Foundations of Computing Series

Logic and Information Flow

edited by

Jan van Eijck
and Albert Visser

Logic and Information Flow

edited by
Jan van Eijck and Albert Visser

The MIT Press
Cambridge, Massachusetts
London, England

This book was set in LaTeX by the authors and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Logic and information flow / edited by Jan van Eijck and Albert Visser.
p. cm. — (Foundations of computing)
Includes bibliographical references.
ISBN 0-262-22047-4
1. Computer science. 2. Natural language processing (Computer science) 3. Logic, Symbolic and mathematical. I. Eijck, J. van (Jan) II. Visser, Albert. III. Series.
QA76.L5662 1994
511.3—dc20 93-40778
CIP

Contents

11 Meeting Some Neighbours

Maarten de Rijke

11.1 Introduction

Over the past several years the computer science community seems to have lost interest in dynamic logic and related systems somewhat. In the philosophical community, on the other hand, more and more people have felt a need for systems in which changes and processes can be modelled. This has lead to the birth of quite a number of systems blessed with the predicate 'dynamic'.

In this chapter one such system, called *DML*, is taken as a starting point, and its connections with alternative dynamic proposals are examined. Specifically, a revision operator is defined in *DML* which can be shown to satisfy most of the postulates such operators are currently supposed to satisfy. Further links are established with terminological logic, Veltman's update semantics, and preferential reasoning. Technical results pertaining strictly to the dynamic modal system of this chapter are given in a companion paper.

The purpose of this chapter is to discuss links between a recent proposal for reasoning about the dynamics of information, called *dynamic modal logic* or *DML*, and other such proposals, as well as connections with some other formalisms in philosophical logic, cognitive science and AI. The key phrases common to most of the systems that come up in this note are (minimal) change and reasoning about information.

As many dynamic-like formalisms have been proposed over the last few years, the danger that several researchers might be re-inventing the wheel is not entirely fictitious. For that reason I think it is important to have occasional comparisons across platforms. As a result of such comparisons results known in one domain may shed light on problems in the other domains, allowing the field at large to benefit. And at a more down to earth level the obvious advantage of such comparisons is that they may serve as partial maps of rapidly changing research areas. Thus, the purpose of this note is to sketch such a partial map by comparing or unifying some related dynamic systems using the *DML* formalism.

What's commonly considered to be the minimal requirement for a system to be called dynamic, is that it has a notion of state, and a notion of change or transition from one state to another. States and transitions are precisely the basic ingredients of the system *DML*; in addition it has various systematic connections between those basics. Although *DML* may at first appear to be a somewhat unorthodox modal system, it can be analyzed using fairly traditional tools from modal logic, yielding results on its expressive power, the hardness of the satisfiability problem for the language, and axiomatic completeness.

The main benefits of using *DML* as a guide-line for linking a number of dynamic proposals are the fact that many dynamic proposals are, so to say, de-mystified by being embedded in a system itself comprising of two well-known components (Boolean algebra and relational algebra); the embedding of such proposals into (a fragment of) *DML* suggests natural additions to, and generalizations of, these proposals. Moreover, the work presented here shows how fairly orthodox dynamic proposals like *DML* can be used fruitfully far beyond their traditional boundaries.

In 11.2 I describe the basics concerning *DML*, including two ways of dealing with the states of *DML* models: one can either take the usual view as states as objects devoid of any structure, or one can endow them with an internal structure and logic of their own.

After that I move on to two connections between *DML* equipped with 'structureless' states and other systems. In 11.3 an example from cognitive science and AI is considered when I model certain postulates for theory change inside *DML*. In 11.4, a link is established between *DML* and terminological logic and knowledge representation. I obtain an exact match between *DML* and a KL-ONE dialect, called the Brink and Schmidt language, plus an axiomatization of the representable algebras underlying this language.

In 11.5 and 11.6 the states of our *DML*-models will be equipped with structure. This is needed in 11.5 to link *DML* to a system of update semantics from the philosophical logic tradition proposed by Frank Veltman, while 11.6 contains some suggestions on how one would have to go about dealing with preferences and other more complex systems in *DML*.

Section 11.7 rounds off the chapter with some conclusions and questions.

11.2 *DML*: A Quick Review

11.2.1 Basics

The system of dynamic modal logic *DML* figuring in this note first appeared in what's more or less its present form in (Van Benthem [4]), but parts of it can be traced back to Van Benthem[2]. The original application of the system was reasoning about the knowledge of a single agent, and the "epistemic moves" this agent makes in some cognitive space to acquire new knowledge. Thus, in *DML* provisions have been made to talk about transitions that represent the acquisition of new knowledge, and about transitions representing the loss or giving up of knowledge. Moreover, these transitions may be structured in a variety of ways. To sum up, the *DML*-language has Boolean ingredients to reason about the static aspects of the agents knowledge, and relational ingredients to reason about the dynamic aspects thereof. In addition there are systematic connections between the two realms, as depicted in Figure 11.1.

After some cleaning up had been performed, a stable version of the language was given in (De Rijke [23]). Here it is:

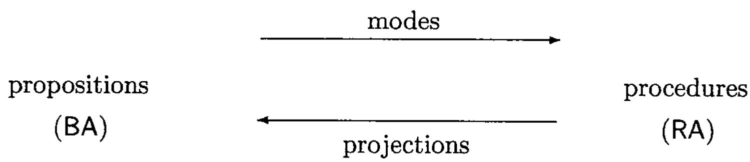

Figure 11.1
DML, the basic ingredients

Atomic formulas:	$p \in \Phi$,
Formulas:	$\varphi \in Form(\Phi)$,
Procedures:	$\alpha \in Proc(\Phi)$.

$\varphi ::= p \mid \bot \mid \top \mid \varphi_1 \to \varphi_2 \mid \mathsf{do}(\alpha) \mid \mathsf{ra}(\alpha) \mid \mathsf{fix}(\alpha)$,
$\alpha ::= \mathsf{exp}(\varphi) \mid \mathsf{con}(\varphi) \mid \alpha_1 \cap \alpha_2 \mid \alpha_1 ; \alpha_2 \mid -\alpha \mid \alpha \;\mid \varphi?$.
I will refer to elements of $Form(\Phi) \cup Proc(\Phi)$ as *expressions*.

The intended interpretation of the above connectives and mappings is the following. A formula $\mathsf{do}(\alpha)$ ($\mathsf{ra}(\alpha)$) is true at a state x iff x is in the domain (range) of α, and $\mathsf{fix}(\alpha)$ is true at x if x is a fixed point of α. The interpretation of $\mathsf{exp}(\varphi)$ (read: expand with φ) in a model $\mathfrak{M}$ is the set of all moves along the "informational ordering" in $\mathfrak{M}$ that take you to a state where φ holds; the interpretation of $\mathsf{con}(\varphi)$ (read: contract with φ) consists of all moves *backwards* along the ordering to states where φ *fails*; $\varphi?$ is the "test-for-φ" relation, while the intended interpretation of the operators left unexplained should be clear.

The models for this language are structures of the form $\mathfrak{M} = (W, \sqsubseteq, [\![\cdot]\!], V)$, where $\sqsubseteq \subseteq W^2$ is a transitive and reflexive relation (the informational ordering), $[\![\cdot]\!] : Proc(\Phi) \to 2^{W \times W}$, and $V : \Phi \to 2^W$. The interpretation of the modes is:

$\mathfrak{M}, x \models \mathsf{do}(\alpha)$ iff $\exists y\, ((x, y) \in [\![\alpha]\!])$,
$\mathfrak{M}, x \models \mathsf{ra}(\alpha)$ iff $\exists y\, ((y, x) \in [\![\alpha]\!])$,
$\mathfrak{M}, x \models \mathsf{fix}(\alpha)$ iff $(x, x) \in [\![\alpha]\!]$,

while the relational part is interpreted using the mapping $[\![\cdot]\!]$:

$$
\begin{aligned}
[\![\mathsf{exp}(\varphi)]\!] &= \lambda xy.\, (x \sqsubseteq y \wedge \mathfrak{M}, y \models \varphi),\\
[\![\mathsf{con}(\varphi)]\!] &= \lambda xy.\, (x \sqsupseteq y \wedge \mathfrak{M}, y \not\models \varphi),\\
[\![\alpha \cap \beta]\!] &= [\![\alpha]\!] \cap [\![\beta]\!],\\
[\![\alpha ; \beta]\!] &= [\![\alpha]\!] ; [\![\beta]\!],\\
[\![-\alpha]\!] &= -[\![\alpha]\!],\\
[\![\alpha\ ]\!] &= \{\, (x, y) : (y, x) \in [\![\alpha]\!] \,\},
\end{aligned}
$$

$$[\![\varphi?]\!] \quad = \quad \{\,(x,x) : \mathfrak{M}, x \models \varphi\,\}.$$

Obviously, ra and fix are definable using the other operators, however, for conceptual and notational convenience they will be part of the official definition of the language. Further examples of operators definable in terms of the others will be given below.

I will refer to this language as the *DML*-language, and in more formal parts of this chapter also as $DML(\sqsubseteq, \Phi)$, where Φ is the set of proposition letters. A natural extension is obtained by considering multiple basic relations $\{\,\sqsubseteq_i\,\}_{i\in I}$ instead of the single relation $\sqsubseteq$; I will write $DML(\{\,\sqsubseteq_i\,\}_{i\in I}, \Phi)$ for the language thus extended. (In this extended language the expansion and contraction operators will be indexed with the relations they are based upon, viz. $\mathsf{exp}(\varphi)_i$ and $\mathsf{con}(\varphi)_i$.)

