam Hussein) as parallel to, or more likely in this context, exemplified by the text fragment 8a (Parliament debating) and 8b (Baghdad rejecting). This would at least serve to bring together the first and the last large text segments. A similar relation, but in this case probably `parallel`, could be found between the exception of Saddam Hussein to those disliking war (CR2) and the exception of 'those who rebelled' (CR3, CR4, CR5). In an editorial text such as this, it is quite likely that the overall meaning is going to be woven out of such patterns, where a relational structure attributing properties to certain participants (e.g., Saddam Hussein, pro-Saddam propaganda) in a very explicit way is used to indirectly attribute similar properties to other participants (e.g., the MP's who rebelled, MP-rebellion and peace protests) *by virtue of the latter's participation in similar coherence structures.* This could then fall under Hobbs' *figure-ground* coherence relation.

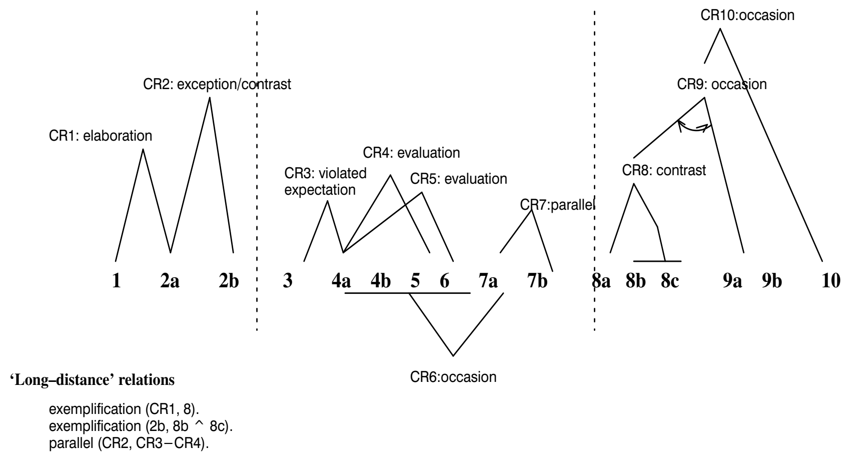

Figure 15: Hobbs' style analysis of Text II

3.2.3 Discussion of individual modes of analysis

3.2.3.1 Lascarides/Asher There is in this mode of analysis rather little linguistic detail; much of naturally occuring texts remain untreatable. The principle contribution would appear to be the mechanism proposed—and so it is this that needs evaluation in comparison with the other approaches available. The main question then is whether the proposed mechanism offers a useful treatment even when a more realistic range of coherence relations are adopted.

The principle components of the mechanism are:

- Asher's [Asher, 1993] SDRT (as introduced in Section 3.1.4.1 above),
- a logic of commonsense entailment (CE),
- a nonmonotonic inferencing system (DICE).

Local discourse contribution-by-contribution construction of discourse structure is attempted analogously to the situation with Martin's conjunctive relation analysis. This is in sharp contrast to the analytic mode proposed for RST.

Alternatively, it would seem, particularly following the introduction of beliefs and attitudes into the account in [Asher and Lascarides, 1994], that the formalisation possiblities may also be appropriate for the kinds of statements required for the definitions of RST. There are also clearly similarities to be drawn here with the rational planning accounts proposed by [Cohen and Perrault, 1979, Cohen and Levesque, 1985, Cohen and Levesque, 1987].

3.2.3.2 Martin What is important about the kind of analysis we see using Conjunctive Relations is the possibility it offers of 'getting more out of the text' in a linguistically principled way. To make full use of this we need to fill out the coverage of the resources proposed by Martin, since there are many cases which were not already to be found in the networks he proposes. This also involves filling out the cases of noncongruent realizations so as to net in further surface forms. The range of realizational possiblities is so large here that we have also to be sure that some constraint on interpretation remains.

Given this, the possible relation to an account such as RST, or Lascarides/Asher, or Hobbs is clear: more sources of constraint are made available. But are these other accounts in a position to benefit?

3.2.3.3 Mann and Thompson There has been considerably broader discussion and computational experimentation on RST than on other approaches to coherence relations. Our discussion here therefore draws on this in addition to our own experiences in applying RST for text analysis.

The main problems that we found in attempting to apply 'standard' RST, i.e., that version of the theory presented in [Mann and Thompson, 1987], were the following:

- 'standard' RST relations are often ambiguous (not enough constraints to choose the "right" relation);
- 'standard' RST supposes that the information unit corresponds minimally to a clause (problem of granularity);
- 'standard' RST relations do not take deep level intentions into account;
- 'standard' RST forces the user to choose one specific relation, even if more than one relation can apply;
- 'standard' RST enforces an isomorphism between relational structure and discourse structure: text 'spans' must be related by single relations and the spans so related must be contiguous;[22]
- 'standard' RST relations are all treated in the same way, although there is reason to believe that in fact rather different relations have been grouped together.

[22] This latter aspect is not actually a component of the original RST definition (cf. [Mann *et al.*, 1989, p11]), although most subsequent work has been done as if it were.

These problems are in fact typical of those reported by researchers working with RST and there is now considerable discussion in the computational linguistics literature concerning them; some of this discussion will be briefly mentioned below. In many respects, the problems encountered can mostly be accounted for by the fact that RST was developed as a descriptive theory for supporting the discourse analyst, and not for generation. As a consequence, the theory is underdefined.

3.2.3.4 Hobbs The Hobbs-style of analysis was generally very revealing—particularly in cases where it was problematic to find a simple conjunctive or rhetorical structure. As with Lascarides and Asher types analyses, this clearly underlines the value of having a well-defined set of specifications and acceptable inferences over those specifications for discourse-related reasoning.

However, the general method for analysing a discourse suffers from the same problem as that for RST: it is not always possible to find a convincing tree-structure for a text. It is still possible within Hobbs' account to apply the definitions of the coherence relations in order to see which apply *independently* of an informal segmentation, but one clearly cannot adopt the position of attempting to check pairwise for all propositions asserted in a text without some further constraints. One might expect, following the structural tradition in discourse analysis, that discourse structure would have a particular role to play here. But any removal of the contiguity constraint means that the discourse 'structure' may not be as straightforward to define. This may also increase the value of of linguistic constraints; Martin's conjunctive relations provide one possible set of constraints, although patterns of anaphora, thematic progressions, and other textually significant lexicogrammatical patternings are also good constraint providers.

It was also important that a way of constructing the assertion made by a complex discourse segment has at least been suggested by Hobbs [Hobbs, 1990, p104]. This may need further specification, however. The problem of ascertaining what is asserted by a discourse contribution or compound discourse segment was glossed over in our analyses above, but in general this will need a much tighter treatment. Hobbs treats the construction of complex-assertions by classifying coherence relations as either *coordinating* or *subordinating*. The former type of relations construct an 'inferrable' proposition common to all of the component segments related; the latter type takes the assertion of the 'dominant' segment [Hobbs, 1990, p104].

Genre, although mentioned, is certainly a second-class citizen in Hobbs' account. Possible properties of genre need to be made accessible; in particular, the constraints of genre that Hobbs also mentions [Hobbs, 1990, p107] also need to be applied. However, the recognition of genre is an important advance.

Related to this is the fact that all, except for one, of Hobbs relations hold with no mention of speaker goals. Only 'evaluation' makes any mention of goals and not much is said about this.

3.3 Preparatory steps towards a synthesis

3.3.1 A comparison of the kinds of information addressed

Our approach to the comparison will be 'problem-driven', drawing both on our experiences in applying the approaches to a body of texts and problems reported in the literature. We will start with the kinds of information present in RST and Hobbs-style coherence relations, and then go on to consider the additional contributions of Asher/Lascarides' and Martin's analyses.

3.3.1.1 RST We can organize both possible ways to build on RST and the current discussion on the content of RST by breaking down our problems according to problems with the content of the RST level of representation itself, with its relations to levels above (intentional goals, domain knowledge, style...), and with its relations below (how syntactically realized? in what kinds of units?). These issues need to be clarified before any real progress in using RST-type descriptions can be made or even to state criticisms accurately. We will address here first the selection and motivation of contents within RST representations, second their anchoring to more abstract representations—particularly, to communicative goals and intentions; and finally, their anchoring to less abstract descriptions—particularly, to the lexicogrammar.

3.3.1.1.1 RST: Internal organization and contents Within 'standard' RST itself there is a distinction drawn between relations that belong to the subject matter and relations that are more concerned with the presentation of that subject matter. Both distinct types are admitted, although in cases of conflict the presentational is to be preferred since it records the changes in the hearer's beliefs — thus making contact once again with the idea of 'rhetoric'. The requirement that *one* relation be selected in the analysis is a very serious constraint on the theory. Indeed, some researchers suggest that it is too strong in that it, firstly, does not work and, secondly, would require composite relations (i.e., cross-product of the two types of relation) to be defined. That the selection of a single relation proves in many practical cases difficult or unconvincing is not surprising given, as we will see in Figure 17 below, the probable conflation of multiple levels of analyses when an 'RST'-level analysis is carried out. Furthermore, even between relations that are arguable at similar levels of abstraction, there is still doubt that it is appropriate to pick a single relation in all cases.

This has recently been argued by [Moore and Pollack, 1992] who divide RST relations similarly into two kinds, presentational and subject matter, and take them to correspond to intentional versus informational 'levels' of discourse respectively.[23] A complete analysis of a text then needs to consist of two, possibly non-isomorphic, RST structures, each drawing its relations from one of the two subtypes of relation. A similar line is taken, although in a completely different theoretical setting, by [Korelsky and Kittredge, 1993], who suggest that the two subtypes of RST relation should be seen in terms of *stratification*, with one being the *means of* expressing the other.

This lack of agreement concerning what types of relations are to be accepted contributes significantly to the difficulties met in agreeing on some set of relations. Even when researchers are loosely agreed on using an RST-based approach, there are differences in the exact sets of relations adopted. In an attempt to achieve some consensus in this area, [Hovy, 1990] began a large-scale collection and classification of the discourse relations proposed in the literature; this has since been taken further in [Maier and Hovy, 1993] where an additional organization is imposed on the classification in terms of the notion of *metafunctions* drawn from systemic-functional linguistics (see e.g. [Halliday, 1978]). The three metafunctions: ideational (propositional content), interpersonal (interactant relationships), and textual (the means by which ideational and interpersonal information is presented as text in context) are considered to be 'simultaneous' in that they simultaneously describe a given linguistic unit from different perspectives. The first attempt to describe RST relations from a metafunctional perspective was that carried out in [Mann and Matthiessen, 1990]. This served as an additional theoretical basis for the presentation/subject-matter division: subject matter was clearly assigned to the ideational metafunction, while presentational, being concerned with the affect on the hearer, was interpersonal. It is then curious that in subsequent versions of RST the constraint of picking a single relation was accepted. An extended use of metafunctional classification has now also been proposed by [Lavid and Maier, 1992] where distinct relations belonging to all three metafunctions are explicitly considered; i.e., purely textual relations are also posited. This goes one step beyond the two-fold structuring suggested by [Moore and Pollack, 1992] and also matches well with the three-fold structuring independently developed in the 'Parallel-Components Model' of [Redeker, 1990,

[23] The by now painful overloading of the word 'levels' is well noted by [Meteer, 1993]; a more finely differentiating technical vocabulary is required. This is one of the reasons why we will be adopting systemic-functional terminology below. For an introduction to all such terminology used in this paper, see [Matthiessen and Bateman, 1991, Chapter 5].

Redeker, 1991, Redeker, 1992].

Concerning the internal organization of resources for coherence relations, there are, then, at least two possible 'kinds' of rhetorical relations proposed and two suggested models of their interaction. They are either related simultaneously, i.e., they occupy the same position in terms of depth of abstraction, in a metafunctional relationship, or stratificationally, i.e., they occupy differing positions in terms of depth of abstraction and one kind is used as a recoding or re-expression of the other. Both theoretical positions can be considered, although the arguments that [Moore and Pollack, 1992] bring to bear showing the non-reducibility of one kind of relation to the other would suggest the metafunctional *layering* view of relational structure to be the more appropriate. [Redeker, 1992] takes a similar stance.[24]

Following this line of argument, however, requires that we address the role of each of the three metafunctions with respect to rhetorical relations more critically. Since rhetorical relations are text organizing, they would seem to be *a priori* part of the textual metafunction. This is probably the most compatible with the original proposal of [Mann and Matthiessen, 1990]: the actual relations would belong to the textual metafunction, and these would be further classifiable according to whether they are concerned with the presentation of ideational or interpersonal information. The textual function is in systemic-functional theory generally seen as such a 'second-order' construct (which is never to be confused with processing order or stages). Moreover, many of the ideational relations seem also to be candidates for organizational principles for *domain* knowledge (e.g., such as temporal sequence, causality, etc.). If domain representations are then also to be subsumed under RST relations, this understandably leads to the further criticism that RST is too unconstrained for any possible agreement on relations to occur. In response to this, [Dale, 1993] argues that all relations that simply "mirror" the underlying knowledge base relations should be removed from the 'rhetorical component' and that attention should be concentrated on those relations that are clearly not domain relations, "such as MOTIVATION and JUSTIFICATION". If both of these positions were accepted, we would be back at the original suggestion that in an RST analysis only a single relation belonging to the 'presentational' class should be considered for relating any two discourse segments. We are now, however, also in a position to say more about the coherence relations excluded from RST, since the referential type is arguably addressed in detail by Hobbs (causality-based) and Lascarides and Asher (temporally-based).[25]

In order to clarify these issues, we now list the original set of 'interpersonal' RST relations proposed by [Mann and Matthiessen, 1990, p15/16], further organized as Mann and Matthiessen suggest. In addition, we propose a specification of the 'effects' field of the relations, drawing on a small list of propositional attitudes. These include: the reader's belief ($\mathcal{B}_R$), regard ($\mathcal{R}$), wanting ($\mathcal{W}$), ability ($\mathcal{A}$), understanding ($\mathcal{U}$), and purposes ($\mathcal{P}$). These attitudes will not be formally defined at this point; their purpose is to suggest what kinds of information would be necessary to formalise the RST relations as given and to bring out more clearly differences and similarities in the definitions. In addition, an operator over such attitudes (deg$\uparrow$) is used to indicate that an increase in the specified attitude is achieved. The expressions me_N and me_S follow Asher and Lascarides in marking the 'main eventuality' associated with either the nucleus or the satellite discourse segment respectively.

