

Modular Algorithms for Heterogeneous Modal Logics via Multi-Sorted Coalgebra

LUTZ SCHRÖDER^{1†} and DIRK PATTINSON^{2‡}

¹ *DFKI Bremen and Dept. of Comput. Sci., Univ. Bremen,
Cartesium, Enrique-Schmidt-Str. 5, 28359 Bremen, Germany*
E-mail: Lutz.Schroeder@dfki.de

² *Dept. of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK*
E-mail: dirk@doc.ic.ac.uk

Received

State-based systems and modal logics for reasoning about them often heterogeneously combine a number of features such as non-determinism and probabilities. Here, we show that the combination of features can be reflected algorithmically and develop modular decision procedures for heterogeneous modal logics. The modularity is achieved by formalizing the underlying state-based systems as multi-sorted coalgebras and associating both a logical and an algorithmic description to a number of basic building blocks. Our main result is that logics arising as combinations of these building blocks can be decided in polynomial space provided that this is the case for the components. By instantiating the general framework to concrete cases, we obtain PSPACE decision procedures for a wide variety of structurally different logics, describing e.g. Segala systems and games with uncertain information.

1. Introduction

Modal logics appear in computer science in a variety of contexts. They are the formalism of choice for reasoning about reactive systems and feature prominently in areas related to artificial intelligence such as knowledge representation and reasoning with uncertainty (Halpern 2003). The semantics of modal logics typically involves a notion of state and transition, which can take a number of different forms. Transitions can be probabilistic or weighted, as in probabilistic modal logic (Larsen and Skou 1991; Heifetz and Mongin 2001) and graded modal logic (Fine 1972; D'Agostino and Visser 2002), induced by joint actions of agents as in coalition logic or alternating-time logic (Pauly 2002; Alur, Henzinger, and Kupferman 2002), or non-monotonically conditioned as in conditional logic (Chellas 1980). An attractive aspect of many of these logics is that they are decidable in comparatively low complexity classes — e.g., in the absence of fixpoint operators and global assumptions typically in PSPACE (Tobies 2001;

[†] Work partially supported by BMBF grant FormalSafe (FKZ 01IW07002)

[‡] Work supported by EPSRC grant EP/F031173/1

Pauly 2002; Schröder and Pattinson 2009), i.e. the same as the standard modal logic K (Blackburn, de Rijke, and Venema 2001) and not dramatically worse than propositional logic.

Features like non-determinism, probabilistic choice or joint actions are often combined, leading to systems that incorporate more than one type of transition. Moreover, features can be combined in different ways: E.g. in the alternating model of probabilistic transition systems (Hansson and Jonsson 1990), states may have either non-deterministic or probabilistic transitions, whereas simple Segala systems (Segala 1995) have a two-layered structure where non-deterministic transitions lead to probability distributions over successor states. Bartels et al. (2003) discuss 12 different types of probabilistic transition systems arising in the literature that come about as such combinations of basic features.

Here, we introduce a simple calculus that formalizes the combination of features and establish that combined logics typically inherit properties of their building blocks in a straightforward way; as an example, we show that shallow models and decidability in PSPACE according to the method of strictly one-step complete rule sets (Schröder and Pattinson 2009) transfer from the components to the combined logic. Our results and algorithms are *generic* and use the same algorithmic template to realize decision procedures at the level of each individual feature. This is achieved by formalising the combined logics in a multi-sorted extension of *coalgebraic modal logic* (Pattinson 2004) whose semantics is parametric in a set functor T ; models then appear as T -coalgebras. This pushes the generic PSPACE-decision procedure of Schröder and Pattinson (2009), which works uniformly for such diverse logics as Hennessy-Milner logic, coalition logic, graded modal logic, and probabilistic modal logic, to the level of combined logics that integrate several features.

Formally, a feature consists of a set of modal operators together with a set of associated proof rules. On the semantic level, a structure for a feature is an endofunctor of type $\mathbf{Set}^n \rightarrow \mathbf{Set}$, where n is the arity of the feature (e.g. choice, fusion, and conditionality are binary features). The notion of *gluing* formalizes specific ways of combining given features. Syntactically, gluings define multi-sorted modal logics. Semantically, gluings induce endofunctors $T : \mathbf{Set}^n \rightarrow \mathbf{Set}^n$ such that T -coalgebras are models of the combined logic. The single sorted case $n = 1$ is of special interest since it captures the standard models of combined systems, including e.g. the ones presented by Bartels et al. (2003), which equip multi-sorted logics with a single-sorted semantics.