In its formulation in Van Benthem [4] the *DML*-language also contained *minimal* versions μ-$\mathsf{exp}(\cdot)$ and μ-$\mathsf{con}(\cdot)$ of the expansion and contraction operators, respectively, where $[\![\mu\text{-}\mathsf{exp}(\varphi)]\!] =$

$$\lambda xy.\,\Big((x,y) \in [\![\mathsf{exp}(\varphi)]\!] \wedge \neg\exists z\,(x \sqsubseteq z \sqsubset y \wedge (x,z) \in [\![\mathsf{exp}(\varphi)]\!])\Big),$$

and likewise for μ-$\mathsf{con}(\varphi)$. However, there is no need to add them explicitly to the language, as both are definable:

$$[\![\mu\text{-}\mathsf{exp}(\varphi)]\!] = [\![\mathsf{exp}(\varphi) \cap -(\mathsf{exp}(\varphi); (\mathsf{exp}(\top) \cap -(\top?)))]\!],$$

and similarly for μ-$\mathsf{con}(\varphi)$.

11.2.2 Some Results

Let me mention some of the work that has been done on *DML*. De Rijke [23] gives an explicit axiomatization of validity in *DML*, comprising of 36 axioms, and 4 derivation rules (including a so-called 'unorthodox' Gabbay-style irreflexivity rule). For future reference let me record this result:

THEOREM 11.1 There exists a complete, finitary axiomatization of validity in the language $DML(\{\,\sqsubseteq_i\,\}_{i\in I}, \Phi)$.

De Rijke [23] uses a difference operator $\mathbf{D}$ ('truth at a different state') to characterize some of the modes and projections, for example

$$p \wedge \neg\mathbf{D}p \rightarrow \Big(\mathsf{fix}(\alpha \cap \beta) \leftrightarrow \mathsf{do}(\alpha \cap \beta; p?)\Big)$$

is an axiom in his axiomatization governing the interaction of fix and $\cap$.

The same paper also establishes the undecidability of satisfiability in *DML*. In addition it gives a number of subsystems and extensions of *DML* whose satisfiability problems are decidable; in particular, deleting (; and) or just $-$ yields decidable fragments again, as does restricting the class of models to those based on trees. Furthermore, exact descriptions, both syntactic, and semantic by means of appropriate bisimulations, are given for the first-order counterpart of DML.

11.2.3 Some Connections

There are obvious connections between DML and *propositional dynamic logic* (PDL, cf. Harel [18]). The 'old diamonds' $\langle\alpha\rangle$ from PDL can be simulated in DML by putting $\langle\alpha\rangle\varphi := \mathsf{do}(\alpha;\varphi?)$. And likewise, the expansion and contraction operators are definable in a particular mutation of PDL where taking converses of program relations is allowed and a name for the informational ordering is available: $[\![\mathsf{exp}(\varphi)]\!] = [\![\sqsubseteq;\varphi?]\!]$ and $[\![\mathsf{con}(\varphi)]\!] = [\![\sqsubseteq\ ;\neg\varphi?]\!]$. The operator $\mathsf{do}(\alpha)$ can be simulated in standard PDL by $\langle\alpha\rangle\top$. An obvious difference between DML and PDL is that (at least in it's more traditional mutations) PDL only has the *regular* program operations $\cup$, ; and *, while DML has the full relational repertoire $\cup, -,$ and ;, but not the Kleene star. Another difference is not a technical one, but one in emphasis; whereas in PDL the Boolean part of the language clearly is the primary component of the language, in DML some effort is made to give the relational part the status of a first-class citizen as well by shifting the notation towards one that more clearly reflects the aspects of relations which we usually consider to be important.

A related formalism whose relational apparatus is more alike that of DML is the *Boolean modal logic* (BML) studied by Gargov and Passy [14]. This system has atomic relations $\rho_1, \rho_2, \ldots$, a constant for the Cartesian product $W \times W$ of the underlying domain W, and relation-forming operators $\cap$, $\cup$ and $-$. Relations are referred to within the BML-language by means of the PDL-like diamonds $\langle\alpha\rangle$. Since BML does *not* allow either ; or as operators on relations, it is a strict subsystem of $DML(\{\rho_1, \rho_2, \ldots\}, \Phi)$.

Further connections between DML and related work have been given in (Van Benthem [4]). These include links with Hoare Logic, and with various styles of non-standard inference.

11.2.4 Adding Structure

Usually no assumptions are made on the nature of the states of modal models. But for some applications of modal or temporal logics it may be necessary to be more specific about their nature. (Cf. (Gabbay, Hodkinson and Reynolds [11]) for a whole array of examples.) In such a structured setting models will have the form $\mathfrak{M} = (W_g, \ldots)$, where the *global* components of the model are given by the... , while the set W_g is a set of models $\{\,\mathfrak{m}\,\}_{i\in I}$ each of which may have further structure. For instance, they may themselves be of the form $\mathfrak{m} = (W_l, R, V_l)$. Clearly, two languages are involved here: a *global* language which talks about global aspects of the structure, but which does not deal with local aspects, and, secondly, there is *local* language used to reason only about the internal structure of the elements of the model $\mathfrak{M}$. Below, in 11.5 and 11.6, I will equip the states of DML-models with structure in different ways, each with an appropriate local language, but in every case DML will be the global language.[1]

11.3 On Postulates for Theory Change

In this section I will first discuss to which extent Gärdenfors' theory on the dynamics of belief and knowledge can be dealt with in the DML language. After that I will discuss two alternative proposals, and finally I will tie up some loose ends.

11.3.1 The Gärdenfors Postulates

Consider a set of beliefs or a knowledge set T.[2] As our perception of the world as described by T changes, the knowledge set may have to be modified. In the literature on theory change or belief revision a number of such modifications have been identified (cf. (Alchourrón, Gärdenfors and Makinson [1]), and (Katsuno and Mendelzon [20])); these include expansions, contractions and revisions. If we acquire information that does not contradict T, we can simply *expand* our knowledge set with this piece of information. When a sentence φ previously believed becomes questionable and has to be abandonded, we *contract* our knowledge with φ. Somewhat intermediate between expansion and contraction is the operation of *revision*, this is the operation of resolving the conflict that arises when the newly acquired information contradicts our old beliefs. The revision of T by a sentence φ, $T * \varphi$, is often thought of as consisting of first making changes to T, so as to then be able to expand with φ. According to general wisdom on theory change, *as little as possible* of the old theory should be given up in order to accommodate for newly acquired information.

Gärdenfors and others have proposed a set of rationality postulates that the revision operation must satisfy. To formulate these, let a *knowledge set* be a deductively closed set of formulas. Given a knowledge set T and a sentence φ, $T * \varphi$ is the revision of T by φ. $T + \varphi$ ("the expansion of T by φ") is the smallest deductively closed set extending $T \cup \{\varphi\}$.

Basic Gärdenfors postulates for revision

(∗1) $T * \varphi$ is a knowledge set.

[1] The essential syntactic restriction corresponding to the above global-local distinction is that operators from the global language are *not* allowed to occur inside the scope of operators from the local language. By results of Finger and Gabbay [9], if both the local and the global language have some nice property P, then so does their composition, provided that the above syntactic restriction is met; here P can be a property like enjoying a complete recursive axiomatization, decidability, or the finite model property.

[2] This subsection was inspired by a reading of (Fuhrmann [10]).

(*2) $\varphi \in T * \varphi$.
(*3) $T * \varphi \subseteq T + \varphi$.
(*4) If $\neg\varphi \notin T$ then $T + \varphi \subseteq T * \varphi$.
(*5) If $\bot \in T * \varphi$ then φ is unsatisfiable.
(*6) If $\varphi \leftrightarrow \psi$ then $T * \varphi = T * \psi$.

Additional Gärdenfors postulates for revision

(*7) $T * (\varphi \wedge \psi) \subseteq (T * \varphi) + \psi$.
(*8) If $\neg\psi \notin T * \varphi$ then $(T * \varphi) + \psi \subseteq T * (\varphi \wedge \psi)$.

For an intuitive explanation of this postulates I refer the reader to (Alchourrón et al [1], Gärdenfors [12]).

To represent the revision operator in DML some choices need to be made. First, we have to agree on some kind of structure in which our theories will be represented, and in which transitions between theories will take place. To keep things simple, and exclude what I consider to be aberrations in this context (like densely ordered sequences of theories), let us assume that our structures are well-founded ones (in addition to being pre-orders, of course).

Next, we have to decide how to represent theories or knowledge sets. The natural option suggested by standard practice in epistemic logic is to do this. Let $\sqsubseteq$ abbreviate $\mathsf{exp}(\top)$, and let $[\sqsubseteq]\varphi$ be short for $\neg\langle\sqsubseteq\rangle\neg\varphi$. Then, I represent theories as sets of the form $w_{\Box} = \{\varphi : \mathfrak{M}, w \models [\sqsubseteq]\varphi\}$,for some w in the model $\mathfrak{M}$. Then "$\varphi \in T$" may be represented as "$[\sqsubseteq]\varphi$," that is, as $\neg\mathsf{do}(\mathsf{exp}(\top); \neg\varphi?)$.