[24] Note that to completely remove any constraint on the number and types of relations that may hold between discourse segments, such as suggested by [Hobbs, 1990, p88], raises significant issues for the supposedly accompanying discourse 'structure' that is recursively constructed from the relations obtaining. Some constraint would appear useful and the approach that we will be adopting is to propose explicit 'levels' and 'layers' of relations, within which relation attribution is very constrained.

[25] Although this argument may work for the ideational relations, the additional organization that appears purely textual (cf. [Lavid and Maier, 1992, Redeker, 1992] and also [Martin, 1992, p260] for some further very clear examples) does not seem so simply subsumed. However, it is still not yet clear that there is no distinction to be drawn between *presenting* something as a circumstance or cause and this information being represented in the domain knowledge. It may be that domain representation relations are simply 're-usable' as rhetorical relations without imposing an identity. Their use would then differ in function according to whether they are viewed from the ideational perspective, in which case they are representing domain relationships, or from the rhetorical perspective, in which case they are holding the text together. We clarify this below when we propose a definition of textual rhetorical relations.

- aids to assertion-like activities:
  - Evidence deg$\uparrow$ $\mathcal{B}_R(me_N)$
  - Antithesis deg$\uparrow$ $\mathcal{R}_R(me_N)$
  - Concession deg$\uparrow$ $\mathcal{R}_R(me_N)$
- aids to various kinds of requesting and offering:
  - Motivation deg$\uparrow$ $\mathcal{W}_R(me_N)$
  - Enablement deg$\uparrow$ $\mathcal{A}_R(me_N)$
- aids to any kind of speech-function, "making it more comprehensible or more acceptable":
  - Justify deg$\uparrow$ $\mathcal{U}_R(\mathcal{P}_w(me_N))$
  - Background deg$\uparrow$ $\mathcal{U}_R(me_N)$

These relations are necessarily 'intentional', they bring additional organization and information to bear that is not to be found in a representation of the information being expressed in a text. [Moore and Paris, 1993] suggest that there is a one-to-one correspondence between the presentational relations and communicative intentions, which is not the case for the subject matter relations.[26] All the effects fields of these relations (and none of the effects field of the others) refer to an increase in the degree of some attitude of the reader.

This discrimination between interpersonal ('intentional') and ideational ('informational', 'subject matter', 'domain') relations would, then, at first seem to suggest there there is a substantial number of subject matter relations to consider. However, on closer examination of the effects fields defined for these relations, there turn out to be very few 'pure' informational relations of the kind that mostly constitutes Hobbs' account.

Attempting to specify the effects field of the following relations reveals supporting relationships and predicates that seem unlikely candidates for inclusion in a knowledge base prior to the contingencies of constructing a text. In these specifications the expression ${}_R\langle p\rangle$ indicates that the proposition or state of affairs p is recognized by the reader R.

- interpretation: ${}_R\langle \text{Framework}(me_S, me_N) \wedge (me_S \cap me_N = 0)\rangle$[27]
- evaluation: ${}_R\langle me_S = \text{assessment}(me_N)\rangle$
- restatement: ${}_R\langle me_S = \text{restatement}(me_N)\rangle$
- summary: ${}_R\langle me_S = \text{restatement}(me_N) \wedge S\ \textit{shorter-than}\ N\rangle$

All of these appear to be essentially textual[28], even though Mann and Matthiessen explicitly reject the existence of textual relations. The relation 'joint' may also be a candidate here. These relations impose constraints on the information related (or induce presuppositions concerning the relation between related information), since it is clear that not everything can stand as an assessment or as a restatement; but these constraints are nevertheless quite general. The main source of constraint in application is the purpose of the speaker in the construction of the text at hand.

A further subset of relations can more clearly be related back to informational specifications that would support their use, although the relations themselves present a textually classified and motivated view of those specifications. The examples here are:

[26] Curiously Moore and Paris include 'concession' as a subject matter relation. The communicative intention for which 'concession' stands as a realization appears to include an implicit *interpersonal* element however: (BEL ?h BEST-METHOD-FOR ?goal ?method)). This may explain the discrepancy.

[27] Recall the RST effect field: R recognizes that S relates the situation presented in N to a framework of ideas not involved in the knowledge presented in N itself.

[28] Interestingly analogous to [Fox, 1987]'s 'floating' relations: 'summary', 'assessment', and 'conclusion'.

- elaboration: the textually motivated selection of almost any informationally linked information;
- circumstance: the textually motivated selection of a temporal, spatial, etc. linked piece of information as presenting a general framework for the current text segment's interpretation;
- sequence: the textually motivated selection of elements along some dimension of ordering (such as size, or age, time of occurrence, etc.).

Again, these need to have very general applicability conditions defined for them with respect to the knowledge available. As also argued by [Asher and Lascarides, 1994, Lagerwerf and Oversteegen, 1994] textual relations do usually impose such informational constraints. But they also need some kind of specification in terms of when a speaker might use them.

Even the few RST relations remaining contain some arguable cases. Is the notion of 'solutionhood' so clear cut as an ideational, subject-matter relations? Or, for that matter, 'purpose'? Certainly one way to achieve this is to recognize in the ideational realm a domain of agents, plans and goals. This does not seem so unlikely and may even serve to net in and define more finely the one exception in Hobbs' account that refers to the 'goals' of the discourse: 'evaluation'.

3.3.1.1.2 Anchoring 'above' RST has been criticized concerning its apparent equivocation of intentions and text structure. As [Meteer, 1993, p83] cites from [Hovy *et al.*, 1992][29]:

> "We consider discourse to be a structured collection of clauses. This structure is expressed by the nesting of segments in the discourse; a discourse can thus be represented as a tree structure, in which each node of the tree governs the segment (subtree) beneath it. At the top level, the discourse is governed by a single root node; at the leaves, the basic segments are single grammatical clauses. Each discourse segment has an associated purpose, which, following [Grosz and Sidner, 1986], we call the Discourse Segment Purpose, i.e., a communicative goal. Each adjacent pair of such purposes is related in the discourse structure by one or more discourse structure relations..." [Hovy *et al.*, 1992, p58; footnote 2]

This indeed, as Meteer argues, seems to identify three distinct kinds of information — structural, intentional, and rhetorical — while forcing them to be isomorphic. Whereas the current RST account has discourse segments corresponding to communicative goals, Meteer proposes that "intentions control the decision making of the text ... and do not necessarily appear explicitly in the structure of the text." [Meteer, 1993, p82] This is no doubt correct, but as an argument against RST it seems to rely on an overly vague notion of 'intention'.

Among the early enhancements to descriptive RST that were made in creating a computational version of the theory was precisely the addition of links to more abstract 'discourse goals' (cf. [Paris, 1991, Hovy, 1991, Moore and Swartout, 1991]), thereby taking 'deep level intentions' into account. These links typically functioned as additional constraints on the selection of rhetorical relations.

They also enabled the rhetorical relations to be used within planning processes, since their preconditions and consequences were formally specified. Thus, while it is clear that there is no necessity that communicative intentions in general themselves appear in the structure of a text, this does not entail that there are no structural consequences of such intentions and it is furthermore natural to interpret these structural consequences in functional terms (as often carried out even in the grammar in dependency approaches or systemic-functional grammar). This leaves at least two kinds/levels of description: rhetorical intentions, at a lower level of abstraction, and discourse purposive intentions,

[29] Here presented in full.

at a higher level of abstraction. The former is seen as the *realization* of the latter; both can be accommodated if they are found to be motivated and necessary for a theory of discourse (cf. also [Hovy, 1993] on these two 'levels' of description).

3.3.1.1.3 Anchoring 'below' With respect to their anchoring to less abstract kinds of description, RST descriptions are in an awkward position. This is due to the following two simultaneous requirements:

- first, for use in computation, it should be possible to exactly specify the consequences of selecting a particular rhetorical relation on the linguistic form;
- however, "*... all of RST is pre-realizational, since it makes statements about how such meanings and intentions are structured and combined, and not about how they are realized.*" [Mann *et al.*, 1989, p8] (emphasis in original)

Thus RST itself has not presented clear realizational consequences for its structures, which is particularly problematic in text generation where that is precisely what we need to know following the selection of any such RST relation. This has led to alternative proposals for recognizing rhetorical relations, such as that of [Knott and Dale, 1992], whereby some straightforwardly recognizable linguistic feature (e.g., cue phrases) is taken as direct evidence for a particular discourse relation. Similar reliance on surface-observable features of sentences has, of course, been used in a wide range of accounts: e.g., [Reichman, 1985]'s use of cue phrases, [Grosz and Sidner, 1986]'s use of anaphora, [Redeker, 1990]'s use of 'discourse operators'. It should be clear that, as long as the information can be organized appropriately, the more such evidence the better. However, it does little good to have a list of unrelated 'surface observable phenomena'. This is one value of a highly systematic account of such phenomena as proposed by Martin's Conjunctive Relations. This information should not, arguably, become part of the RST-type of description since a close binding of coherence relation and linguistic realization of this kind would clearly move the relations discussed away from those originally intended by RST, to a lower level of abstraction more closely bound to lexicogrammar, which brings us back to the consideration above of the internal content and organization of an RST account.

Finally, there is the issue of the granularity of RST analyses. This has repercussion for the content of RST analyses, but also relates crucially to the way in which RST analyses bottom out. It is now often stated that the RST selection of *grammatical clause* as the unit of analysis is multiply problematic. [Meteer, 1993, Rösner, 1993] argue that this is too large a unit at which to stop analysis; [Dale, 1993] argues that it is sometimes too small a unit to require an RST functional analysis. However, it is clearly stated in [Mann *et al.*, 1989] that grammatical clauses are not the chunks of discourse between which RST relations hold.[30]

> "Strictly speaking, the relations of a text do not hold across the various word sequences of which the text consists. Rather, the word sequences are realizations of more abstract entities: meanings and intentions that are represented by those word sequences. ...It is inconvenient to acknowledge the abstraction on every mention, so we will generally say that relations hold between spans of text, but the distinction between the abstract entity and its realization always remains. [p8]
>
> A *text span* is any portion of text that has an RST structure ..., or that is realized by a *unit*. [p11]
>
> Unit size is arbitrary in RST; in principle the units can be of any size from typical lexical items to entire paragraphs or larger." [p16]

[30] Although it is also clear that certain of the expressions of "convenience" adopted in [Mann *et al.*, 1989] help confuse this point.

The crucial distinction of the spans related by relations and the realizations of those spans appears to have been lost in much current discussion. The 'leaves' of an RST structure tree *cannot* be clauses: they can only be discourse-level spans *realized* by clauses (or by many other kinds of linguistic entity). This difference opens up a wealth of realizational possibilities that go far towards providing the space required for answering criticisms of granularity. This is also compatible with other accounts, and Martin in particular fills in significant parts of this theoretical space in detail.

3.3.1.2 Mechanisms It is also necessary to specify how the kinds of resources described in Lascarides/Asher, RST, Hobbs, or Martin are to be applied. For this, it is crucial that one can build up the meaning of larger discourse segments. This is a common feature of all the accounts investigated. To do this most effectively, however, requires a constructive semantics based on combinations of the meanings of discourse segments plus the coherence relation. Although RST specifications are weaker from this perspective, there is a basis given in the specifications of their 'effects' fields. This needs to be examined more closely therefore. Such 'semantics' for relations are also given somewhat intermittently for Martin's conjunctive relations—the most detailed account that he gives is for external consequentual relations [Martin, 1992, §4.3.2, pp193-202]; we will build on this below in our computational pre-specification.

Hobbs and RST have no formal specification of how to build up an analysis of a text; the analysis tends to rely first on the analysist recognizing what is important for the text and producing a segmentation. Asher/Lascarides have a more formal account of discourse structure construction with specification of attachment points, similar to that of Webber and others (the 'right-frontier') and extensions of DRT. In more recent work, e.g., [Hobbs *et al.*, 1993], Hobbs suggests that abduction over the definitions of coherence relations may succeed in producing an analysis. Martin's conjunctive relations are also intended to be applied to a text, left-to-right as it is produced and so is more similar to the Asher/Lascarides account; Martin criticizes the view of a text as a static, pre-formed entity assumed by RST [Martin, 1992, p262]. This interpretive mechanism is not specified formally in any way, however. Moreover, Martin's conjunctive relations make available more information when it is available in the text, otherwise—i.e., for implicit relations—a similar problem obtains to Hobbs and RST.