The central technical contribution of this work is the construction of a logically equivalent *flattening* of a given gluing, where flat gluings assign to each occurrence of a feature an individual sort in the semantics. Flat gluings are technically more tractable than general gluings, and in fact the extension of generic results in coalgebraic modal logic from the single-sorted case to flat gluings typically requires no more than accommodating some notational overhead for the sorts. We demonstrate this by establishing the shallow model property and a generic PSPACE algorithm for flat gluings, in generalization of previous results for the single-sorted case (Schröder and Pattinson 2009). Together, these results imply PSPACE upper bounds for satisfiability over general gluings, including the standard single-sorted semantics of multi-sorted modal logics such as the logic of simple Segala systems (Jonsson, Yi, and Larsen 2001).

Related Work. Our work is closely related to several previous approaches (Cîrstea and Pattinson 2007; Jacobs 2001), which focus on completeness issues, with the main difference that our re-

sults make the multi-sorted nature of heterogeneous logics explicit by considering multi-sorted models. Our treatment of typed formulas resembles the concept of *ingredients* (Jacobs 2001), but the multi-sorted semantics avoids the use of the next-operator of *op.cit.* One advantage of the new framework is that constructions such as Cartesian product or disjoint union no longer have to be treated separately but are covered as special cases by the overall framework, which needs to talk only about polyadic set functors and their predicate liftings. Moreover, generic results in coalgebraic modal logic such as the decision procedures of (Schröder 2007; Schröder and Pattinson 2009), and many others as listed in more detail in Section 6, generalize straightforwardly to the multi-sorted case. The multi-sorted approach to the complexity of composite modal logics complements transfer results obtained for the fusion of modal logics (Hemaspaandra 1994; Wolter 1998) in the sense that our framework currently covers logics with iterative axioms (Lewis 1975) (i.e. axioms that nest modalities) only in a somewhat limited way via a multi-sorted extension of the results of (Schröder and Pattinson 2008a), but allows more flexible logic composition and especially applies also to logics that are not amenable to Kripke semantics.

Outline We begin with a brief informal discussion of a few examples of multi-sorted modal logics (Section 2) before we proceed to give a formal syntactic definition of our logic composition mechanism in Section 3. We then define the multi-sorted coalgebraic semantics of composite modal logics in Section 4, where we also present the main result of this work, the flattening construction which reduces functor composition to multi-sortedness. Finally, we demonstrate the straightforward extension of single-sorted results to the multi-sorted setting in Section 5, using the above-mentioned generic PSPACE-algorithm as an example. This work is an extended and revised version of (Schröder and Pattinson 2007).


2. Multi-Sorted Modal Logics by Example

We proceed to present examples of logics that combine reasoning mechanisms in various ways, to motivate the technical definitions that follow. We discuss both the standard definitions of system types and their coalgebraic rephrasing. Readers unfamiliar with coalgebraic terminology can safely ignore this coalgebraic aspect for now, as we will recall the basic notions of coalgebra in Section 4. The coalgebraic treatment of probabilistic systems follows (Bartels, Sokolova, and de Vink 2003).

2.1. Logics for Probabilistic Systems

Segala systems (Segala 1995) and alternating systems (Hansson and Jonsson 1990) both combine probabilistic transitions and non-determinism. In Segala systems, each system state can non-deterministically perform actions that lead to probability distributions over states. Contrastingly, alternating systems have two kinds of states engaging in purely probabilistic transitions and non-deterministic actions, respectively, that may end up in either kind of state. In the example below, the non-deterministic states are represented by solid circles ●, and the probabilistic states by hollow circles ○. Note that we have also represented the intermediate probability distributions occurring in Segala systems as hollow circles, although these do not as yet correspond to actual states of the system — it is the main message of the current work, to be elaborated in the

following sections, that nothing is changed if we do eventually promote them to states in their own right.