A third choice needs to be made to represent the *expansion* operator $[+\varphi]\psi$ ("ψ belongs to every theory resulting from expanding with φ"). Here I opt for:

$$[+\varphi]\psi := \neg\mathsf{do}(\mu\text{-}\mathsf{exp}([\sqsubseteq]\varphi); \neg[\sqsubseteq]\psi?).$$

So, a formula $[+\varphi]\psi$ is true at some point x if in every 'minimal' $\sqsubseteq$-successor y of x where $[\sqsubseteq]\varphi$ holds (i.e. where φ has been added to the theory), the formula $[\sqsubseteq]\psi$ is true (i.e. ψ is in the theory). Obviously, $[\sqsubseteq]\psi$ may be viewed as the special case of $[+\varphi]\psi$, where one expands with $\varphi = \top$.

Representing the revision operator $[*\varphi]\psi$ ("ψ belongs to every theory resulting from revision by φ") is a slightly more complex matter. Recall that revision of T by φ is explained as removing from T all (and only those) sentences that are inconsistent with φ, and subsequently expanding T by φ.[3] Mimicking the removal from T of the formula that causes the inconsistency with φ by $\mu\text{-}\mathsf{con}([\sqsubseteq]\neg\varphi)$, and the subsequent expansion with φ as before, I end up with the following definition:

[3]Isaac Levi has in fact suggested that revisions should be *defined* in terms of such contractions and revisions.

$$[*\varphi]\psi := \neg \mathsf{do}\Big([\mu\text{-}\mathsf{con}([\sqsubseteq]\neg\varphi); \mu\text{-}\mathsf{exp}([\sqsubseteq]\varphi)]; \neg[\sqsubseteq]\psi?\Big).$$ [4]

Before actually translating the revision postulates into DML, let me mention a possible point of discussion here. In my approach the expansion and revision operators lack the functional character they have in the Gärdenfors approach. This is due, of course, to the fact that the underlying $\sqsubseteq$-paths to points where "$\varphi \in T$" holds or fails for the first time, need not be uniquely determined. I don't see this as a shortcoming of the way I've set up things. On the contrary, one can view this as an attempt to take the non-deterministic character of everyday expansions and revisions seriously, instead of dismissing it as being "non-logical".

Another source of indeterminism is that, starting from a given node/theory and a formula φ that you want to expand with, you may have to pass several other nodes/theories before ending up at an outcome of the expansion, while a move to contract by φ at this outcome need not take you all the way back to your starting point.[5]

Finally, despite the fact that expansions and revisions may have multiple outcomes in my setup, they need not have a single one, i.e., expansions and revisions need not be defined in every situation.

Given the above points some of the postulates $(*1)$–$(*8)$ are bound to come out invalid when translated into DML. But on the other hand, they also allow for some choices when doing the translation. The statement $\psi \notin T * \varphi$ may be read as "ψ does not belong to *any* theory resulting from revision by φ," or as "for *some* outcome T' of revising T by φ, $\psi \notin T'$." The modal counterparts of these options are

$$\neg \mathsf{do}\Big([\mu\text{-}\mathsf{con}([\sqsubseteq]\neg\varphi); \mu\text{-}\mathsf{exp}([\sqsubseteq]\varphi)]; [\sqsubseteq]\psi?\Big),$$

or $[\dagger\varphi]\psi$ for short, and $\neg[*\varphi]\psi$, or

$$\mathsf{do}\Big([\mu\text{-}\mathsf{con}([\sqsubseteq]\neg\varphi); \mu\text{-}\mathsf{exp}([\sqsubseteq]\varphi]; \neg[\sqsubseteq]\psi?\Big),$$

respectively. These subtleties will make some difference for postulate $(*8)$.

On a similar note, as expansions and revisions need not be defined in every situation, one might consider adding a clause $\neg[+\varphi]\bot$ ($\neg[*\varphi]\bot$) saying "and if expansion (revision) with φ is at all possible" to some of the Gärdenfors postulates. However, for none of the postulates this has any visible effects.

---

[4] This definition is clearly in accordance with the earlier maxim "change as little as possible of the old theory."

[5] In other words: it may be that you need to expand with some formulas $\psi_1, \ldots, \psi_n$ before you can expand with φ. Admittedly, this kind of interference may be undesirable, especially when $\psi_1, \ldots, \psi_n$ and φ are logically independent; on the other hand, this interference might be useful to model various kinds of *non*-logical dependencies between formulas.

(G2)	$[*\varphi]\varphi$
(G3)	$[*\varphi]\psi \to [+\varphi]\psi,$
(G4)	$\neg[\sqsubseteq]\neg\varphi \wedge [+\varphi]\psi \to [*\varphi]\psi,$
(G5)	$[*\varphi]\bot \to [*\psi]\neg\varphi,$
(G6)	$\varphi \leftrightarrow \psi \,/\, [*\varphi]\chi \leftrightarrow [*\psi]\chi,$
(G7)	$[*(\varphi \wedge \psi)]\chi \to [*\varphi][+\psi]\chi,$
(G8a)	$\neg[*\varphi]\neg\psi \wedge [*\varphi][+\psi]\chi \to [*(\varphi \wedge \psi)]\chi,$
(G8b)	$\neg[\dagger\varphi]\psi \wedge [*\varphi][+\psi]\chi \to [*(\varphi \wedge \psi)]\chi.$

Table 11.1
Translating the Gärdenfors postulates.

Which translations does this give, then? Translating $\chi \in T + \varphi$ as $[+\varphi]\chi$, with $[\sqsubseteq]\chi$ as the limiting case where $T + \varphi$ is in fact T (or $T + \top$), and, likewise, translating $\chi \in T * \varphi$ as $[*\varphi]\psi$, I arrive at Table 11.1, where Gn is the translation of postulate (*n). Observe that there is no schema corresponding to postulate (∗1) in Table 11.1; this one seems to resist a direct translation, but its validity is guaranteed given the choices I have made.

Which of the schemata G2–G8b is valid on the well-founded DML-models we are considering here? First, the translation G2 of (∗2) comes out valid, as an easy calculation shows. To see that G3 is valid, assume that in some model we have $x \not\models [+\varphi]\psi$. So there is a minimal $\sqsubseteq$-successor y of x with $y \models [\sqsubseteq]\varphi$, $\neg[\sqsubseteq]\psi$. Let us verify that $x \not\models [*\varphi]\psi$. Clearly, $y \models [\sqsubseteq]\varphi$ implies $x \not\models [\sqsubseteq]\neg\varphi$, so $(x, x) \in [\![\mu\text{-con}([\sqsubseteq]\neg\varphi)]\!]$. In addition $(x, y) \in [\![\mu\text{-exp}([\sqsubseteq]\varphi)]\!]$. Hence, as $y \models \neg[\sqsubseteq]\psi$, we must have

$$x \models \mathsf{do}\Big([\mu\text{-con}([\sqsubseteq]\neg\varphi; \mu\text{-exp}([\sqsubseteq]\varphi)]; \neg[\sqsubseteq]\psi?\Big),$$

which is what we were after. Ergo, G3 is valid on all DML-models.

Next comes G4. Suppose that $x \models \neg[\sqsubseteq]\neg\varphi$, $[+\varphi]\psi$, but that $x \not\models [*\varphi]\psi$. We derive a contradiction. By $x \not\models [*\varphi]\psi$ there is a minimal $\sqsubseteq$-predecessor y of x with $y \models \neg[\sqsubseteq]\neg\varphi$. But as $x \models \neg[\sqsubseteq]\neg\varphi$, x itself must be this y. But then, by assumption, $x \models \mathsf{do}(\mu\text{-exp}([\sqsubseteq]\varphi; \neg[\sqsubseteq]\psi?)$, that is: for some minimal $\sqsubseteq$-successor z of x, $z \models [\sqsubseteq]\varphi$, $\neg[\sqsubseteq]\psi$. But by $x \models [+\varphi]\psi$, we must also have $z \models [\sqsubseteq]\psi$, yielding the required contradiction. Hence G4 is valid.

G5 is trivially valid, as its antecedent can never be satisfied. The validity of G6 is also obvious, so let us consider G7. Seeing that it is valid requires a small argument. Assume that in some model we have $x \not\models [*(p \wedge q)]r \to [*p][+q]r$. Then there are y, z, u such that

1. y is a minimal $\sqsubseteq$-predecessor of x with $y \not\models [\sqsubseteq]\neg p$,
2. z is a minimal $\sqsubseteq$-successor of y with $z \models [\sqsubseteq]p$,
3. u is a minimal $\sqsubseteq$-successor of z with $u \models [\sqsubseteq]q, \neg[\sqsubseteq]r$.

To arrive at a contradiction assume that

4. $x \models [*(p \wedge q)]r$.

Then, by (1) and an easy argument, y must be a minimal $\sqsubseteq$-predecessor of x with

5. $y \not\models [\sqsubseteq]\neg(p \wedge q)$.

To arrive at a contradiction, we will show that $u \models [\sqsubseteq]r$ — conflicting with (3). By (4) and (5), if u is a minimal $\sqsubseteq$-successor of y with $u \models [\sqsubseteq](p \wedge q)$, we must have $u \models [\sqsubseteq]r$. If, on the other hand, u is not such a successor, then, as $u \models [\sqsubseteq](p \wedge q)$ by (2) and (3), there must be a v such that

6. v is a minimal $\sqsubseteq$-successor of y with $v \models [\sqsubseteq](p \wedge q)$ and $v \sqsubseteq u$,

because we have assumed our structures to be well-founded. But then, by (4) and (5), $v \models [\sqsubseteq]r$, and by (6), $u \models [\sqsubseteq]r$, and we have reached our contradiction. This implies that G7 is valid.