3.3.1.3 Summary As far as the content of the levels of description the accounts propose, Hobbs has a very weak account of time—perhaps it is not necessary to have more detail than precedence and overlapping at the level of abstraction with which he is concerned. This is clearly different from the somewhat detailed view of time and its discourse import offered by Lascarides and Asher [Lascarides and Asher, 1993]. Hobbs' account is primarily referentially based (apart from 'evaluation'), whereas RST has closest connection to speaker 'goals' and beliefs. Speaker intentions, attitudes and belief do find their way into the Asher and Lascarides account in the more recent [Asher and Lascarides, 1994], however. Finally, in addition to an organization of resources, we also need an account of how those resources are to be processed. Here the abductive inferential mechanisms adopted by both Hobbs and Asher are the most powerful. The extent to which this power is required—or in what cases is it required—must also be clarified before a computational specification can be completed. To summarize, then, the approaches examined show the relative strengths and weaknesses outlined in Figure 16.

Some of these problems and weaknesses have been discussed within the respective literature of the individual perspectives; however, there has been no proposal which addresses all weaknesses, maintaining the strengths, within a single unified account. In general, we can say that approaches rooted in analysis develop insufficient detail of linguistic information for supporting text generation, while approaches rooted in text generation develop insufficient formalisations and recognition criteria. Analogously, those approaches that posit a well-defined discourse structure are in a better position to specify a

Property	Asher/Lascarides	RST	Hobbs	Martin
breadth of coverage necessary for addressing natural texts	$\times$	$\surd$	$\surd$	$\surd$
degree of formalization	$\surd$	$\times$	$\surd$	$\times$
construction algorithm for interpretation	$\surd$	$\times$	$\times$	$\times$
degree of modularization of information	$\times$	$\times$	$\times$	$\surd$
contact with grammatical realization	$\times$	$\times$	$\times$	$\surd$
density of detail for generation	$\times$	$\surd$	$\times$	$\surd$

Figure 16: Relative strengths ($\surd$) and weaknesses ($\times$) of the approaches

formal and computationally specifiable interpretation procedure. This imbalance must be rectified in any more appropriate framework.

The example approaches and analyses now discussed, which span the major distinctive approaches to coherence and coherence relations available at this time, can be organized in terms of their 'vertical' positioning within a continuum of linguistic abstraction—this begins to impose the Dandelion 'vertical methodology' on the wide range of information covered and refers to the kinds of information that the approaches manipulate and the kinds of motivations for the organizations of that information that they require. This is depicted graphically in Figure 17. As can be seen,[31] there is a large degree of overlap between some accounts, but also some differences.

3.3.2 An initial critical synthesis and unresolved issues

There is often argument against further discriminations in types of information beyond 'semantics'. The principle reasons for this reluctance are economy of representation, interactions among all types of information, and the reoccurence of similar appearing information at levels posited to be distinct.[32] The negative consequences are, however, that modularity is weakened and some information, although recognized, is accorded second-class status (e.g., genre in Hobbs' account). We will make a hypothesis of strong modularity and ask whether it is possible to find a finer discrimination of types of information that simplifies the account as a whole—particularly in terms of maintainability of resources and the specification of interactions. Towards this we therefore expand the latter of our above initial questions into a number of investigative goals. These may also be organized roughly 'vertically' according to the Dandelion vertical methodology and related to our assumed boundaries for coherence relations as follows:

1. to ascertain the utility of an upper bound on the level of abstraction covered reasonably by coherence relations.

   The resources constituting coherence relations can then be 'anchored above' with respect to resources at higher levels of abstraction. Our first hypothesis here will be to use the discourse purposes, or *generic purposes*, suggested in Dandelion deliverable [Lavid, 1993] drawing on work of [Martin, 1992] and others. For each mode of analysis considered, we will investigate particularly the role of genre and generic purposes for that mode.

[31] The figure also shows the approaches of [Knott and Dale, 1992] and [Sanders *et al.*, 1992] that have been omitted from our presentation here.

[32] Examples here are the debate over the (non)distinction between semantic-conceptual information, or between register and genre [Martin, 1992].

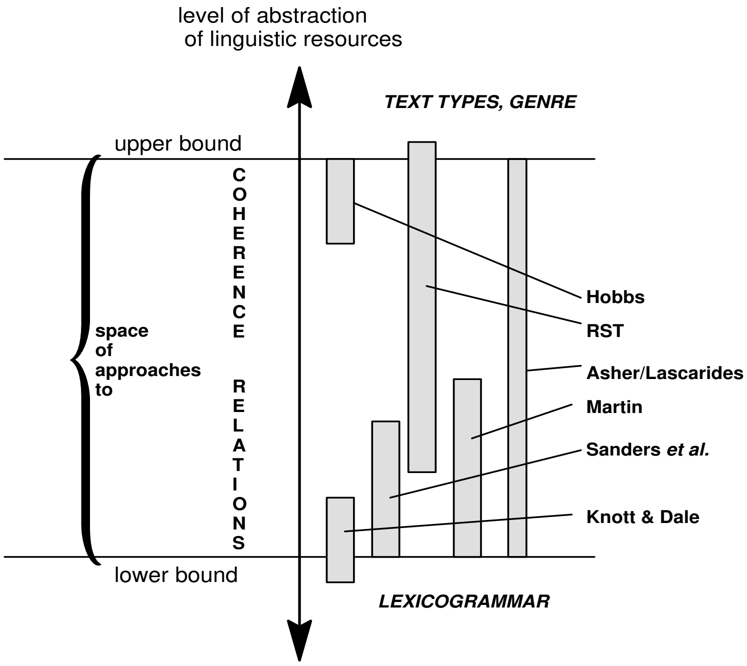

Figure 17: Overview of approaches to coherence

2. to refine the continuum of levels of abstraction lying between the upper and lower bounds into a discrete number of linguistic *strata* and to provide details of the contents of each such stratum.

3. to find a lower bound on the level of abstraction covered reasonably by coherence relations.

   The resources constituting coherence relations then become 'anchored below'. A first hypothesis here could reasonably be the lexicogrammar: that is, coherence relations would be *realized by* lexicogrammatical features (conjunctions, anaphora, tense, modification, clause combining, etc.) but would not themselves *be* lexicogrammatical resources.

A good candidate for a set of resources firmly based on language-related criteria is clearly that of Martin: here the primary motivation is in terms of patterns of partially shared linguistic realizations.

In terms of what happens above the lower bound for coherence relations, it is clear that RST in particular is being stretched very thin. It includes both attempts to constrain the linguistic form that occurs in fine detail (although this is not part of its original brief: see above) and to state connections with communicative goals and intentions. It thus tends to conflate at least.

- language related criteria for relation discrimination,
- 'logical', world-related, or natural ontological criteria of the commonsense world,
- rhetorical goals for constructing texts concerning speakers and hearers and their interactions in the world.

It has already been observed that in lists of multiple levels of description for coherence relations, such levels are typically not of equal status. Examples of this include the division between rhetorical goals and intentional goals pointed out above (cf. [Hovy, 1993]) and [Redeker, 1991]'s distribution of the five 'planes of talk' proposed by [Schiffrin, 1987] to cope with the multifunctionality of discourse markers over two, rather different statuses, instead of the flat list given by Schiffrin. Similarly, [Fox, 1987] makes a clear distinction between 'higher-level' rhetorical relation patterns (including 'issue', 'list' and 'narrate' [Fox, 1987, p79]) and 'floating' relations (such as 'summary' [Fox, 1987, p88]). Opening up the discussion to allow for descriptions at varying levels of abstraction raises, of course, some rather important further issues of its own. It appears that many proposals made involving coherence relations and RST for computational purposes have not sufficiently separated out the distinct levels of representation involved. Thus, the goal hierarchies found in the text planning process often contain rather different kinds of information; an example is the hierarchy given in [Moore and Paris, 1992, p318], where not all goals correspond directly to RST relations (construed as operator plans), although some of them do. This can be criticized on the grounds given by Meteer above. It is also likely to make a larger-scale declarative statement of the resources involved, as sought by [Paris and Maier, 1991], more difficult. We take this as general supportive evidence that a more discriminating set of definitions might separate out these distinct motivations for distinct sets of resources.

A similar fate would seem to await Asher's account once it becomes more developed. There is no distinction drawn there for coherence relations that apply in a local fashion and relations motivated by virtue of their being an established genre, or text type, in the culture—e.g., 'explanation'. This may be appropriate for Asher's concerns—that is, if the same formal processing mechanisms apply regardless of the nature, degree of abstraction, and granularity of a 'coherence' relation—but it would seem of less help in the construction of richly detailed sets of linguistic resources. More concrete evidence here may be obtained by considering the *nature* of the discourse structure that applies at different levels. If evidence for differing organizations is to be found, the case for a single undifferentiated account is weakened.[33] Although this goes beyond the initial brief for this deliverable, it is an essential issue for further formalisation.

In constrast to these mixed accounts, a good candidate for a set of resources motivated on 'logical', referential grounds is that of Hobbs. The clean separation of Hobbs' account from actual linguistic realization and his concentration on logical conditions on states of affairs in the 'world' as primary criteria for establishing coherence relations suggests that this type of description could provide an appropriate most abstract account within this level. Interesting here is the rather weak acknowledgement of speaker's goals that Hobbs offers, with only one relation referring to goals in its definition and, even then, in a form which either seems to be far too general (e.g., the communicative goal of 'being understood' [Hobbs, 1990, p89] or does not seem widely applicable enough (see Section 3.2.2.4). Removing this aspect from the Hobbs-like specifications altogether may, therefore, improve its internal consistency.

We thus have a first breakdown of the overall space of coherence relations. However, the relations between conjunctive relations and rhetorical coherence relations have not yet been sufficiently clarified. This is addressed in a second phase of analysis.

4 Analysis 2

In the second phase of analyses the conjunctive relations and the rhetorical coherence relations have been investigated more closely. How do they relate to each other? The hypothesis we have started out with is that conjunctive relations are merely a subset of rhetorical relations. Progressive decomposition of a text would produce a tree of RST relations whose leaves are formed by Conjunctive Relations. In other words: there is a bottom layer of Conjunctive Relations on top of which there are layers of RST relations. This view implies that Martin's analysis is best suited for describing coherence

[33] Analogous also is [Asher and Lascarides, 1994]'s argument for a separation of intention, information, and discourse structure.

relations between small units of texts and that it would run into difficulties if applied to large text units. The opposite holds good for RST relations. For testing our hypothesis we have used two letters from the fund-raising letter corpus of [Abelen *et al.*, 1993]. The organisations which seek funds are the World Wildlife Fund (WWF) and the Nature Conservancy (NC). All examples used are marked by either WWF or NC to indicate their origin. The numbers in front of each relation refer to the unit numbering of [Abelen *et al.*, 1993]. In some cases relations are distinguished within a unit. The unit size then no longer corresponds to the size adopted by 'classical' Rhetorical Structure Theory. These cases are marked by the title non classical RST unit. Generally speaking, if the unit size is reduced to a single clause, then RST is descriptively less tight than Martin's conjunctive relations. Multiple RST relations can apply. This is shown by the following example:

WWF(non classical RST unit)
31 "I'm hoping that with your help we can bring many forms of life on earth back from the brink of extinction."
RST relation: Volitional Cause
Martin:
consequential:external:cause:contingency:condition
realisation: *with your help* (explicit, **within process:as circumstance**)

In this example not only the Volitional Cause relation can apply, but the Condition relation as well. The problem of multiple RST relations seems to occur most frequently when the relation is realised as a single process. A very small unit seems to introduce a multiplicity of RST relations. Martin is much tighter in this respect. However, RST does not always fail to provide a tight description when the unit size is very small:
WWF(non classical RST unit)
30a "But in spite of all this, we need many more committed members, many more dollars, and a vastly accelerated timetable
30b if we are to win the race for wildlife survival." RST relation: Volitional Cause
Martin:
consequential:external:cause:contingency:condition:inclusive:factual:mod.possible
realisation: *if* (explicit, **between processes:taxis:hypotaxis:finite**)

RST and Martin perform equally well in this case. On the basis of the earlier hypothesis Martin should offer the tightest description here. Note that the relation is now realised as being between two processes.

It is also possible that multiple conjunctive relations apply if the unit size is limited to a single clause.
WWF(non classical RST unit)
9 "Pleading ignorance of these vital and fragile ecosystems can only spell global disaster." RST relation: Non-Volitional Cause
Martin:
consequential:external:cause:consequence
consequential:external:cause:contingency:condition:inclusive:factual:mod.certain
realisation: *spell* (explicit,**within process:as relation:within clause:identifying:as process**

In the first reading *spell* corresponds to *lead to*, in the second to *ensures*.

The above examples have made clear that Martin does not automatically offer the tightest description if the unit size is small (limited to a single clause). This can be taken as proof to the contrary for our starting hypothesis. We have also seen that small units can result in multiple application of rhetorical relations. This can also happen if the unit size is larger than a single clause:
WWF:
51-52 "Runaway destruction of our planet's rainforests —25 to 50 acres every single minute— is alarm-

ing. Especially when you consider that half of our prescription medicines come from natural sources like plants and nearly half of all plant species on earth are found in these rainforests!
RST relation: Evidence
Martin:
comparative:internal:similarity:reformulation:rework:generality:local:particularise
realisation: *especially* (explicit, **between processes:cohesion**)

Both the Evidence and the Interpretation relation apply here. The difficulty in assigning a single relation may reside in the fact that 51 and 52 form a clause complex together. The unit size may then be too small for any RST relation to apply with ease and this can be taken as an indication that RST is being applied outside its borders. The cue phrase *especially* does not occur in any of Martin's networks, but can be paraphrased by *particularly so*. The problem is now that the combination of *particularly* and *so* doesn't occur in the networks. We decided to class it in the same category as *particularly* until a better solution is available.

We will examine how both analyses will perform if applied to large(r) units. RST should offer the tightest description here, according to our hypothesis.