Formally, a simple Segala system over a set A of actions consists of a set X of states and for each state $x \in X$ and each action $a \in A$, a set $\xi(x)(a)$ of probability distributions on X . In coalgebraic terms, this means that we have a coalgebra

$$\xi : X \rightarrow (A \rightarrow \mathcal{P}(\mathcal{D}(X)))$$

for the functor T defined by $TX = (A \rightarrow (\mathcal{P}(\mathcal{D}(X))))$, where \rightarrow denotes function space, \mathcal{P} denotes powerset, and $\mathcal{D}(X)$ is the set of discrete probability distributions on X . In contrast, an alternating system consists of a set X of states and, for each state $x \in X$, either a probability distribution $\xi(x)$ on X or a map $\xi(x)$ assigning to each action $a \in A$ a set $\xi(x)(a)$ of successor states. In other words, we have a coalgebra

$$\xi : X \rightarrow \mathcal{D}(X) + (A \rightarrow \mathcal{P}(X))$$

where $+$ denotes disjoint sum.

It has been shown that a suitable variant of *probabilistic modal logic* (Larsen and Skou 1991) over a set A of actions characterizes states of image-finite Segala systems up to bisimilarity (Jonsson, Yi, and Larsen 2001, Theorem 8). This logic has two sorts n and u of non-deterministic and probabilistic (‘uncertain’) formulas, respectively, and two families of modal operators

$$\Box_a : u \rightarrow n \quad (a \in A) \quad \text{and} \quad L_p : n \rightarrow u \quad (p \in [0, 1] \cap \mathbb{Q}),$$

where L_p reads ‘with probability at least p ’. The sets \mathcal{L}_n and \mathcal{L}_u of non-deterministic and probabilistic formulas, respectively, are thus defined by the grammar

$$\begin{aligned} \mathcal{L}_n \ni \phi &::= \top \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid \Box_a \psi & (\psi \in \mathcal{L}_u, a \in A) \\ \mathcal{L}_u \ni \psi &::= \top \mid \psi_1 \wedge \psi_2 \mid \neg\psi \mid L_p \phi & (\phi \in \mathcal{L}_n, p \in [0, 1] \cap \mathbb{Q}). \end{aligned}$$

Given a Segala system (X, ξ) as above, the semantics of the two sorts of formulas is given by satisfaction relations \models between distributions P on X and probabilistic formulas and between states $x \in X$ and non-deterministic formulas, respectively, defined by mutual recursion with the obvious clauses for Boolean operators and

$$\begin{aligned} x \models \Box_a \psi &\text{ iff } \forall P \in \xi(x)(a). P \models \psi \\ P \models L_p \phi &\text{ iff } P(\{x \mid x \models \phi\}) \geq p \end{aligned}$$

for the modal operators. Alternating systems, on the other hand, can be captured by a logic

comprising three sorts n , u , and o of non-deterministic, probabilistic, and alternating formulas, respectively, and modal operators

$$+ : u, n \rightarrow o \quad L_p : o \rightarrow u \quad \Box_a : o \rightarrow n$$

that induce the three-sorted grammar

$$\begin{aligned} \mathcal{L}_u \ni \phi &::= \top \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid L_p\chi & (\chi \in \mathcal{L}_a, p \in [0, 1] \cap \mathbb{Q}) \\ \mathcal{L}_n \ni \psi &::= \top \mid \psi_1 \wedge \psi_2 \mid \neg\psi \mid \Box_a\chi & (\chi \in \mathcal{L}_a, a \in A) \\ \mathcal{L}_a \ni \chi &::= \top \mid \chi_1 \wedge \chi_2 \mid \neg\chi \mid \phi + \psi & (\phi \in \mathcal{L}_u, \psi \in \mathcal{L}_n) \end{aligned}$$

which defines the sets \mathcal{L}_n , \mathcal{L}_u , and \mathcal{L}_a of non-deterministic, probabilistic, and alternating formulas, respectively. The binary modal operator $+$ implements the choice between probabilistic and non-deterministic transitions, which is a case statement: $\phi + \psi$ demands that ϕ holds if the present state is probabilistic whereas ψ holds if present state is non-deterministic. Formally, the semantics of the logic over an alternating system (X, ξ) is given by satisfaction relations between maps $f : A \rightarrow \mathcal{P}(X)$ and non-deterministic formulas, between distributions $P \in \mathcal{D}(X)$ and probabilistic formulas, and between states $x \in X$ and alternating formulas, respectively, defined by mutual recursion, with the clauses for the modal operators being

$$\begin{aligned} P \models L_p\chi &\text{ iff } P(\{x \mid x \models \chi\}) \geq p \\ f \models \Box_a\psi &\text{ iff } \forall x \in f(a). x \models \psi \\ x \models \phi + \psi &\text{ iff } \begin{cases} \xi(x) \models \phi & \text{if } \xi(x) \in \mathcal{D}(X) \\ \xi(x) \models \psi & \text{if } \xi(x) \in (A \rightarrow \mathcal{P}(X)). \end{cases} \end{aligned}$$