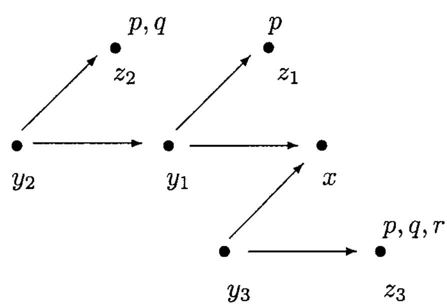

Figure 11.2
Refuting G8a

In G8a the antecedent $\neg\psi \notin T * \varphi$ of (*8) is translated as $\neg[*\varphi]\neg\psi$. The instance $\neg[*p]\neg q \wedge [*p][+q]r \to [*(p \wedge q)]r$ of G8a is refuted at x in the model depicted in Figure 11.2. To see this, notice first of all that $[*(p \wedge q)]r$ is refuted at x because

$(x, z_2) \in [\![\mu\text{-con}([\sqsubseteq]\neg(p \wedge q)); \mu\text{-exp}([\sqsubseteq](p \wedge q)); \neg[\sqsubseteq]r?]\!]$.

Second, $\neg[*p]\neg q$ holds at x as

$(x, z_3) \in [\![\mu\text{-con}([\sqsubseteq]\neg p); \mu\text{-exp}([\sqsubseteq]p); [\sqsubseteq]q?]\!]$.

Third, $[*p][+q]r$ holds at x because there's only one "revise by p, expand by q" path leading from x, notably (x, z_3), and at the end of that path $[\sqsubseteq]r$ holds. (In particular, (x, z_2) is not a "revise by p, expand by q" path since $(x, y_2) \notin [\![\mu\text{-con}([\sqsubseteq]\neg p)]\!]$.)

There are several aspects to the invalidity of G8a, and it's worth identifying them. For a start, we are able to perform a contraction with $\neg p$ (moving from x to y_1) *before* we can contract with $\neg(p \wedge q)$ (move from x to y_1 to y_2). As a consequence it is consistent to have $(x, y_2) \in [\![\mu\text{-con}([\sqsubseteq]\neg(p \wedge q))]\!] \setminus [\![\mu\text{-con}([\sqsubseteq]\neg p)]\!]$.[6] A related point is this. Since in our indeterministic set-up we have interpreted $\neg\psi \notin T * \varphi$ as "for *some* result of revising by φ, $\neg\psi$ is not in that result," we are able to have a revision with p that contains $\neg q$, notably z_1, while at the same time having one that does contain q. And as expansions need not always be defined in my set-up a revision with p (the move from x to y_1 to z_1) need not be a revision with $p \wedge q$.

Some of the causes underlying the invalidity of G8a can be eliminated. For example, reading $\neg\psi \notin T * \varphi$ as "for *all* results T' of revising T by φ" as in G8b, some of the indeterminism can be lifted. In particular, points like z_1 in Figure 11.2 will then be forbidden. Nevertheless, G8b is still not valid, as the reader may verify. Although one might go still furhter towards ensuring that expansions and revisions are defined when needed, I don't think that all aspects of indeterminism can be done away with. Specifically, I don't think that the kind of dependencies noted in footnotes 5 and 6 can be removed. In conclusion: there is no reasonable translation of (∗8) into DML that will make it come out valid.

I have so far tried to give a modal analysis of the Gärdenfors postulates inside DML, yielding a formal machinery for reasoning about Theory Change. The surplus value of having the full relation algebraic repertoire available in conjunction with Gärdenfors style expansion and revision operators will be discussed towards the end of this section. At this point I want to pursue the fact that one postulate, viz. (∗8), did not come out valid despite some alterations to its initial translation. This failure may prompt three reactions. One can leave things as they are, and not be bothered by the invalidity of (∗8); as (∗8) has been criticized extensively in the literature, this choice could be well argued for (cf. for example (Ryan [25])). Alternatively, one can change the rules of the game somewhat by changing the relevant postulate to one that no longer rests on the assumptions that expansions and revisions be functional and always defined. A third possibility would be to look for an alternative (modal) modelling of the postulates in DML or some other formalism. Two proposals pursuing the second option will be discussed in the following two subsections. Readers interested in alternative (modal) modellings of the Gärdenfors postulates and of postulates proposed by others are referred to (Fuhrmann [10]) and (Grahne [15]).

[6] As another consequence, the so-called *recovery* postulate for contraction ($T \subseteq T - \varphi + \varphi$, or in modal terms $[\sqsubseteq]\psi \to [-\varphi][+\varphi]\psi$, where $[-\varphi]$ has the obvious interpretation) is not valid in my set-up. This may not be such a bad thing as the recovery postulate is commonly considered to be the intuitively least compelling of the Gärdenfors postulates for contracting, cf. (Hansson [17]).

11.3.2 The Lindström-Rabinowitz postulates

While discussing the indeterminacy arising in the context of revision of probabilistic functions modelling belief states, one proposal Lindström and Rabinowitz [22] come up with, is letting belief revision be a relation rather than a function. They argue that this way of looking at belief revision is natural if one thinks that an agent's policies for belief change may not always yield a *unique* belief set as the result of a revision. Let a *belief revision relation* be a ternary relation $\mathcal{R}$ between knowledge sets, (consistent) formulas and knowledge sets. Lindström and Rabinowitz propose postulates ($\mathcal{R}0$)–($\mathcal{R}4$) below for all T, S, U and φ, ψ.

Lindström-Rabinowitz postulates for revision as a relation

($\mathcal{R}0$) There exists a T' such that $T\mathcal{R}_\varphi T'$.
($\mathcal{R}1$) If $T\mathcal{R}_\varphi S$ then $\varphi \in S$.
($\mathcal{R}2$) If $T \cup \{\varphi\}$ is consistent and $T\mathcal{R}_\varphi S$, then $S = T + \varphi$.
($\mathcal{R}3$) If $\varphi \leftrightarrow \psi$ and $T\mathcal{R}_\varphi S$, then $T\mathcal{R}_\psi S$.
($\mathcal{R}4$) If $T\mathcal{R}_\varphi S$, $S\mathcal{R}_\psi U$ and $S \cup \{\psi\}$ is consistent, then $T\mathcal{R}_{\varphi\wedge\psi} U$.

The intuitive reading of $T\mathcal{R}_\varphi S$ is: S is a (possible) outcome of revising T by φ. Postulate ($\mathcal{R}0$) corresponds to the requirement that revision should be defined for all T and (consistent) φ. Postulates ($\mathcal{R}1$)–($\mathcal{R}3$) are the relational counterparts to the Gärdenfors postulates (∗2), (∗3) and (∗4), (∗6), and (∗8), respectively. Lindström and Rabinowitz don't give relational counterparts to (∗5) and (∗7). ($\mathcal{R}4$) is new.

How can the Lindström-Rabinowitz postulates be accounted for in *DML*? As before we let knowledge sets be represented as sets of the form $w_\Box = \{\varphi : \mathfrak{M}, w \models [\sqsubseteq]\varphi\}$. And following the definition of $[*\varphi]\psi$, the obvious choice for the relation $\mathcal{R}_\varphi$ seems to be

$$\mathcal{R}_\varphi = [\![\mu\text{-con}([\sqsubseteq]\neg\varphi); \mu\text{-exp}([\sqsubseteq]\varphi)]\!].$$

So $T\mathcal{R}_\varphi S$ iff $\exists t, s\, (T = t_\Box \wedge S = s_\Box \wedge (t,s) \in \mathcal{R}_\varphi)$.

Given this representation, one can reason about the revision relation $\mathcal{R}$ and its properties using the *DML* apparatus. For instance, idempotency properties like

$$\mathsf{fix}(\mathcal{R}_\varphi; \mathcal{R}_\varphi)$$

can now be tested for. I leave it to the reader to check that ($\mathcal{R}0$) fails under this representation, and that ($\mathcal{R}1$)–($\mathcal{R}3$) are all valid. As to ($\mathcal{R}4$), in order to make sense of it in *DML* we have to decide how to represent "$S \cup \{\psi\}$ is consistent" in *DML*. One natural candidate is "$[\sqsubseteq]\neg\psi \notin s_\Box$," where $s_\Box$ represents S. But this reading does not make ($\mathcal{R}4$) come out valid in *DML*. An easy counter model is given in Figure 11.3, with $T = t_\Box, S = s_\Box, U = u_\Box, \varphi = p$ and $\psi = q$. In Figure 11.3 $(s,t) \in \mathcal{R}_p$, $(t,u) \in \mathcal{R}_q$, $s \not\models [\sqsubseteq]\neg q$, but $(s,u) \notin \mathcal{R}_{p\wedge q}$. Hence, in *DML* an agent has the possibility to distinguish between

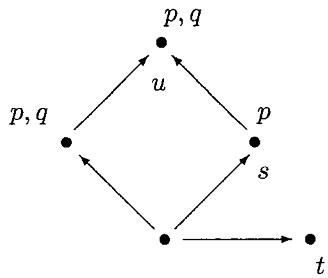

Figure 11.3
A counter model for ($\mathcal{R}4$)

revising his knowledge by φ (without excluding ψ as an unacceptable proposition) and subsequently revising by ψ on the one hand, and revising by the conjunction $\varphi \wedge \psi$ on the other hand. Thus, in DML there's still more (room for) indeterminism than is allowed for by the Lindström-Rabinowitz postulates.