WWF
1,18-19,43 "[..] But you may be wondering if your help [..]"
RST relation: Justify
Martin:
consequential:internal:concessive
realisation: *but* (explicit, **between processes:cohesion**)

The above relation seems to provide evidence to the contrary of the view that Martin's analysis is more suited to deal with small text units. Despite the fact that the text units are very large, there is an explicit cue phrase present which provides the coherence of the text units. Martin's Conjunctive Relations have no problems dealing with these large units.

WWF
31,32-33,36 "I'm hoping that with your help we can bring many forms of life back from the brink of extinction. We can and must make sure that "extinct" does not become their final epitaph. How can we do this? By making our commitment move beyond "just talk". Beyond merely marvelling at the perfect flying formation of migratory geese cross our bit of sky. Beyond delighting in the majestic movements of the great whales as they traverse the global oceans."
RST relation: Justify
Martin:
additive:internal:turnbuilding:staging:framing:open
realisation: *well* (implicit, **between processes:cohesion**)

NC
12-13 "We have a staggering difficult job to do and not much time to do it. So we urgently need all the support we can get."
RST relation: Non-Volitional Result
Martin:
consequential:external:cause:consequence
realisation: *so* (explicit, **between processes:cohesion**)

In the above example two RST relations can apply: Non-Volitional Cause and Non-Volitional Result. Their application depends on what one considers to be the nucleus. Here the genre comes into play. Since the letter is for fund-raising it seems justified to consider the urge for support as the nucleus (i.e. unit 13). In that case the cause is in the satellite and Non-Volitional Cause is the relation that holds.

NC
45,48-49,51 "It can be done. Our 39-year record confirms it. And one more thing confirms it — al that acreage already protected by the Conservancy. For life. Think of the real estate developments, waste dumps, and amusement parks that weren't built in ecologically important and beautiful areas in your state or maybe even in your neighborhood. Things that don't happen seldom make the headlines. But they can be vastly more important to you and your family — and your quality of life — just the same."
RST relation: Evidence
Martin:
comparative:internal:similarity:reformulation:rework:abstraction:exhaust
realisation: *in other words* (implicit, **between processes:cohesion**)
Unit 49,51 can be seen as a kind of restatement or paraphrase of the unit 45,48 (i.e. buy land for protection = prevent waste dumps etc. from being built).

The examples above show that large units do not provide difficulties for the Martin's analysis. As one might expect, the large text units are related by internal conjunctive relations. Although there doesn't seem to be a difference in descriptive tightness between the two analyses at this level, other examples make clear that it can sometimes be difficult to find an appropriate conjunctive relation for very large text units. So generally speaking RST does a better job here. Nevertheless the examples have shown that our hypothesis can no longer upheld. Conjunctive relations are not merely a subset of rhetorical relations. Their actual relation needs further investigation.

Subsequently a comparison was made between the conjunctive relations and the rhetorical relations to see how they would map onto one another. Figure 18 shows the mapping for the two fund-raising letters described above, whereas figure 19 shows it for three articles from the Gulf War corpus (among them Texts I and II). The dotted lines indicate mappings of relations which hold between nonclassical RST units. Starting from the conjunctive relations side, we see that to every relation corresponds a RST relation, but that in some cases a conjunctive relation maps onto multiple RST relations. This provides further proof for the falsification of our starting hypothesis. If conjunctive relations were a mere subset of rhetorical relations then a one onto many relationship would exist. The figures show that this is not (always) the case. Another interesting phenomenon which requires further investigation is how genre influences the way (RST) relations are mapped onto different conjunctive relations. For instance Volitional Result is mapped onto **additive:external:addition:positive** and **additive:internal:turnbuilding:developing:addition:simple** in the Gulf War articles, whereas it is mapped onto **consequential:external:manner** and **consequential:internal:concessive:objection:concede:mod.certain** in the fund-raising texts.

In short, a straightforward combination of approaches was not supported. This would rest on it being possible to place analyses reliably in a one-to-one relationship so that when as analysis for one perspective obtains, we know that a particular analysis at each other combinable perspective also obtains. This did not seem to hold for RST and conjunctive relations. A more complex relationship needs to be adopted.

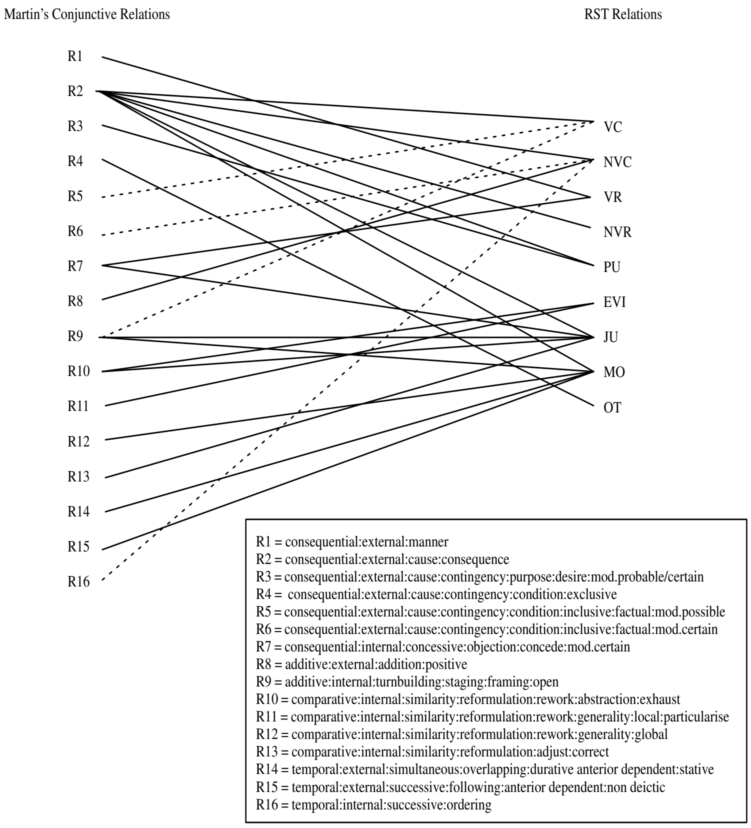

Figure 18: Mapping of Martin's Conjunctive Relations on RST relations: Fund-raising texts

5 Prespecification of a comprehensive coherence relation account

What happens between the lower and upper bounds has been investigated empirically above. We have asked for a range of approaches where they are situated 'vertically' in the abstraction continuum between our lower bound (initially, the lexicogrammar) and our upper bound (initially, generic discourse purposes). Where there is overlap, we can expect the accounts concerned to be reconcilable and to provide mutual input. Where there are divisions, we have evidence for finer stratification. Finally, building on the strata uncovered will provide specifications for resources usable in text generation for progressively mapping from the most abstract discourse purposes down to lexicogrammatical realization and in text analysis by constraining inferences.

Our intention is that the increased detail of the linguistic resources will significantly constrain the necessary access to real world knowledge for generating and interpreting texts. The tighter modularity is intended to provide for a specification with similar properties to that aimed at by Sanders: i.e.,

> "...we have limited the appeal to world knowledge to the absolute minimum; as far as it is made use of, it has been restricted to dictionary-like knowledge. At what cost we have done this, has been made clear by comparing the computed structures with the ones assigned by expert analysts on the basis of their intuitions ...This comparison can reveal what knowledge is necessary at what point to improve the analysis. It turns the inclusion

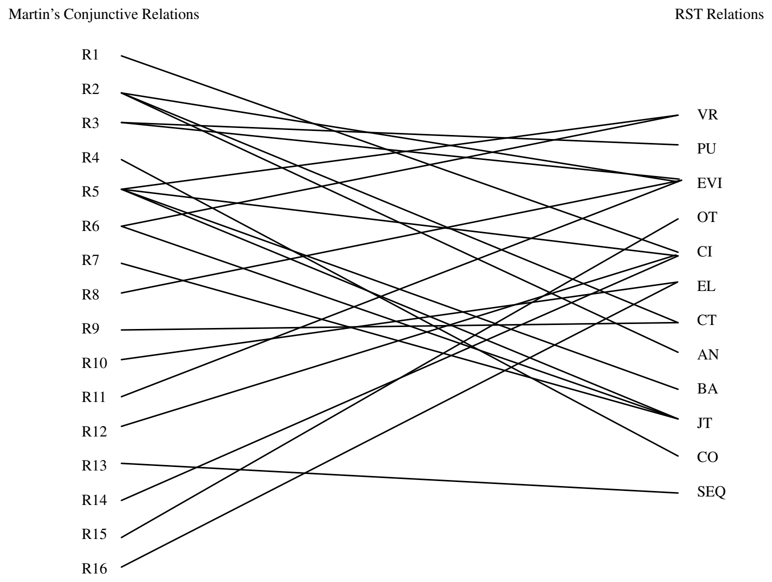

R1 = consequential:external:manner
R2 = consequential:external:cause:consequence
R3 = consequential:external:cause:contingency:purpose:desire:mod.probable/certain
R4 = consequential:external:cause:contingency:+concession:universal
R5 = additive:external:addition:positive
R6 = additive:internal:turnbuilding:developing:addition:simple
R7 = additive:internal:turnbuilding:developing:addition:cumulative:amplifying
R8 = comparative:external:similarity
R9 = comparative:internal:difference:opposition:contrast
R10 = comparative:internal:similarity:reformulation:adjust:correct
R11 = temporal:external:simultaneous:coextensive:punctiliar
R12 = temporal:external:simultaneous:overlapping:durative anterior dependent: non stative
R13 = temporal:external:successive:following:anterior dependent:non deictic
R14 = temporal:external:successive:following:anterior dependent:deictic:extension to present
R15 = temporal:external:successive:following:posterior dependent
R16 = temporal:internal:successive:ordering

Figure 19: Mapping of Martin's Conjunctive Relations on RST relations: Gulf War texts

of world knowledge from an article of faith into an empirical affair." [Sanders, 1992, p110]

On the basis of our discussions and analyses above, we now adopt the following vertically and horizontally organized set of general types of coherence relation resources and their boundaries:

- Lexicogrammar (external to coherence relation model)
- Conjunctive Semantic Relations (based on Martin)

  The close linking of Martin's approach with the lexicogrammar, suggests that we build on this as an appropriate lower-bound for coherence relation accounts. Problems to be addressed here in order for a more complete evaluation of this possibility are to make fully explicit the realizations of conjunctive relations in terms of concrete lexicogrammars and to investigate how the 'semantics' of the distinct alternatives that Martin describes can be specified. The firm link to lexicogrammar is aimed at enabling us to accumulate as much surface evidence as possible in a systematic way.
- Rhetorical Coherence Relations

  Here particularly, we also adopt the 'horizontal' metafunctional organization argued for above with respect to RST. This allows us to pinpoint the various types of information involved more accurately. In particular, we have the following organization of definitions and information.

- ideational: domain knowledge, world knowledge, including intentions of agents, their plans, goals, and problems.
  This will be based on general referentially-based relations (which is now covered by Hobbs-style coherence relations for causality and Lascarides/Asher-style relations for temporality).
- interpersonal: the enaction of interaction, with direct alterations of speaker/hearer states, beliefs, and attitudes.
  This will be based on an improved specification of relations of the kind seen in interpersonal rhetorical relations.
- textual: the means of putting text together using generic mode notions such as 'assessment', 'restatement', etc. plus the central textual notion from RST of nuclearity.

All of the relations at this level are pruned of imperfect indications of surface form (which is now covered by conjunctive relations).

- Generic Purposes (based on Hasan/Martin, but again external to coherence relation model proper).

In the remainder of this section, we will present examples of how this specification is to be made concrete.

5.1 Conjunctive relations

5.1.1 General Organization

The units involved at this level are 'messages'; messages are realized by grammatical units. However, internal conjunctive relations can organize longer chunks of discourse directly with explicit structuring devices (e.g., a text such as *First*...*Second*......*And finally,* ...). Such chunks can be quite lengthy, but we will hypothesise here that they remain describable as message complexes (i.e., messages combined by conjunctive relations).

The functions of the relations between units are classified in a type subsumption lattice with the following properties. Each type specifies an applicability constraint defined in terms of semantic categories and relations, and a realization constraint that restricts the expression of the related units in the lexicogrammar. For the former, we will take unresolved, metafunctionally broad specifications of the semantics of grammatical units—analogous to the Penman Sentence Plan Language (SPL) [Kasper, 1989], or Quasi-Logical Form (QLF) [van Eijck and Alshawi, 1992, Geurts and Rentier, 1993], as the basis. For the latter, it may in some cases be appropriate to further conditionalize the constraints according to the 'text type'. In general it should be made as straightforward as possible for the discourse analyst to add to and change these constraints.

Since a conjunctive relation analysis prefers, following [Martin, 1992, p253], an $\alpha\beta\gamma$ dependency analysis (in contrast to RST's preference for $\alpha\beta(\alpha\beta)$),[34] interpretation is attempted left-to-right on successive partial semantic specifications obtained from the source text lexicogrammatical units. Generation attempts similarly to construct left-to-right related dependency sequences of messages.

The relations to be hypothesized as holding follow the subsumption/classification network.[35]

[34] Using [Halliday, 1985b]'s dependency notation.

[35] As will be taken further in the computational specification, we accept that it will not always be formally and fully automatically possible to find the desired relations. In such cases the Dandelion discourse workbench should be able to offer options to the analyst for direct selection—one way of achieving this is on the model of a Penman-style user interface (e.g., [Penman Project, 1989, Bateman, 1994]), where insufficient semantic constraints for stepwise classification may be resolved by querying the user.

A possible complication of the structural configurations required is the ability of conjunctive relations to hold over sequences of grammatical units. We will also assume for the present that such sequences are always realizing message complexes—but this needs further confirmation.