E.g., a state x in an alternating system satisfies the formula

$$L_{1/2}(\perp + \Box_a\perp) + \Box_a\perp$$

(where $\perp \equiv \neg\top$) if either x behaves probabilistically, and with probability at least 1/2 leads to a non-deterministic state where a is blocked, or x behaves non-deterministically and blocks a .

2.2. Fusion of Modal Logics

Both logics described above wire up the component logics in a restricted way, by imposing layering and choice, respectively. The unrestricted combination of logics \mathcal{L}_a and \mathcal{L}_b can be modelled by a logic with sorts a , b , f and modal operators with associated source and target sorts

$$[\pi_1] : a \rightarrow f \quad [\pi_2] : b \rightarrow f \quad \Box : f \rightarrow a \quad \heartsuit : f \rightarrow b$$

where \Box ranges over the modal operators of \mathcal{L}_a and \heartsuit over those of \mathcal{L}_b . This gives rise to an obvious three-sorted grammar. The semantics of $[\pi_1]$ and $[\pi_2]$ (given in Section 4) will ensure that both operators commute with all Boolean operations. Taking into account the typing information for the modal operators, this means that formulas of sort f are essentially just the well-known *fusion* $\mathcal{L}_a \otimes \mathcal{L}_b$ (see, e.g., (Kurucz 2006)). Recall that the fusion disjointly combines the axioms and modalities of \mathcal{L}_a and \mathcal{L}_b , giving rise to a one-sorted grammar. One can translate back and forth between the fusion and formulas of sort f by translating the operator \Box of the fusion to the composite operator $[\pi_1]\Box$, and the operator \heartsuit to $[\pi_2]\heartsuit$; a converse translation works by first

distributing $[\pi_1]$ and $[\pi_2]$ over Boolean connectives and then replacing $[\pi_2]\heartsuit$ with \heartsuit and $[\pi_1]\square$ with \square . Thus, fusion is an instance of the multi-sorted combination of modal logics.

As fusion does not impose any well-typedness constraints on formulas, it can be regarded as the maximally permissive way of combining two modal logics. However, as shown by the previous example, formulas of the fusion do not in general have an interpretation over the intended type of systems, so that it is for many purposes preferable to work with the more restrictive well-typed combinations considered here. Moreover, although the theoretical complexity of the fusion typically does not exceed that of well-typed fragments, excluding ill-typed formulas still does reduce the search space for purposes of practical implementations.

2.3. Conditional Logic

The standard conditional logic CK (Chellas 1980) has a binary modal operator \Rightarrow , where $\phi \Rightarrow \psi$ is read as a non-monotonic conditional ‘if ... then normally ...’. In the right hand argument, \Rightarrow behaves essentially like the \square of the normal modal logic K , and in particular obeys the usual K -axiom when the left argument is fixed. Indeed we can embed CK into a two-sorted *extended conditional logic* with sorts c, k and modal operators

$$\overset{\bullet}{\Rightarrow} : c, k \rightarrow c \quad \square : c \rightarrow k$$

by translating $\alpha \Rightarrow \beta$ to $\alpha \overset{\bullet}{\Rightarrow} \square\beta$. Here, $\overset{\bullet}{\Rightarrow}$ represents a rudimentary conditional, and \square is the standard box modality of K . The semantic and syntactic details of this combination will be made explicit in Sections 3 and 4. This shows how a given complex logic can be broken down into simpler building blocks.

3. Compositional Syntax of Multi-Sorted Modal Logic

For our purposes, it is convenient to present the syntax of multi-sorted modal logic in a way that provides explicitly for a decomposition into building blocks. The building blocks, which we call *features*, are collections of (possibly polyadic) modal operators and associated proof rules that capture specific properties of a logic, such as the ability to describe choice, non-determinism, or uncertainty.