11.3.3 The Katsuno-Mendelzon Postulates for Indeterministic Revision

Katsuno and Mendelzon [21] give a model-theoretic characterization of all revision operators that satisfy the Gärdenfors postulates (*1)–(*8). They show that these operators are precisely the ones that accomplish a revision with minimal change to the class of models of the knowledge set. This minimality is measured in terms of total pre-orders among models of the "initial" knowledge set. Katsuno and Mendelzon also study variations on the ordering notions and the corresponding postulates; in one of their variations they change the above *total* pre-orders to *partial* ones, and formulate postulates characterizing the corresponding *indeterministic* revision operators. Below I will translate these postulates into DML.

The Katsuno-Mendelzon postulates are formulated for knowledge sets T that are assumed to be represented by a propositional formula ψ_T such that $T = \{\varphi : \psi_T \vdash \varphi\}$. The notation $\psi \circ \mu$ is used to denote the revision of (the knowledge set represented by) ψ with (the formula) μ. Katsuno and Mendelzon propose seven postulates for indeterministic revision, the first five of which are in fact equivalent to the Gärdenfors postulates (*1)–(*7), and thus valid (when translated) in DML. Here are the remaining two.

Katsuno-Mendelzon postulates for indeterministic revision

(R7) If $\psi \circ \mu_1$ implies μ_2 and $\psi \circ \mu_2$ implies μ_1, then $\psi \circ \mu_1$ is equivalent to $\psi \circ \mu_2$.
(R8) $(\psi \circ \mu_1) \wedge (\psi \circ \mu_2)$ implies $\psi \circ (\mu_1 \vee \mu_2)$.

Intuitively, (R7) says that if μ_2 holds in every result of revising with μ_1, and μ_1 holds in every result of revising with μ_2, then the revision with μ_1 and the revision with μ_2 have the same effect. Postulate (R8) says that every knowledge set that may be arrived at after revising with μ_1, and also after revising with μ_2, must be among the knowledge sets obtained after revising with $\mu_1 \vee \mu_2$.

Given these intuitive readings of (R7) and (R8) the following seem to be the natural translations of these postulates into *DML*. ($\mathcal{R}$ is the revision relation defined in the previous subsection.)

(KM7) $[*\varphi]\psi \wedge [*\psi]\varphi \rightarrow ([*\varphi]\chi \leftrightarrow [*\psi]\chi)$.
(KM8) $[*(\varphi \vee \psi)]\chi \rightarrow \neg \mathsf{do}((\mathcal{R}_\varphi \cap \mathcal{R}_\psi); \neg[\sqsubseteq]\chi?)$.

Although (∗8) or G8a has now been weakened to (R7) ∧ (R8) or (KM7) ∧ (KM8), this weaker version is still not valid in *DML*. In Figure 11.4 the instance $[*p]q \wedge [*q]p \rightarrow$

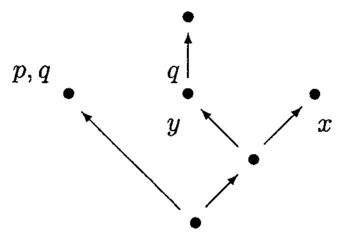

Figure 11.4
Refuting (KM7)

$([*p]r \leftrightarrow [*q]r)$ of (KM7) fails at x. As before, one thing that makes the model depicted there a counter model for (KM7) is the fact that expansions and revisions need not always be defined in my set-up. In particular, (KM7) would not fail at x in Figure 11.4 if it were possible to expand with q at y. Furthermore, in Figure 11.5 the instance $[*(p \vee q)]r \rightarrow \neg \mathsf{do}((\mathcal{R}_p \cap \mathcal{R}_q); \neg[\sqsubseteq]r?)$ of (KM8) fails at x. What this seems to amount to is that in *DML* an agent can get to know a (non-trivial) disjunction without having to know either disjunct. Apparently this possibility is excluded by the Katsuno-Mendelzon postulates.

A Look Back

Let's step back and review some points made in this section. One of the main features of the revision and expansion operators defined in this section as opposed to other formalisms for theory change, is that in my set-up revisions and expansions need not always be defined. Just as one can argue for giving up the functionality or determinism implied by the Gärdenfors postulates by saying that an agent's strategies for belief revision may

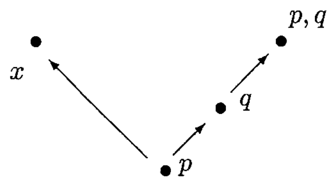

Figure 11.5
Refuting (KM8)

not always tell him how to choose between possible outcomes,—one can also argue for the possibility of revisions and expansions not being defined at all by pointing out that an agent's strategy for belief revision may not always tell him how to revise or expand. Everyday life examples to this effect are easily found.

As was pointed out before, revisions and expansions as defined in this section lack the *total* independence of sentences implicitly assumed by, for instance, the Gärdenfors postulates for belief contraction (cf. footnotes 5, 6). This lack of independence might be useful for modelling non-logical relations between beliefs.

Apart from the above two deviations this section shows that it is possible to define revision and expansion operators in a fairly standard dynamic modal formalism like DML that satisfy most of the postulates given by Gärdenfors, Lindström-Rabinowitz, or Katsuno-Mendelzon.

There are several advantages to having revision and expansion operators satisfying those postulates defined using the well-known Boolean and relation algebraic repertoire. To a large extent this embedding de-mystifies the enterprise of theory change. Next, in this larger repertoire one is no longer restricted to classical combinations of expansions and revisions, but further operations become visible as well. One can think of sequential composition of revisions, of reversals or 'un-doings' of revisions, and given that revisions and expansions need not always be defined in my set-up, one might introduce *conditional* revisions or expansions, where the conditions could read something like "after having contracted with $\neg\varphi$ you should always be able to expand with φ."

Having the revision and expansion operations embedded in a Boolean and relation algebraic setting also reveals possible generalizations. One might consider weaker forms of revision in which some of the minimality requirements are weakened. Second, this section discussed revision, that is, changing beliefs as a result of newly obtained information about a static world; one could also try and define so-called *updates* in DML; an update is a theory change reflecting a change in the world. As shown in (Katsuno and Mendelzon [20]) updates can be characterized by a set of postulates similar to the Gärdenfors postulates. Another obvious generalization is to allow for several copies of these operators, possibly interacting in certain prescribed ways, to model not only the belief change of several agents simultaneously but also the belief changes resulting from interaction between the agents.

Below the states of DML-models will be equipped with structure, a move that could be made here as well, allowing the theories that are being revised to explicitly have structure. One can think here of a hard core of sentences not admitting revisions, surrounded by sentences which do admit revisions but which need not all have the same epistemic status; the latter kind of sentences would then be ranked according to their "epistemic entrenchment", and the revision process would need to take this into account (compare (Gärdenfors and Makinson [13])).

11.4 Terminological Languages

As Blackburn and Spaan [5] put it, in recent years modal logicians have considered a number of enriched modal systems that bear on issues of knowledge representation. One example is Schild's [26] in which the correspondence between terminological languages and modal logic is used to obtain complexity results for terminological reasoning. In this section the correspondence between DML and one particular terminological proposal will be described.

Recall that terminological languages provide a means for expressing knowledge about hierarchies of concepts. They allow the definition of concepts and roles built out of primitive concepts and roles. *Concepts* are interpreted as sets (of individuals) and *roles* are interpreted as binary relations between individuals. For instance, `traveler` and `Amsterdam` may be concepts, and `has-flown-to` may be a role. Compound expressions are built up using various language constructs. Quite a number of proposals for such constructs have been and still are being put forward (cf. (Schmidt [27]) for a comprehensive survey). Here I will link DML to a KL-ONE dialect discussed by Brink and Schmidt [7]; I will refer to this language as the Brink and Schmidt language.

The operations considered by Brink and Schmidt are the usual Boolean ones for the concepts plus the usual RA-operations for the roles. In addition they consider a binary operator $\Diamond$ taking a role and a concept, and returning a concept: $\Diamond(R, C) = \{x : \exists y\,((x, y) \in R \wedge y \in C)\}$, and a mapping $(\cdot)^c$ called *(left) cylindrification* taking concepts to roles: $C^c = \{(x, y) : x \in C\}$. Other operations usually considered in terminological languages are *role quantifications* of the form (`SOME has-flown-to amsterdam`) and (`ALL has-flown-to amsterdam`). These expressions can be read as "objects having flown (at least once) to Amsterdam" and "objects all of whose flying trips went to Amsterdam". The quantifications (`SOME R C`) and (`ALL R C`) can be defined in Brink and Schmidt's language as $\Diamond(R, C)$ and $-\Diamond(R, -C)$, respectively.

Here's an example; while the present author is abroad one thing he may try to achieve is "writing a paper and not phoning to a Dutch person", or:

$\Diamond(\texttt{writing}, \texttt{paper}) \wedge \neg\Diamond(\texttt{phone} \cap (\texttt{dutch} \wedge \texttt{human})^c, \top)$,

where $\top$ is the Boolean 1.

The main questions in terminological reasoning are *satisfiability problems* (does a concept (role) have a non-empty denotation in some interpretation), and the *subsumption problem* (a concept (role) C subsumes a concept (role) D iff in every interpretation the denotation of C is a superset of the denotation of D). For example, on the understanding that `amsterdam` is in `europe` the concept

```
(ALL has-traveled-to amsterdam) ∩
   (SOME has-flown-to north-of-paris)
```

is subsumed by

```
(ALL has-flown-to europe).
```

As conjunction and negation are available in this languae, the subsumption problem can be reduced to the satisfiability problem.