5.1.2 Definitions

For illustrative purposes, we will propose definitions here for 'external consequential conjunctive relations'. For the discourse grammar, we will also address internal consequential conjunctive relations, temporal relations (to make contact with the account of temporal connective semantics given in Dandelion deliverable [Oversteegen, 1993]), and internal additive conjunctive relations. These provide a good range of rather diverse relation types. The starting point in each case is the classification offered in [Martin, 1992].

Martin usually presents his realizations in the form of example congruent grammatical conjunctions (e.g., examples for the relation [extending] are *further, furthermore*) or as example incongruent sentences/phrases (cf. examples in Section 3.1.3). We can specify this more concretely using semantic categories such as those proposed for the Nigel grammar [Matthiessen, 1992b]. This grammar already has extensive grammatical resources for connectives and so is a convenient choice; as long as a grammatical treatment includes a sufficient range of such connectives, however, nothing in our account particularly calls for one form of grammar rather than another.

In order to generate a sentence with a conjunction such as the following, for example:

However, a farmer kills a duckling.

the following semantic specification would be specified in Penman SPL notation (the input specification used for driving the Nigel grammar):

```
(e1 / kill
    :actor (p1 / farmer)
    :actee (o1 / duckling)

    :conjunctive-relation-q conjunctive
    :conjunctive-relation-id
          (rr  / rhetorical-relation
                :process-regulated-q notprocessregulated
                :joint-regard-q notjoint
                :presentational-relation-q notpresentational
                :hypothesis-opposition-q hypothesisopposition ))
```

As a first stage of specification, we can simply use the feature values marked in bold in this example as the realization for the set of conjunctive relation features corresponding to 'however'. The next stage is to replace the implementations of the 'inquiries' in the semantic specification (`:joint-regard-q`, etc.) so that they directly refer to the conjunctive relation hierarchy. This makes it possible to replace them in a declarative fashion just as has been done for the ideational inquiries in [Bateman *et al.*, 1992]. The realizational relationship itself will be expressed by a hierarchy of relations with domains drawn from the set of conjunctive relation features and ranges from the partial SPL- or QLF-like structures needed to motivate lexicogrammatical expression.

Martin's description of the semantics of consequentials is quite detailed, although informally expressed. Whereas temporal relations construct a connnection between events along a temporal dimension, consequential relations construct a connection either of *enablement* or *determination* [Martin, 1992, p193]. The generality of these notions corresponds to the underspecified nature of the kind of 'cause', or 'entailment', for natural language consequentials often refered to in the literature (for some discussion of this, see, e.g., the discussion of Mackie in [Lagerwerf and Oversteegen, 1994]).

The corresponding semantics at issue, Martin describes in terms of two kinds of 'modulation', modulation through *potentiality* and modulation through *obligation*.[36] Martin suggests that this is "the 'natural logic' of the distinction between sufficient and necessary conditions", but we will not make any use of this here as it does not relate particularly to the account developed by Martin. In addition to obligation, a relation can be modalized in terms of *probability* whereby the effect becomes non-actual, or potential ("Effect irrealis" [Martin, 1992, p194]).

Modalized relations can also indicate *inclination* (i.e., purpose). A further polarity opposition then has different meanings depending whether inclination is involved or not. Without inclination, the polarity opposition relates to the 'exclusive/non-exclusive'-ness of the relation; with inclination the opposition relates to whether the *inclination* is positive ('desire') or negative ('fear'). An example for exclusiveness is the following [Martin, 1992, p196]:

(2) a. **Unless** you go that way, you'll be there by six.
(i.e., As long as you don't go that way, you'll be there by six.)

b. **If** you go that way, you'll be there by six.

In (2a), a negative condition is expressed, which has the exclusive meaning that 'that way is the only way you can go wrong'; in (2b), there is no statement as to whether there are other ways of 'going there and getting there by six'.

As a prepatory stage to constructing a computational specification, these conditions can be straightforwardly represented in a quasi-logical notation. At present this will serve simply to clarify what terms need further definition. They are not intended as definitions in their own right. It must also be remembered that the statements made by these relations are in any case *assertions of the speaker*. That is, they are statements of what the speaker intends the hearer to believe; this may well differ from any 'actual' state of affairs. Some embedding of the account within one involving knowledge and belief is also therefore presumed.

The individual conditions propose a semantic contribution for each feature from the conjunctive relation classification hierarchy. They provide the 'semantics'—i.e., the reasons for choice—for the alternatives in the hierarchy in precisely the same way as choosers and inquiries standardly do for a systemic-functional grammar in the Penman-style. The combined semantics is at present simply the conjunction of the individual additions. The behaviour of embedded conjunctive relations will be specified more completely when the computational context of the discourse grammar is available. In general, a consequential relation between a and b, written `consequential (a, b)`, indicates that a 'generalized cause' relation holds between events a and b. This we will write $a \leadsto b$, although the precise relation is specialized by further conditions on the classification network.

Drawing on Martin, the following informational conditions for the features of the network shown in Figure 3 above can be given.[37]

- `manner (a, b)` :- $\leadsto\ <$ *generalized-potential.*
- `cause (a, b)` :- $\leadsto\ <$ *generalized-obligation.*
- `contingency (a, b)` :- *irreal* (b).
- `consequence (a, b)` :- *real* (b).

[36] As with all labels for abstract semantic categories, these must be understood in a somewhat 'generalized' sense. For appropriate warnings about such labels in general, see [Halliday, 1984].

[37] This listing does not capture the inter-dependencies and constraints on co-selection between features that the network represents. The $<$ represents subsumption.

- `positive-relation` (a, b) :- a $\rightsquigarrow$ b.
  (including conjunctive relation features: [desire], [inclusive])
- `negative-relation (a, b)` :- a $\rightsquigarrow \neg b$.
  (including conjunctive relation features: [fear], [exclusive])
- `inclination (a, b)` :- $\rightsquigarrow <$ *'is in order to achieve'*
- `counterfactual (a, b)` :- $\neg b$.
- `modality (possibility, probability, or certainty)` :- b follows from a with indicated degree of probability; notated as one of: $\rightsquigarrow_{\text{possible}}$, $\rightsquigarrow_{\text{probable}}$, or $\rightsquigarrow_{\text{certain}}$.
- `negative-relation (a, b)` $\wedge\neg$`inclination (a,b)` $\wedge$ `contingency (a,b)` :-
  $(\forall\ T : a', \quad \text{s.t.} \quad T : a' \neq T : a)(T : a' \rightsquigarrow b)$.
  where $T : a$ is an event/state of some specified type. This says that if any other appropriately typed event or state is adopted than that specified by a, then this will generalized cause the state/event b.

Depending on what semantic conditions hold, different paths through the conjunctive relation network are motivated (and accordingly different constraints on lexicogrammatical form are established). Conversely, different lexicogrammatical selections can constrain the semantic conditions that are entailed.

An example from the fund-raising corpus is:

(3) Unless we act immediately to identify, protect and manage these natural areas and the wildlife they harbor, the fast escalating pressure of civilization will triumph. *The Nature Conservancy* (41)

This has a lexicogrammatical realization (*Unless* α, β) which receives a conjunctive relation classification of: [contingency, condition, exclusive]. Grouping the semantic constraints entailed by these features together gives the following:

- `contingency`:
  - $irreal\,(\beta)\ \wedge$
- `exclusive`: (which entails `negative-relation`):
  - $\alpha \rightsquigarrow \neg\beta \quad \wedge.$
- `condition` (which rules out [purpose] and hence `inclination`, and because `negative-relation` also holds):
  - $(\forall\ T : a', \quad \text{s.t.} \quad T : a' \neq T : \alpha)(T : a' \rightsquigarrow \beta)$

That is: 'the fast escalating pressure of civilisation' triumphing is not yet a real situation (*irreal*); any course of action of the general type of 'acting' which is not the particular course of action 'acting immediately to identify, protect and manage these natural areas and the wildlife they harbor' will cause 'the fast escalating pressure of civilisation triumphing' to become real; the course of action 'acting immediately to identify, protect and manage these natural areas and the wildlife they harbor' will prevent this. Which is, arguably, part of what this sentence is attempting to say.

These definitions also support the paraphrases that Martin provides as examples of various conjunctive relations. For example, the two versions of example (2a) have the following consequential relation classifications respectively:

- [contingency:condition:exclusive](a,b).
- [contingency:condition:inclusive:factual:modality-certain]($\neg a, b$).

where

a: go that way.
b: get there by six.

Grouping the conditions for the first alternative gives the same case as above, i.e.:

`irreal (b)` and
$(\forall T : a', \text{s.t.} T : a' \neq T : a)(T : a' \rightsquigarrow b)$.

Grouping those for the second alternative gives:

`irreal (b)` and
$\neg a \quad \rightsquigarrow_{\text{certain}} \quad b$.

This shows the two to be largely equivalent, but with rather different perspectives or saliences.

In addition to modulation, Martin also describes a dimension of *concession*. As shown in Figure 3, this cross-classifies the other consequential relations, rather than standing as an additional type. Martin's motivation for this is that "concessives function as a kind of 'anti-modulation', cancelling the potentiality or obligation which would otherwise enable or determine the consequential relation between events." [Martin, 1992, p198] Here we expect a further area of interaction with the results of Dandelion deliverable [Oversteegen, in preparation]. In a related paper, [Lagerwerf and Oversteegen, 1994], show that concessive generally involve a notion of expectation, since in a sentence such as:

(4) Ben didn't improve his time even by training hard.

there is some expectation that there would normally be a connection between training and improving time. In addition, as [Lagerwerf and Oversteegen, 1994] show at length, the expectation is of a more general nature than the concrete cases refered to in a given sentence. That is, there does not need to be an expectation that Ben's training hard will improve his time, just that, in general, someone training hard should improve their time. The precise generalization that is required is, however, not straightforward to determine.

Here the two-levels introduced by a level of conjunctive relations and a level of 'rhetorical' relations appears also to be beneficial. We can adopt the same notion of an 'unresolved' piece of information as has been used for other relations and predicates in quasi-logical forms. That is, we will assume that the more general expectation, and the precise form of that expectation, is in fact generated in context on the basis of the direct expectation. That this is likely, and perhaps necessary, is supported by the fact that the degree of generalization required is very variable, sometimes not going beyond the particular circumstances set up in the text itself. This can readily be seen in the following constructed example.

> The sprinter Captain Marvel was unique. He could improve his time by eating sausages. But today something was wrong. He hadn't improved his time although he had eaten nothing but sausages for days.

The only expectation here is that given explicitly in the text, i.e., for Captain Marvel, eating sausages usually improves his time. The essential meaning of concession at this level is recognition of violated expectation. Thus we may directly relate the conjunctive relation `concession` with two non-generalized inferences:

- `concession (a, b)` :- $a > \neg P \quad \wedge \quad b > P$

where P is some expectation (often given explicitly by b), but relate only indirectly the relation at the conjunctive relation level with some defeasible inference at the contextualized situation level such as those argued for by [Lagerwerf and Oversteegen, 1994].

This gives for a `concession:manner (a, b)` interpretation of the sequence:

a: Ben trained hard.
b: Ben didn't improve his time.

an analysis of:

Ben trained hard $> \neg P \quad \wedge$
Ben improving his time $> P \quad \wedge$
Ben trained hard $\rightsquigarrow_{\text{potential}}$ Ben didn't improve his time.

Letting P become identical with b, gives the defeasible inference that Ben improved his time by training hard and the generalized cause of him in fact not improving his time in this case.

There is clearly very much more here to discuss; this will be taken up in the integration work in the context of the discourse grammar deliverables.

5.2 Ideational, Interpersonal, and Textual rhetorical relations

5.2.1 General organization

The units involved at this level are 'discourse segments'; discourse segments are arranged in hierarchical structures, and their leaf discourse segments are realized by (complexes of) messages. The root of a hierarchically organized set of discourse segments will be presumed to be a generic stage in the sense described in Dandelion deliverable [Lavid, 1993]. The relation of such stages to units suggested in deliverable [Rothkegel and Villiger, 1993] is also to be assessed.

One principle textual contribution of an RST-analysis and of the definitions on which it relies is the attribution of *nuclearity*. Differences in nuclearity are sufficient to define distinct RST relations (e.g. (non-)volitional-cause *vs.* (non-)volitional-effect). This combined with presumed subject-matter relations conflates a statement of ideational import and its textual presentation. For example, the definition of the subject-matter RST relation PURPOSE requires that:

$$me_N \ \textit{in-order-to} \ me_S$$

It is not then possible to recognize a PURPOSE relationship when the text for some reason wants to assign nuclearity to the second parameter.

It is also now necessary to begin refining the structure/relation issue for rhetorical organizations. We attempt in our definitions below

1. to separate out the notion of the set of possible coherence relations that may be selected, particularly how such a set is to be organized (paradigmatic organization), and
2. to separate out the discoursal structural consequences of selecting particular (types of) coherence relations (syntagmatic organization).

At present, the first of these is often collapsed together with a classification of the dependency relations, whereas taking the analogy with the grammar again, it would be more usual for such dependency relations to be considered part of the syntagmatic organization. This needs to be clarified.

5.2.2 Definitions

We classify possible rhetorical relations in terms of the three simultaneous metafunctional views introduced above.

Ideational rhetorical relations select some informational configuration for presentation and assign nuclearity status to some element of that configuration. In order to avoid rhetorical relations for very specific informational specifications—for which there is no support—the informational specifications are necessarily general. The terms for expressing these that we adopt will essentially be those of Hobbs. This offers a classification system that is ideationally based but discourse coherence motivated. This also means that we do not need to assume that the particular ideational specifications are 'already present' in some knowledge base, only that they may be inferred if necessary (i.e., if we wish to attempt to use some particular rhetorical relation and need to check applicability). We take 'occasion' and 'expansion' as basic since, as we have argued, Hobbs' remaining relations are less centrally informationally. This provides the initial ideational component of the rhetorical relation classification hierarchy.

The initial interpersonal component of the classification hierarchy is drawn largely from the effect fields of the relations so classified above.

The textual organization is similarly drawn from the effect fields of the textual rhetorical relations. It is probably theoretically desirable (but implementationally somewhat tricky) to have the assignment of nuclearity in this component also.

The complete lattice offers a three-way parallel classification of possible rhetorical relations—it thus represents the *functions* (discourse purposes) of such relations. This is the paradigmatic axis of the rhetorical relation resource description. The realizations of the classifications supported are partial discourse structures consisting of two discourse units related by a 'rhetorical relation'.[38] Each discourse unit has associated with it a set of constraints concerning possible further rhetorical decompositions. This is the syntagmatic axis of the rhetorical relation resource description.

This classification is shown in the form of a systemic network in Figure 20.

Each feature in the network can, in the normal fashion, set constraints for their realization. The realization here must be specified in terms of text segments, text segment constituency, and specified relations between text segments. Simple examples for RST-style relations would be:

- [cause/S_0-nucleus] $\searrow$ (Non-)Volitional-Result (TS_1, TS_2).
- [cause/S_1-nucleus] $\searrow$ (Non-)Volitional-Cause (TS_1, TS_2).
- [plan-step/S_0-nucleus] $\searrow$ Purpose (TS_1, TS_2).

[38] Note that the co-classification in effect implicitly provides the $m \times n$ rhetorical relations that [Moore and Pollack, 1992] reject as unreasonable.

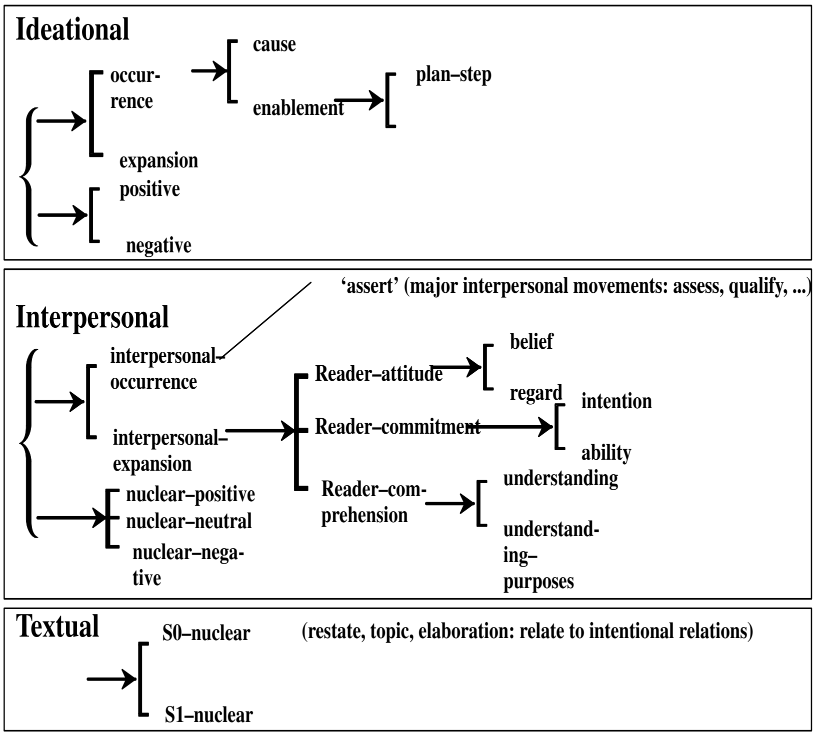

Figure 20: Initial Rhetorical relation paradigmatic network

- [occurence/S_1-nucleus] $\searrow$ Condition (TS_1, TS_2).
- [occurence/S_1-nucleus/negative] $\searrow$ Otherwise (TS_1, TS_2).
- [expansion] $\searrow$ Elaboration (TS_1, TS_2).

Just as reported in [Sanders, 1992], it is sometimes necessary to add additional information to that of the relation in order to obtain a given RST-relation (e.g., volition). In our further specification work, we will consider the features that Sanders proposes more closely.[39]

5.3 Relations between conjunctive relations and rhetorical relations

5.3.1 General organization

Our first proposal for relating the rhetorical relation resources and the conjunctive relation resources is based on the model of 'relational constraints' described by [Zajac, 1992]. The applicability of this approach for parallel classification hierarchies representing distinct linguistic strata has been shown in [Bateman *et al.*, 1992].

The form of these constraints is simply a subsumption lattice of relations, whose domains are drawn from the rhetorical relation resources and whose ranges are drawn from the conjunctive relation

[39] We will also consider some other proposals for the interpersonal part of the network, for example, those of [Lavid and Maier, 1992] and of [Maier and Sitter, 1992].

resources. This organization has as an advantage that it is nondirectional and straightforward to represent/implement. Whether it is sufficiently powerful must be evaluated by more detailed studies.

One complexity that we can only address in a very limited way here is that the conjunctive relations are not the only realizational possibilities for the rhetorical relations. As noted above, thematic progressions, temporal perspectives, and reference/lexical chains are also often involved. These co-constrain conjunctive relation choice and *vice versa*. The interaction with theme particularly needs to be addressed in the discourse grammar design deliverables, since thematic progression is one area that has received empirical investigation in Dandelion.

A further, related complexity is the different natures of the structures being related. Conjunctive relation structures are sequential, left-to-right oriented, while rhetorical relation structures are organized around the importance assigned by a speaker to text segments according to his/her communicative purposes. the reconciliation of these modes will no doubt require considerable further work.

5.3.2 Definitions

As examples here we present codings of some of the realizations discussed in Section 4 above. There we saw, for example, that the RST relation Volitional Result is mapped onto **additive:external:addition:positive** and **additive:internal:turnbuilding:developing:addition:simple** in the Gulf War articles, whereas it is mapped onto **consequential:external:manner** and **consequential:internal:concessive:objection:concede:mod.certain** in the fund-raising texts. If this pattern were supported by more data, then it would be straightforward to state disjunctions of relations between the corresponding members of the RST-conjunctive relation pairs. Then, more interesting, is the specification of constraints concerning which of the theoretically possible pairs is to be adopted in a particular case. This will probably need to consider information internal to the items being related as well as text type considerations.

5.4 Mechanisms

There are a number of possible mechanisms for managing the above resources during generation and analysis and the specification of this processing will require considerable further work. This will also impact directly on any attempts to formulate 'discourse grammars', since that is where the necessary mechanisms must be situated. In general, however, we will want to achieve a model where interpretation/production is free to proceed incrementally across all levels simultaneously. The following questions must be answerable:

- what is the conjunctive relation status of a contribution and where does it attach in the conjunctive relation structure?
- what is the contribution to field? (external conjunctive relations)
- what is the contribution to genre (to which generic stage does it make a contribution) and mode (and where does it attach in the discourse structure)? (internal conjunctive relations)
- what are the constraints between the rhetorical specification and the conjunctive specification?

It is also possible that each of the proposed levels has its own notion of 'discourse structure'. The nature of the units related at each level must therefore be clarified. Relations between units and relations at distinct levels must also be more tightly specified before a computational specification can be finalized.

There are a number of mechanisms proposed in the literature that need to be investigated for their adequacy with respect to our requirements for an account of coherence relations. In this section, we will simply list the most promising. These will be taken up again in the appropriate Discourse Grammar specification.

- The adequacy of simple constraint resolution using a system such as TFS and CUF needs to be evaluated in order to concretely assess whether, and precisely where, more complicated mechanisms might be necessary.
- The Asher/Lascarides/Oberlander algorithm offers a method for the left-to-right construction of a discourse structure. This needs to be investigated for a conjunctive relation analysis also.
- The LRE 'Towards a Declarative Representation of Discourse' is investigating (distinct) methods for generation and analysis. These need to be evaluated with respect to our needs.
- The discourse planning system DPOCL [Young and Moore, 1994] intended as a formally adequate planner for non-hierarchical rhetorical structures using rhetorical relations of the kind defined in RST. The system is claimed to be "sound, and for some class of plans, complete." This needs to be investigated for possible application to our rhetorical level.
- Possibilities of the kind of generation by abduction discussed in [Lascarides and Oberlander, 1992] need also to be investigated.

6 Applications of the account to noted problems in the literature

6.1 The benefit of nuclearity assignment

It would seem that nuclearity assignment is a clearly textual decision that should not be 'hardcoded' as combining in particular fixed orientations with ideational content. In then considering possible motivations for nuclearity assignment, it is natural to consider more global text organizations such as thematic development (cf. Dandelion deliverable R1.2.1), generic structure (cf. Dandelion deliverable R1.1.1b), and topic as discussed in, e.g., [Asher, 1993].

The following example from the fund-raising letter corpus demonstrates this well.

> (12) We have a staggeringly difficult job to do and not much time to do it. (13) So we urgently need all the support we can get.

The issue is whether this should receive a non-volitional cause analysis (as [Abelen *et al.*, 1993] assign it) or a non-volitional result analysis (as [Julia Lavid and Alonso, 1994] assign it).

Independently motivated nuclearity assignment then also dispenses with the problem of conflating nuclei raised for RST by [Asher and Lascarides, 1994]. We maintain nuclearity as a textual status but do not need to attribute it definitionally to ideational relations independently of their use in text.

Separating out nuclearity from informational specifications also provides a more satisfactory resolution of the problems raised by [Moore and Pollack, 1992]. The important contribution of an informational RST relation is *not* the informational specification, but the *assignment of nuclearity to that information.* This then fully supports Mann and Thompson's method for reconciling possibly conflicting RST analyses that is criticised by Moore and Pollack. The relevant quotation is as follows.

> "Imagine that a satellite provides evidence for a particular proposition expressed in its nucleus, and happens to do so by citing an attribute of some element expressed in the nucleus. Then. . . the conditions for both EVIDENCE and ELABORATION are fulfilled. If the analyst sees the speaker's purpose as increasing the hearer's belief of the nuclear propositions, and not as getting the hearer to recognize the `object:attribute` relationship, then the only analysis is the one with the EVIDENCE relation." [Mann and Thompson, 1987, p30]

Against this, Moore and Pollack argue that "Even if the speaker aims to get the hearer to recognize some `object:attribute` relationship, she has some underlying intention for doing that: she wants to enable the hearer to perform some action, or to increase the hearer's belief in some proposition, etc." [Moore and Pollack, 1992, p539] This appears to somewhat arbitrarily remove:

> *inform the hearer that the nucleus has an attribute as given in the satellite*

from the 'proper' intentions of a speaker. Both this *and* the intention admitted by Moore and Pollack, i.e.,

> *increase the hearer's belief in the nuclear proposition*

are perfectly viable intentions. And, morever, *both* are crucially distinct from the communicative intention-free, informational specification that an `object:attribute` relationship obtains.

This does, however, predict gaps in the RST-relation list standardly given—i.e., there should be relations where the informational specification is unchanged but nuclearity assignment is different; one candidate here was the PURPOSE relation mentioned above. Mann and Thompson only offer a complete paradigm for the causal relationships.

These gaps can lead to misleading (or mistaken) analyses. Moore and Pollack's second main example, which claims that informational and intentional RST analyses can have non-isomorphic RST structures, illustrates this. The text fragment at issue is as follows.

> (a) Come home by 5:00. (b) Then we can go to the hardware store before it closes. (c) That way we can finish the bookshelves tonight. [Moore and Pollack, 1992, p542]

Moore and Pollack's competing analyses are:

- intentional: motivation (N: a, S: motivation (N: b, S: c))
- informational: condition (N: c, S: (condition (N: b, S: a))

But the second analysis should only apply if the analyst is really prepared to admit (c) as the main nucleus. If that were the case, then the motivation analysis would not be possible. Moore and Pollack's argument is then in fact equivalent to stating that if the analyst interprets the text in two different ways, then these may be non-isomorphic—which is not surprising. What must be strongly emphasized in relation application is, again, the crucial nature of the *assignment of nuclearity*. This cannot be considered an optional add-on to be done after selecting relations. One cannot reason from the applicability of an RST-relation on informational grounds to the allocation of nuclearity; this would be to claim that one can map from informational accounts to intentional accounts, a position that Moore and Pollack themselves explicitly reject.[40] This also includes [Asher and Lascarides,

[40] Deciding on nuclearity—which in the present example would appear to be linguistically well motivated to fall on (a)—would here lead to the isomorphic informational selection of PURPOSE as the appropriate analysis. This is arguably a better analysis of this text in any case.

1994] solution, which rests on allowing both subordinating and non-subordinating relations (thereby escaping conflicting nuclearity assignments), but improves on their account by making it clear why informational specifications do not impose nuclearity—they are not textual.

We in fact find a similar, but much more restricted, conflation of functions in Hobbs' account. On closer analysis, the Evaluation, Ground-Figure, and Explanation relations combine both an intentional communicative goal and an informational applicability condition in a manner which is probably too rigid. Hobbs claims that:

> If S_1 *causes* S_0,
> then S_1 *explains* S_0.
> If S_1 *causes the speaker to say* S_0, then S_1 *evaluates* S_0.
>
> [Hobbs, 1990, p90]

Giving a cause may in some cases serve as an explanation, but the simple stating of this condition hardly does justice to the complexity and contingency of this relation; similarly for evaluation. Accordingly, we will henceforeth separate out the purer ideational relations of Hobbs account—occasion and expansion—from those of a mixed status.