What's the connection between Brink and Schmidt's terminological language and DML? Clearly, the terminological concepts are simply the propositions of DML, and the roles have their counterparts in the extension $DML(\{\sqsubseteq_i\}_{i\in I}, \Phi)$ of DML in which multiple 'primitive' relations $\sqsubseteq_i$ are available. So the two systems have the same basic ingredients. But what about their operators? Are they interdefinable, for example? Tables 11.2 and 11.3 show that in fact they are.

$$\begin{array}{rclcl} \Diamond(\alpha,\varphi) & = & \{\, x : \exists y\, ((x,y) \in [\![\alpha]\!] \wedge \mathfrak{M}, y \models \varphi) \,\} & = & \mathsf{do}(\alpha; \varphi?), \\ \varphi^c & = & \{\, (x,y) : \mathfrak{M}, x \models \varphi \,\} & = & \varphi?; (\delta \cup -\delta), \end{array}$$

Table 11.2
From terminological logic to DML ....

$$\begin{array}{rclcrcl} \mathsf{do}(\alpha) & = & \Diamond(\alpha, 1), & \quad & \mathsf{exp}(\varphi) & = & \sqsubseteq \cap \varphi^c\ , \\ \mathsf{ra}(\alpha) & = & \Diamond(\alpha\ , 1), & & \mathsf{con}(\varphi) & = & \sqsubseteq\ \ \cap (\neg\varphi)^c\ , \\ \mathsf{fix}(\alpha) & = & \Diamond((\alpha \cap \delta), 1), & & \varphi? & = & \delta \cap \varphi^c. \end{array}$$

Table 11.3
... and conversely.

To illustrate this connection, here's an example expressing the concept "people having flown only to cities called Amsterdam" in DML:

$\texttt{human} \wedge \mathsf{do}(\texttt{has-flown-to}; (\texttt{city} \wedge \texttt{amsterdam})?) \wedge$
$\quad \neg\mathsf{do}(\texttt{has-flown-to}; -(\texttt{city} \wedge \texttt{amsterdam}?))$.

The above connection may be formulated 'officially' by means of two mappings between the two languages, thus establishing the following result.

PROPOSITION 11.1 Brink and Schmidt's language for terminological reasoning with primitive concepts Φ and primitive roles $\{\sqsubseteq_i\}_{i\in I}$ is a notational variant of the modal language $DML(\{\sqsubseteq_i\}_{i\in I},\Phi)$.

Thus, the main issues in terminological reasoning, viz. satisfiability and subsumption, may be re-formulated as satisfiability problems in (an extension of) DML, and results and topics from the modal domain can be transferred to the terminological domain, and vice versa. To substantiate this claim, let me give some examples.

COROLLARY 11.1 Modulo the translation induced by Table 11.3, the axioms and inference rules of $DML(\{\sqsubseteq_i\}_{i\in I},\Phi)$ are a sound and complete axiomatization of subsumption of concepts in the Brink and Schmidt language.

We can be very brief about the proof of Corollary 11.1: apply 11.1 and 11.1. Indirectly, the axioms and rules also of $DML(\{\sqsubseteq_i\}_{i\in I},\Phi)$ also axiomatize subsumption of *roles* in the Brink and Schmidt language; this is because any equation $\alpha=\beta$ between roles (relations) can be mimicked at the level of concepts (propositions) by

$$EQUAL(\alpha,\beta) := \mathbf{A}(\neg\mathsf{do}(\alpha\cap-\beta)\wedge\neg\mathsf{do}(-\alpha\cap\beta)).$$

Although the following result is not new (cf. Schmidt-SchauS [28]), its proof too comes very easy given Proposition 11.1, and the fact that satisfiability in DML is undecidable (by De Rijke [23, Theorem 5.1]).

COROLLARY 11.2 Satisfiability and subsumption in the Brink and Schmidt language are undecidable.

As is well known, part of the Knowledge Representation community is concerned with finding *tractable* terminological systems, either by limiting the expressive power of the representation language, or by limiting the inference capabilities of the formalisms. This has resulted in the description of quite a number of decidable or even tractable systems, many of which can be seen as fragments of the Brink and Schmidt system. By 11.1, this work is relevant to the search for decidable or tractable fragments of DML.

Here's a final possibility for exchange between the modal and terminological domain. Terminological reasoning often deals with *number restrictions* like ($\geq$ 2 `has-flown-to` `amsterdam`) (which can be read as "objects having flown to Amsterdam at least twice") to perform numerical comparisons. The modal logic of these counting expressions (by themselves) has been analyzed by Van der Hoek and De Rijke [19]. The link between terminological languages and DML established in 11.1 suggests that it may be worth the effort to add the counting quantifiers to DML, and examine the resulting language.

To finish this section let me cast the connection between the Brink and Schmidt language and DML in algebraic terms. Schmidt [27] equips the Brink and Schmidt language with an algebraic semantics called Peirce algebras. To understand what these are we have to go through one or two definitions. First of all, a *Boolean module* is a structure $\mathfrak{M} = (\mathfrak{B}, \mathfrak{R}, \Diamond)$, where $\mathfrak{B}$ is a Boolean algebra, $\mathfrak{R}$ is a relation algebra and $\Diamond$ is a mapping $\mathfrak{R} \times \mathfrak{B} \to \mathfrak{B}$ such that

M1	$\Diamond(r, a+b) = \Diamond(r,a) + \Diamond(r,b)$	M4	$\Diamond(\delta, a) = a,$
M2	$\Diamond(r+s, a) = \Diamond(r,a) + \Diamond(s,a)$	M5	$\Diamond(0, a) = 0,$
M3	$\Diamond(r, \Diamond(s,a)) = \Diamond((r;s), a)$	M6	$\Diamond(r\ , \Diamond(r,a)') \leq a'.$

Just as Boolean algebras formalize reasoning about sets, and relation algebras formalize reasoning about relations, Boolean modules formalize reasoning about sets interacting with relations through $\Diamond$. In the full Boolean module $\mathfrak{M}(U) = (\mathfrak{B}(U), \mathfrak{R}(U), \Diamond)$ over a set $U \neq \emptyset$ the operation $\Diamond$ is defined as described earlier, by

$$\Diamond(R, C) = \{\, x : \exists y\, ((x,y) \in R \wedge y \in C)\,\}.$$

(See (Brink [6]) for a formal definition of Boolean modules and some examples.) Boolean modules are almost, but not quite, the right algebraic semantics for Brink and Schmidt's terminological language. To obtain a perfect match, what we need in addition to the *set* forming operation $\Diamond$, is an operation that forms new *relations* out of sets. This yields the notion of a *Peirce algebra*, which is a two-sorted algebra $\mathfrak{P} = (\mathfrak{B}, \mathfrak{R}, \Diamond, (\cdot)^c)$ with $(\mathfrak{B}, \mathfrak{R}, \Diamond)$ a Boolean module, and $(\cdot)^c : \mathfrak{B} \to \mathfrak{R}$ a mapping such that for every $a \in \mathfrak{B}$, $r \in \mathfrak{R}$ we have

P1 $\Diamond(a^c, 1) = a,$

P2 $\Diamond(r, 1)^c = r; 1.$

In the full Peirce algebra $\mathfrak{P}(U)$ over a set $U \neq \emptyset$, $(\cdot)^c$ is defined, as before, as $C^c = \{\,(x,y) : x \in C\,\}$. The algebraic apparatus of Peirce algebras has been used by Brink and Schmidt [7] as an inference mechanism in terminological representation

Where does DML come in here? Because of Proposition 11.1, the modal algebras for the DML-language $L(\{\sqsubseteq_i\}_{i\in I}, \Phi)$ are the Peirce algebras generated by the relations $\{\sqsubseteq_i\}_{i\in I}$ and the propositions Φ. Let a *set identity* be one of the form $a = b$ where a, b are terms living in the Boolean reduct of a Peirce algebra (observe that a, b may contain the relational operations as well as $\Diamond$ and $(\cdot)^c$). Then, the completeness result 11.1 may be interpreted as follows.

PROPOSITION 11.2 The set identities valid in all representable Peirce algebras are completely axiomatized by the algebraic counterpart of the modal axiom system for the language $DML(\{\sqsubseteq_i\}_{i\in I}, \Phi)$.[7]

[7]Although, strictly speaking, the completeness result 11.1 only axiomatizes validity on pre-ordered DML-structures, the construction does not depend in an essential way on these structural assumptions.

Axiomatic aspects of the full two-sorted language over Peirce algebras, invlobing both set identities and identities between terms denoting relations, are studied by De Rijke [24].

11.5 Structured States: Update Semantics

In both this and the next section I will equip the states of DML-models with additional structure to be able to link DML with other dynamic proposals. The formalism I will consider in the present section is Veltman's *update semantics* [30]. In this system the standard explanation of the meaning of a sentence being its truth-conditions, is replaced by: "you know the meaning of a sentence if you know the changes it brings about in the information state of anyone who accepts the information conveyed by the sentence." According to this point of view the meaning of a sentence becomes a dynamic notion, an operation on information states. In this dynamic approach phenomena surrounding the instability and changing of information caused by modal qualifications like 'might,' 'presumably' and 'normally' can be adequately accounted for, as is shown by Veltman using a number of systems. The simplest one, called US_1 here, has in its vocabulary a unary operator `might` and a connective '$\bullet$', in addition to the usual Boolean connectives; in US_1 one can reason about an agent acquiring new information about the actual facts.

DEFINITION 11.1 The language of $US_1(\Phi)$ is given by the following definition.