6.2 Relation to text type

We can now also begin to make the relationship with text types more explicit. One can see in each of RST, Hobbs, and Asher/Lascarides the somewhat ambivalent influence of text types. As we have noted, Hobbs has 'genres' as constrained discourse structures [Hobbs, 1990, p107], RST as precompiled RST analyses, Asher/Lascarides as defaults for relation interpretation. There is no clear line drawn between such generic structures and discourse structures—particularly when the generic structure has only two elements. Examples here are Hobbs' *function-structure* elaboration [Hobbs, 1990, p97] and the "convention of argumentation for there to be just three [parallel arguments], and those ordered by increasing strength." [Hobbs, 1990, p107] A further candidate may well be both Hobbs and RST's *problem-solution* relation. In Mann and Thompson's analysis of American texts this was a very common organizing device for larger segments of discourse—i.e., as well-motivated a generic structure as one is likely to find.[41]

Finally, Asher and Lascarides' use of 'narrative' as a relation is probably best seen as conflating two notions: first, that one way of relating discourse units is by temporal sequences (i.e., RST sequence along a temporal dimension); and, second, that such a discourse relation is frequently employed in narrative 'text types' (or, to be more specific, in *field-structured* texts, where the text (re-)constructs an activity sequence [Martin, 1992, p517]). The value of separating these is that sequences occur in texts that would not be classified in any normal way as 'narratives', and 'narrative' texts have many more features than simple sequence in common (cf. [Polanyi, 1978, Labov and Waletzky, 1978]). 'Narrative' as a relation is thus best avoided, contributing as it does to a weakening, or overloading, of the technical vocabulary.

7 Conclusions

We have now provided a brief appraisal of the theoretical problems of coherence relations and their practical use in text analysis. We can sum up the weaknesses and strengths of previous approaches as follows:

[41] There is no doubt going to be diachronic slippage here. A context-restricted generic structure may, with frequent use, slip over into a coherence relation with wider applicability.

- under-differentiation above syntax in semantic/pragmatics,
- under-differentiation within syntax (structurally based and functionally lean).

We build on this in order to construct an initial framework and to set out some further investigative goals to be achieved by continuing analysis of texts in terms of coherence relations. Achieving these goals should directly contribute to the development of a body of linguistic *resources for achieving coherence* that is usable for text generation and which also significantly improves our theoretical understanding of the nature of coherence relations. We also need to specify appropriate processing mechanisms for these resources, although at present this can only be done in terms of suggested design constraints on the representation to be constructed for the Dandelion workbench (Task 2.2.1).

The preliminary organization of resources suggested can be set out in terms of the following basic properties. There exist several levels of description with correspondence mappings between these; the levels are characterized by the kinds of information they contain. Each level is further defined as a paradigmatic set of options, with each alternation constrained by configurations drawn from the level above. The particular alternations also have specified constraints for the level below.

In addition, we have posed and offered affirmative answers for the following questions:

- **Is it possible to obtain more information useful for coherence relation recognition/use from further differentiation within grammar?**—i.e., widening evidence.
- **Is it possible/useful to provide further differentiation in semantics/pragmatics?**

We therefore hope that the further specification of coherence relations along the lines began here will provide many additional sources of constraint for discourse processing mechanisms being proposed elsewhere.

Acknowledgments

Thanks are due to the many analyses carried out by Julia Lavid and colleagues (Universidad Complutense Madrid) and the detailed discussions and input of Wilbert Spooren and Leonoor Oversteegen (CLS) and Gisela Redeker (Vrije Universiteit Amsterdam) for making the Fund Raising Letter corpus and analyses available to us and for detailed criticisms of our RST analyses.

References

[Abelen *et al.*, 1993] Eric Abelen, Gisela Redeker, and Sandra A. Thompson. The rhetorical structure of US-American and Dutch fund-raising letters. *Text*, 13(3):323 – 350, 1993.

[André and Rist, 1993] Elisabeth André and Thomas Rist. The design of illustrated documents as a planning task. In Mark T. Maybury, editor, *Intelligent Multimedia Interfaces*, pages 94 – 116. AAAI Press/The MIT Press, Menlo Park (CA), Cambridge (MA), London (England), 1993.

[Asher and Lascarides, 1994] Nicholas Asher and Alex Lascarides. Intentions and information in discourse. In *32nd. Annual Meeting of the Association for Computational Linguistics*, pages 34 – 41, New Mexico State University, Las Cruces, New Mexico, 1994. Also available as `CMP-LG` E-print Archive paper: `cmp-lg/94060001`.

[Asher, 1993] Nicholas Asher. *Reference to abstract objects in discourse*. Kluwer Academic Publishers, Dordrecht, 1993.

[Bateman and Matthiessen, 1990] John A. Bateman and Christian M.I.M. Matthiessen. Uncovering the text base, 1990. (To appear in *Selected Papers from the International Conference on Research in Text and Language, 29-31 March 1989, Xi'an Jiaotong University*, edited by Hao Keqi and Hermann Bluhme, Xi'an Jiaotong University Press, P.R. China, 1993).

[Bateman *et al.*, 1991] John A. Bateman, Elisabeth A. Maier, Elke Teich, and Leo Wanner. Towards an architecture for situated text generation. In *International Conference on Current Issues in Computational Linguistics*, Penang, Malaysia, 1991. Also available as technical report of GMD/Institut für Integrierte Publikations- und Informationssysteme, Darmstadt, Germany.

[Bateman *et al.*, 1992] John A. Bateman, Martin Emele, and Stefan Momma. The nondirectional representation of Systemic Functional Grammars and Semantics as Typed Feature Structures. In *Proceedings of COLING-92*, volume III, pages 916 – 920, 1992.

[Bateman, 1994] John A. Bateman. KPML: The KOMET-Penman (Multilingual) Development Environment. Technical report, Institut für Integrierte Publikations- und Informationssysteme (IPSI), GMD, Darmstadt, September 1994. Release 0.6.

[Bloembergen and Gardent, 1994] Siebe Bloembergen and Claire Gardent. Parsing discouse. Technical report, Amsterdam, April 1994. [DRAFT].

[Brew *et al.*, 1993] Chris Brew, Joke Dorrepaal, Claire Gardent, Claire Grover, Suresh Manandhar, Marc Moens, Hub Prüst, and Andreas Schöter. Representation of discourse information. Technical report, University of Edinburgh, December 1993.

[Cohen and Levesque, 1985] Philip R. Cohen and Hector J. Levesque. Speech acts and rationality. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 49–59, Chicago, 1985.

[Cohen and Levesque, 1987] Philip R. Cohen and Hector J. Levesque. Persistence, intention and commitment. Technical Report CSLI-87-88, Center for the Study of Language and Information, Stanford, CA, March 1987.

[Cohen and Perrault, 1979] P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:pages 177 – 212, 1979.

[Cross, 1992] Marilyn Cross. *Choice in text: a systemic approach to computer modelling of variant text production.* PhD thesis, School of English and Linguistics, Macquarie University, Sydney, Australia, 1992.

[Dale, 1993] Robert Dale. Rhetoric and intentions in discourse. In Owen Rambow, editor, *Intentionality and structure in discourse relations*, pages 5 – 6. Association for Computational Linguistics, 1993. (Proceedings of a Workshop sponsored by the Special Interest Group on Generation, 21 June, 1993, Columbus, Ohio).

[Daneš, 1974] František Daneš. Functional sentence perspective and the organisation of the text. In František Daneš, editor, *Papers on Functional Sentence Perspective*, pages 106–128. Academia, Prague, 1974.

[Degand and Bateman, 1993] Liesbeth Degand and John Bateman. Analysis of coherence relations in texts: starting points. Technical report, GMD-IPSI, Darmstadt, FRG, September 1993. (Preliminary version of ESPRIT Basic Research Action: Dandelion, EP6665; Deliverable R1.1.2a).

[Dobeš and Novak, 1991] Zuzana Dobeš and Hans-Joachim Novak. From knowledge structures to text structures. In O. Herzog and C.-R. Rollinger, editors, *Text understanding in LILOG: integrating computational linguistics and artificial intelligence, Final report on the IBM Germany LILOG-Project*, pages 670 – 684. Springer-Verlag, Berlin, 1991. Lecture notes in artificial intelligence, 546.

[Fawcett and Davies, 1992] Robin P. Fawcett and Bethan L. Davies. Monologue as a turn in dialogue: towards an integration of exchange structure theory and rhetorical structure theory. In Robert Dale, Eduard Hovy, Dietmar Rösner, and Olivero Stock, editors, *Aspects of automated natural language generation*, pages 151 – 166. Springer-Verlag, 1992. (Proceedings of the 6th International Workshop on Natural Language Generation, Trento, Italy, April 1992).

[Fox, 1987] Barbara A. Fox. *Discourse Structure and Anaphora in Written and Conversational English.* Cambridge University Press, Cambridge, 1987.

[Gardent, 1994] Claire Gardent. Parsing discouse. Technical report, University of Clermont-Ferrand, April 1994. [DRAFT].

[Geurts and Rentier, 1993] Bart Geurts and Gerrit Rentier. Quasi Logical Form in PLUS, 1993. Deliverable to WP3 of the PLUS Project: *A pragmatics-based language understanding system*, ESPRIT Project P5254.

[Groenendijk and Stokhof, 1991] J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39 – 100, 1991.

[Grosz and Sidner, 1986] Barbara J. Grosz and Candace L. Sidner. Attention, intentions and the structure of discourse. *Computational Linguistics Journal*, 12(3):175–204, 1986.

[Halliday and Hasan, 1976] Michael A.K. Halliday and Ruqaiya Hasan. *Cohesion in English.* Longman, London, 1976.

[Halliday, 1978] Michael A.K. Halliday. *Language as social semiotic.* Edward Arnold, London, 1978.

[Halliday, 1984] Michael A.K. Halliday. On the ineffability of grammatical categories. In *Tenth LACUS Forum*, Columbia, 1984. Hornbeam Press.

[Halliday, 1985a] Michael A.K. Halliday. *An Introduction to Functional Grammar.* Edward Arnold, London, 1985.

[Halliday, 1985b] Michael A.K. Halliday. *An Introduction to Functional Grammar.* Edward Arnold, London, 1985.

[Heim, 1983] I. Heim. File change semantics and the familiarity theory of definiteness. In C. Schwarze and A. von Stechow, editors, *Meaning, Use and Interpretation of Language*, pages 164 – 178. De Gruyter, Berlin, 1983.

[Hobbs *et al.*, 1992] Jerry R. Hobbs, Douglas E. Appelt, John Bear, and Mabry Tyson. Robust processing of real-world natural-language texts. In *Proceedings of the Third Conference on Applied Natural Language Processing*, pages 186 – 192. Association for Computational Linguistics, 1992. 31 March - 3 April.

[Hobbs *et al.*, 1993] Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69 – 142, 1993. (ARTINT 1059).

[Hobbs, 1986] Jerry R. Hobbs. Sublanguage and knowledge. In Ralph Grishman and Richard Kittredge, editors, *Analyzing language in restricted domains: sublanguage description and processing*, pages 53 – 68. Lawrence Erlbaum Associates, Hillsdale, New Jersey and London, 1986.

[Hobbs, 1990] J.R. Hobbs. *Literature and cognition.* CSLI, Stanford, California, 1990. CSLI Lecture Notes.

[Hovy *et al.*, 1992] Eduard Hovy, Julia Lavid, Elisabeth Maier, Vibhu Mittal, and Cécile Paris. Employing knowledge resources in a new text planner architecture. In *Aspects of automated natural language generation*, pages 57 – 72, Berlin, 1992. Springer-Verlag. (Proceedings of the 6th International Workshop on Natural Language Generation, Trento, Italy, April 1992).

[Hovy, 1988] Eduard H. Hovy. Planning coherent multisentential texts. In *The Proceedings of the 26th. Annual Meeting of the Association of Computational Linguistics*, pages 163 – 169. Association for Computational Linguistics, June 1988.

[Hovy, 1990] Eduard Hovy. Parsimonious and profligate approaches to the question of discourse structure relations. In *5th. International Workshop on Natural Language Generation, 3-6 June 1990*, Pittsburgh, PA., 1990. Organized by Kathleen R. McKeown (Columbia University), Johanna D. Moore (University of Pittsburgh) and Sergei Nirenburg (Carnegie Mellon University).

[Hovy, 1991] Eduard H. Hovy. Approaches to the planning of coherent text. In Cécile L. Paris, William R. Swartout, and William C. Mann, editors, *Natural language generation in artificial intelligence and computational linguistics.* Kluwer Academic Publishers, July 1991. Presented at the Fourth International Workshop on Natural Language Generation. Santa Catalina Island, California, July, 1988.

[Hovy, 1993] Eduard Hovy. In defense of syntax: informational, intentional, and rhetorical structures in discourse. In Owen Rambow, editor, *Intentionality and structure in discourse relations*, pages 35 – 39. Association for Computational Linguistics, 1993. (Proceedings of a Workshop sponsored by the Special Interest Group on Generation, 21 June, 1993, Columbus, Ohio).

[Jordan, 1988] Michael P. Jordan. Some advances in clause-relational theory. In James D. Benson and William S. Greaves, editors, *Systemic Functional Approaches to Discourse: Selected Papers from the Twelfth International Systemic Workshop*, pages 282 – 301. Ablex, Norwood, New Jersey, 1988.

[Julia Lavid and Alonso, 1994] Francisco José Ballesteros Julia Lavid and Isabel Alonso. Conjunctive relations analysis of the fund-raising texts corpus, 1994.

[Kamp and Reyle, 1993] Hans Kamp and Uwe Reyle. *From discourse to logic: introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory.* Kluwer Academic Publishers, London, Boston, Dordrecht, 1993. Studies in Linguistics and Philosophy, Volume 42.

[Kamp, 1981] H. Kamp. A theory of truth and semantic interpretation. In *Truth, Interpretation and Information.* Foris Publications, Dordrecht, 1981.

[Kasper, 1989] Robert T. Kasper. A flexible interface for linking applications to PENMAN's sentence generator. In *Proceedings of the DARPA Workshop on Speech and Natural Language*, 1989. Available from USC/Information Sciences Institute, Marina del Rey, CA.

[Knott and Dale, 1992] Alistair Knott and Robert Dale. Using linguistic phenomena to motivate a set of rhetorical relations. Technical Report HCRC/RP-39, Human Communication Research Centre, University of Edinburgh, Edinburgh, Scotland, 1992.

[Korelsky and Kittredge, 1993] Tania Korelsky and Richard Kittredge. Towards stratification of RST. In Owen Rambow, editor, *Intentionality and structure in discourse relations*, pages 52 – 55. Association for Computational Linguistics, 1993. (Proceedings of a Workshop sponsored by the Special Interest Group on Generation, 21 June, 1993, Columbus, Ohio).

[Labov and Waletzky, 1978] W. Labov and J. Waletzky. Narrative analysis. In J. Helm, editor, *Essays on the verbal and visual arts*, pages 12 – 44. University of Washington Press, Seattle, 1978. (Proceedings of the 1966 Spring Meeting of the American Ethnological Society).

[Lagerwerf and Oversteegen, 1994] Luuk Lagerwerf and Leonoor Oversteegen. On causal contrastive relations. Paper in preparation; Tilburg University, May 1994.

[Lascarides and Asher, 1991] Alex Lascarides and N. Asher. Discourse relations and defeasible knowledge. In *Proceedings of the 29th. Annual Meeting of the Association for Computational Linguistics*, pages 55 – 63, Berkeley, California, 1991. Association for Computational Linguistics.

[Lascarides and Asher, 1993] Alex Lascarides and Nicholas Asher. A semantics and pragmatics for the pluperfect. In *Proceedings ofthe Seventh Conference of the European Chapter of the Association for Computational Linguistics*, pages 250 – 259, Utrecht, the Netherlands, 1993.

[Lascarides and Oberlander, 1992] Alex Lascarides and Jon Oberlander. Abducing temporal discourse. In *Aspects of automated natural language generation*, pages 167 – 182, Berlin, 1992. Springer-Verlag. (Proceedings of the 6th International Workshop on Natural Language Generation, Trento, Italy, April 1992).

[Lascarides and Oberlander, 1993] Alex Lascarides and Jon Oberlander. Temporal connectives in a discourse context. In *Proceedings of the Seventh Conference of the European Chapter of the Association for Computational Linguistics*, pages 260 – 268, Utrecht, the Netherlands, 1993.

[Lavid and Maier, 1992] Julia Lavid and Elisabeth Maier. Text relations: their usefulness for a dynamic account of discourse structure in a text planning system. Technical report, GMD/Institut für Integrierte Publikations- und Informationssysteme, Darmstadt, Germany, 1992.

[Lavid, 1993] Julia Lavid. Generic structure potential: a functional characterization of global discourse structures. Technical report, Universidad Complutense Madrid, Madrid, Spain, September 1993. (ESPRIT Basic Research Action: Dandelion, EP6665; Deliverable R1.1.1b).

[Maier and Hovy, 1993] Elisabeth A. Maier and Eduard H. Hovy. Organizing discourse structure relations using metafunctions. In Helmut Horacek and Michael Zock, editors, *New concepts in natural language generation*, pages 69 – 86. Pinter Publishers, London, 1993.

[Maier and Sitter, 1992] Elisabeth Maier and Stefan Sitter. An extension of rhetorical structure theory for the treatment of retrieval dialogues. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 968 – 973. Lawrence Erlbaum Associates, 1992. July 29 to August 1.

[Mann and Matthiessen, 1990] William C. Mann and Christian M. I. M. Matthiessen. Functions of language in two frameworks. Technical Report ISI/RR-90-290, Information Sciences Institute, Los Angeles, California, 1990. (Presented at the 1987 International Systemic Workshop).

[Mann and Thompson, 1987] William C. Mann and Sandra A. Thompson. Rhetorical structure theory: a theory of text organization. Technical Report RS-87-190, USC/Information Sciences Institute, 1987. Reprint series.

[Mann *et al.*, 1989] William C. Mann, Christian M. I. M. Matthiessen, and Sandra A. Thompson. Rhetorical structure theory and text analysis. Technical Report ISI/RR-89-242, University of Southern California, Information Sciences Institute, Marina del Rey, California, 1989. (A revised version of this paper is published in: Mann, William C. and Sandra A. Thompson (eds.) *Text Description: Diverse Analyses of a Fund Raising Text* (Amsterdam: John Benjamins), 1993).

[Martin, 1992] James R. Martin. *English text: systems and structure*. Benjamins, Amsterdam, 1992.

[Matthiessen and Bateman, 1991] Christian M.I.M. Matthiessen and John A. Bateman. *Text generation and systemic-functional linguistics: experiences from English and Japanese*. Frances Pinter Publishers and St. Martin's Press, London and New York, 1991.

[Matthiessen and Thompson, 1989] Christian Matthiessen and Sandra A. Thompson. The structure of discourse and 'subordination'. In Haiman and Thompson, editors, *Clause Combining in Grammar and Discourse*. Benjamins, Amsterdam, 1989. Also available as ISI/RS-87-183.

[Matthiessen, 1992a] Christian M.I.M. Matthiessen. Lexicogrammatical cartography: English systems. Technical report, University of Sydney, Linguistics Department, 1992. Ongoing expanding draft.

[Matthiessen, 1992b] Christian M.I.M. Matthiessen. Lexicogrammatical cartography: English systems. Technical report, University of Sydney, Linguistics Department, 1992. Draft 5.

[Meteer, 1993] Marie Meteer. Assumptions underlying discourse relations: which ones are there and where are they? In Owen Rambow, editor, *Intentionality and structure in discourse relations*, pages 82 – 85. Association for Computational Linguistics, 1993. (Proceedings of a Workshop sponsored by the Special Interest Group on Generation, 21 June, 1993, Columbus, Ohio).

[Moore and Paris, 1988] Johanna D. Moore and Cécile L. Paris. Constructing coherent texts using rhetorical relations. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Cognitive Science Society, August 1988.

[Moore and Paris, 1992] Johanna D. Moore and Cécile L. Paris. Exploiting user feedback to compensate for the unreliability of user models. *User Modeling and User-Adapted Interaction*, 2:287 – 330, 1992.

[Moore and Paris, 1993] Johanna D. Moore and Cécile L. Paris. Planning texts for advisory dialogs: capturing intentional and rhetorical information. *Computational Linguistics*, 19(4):651 – 694, December 1993.

[Moore and Pollack, 1992] Johanna D. Moore and Martha E. Pollack. A problem for RST: the need for multi-level discourse analysis. *Computational Linguistics*, 18(4):537 – 544, December 1992.

[Moore and Swartout, 1991] Johanna D. Moore and William R. Swartout. A reactive approach to explanation. In Cécile L. Paris, William R. Swartout, and William C. Mann, editors, *Natural language generation in artificial intelligence and computational linguistics*. Kluwer Academic Publishers, July 1991. Presented at the Fourth International Workshop on Natural Language Generation. Santa Catalina Island, California, July, 1988.

[Oversteegen, 1993] Leonoor Oversteegen. Dutch temporal connectives. Technical report, Center for Language Studies, Tilburg, Holland, September 1993. (ESPRIT Basic Research Action: Dandelion, EP6665; Deliverable R1.3.2a).

[Oversteegen, in preparation] Leonoor Oversteegen. Causal and contrastive connectives. Technical report, Center for Language Studies, Tilburg, Holland, in preparation. (ESPRIT Basic Research Action: Dandelion, EP6665; Deliverable R1.1.3a).

[Paris and Maier, 1991] Cécile L. Paris and Elisabeth Maier. Knowledge resources or decisions? In *Proceedings of the 1991 IJCAI Workshop on Decision making in the generation process*, Sydney, Australia, 1991.

[Paris, 1991] Cécile L. Paris. Generation and explanation: Building an explanation facility for the explainable expert systems framework. In Cécile L. Paris, William R. Swartout, and William C. Mann, editors, *Natural language generation in artificial intelligence and computational linguistics*. Kluwer Academic Publishers, July 1991. Presented at the Fourth International Workshop on Natural Language Generation. Santa Catalina Island, California, July, 1988.

[Penman Project, 1989] Penman Project. The Nigel manual. Technical report, USC/Information Sciences Institute, Marina del Rey, California, 1989.

[Polanyi, 1978] Livia Polanyi. *The American Story: Cultural constraints on the meaning and structure of stories in conversation*. PhD thesis, University of Michigan, 1978.

[Polanyi, 1985] Livia Polanyi. A theory of discourse structure and discourse coherence. In W. H. Elifor, P.D. Kroeber, and K.L. Peterson, editors, *Papers from the General Session of the 21st. Regional Meeting of the Chicago Linguistics Society*. Chicago, 1985.

[Polanyi, 1986] Livia Polanyi. The linguistic discourse model: Towards a formal theory of discourse s tructure. Technical Report Technical Report 6409, BBN Laboratories Inc, Cambridge, MA, 1986.

[Redeker, 1990] G. Redeker. Ideational and pragmatic markers of discourse structure. *Journal of Pragmatics*, 14:367–381, 1990.

[Redeker, 1991] Gisela Redeker. Linguistic markers of discourse structure. *Linguistics*, 29:1139 – 1172, 1991.

[Redeker, 1992] Gisela Redeker. Coherence and structure in text and discourse, 1992. Manuscript. Tilburg University.

[Reichman, 1978] Rachel Reichman. Conversational coherency. *Cognitive Science*, 2, 1978.

[Reichman, 1985] Rachael Reichman. *Making computers speak like you and me*. The M.I.T. Press, Cambridge, MA, 1985.

[Rösner and Stede, 1992] Dietmar Rösner and Manfred Stede. Customizing RST for the automatic production of technical manuals. In Robert Dale, Eduard Hovy, Dietmar Rösner, and Olivero Stock, editors, *Aspects of automated natural language generation*, pages 199–214. Springer-Verlag, 1992. (Proceedings of the 6th International Workshop on Natural Language Generation, Trento, Italy, April 1992).

[Rösner, 1993] Dietmar Rösner. Intentions, rhetoric, or discourse relations? a case from multilingual document generation. In Owen Rambow, editor, *Intentionality and structure in discourse relations*, pages 106 – 109. Association for Computational Linguistics, 1993. (Proceedings of a Workshop sponsored by the Special Interest Group on Generation, 21 June, 1993, Columbus, Ohio).

[Rothkegel and Villiger, 1993] Annely Rothkegel and Claudia Villiger. Chaining. Technical report, Universität des Saarlandes, Saarbrücken, FRG, September 1993. (ESPRIT Basic Research Action: Dandelion, EP6665; Deliverable R1.2.2a).

[Sanders *et al.*, 1992] Ted J.M. Sanders, Wilbert P.M. Spooren, and Leo G.M. Noordman. Towards a taxonomy of coherence relations. *Discourse Processes*, 15(1):1 – 36, January – March 1992.

[Sanders, 1992] Ted Sanders. *Discourse Structure and Coherence; Aspects of a cognitive theory of discourse representation*. Nevelland, Landegem, Belgium, 1992.

[Scha and Polanyi, 1988] Remko Scha and Livia Polanyi. An augmented context free grammar for discourse. In *Proceedings of the 12th. COLING*, pages 573 – 577, Budapest, 1988.

[Schiffrin, 1987] Deborah Schiffrin. *Discourse Markers*. Cambridge University Press, Cambridge, 1987.

[Scott and de Souza, 1990] Donia R. Scott and Clarisse Sieckenius de Souza. Getting the message across in RST-based text generation. In Robert Dale, Christopher S. Mellish, and M. Zock, editors, *Current Research in Natural Language Generation*. Academic Press, 1990. Paper presented at the 1989 European Natural Language Generation Workshop, Edinburgh, April.

[van Dijk, 1977] T. van Dijk. *Text and Context: Explorations in the Semantics and Pragmatics of Discourse*. Longman, 1977.

[van Eijck and Alshawi, 1992] Jan van Eijck and Hiyan Alshawi. Logical forms. In Hiyan Alshawi, editor, *The Core Language Engine*, chapter 2, pages 11 – 40. MIT Press, Cambridge, Massachusetts and London, England, 1992.

[Vander Linden, 1993] Keith Vander Linden. *Speaking of actions: choosing rhetorical status and grammatical form in instructional text generation*. PhD thesis, University of Colorado at Boulder, Department of Computer Science, July 1993. Appears as technical report CU-CS-654-93.

[Ventola, 1984] Eija Ventola. The dynamics of genre. *Nottingham Linguistics Circular*, 14, 1984.

[Ventola, 1987] Eija Ventola. *The Structure of Social Interaction: A Systemic Approach to the Semiotics of Service Encounters*. Frances Pinter, London, 1987.

[Young and Moore, 1994] R. Michael Young and Johanna D. Moore. DPOCL: a principled approach to discourse planning. In *Proceedings of the Seventh International Workshop on Natural Language Generation, Kennebunkport, Maine, USA, June 21-24, 1994*, Kennebunkport, Maine, USA, 1994.

[Zajac, 1992] Rémi Zajac. Inheritance and constraint-based grammar formalisms. *Computational Linguistics*, 18(2):159 – 182, June 1992. (Special issue on inheritance: 1).