Atomic formulas:	$p \in \Phi$,
Simple formulas:	$\varphi \in Form_1(\Phi)$,
Formulas:	$\psi \in US_1(\Phi)$.

$$\varphi ::= p \,|\, \bot \,|\, \neg\varphi \,|\, \varphi_1 \vee \varphi_2 \,|\, \mathtt{might}\,\varphi,$$
$$\psi ::= \varphi \,|\, \psi_1 \bullet \psi_2.$$

The important restriction is that no $\bullet$ is allowed to occur in the scope of a `might`.

The intuitive reading of $\mathtt{might}\,\varphi$ is that one has to agree to $\mathtt{might}\,\varphi$ if φ is consistent with one's knowledge; otherwise $\mathtt{might}\,\varphi$ is to be rejected. The operator $\bullet$ is simply the composition of (the functions expressed by) formulas.

DEFINITION 11.2 The semantics of the update system US_1 is as follows. Let $W \subseteq 2^{\Phi}$; a subset of W is an *information state*. Formulas are interpreted as functions in $2^W \longrightarrow 2^W$, that is, as functions from information states to information states. Let $\sigma \subseteq W$. I write $[\varphi]\sigma$ for the result of updating σ with φ.

$$\begin{array}{rcl} [p]\sigma & = & \sigma \cap \{ w : p \in w \}, \\ [\neg\varphi]\sigma & = & \sigma \setminus [\varphi]\sigma, \\ [\varphi \vee \psi]\sigma & = & [\varphi]\sigma \cup [\psi]\sigma, \\ [\mathtt{might}\,\varphi]\sigma & = & \begin{cases} \sigma, & \text{if } [\varphi]\sigma \neq \emptyset \\ \emptyset, & \text{otherwise,} \end{cases} \\ [\varphi \bullet \psi]\sigma & = & [\psi]([\varphi]\sigma). \end{array}$$

Veltman discusses several notions of valid consequence. Since these are not my prime concern here (but see below), I will confine myself to explaining the notion "$US_1 \models \varphi$" as "for all information states σ, $[\neg\varphi]\sigma = \emptyset$."

Van Eijck and De Vries [8] have established a connection between US_1 and the modal system $S5$ (see also [16]). This construction underlies the embedding of US_1 into DML presented below; what it amounts to is that US_1 is a formalism for reasoning about $S5$-models and certain transitions between them. This inspires the following definition.

DEFINITION 11.3 A *structured DML-model* is a tuple $\mathfrak{M} = (W_g, \sqsubseteq, [\![\cdot]\!])$ where $\sqsubseteq$ is a global relation on W_g, and W_g is a set of (finite) *pointed* $S5$-models of the form $\mathfrak{m} = (W, R, w, V)$ such that $w \in W$, $R = W \times W$, and V is a valuation. Moreover, the following conditions should be satisfied:

- $\mathfrak{m} \sqsubseteq \mathfrak{n}$ iff $\mathfrak{m}$ is a submodel of $\mathfrak{n}$,
- if $\mathfrak{m} = (W, R, w, V) \in \mathfrak{M}$ then $(W, R, v, V) \in \mathfrak{M}$ for all $v \in W$, and $\mathfrak{n} \in \mathfrak{M}$ for all $\mathfrak{n} \sqsubseteq \mathfrak{m}$.

The formal language appropriate for reasoning about such 2-level structures, $DML(S5)$, is defined as follows. Starting from a set of proposition letters Φ, $S5$-formulas are built using the operators L, M in the usual way. Let Φ' be the resulting set; this is the set of *local* formulas. They serve as 'proposition letters' for the global language DML; that is: $DML(S5)$ formulas are obtained from Φ' by applying the usual DML-connectives to its elements.

The important semantic clauses then read as follows, for $\mathfrak{m} = (W, R, w, V)$:

$\mathfrak{M}, \mathfrak{m} \models p$ iff $w \in V(p)$

$\mathfrak{M}, \mathfrak{m} \models M\varphi$ iff for some $v \in W$ with wRv, $(W, R, v, V) \models \varphi$,

that is, the value of such formulas is computed locally. For 'purely global' formulas, on the other hand, the value is computed globally, as in the following example:

$\mathfrak{M}, \mathfrak{m} \models \mathsf{do}(\sqsubseteq)$ iff for some $\mathfrak{n} \in \mathfrak{M}$, $\mathfrak{m} \sqsubseteq \mathfrak{n}$.

DEFINITION 11.4 Define a translation $(\cdot)^\dagger$ of the US_1-language into $DML(S5)$ as follows:

$$\begin{array}{rcl}
(p)^\dagger & = & p \\
(\neg\varphi)^\dagger & = & \neg\varphi^\dagger \\
(\varphi \vee \psi)^\dagger & = & \varphi^\dagger \vee \psi^\dagger \\
(\mathtt{might}\,\varphi)^\dagger & = & M\varphi^\dagger \\
(\varphi \bullet \psi)^\dagger & = & \varphi^\dagger \wedge \mathsf{do}\Big([\sqsupseteq; (L\varphi^\dagger \wedge \psi^\dagger)?] \\
& & \cap - [(\sqsupseteq; (L\varphi^\dagger \wedge \psi^\dagger)?); ((\sqsupseteq \cap - \delta); (L\varphi^\dagger \wedge \psi^\dagger)?)]\Big).
\end{array}$$

The intuitive reading of $(\mathtt{might}\,\varphi)^\dagger$ is that we *locally* check whether there is a point verifying $\varphi^\dagger$. The intuitive interpretation of $(\varphi \bullet \psi)^\dagger$ is that $(\varphi \bullet \psi)^\dagger$ holds at $\mathfrak{m} = (W, R, w, V)$ iff $\mathfrak{m} \models \varphi^\dagger$ and for $S = \{x \in \mathfrak{m} : (W, R, x, V) \models \varphi^\dagger\}$ we have that $(S, S^2, w, V \upharpoonright S) \models \psi^\dagger$. Notice that $(\cdot)^\dagger$ takes US_1-formulas into a *decidable* fragment of $DML(S5)$.

PROPOSITION 11.3 Let φ be a formula in $US_1(\Phi)$. Then we have $US_1 \models \varphi$ iff $DML(S5) \models \varphi^\dagger$.

Proof Suppose $US_1 \not\models \varphi$. Then for some $W \subseteq 2^\Phi$, and $\sigma \subseteq W$, $[\neg\varphi]\sigma \neq \emptyset$. Define $\mathfrak{M} = (W_g, \sqsubseteq, [\![\cdot]\!])$, where $W_g =$

$$\{(\sigma, \sigma^2, x, V_\sigma) : \sigma \subseteq W, x \in \sigma, (p \in V_\sigma(y) \text{ iff } p \in y, \text{ for } y \in \sigma)\},$$

and $\sqsubseteq$ and $[\![\cdot]\!]$ have their standard interpretation. Then, by a simple formula induction, we have that $\forall\psi\forall\sigma \subseteq W \forall j \in \sigma\, (j \in [\psi]\sigma \text{ iff } (\sigma, \sigma^2, j, V_\sigma) \models \psi)$. It follows that $DML(S5) \not\models \varphi^\dagger$.

To prove the opposite direction we proceed as follows. Assume that for some $\mathfrak{M}$ and $\mathfrak{m} \in \mathfrak{M}$ we have $\mathfrak{M}, \mathfrak{m} \not\models \varphi^\dagger$. Let $\mathfrak{m} = (W, W^2, w, V_\mathfrak{m})$. By standard modal logic we may assume that in $\mathfrak{m}$ every 'relevant' state description of the form $(\neg)p_0 \wedge (\neg)p_1 \wedge \ldots \wedge (\neg)p_n$ (where $p_0, \ldots, p_n$ are all the proposition letters occurring in φ), occurs only once in $\mathfrak{m}$. We may also assume that for every $\mathfrak{n} = (W_1, R_1, w_1, V_1) \in \mathfrak{M}$, (W_1, R_1, V_1) is a substructure of (W, W^2, V). Now, let $W' = 2^{\{p_0, \ldots, p_n\}}$, and for $\mathfrak{n} \in \mathfrak{M}$ let $\sigma_\mathfrak{n} \subseteq W$ be the set of state descriptions realized in $\mathfrak{n}$. Then, by a simple inductive proof, we have for all formulas ψ containing at most the proposition letters $p_0, \ldots, p_n$, and all $\mathfrak{n} = (W_1, R_1, w_1, V_1) \in \mathfrak{M}$, $\mathfrak{n} \models \psi^\dagger$ iff $w_1 \in [\psi]\sigma_\mathfrak{n}$, which completes the proof.

Proposition 11.3 may be interpreted as saying that the 'internal' notions of US_1 can be turned into internal notions of DML. But some of the 'external' or meta-notions of

US_1 can also be turned into internal notions of DML. Veltman [30] discusses various notions of valid consequence for his update systems, including

$$\varphi_1, \ldots, \varphi_n \models_1 \psi \quad \text{iff} \quad \text{for all } \sigma \text{ such that } [\varphi_i]\sigma = \sigma, \text{ we have } [\psi]\sigma = \sigma \ \ (1 \leq i \leq n), \text{ and}$$

$$\varphi_1, \ldots, \varphi_n \models_2 \psi \quad \text{iff} \quad \text{iff for all } \sigma, \text{ we have that } [\psi]([\varphi_n](\ldots([\varphi_1]\sigma)\ldots)) = [\varphi_n](\ldots([\varphi_1]\sigma)\ldots).$$

In $DML(S5)$ these notions become

$$\varphi_1, \ldots, \varphi_n \models_1 \psi \quad \text{iff} \quad \mathbf{A}(L\varphi^\dagger \wedge \ldots \wedge L\varphi_n^\dagger \to L\psi^\dagger), \quad \text{and}$$
$$\varphi_1, \ldots, \varphi_n \models_2 \psi \quad \text{iff} \quad \mathbf{A}L\Big((\varphi_1 \bullet \cdots \bullet \varphi_n)^\dagger \to \psi^\dagger\Big),$$

respectively, where $\mathfrak{M}, \mathfrak{m} \models \mathbf{A}\chi$ iff for all $\mathfrak{n} \in \mathfrak{M}$ we have $\mathfrak{M}, \mathfrak{n} \models \chi$.

Given the embedding of US_1 into DML, some natural extensions and generalizations become visible. Besides `might` one can consider other tests, the most obvious of which are definable in DML, like an operator testing whether updating with a formula φ will change the current state, or one testing whether the current state is at all reachable via an update with φ, or whether a pre-given goal state may be reached by performing certain updates.

11.6 Structured States: Preferences

Among the structures of logic L there may be some models that are *preferred* for one reason or another. Preferences may differ between applications, thus giving rise to different notions of preferential inference. Shoham [29] offers a general to preferential reasoning in which there is a (strict partial) order of preference $<$ on L-models on top which minimal consequence is defined as "truth of the conclusion in all $<$-minimal or most-preferred L-models of the premisses." By specifying the relation $<$ in alternative ways many formalisms with non-monotonic aspects can be shown to fit this general preferential scheme.

Given the embedding of US_1 into DML of 11.5 as an example, it should be obvious how preferential reasoning can be mimicked in DML: let $\sqsubseteq$ be the preferential ordering, and let the states of our DML-models simply *be* L-models. Then 'φ preferentially entails ψ' is true in the global structure $\mathfrak{M}$ iff $\mathfrak{M} \models \mathbf{A}(\varphi \wedge \neg \mathsf{do}(\sqsubseteq \ ; \varphi?) \to \psi)$. Via this equivalence all preferential reasoning can be performed inside DML.

Just as the preference relation embodies certain dynamic aspects of the underlying L, it itself could also be subjected to change. This point may be illustrated with a system US_2 which is slightly more complex than US_1, and which has also been introduced by

Veltman [30]. In US_2 one is not only able to reason about changing knowledge as new information comes in, but also about changing expectations; the latter are modelled using a notion of *optimality* with respect to a pre-order. Modelling this system in DML requires adding a separate $S4$-like component for expectations to the structured states of 11.5, in addition to the $S5$-like component for knowledge. An agents refinement or revision of his expectations can then be modelled inside such DML-structures by making moves to points with a suitably altered 'expectations' component.

11.7 Final Remarks

Let me point out what I consider to be the main points of this chapter. It has brought out connections and analogies between dynamic formalisms from cognitive science, philosophy and computer science by using a fairly traditional dynamic modal system (DML) in a flexible way, far beyond its traditional boundaries.

Putting DML to work in this manner had the surplus advantage of de-mystifying some of those formalisms, and through these applications natural alternatives and generalizations of formalisms in those areas became visible.

Finally, structuring states as in 11.5 and 11.6 of this note may be seen as initial steps of a larger program of adding structure to objects. As to adding structuring the transitions between states, rather than or in addition to structuring the states, there seems to be a problem. When transitions are equated with pairs of objects rather than treated as first-class citizens in their own right, there does not seem to be an obvious way to structure them. But Van Benthem [3] proposes a system of *arrow logic* in which the transitions or arrows have a primary status in the ontology, without necessarily being identified with pairs of states. Eventually this might be the way to go if one wants to be able to structure transitions as well as objects.

Acknowledgments

This chapter was conceived of during a visit to the Center for the Study of Language and Information, Stanford University, on NWO-grant SIR 11-596. I would like to thank David Israel for valuable suggestions. Back in the Netherlands Wiebe van der Hoek, Frank Veltman and especially Johan van Benthem provided useful comments.
The investigations were supported by the Foundation for Philosophical Research (SWON), which is subsidized by the Netherlands Organization for Scientific Research (NWO).

References

[1] Alchourrón, C., Gärdenfors, P. & Makinson, D. (1985), 'On the logic of theory change: partial meet contraction and revision functions', *Journal of Symbolic Logic* **50**, 510–530.

[2] van Benthem, J. (1989), Semantic parallels in natural language and computation, *in* H.-D. Ebbinghaus et al., eds, 'Logic Colloquium. Granada 1987.', North-Holland, Amsterdam, pp. 331–375.

[3] van Benthem, J. (1991), *Language in Action*, North-Holland, Amsterdam.

[4] van Benthem, J. (1991), Logic and the flow of information, *in* D. Prawitz et al., eds, 'Proc. 9th Intern. Congress of Logic, Method. and Philos. of Science', North-Holland, Amsterdam.

[5] Blackburn, P. & Spaan, E. (1993), 'A modal perspective on the computational complexity of attribute value grammar', *Journal of Logic, Language and Information.* To appear.

[6] Brink, C. (1981), 'Boolean modules', *Journal of Algebra* **71**, 291–313.

[7] Brink, C. & Schmidt, R. (1992), 'Subsumption computed algebraically', *Computers Math. Applic.* **23**, 329–342.

[8] van Eijck, J. & de Vries, F.-J. (1993), Reasoning about update logic, Technical Report CS-R9312, CWI, Amsterdam. To appear in the *Journal of Philosophical Logic.*

[9] Finger, M. & Gabbay, D. (1992), 'Adding a temporal dimension to a logic system', *Journal of Logic, Language and Information* **1**, 203–233.

[10] Fuhrmann, A. (1990), On the modal logic of theory change, *in* A. Fuhrmann & M. Morreau, eds, 'Lecture Notes in AI, 465', Springer, pp. 259–281.

[11] Gabbay, D., Hodkinson, I. & Reynolds, M. (1993), *Temporal Logic: Mathematical Foundations and Computational Aspects*, Oxford University Press. To appear.

[12] Gärdenfors, P. (1988), *Knowledge in Flux*, The MIT Press, Cambridge, MA.

[13] Gärdenfors, P. & Makinson, D. (1988), Revisions of knowledge systems using epistemic entrenchment, *in* 'Proc. 2nd. Conf. Theoret. Aspects of Reasoning about Knowledge', Pacific Grove, pp. 83–95.

[14] Gargov, G. & Passy, S. (1990), A note on Boolean modal logic, *in* P. Petkov, ed., 'Mathematical Logic. Proceedings of the 1988 Heyting Summerschool', Plenum Press, New York, pp. 311–321.

[15] Grahne, G. (1991), Updates and counterfactuals, *in* J. Allen, R. Fikes & E. Sandewall, eds, 'Principles of Knowledge Representation and Reasoning: Proc. 2nd. Intern. Conf.', Morgan Kaufman, San Mateo, pp. 269–276.

[16] Groeneveld, W. (1992), Contextual *S*5: a 'static' renotation of Update Semantics (Draft), Manuscript.

[17] Hansson, S. (1991), 'Belief contraction without recovery', *Studia Logica* **50**, 251–260.

[18] Harel, D. (1984), Dynamic logic, *in* D. Gabbay & F. Guenthner, eds, 'Handbook of Philosophical Logic', Vol. 2, Reidel, Dordrecht, pp. 497–604.

[19] van der Hoek, W. & de Rijke, M. (1993), 'Generalized quantifiers and modal logic', *Journal of Logic, Language and Information* **2**, 19–58.

[20] Katsuno, H. & Mendelzon, A. (1991), On the difference between updating a knowledge base and revising it, *in* J. Allen, R. Fikes & E. Sandewall, eds, 'Principles of Knowledge Representation and Reasoning: Proc. 2nd Intern. Conf.', Morgan Kaufman, San Mateo, pp. 387–394.

[21] Katsuno, H. & Mendelzon, A. (1992), 'Propositional knowledge base revision and minimal change', *Artificial Intelligence* **52**, 263–294.

[22] Lindström, S. & Rabinowitz, W. (1989), 'On probabilistic representation of non-probabilistic belief

revision', *Journal of Philosophical Logic* **18**, 69–101.

[23] de Rijke, M. (1992), A system of dynamic modal logic, Technical report # CSLI-92-170, Stanford University. Also appeared as ILLC Report LP-92-08, University of Amsterdam.

[24] de Rijke, M. (1993), Modal axioms for Peirce algebras, Manuscript, ILLC, University of Amsterdam.

[25] Ryan, M. (1992), Ordered Theory Presentations, PhD thesis, Imperial College.

[26] Schild, K. (1990), A correspondence theory for terminological logics, *in* 'Proc. 12th IJCAI', Sydney, Australia, pp. 466–471.

[27] Schmidt, R. (1991), Algebraic terminological representation, Master's thesis, Department of Mathematics, University of Cape Town.

[28] Schmidt-SchauS, M. (1989), Subsumption in KL-ONE is undecidable, *in* 'Proc. 1st Intern. Conf. Princ. of Knowledge Representation and Reasoning', Morgan Kaufman, San Mateo, pp. 421–431.

[29] Shoham, Y. (1987), Nonmonotonic logics: meaning and utility, *in* 'Proc. IJCAI-87', Milan, Italy, pp. 388–393.

[30] Veltman, F. (1992), 'Defaults in update semantics', *Journal of Philosophical Logic*. To appear.