Definition 3.1. An n -ary *feature* is a pair $F = (\Lambda, \mathcal{R})$ consisting of a (*modal*) *similarity type* Λ , i.e. a set of modal operators \heartsuit with profiles $\heartsuit : i_1, \dots, i_k \rightarrow *$, where $1 \leq i_1, \dots, i_k \leq n$ are formal argument sorts and $*$ is a formal target sort, and a set \mathcal{R} of *one-step rules* of the form

$$R = (\phi_1; \dots; \phi_n) / \psi,$$

where for $i = 1, \dots, n$, ϕ_i is a propositional formula over a set V_i of propositional variables, and ψ is a clause over atoms of the form $\heartsuit(a_1, \dots, a_k)$ with $\heartsuit : i_1, \dots, i_k \rightarrow *$ in Λ and $a_j \in V_{i_j}$, $j = 1, \dots, k$. Here, we use the standard terminology for propositional logic: a *literal* over a set Z of atoms is either an element of Z or the negation of such an element; a *clause* over Z is a finite disjunction of literals over Z ; and a *conjunctive clause* over Z is a finite conjunction of literals over Z . Given a substitution σ of the variables in R by elements of Z , we refer to the substituted rule $R\sigma = (\phi_1\sigma; \dots; \phi_n\sigma) / \psi\sigma$ as a Z -instance of R . Finally, by phrases such as *propositional*

reasoning, propositional entailment, propositional equivalence etc. we refer to reasoning using propositional tautologies and modus ponens.

Although this does not play a crucial role in the present development, we moreover impose the technical condition needed in other contexts that rules be *injective*, which means that every propositional variable can occur at most once in the conclusion of a rule (Schröder and Pattinson 2009). This is not an actual restriction, since different variables a, b can be made to carry the same interpretation by introducing a premise $a \leftrightarrow b$, and keeps the rules in synchrony with the presentation elsewhere.

Note that the rule format disallows nested modalities in the conclusion, so that rules describe the one-step behaviour of a system. As in the single sorted case (Schröder 2007), this format always suffices to completely axiomatize the features of interest, as long as no global conditions (such as transitivity) are imposed on the coalgebraic models.

Example 3.2. We describe the features used in the Examples of Section 2, as well as the standard features that implicitly occur in previous treatments of compositional coalgebraic logics (Jacobs 2001; Cirstea and Pattinson 2007). The similarity types are as follows (with intuitive meanings supported by the respective semantics to be given in Section 4).

Non-Determinism: Given a set A of actions, the unary feature N_A has modal operators $\Box_a : 1 \rightarrow *$ for $a \in A$, read ‘all a -successors satisfy ...’. If A is a singleton, we write K instead of N_A , and \Box instead of \Box_a .

Uncertainty: The unary feature U has modal operators $L_p : 1 \rightarrow *$ for $p \in [0, 1] \cap \mathbb{Q}$, read ‘with probability at least p , the next state satisfies ...’.

Coalition: Given a set N of *agents*, the unary feature G_N has modal operators $[C] : 1 \rightarrow *$ for all *coalitions* C , i.e. all subsets of N , read ‘coalition C has a collaborative strategy to enforce ...’.

Choice: The binary feature S has a single binary modal operator $+ : 1, 2 \rightarrow *$, where $\phi + \psi$ reads ‘if the next state is in the left alternative, then it satisfies ϕ , and if the next state is in the right alternative, then it satisfies ψ ’.

Fusion: The binary feature P has two modal operators $[\pi_i] : i \rightarrow *$, $i = 1, 2$, read ‘the i -th component satisfies ...’.

Conditionality: The binary feature C has a binary modal operator $\overset{\bullet}{\Rightarrow} : 1, 2 \rightarrow *$, read ‘if ... then normally ...’.

Input: Given a set I of inputs, the unary feature I_I has modal operators $(i) : 1 \rightarrow *$ for $i \in I$, read ‘upon input i ’.

Output: Given a finite set O of outputs (infinite sets can be handled as well but have a slightly different rule set), the nullary feature O_O has nullary modal operators (i.e. flexible constants) $o : \rightarrow *$ for $o \in O$, read ‘the current state outputs o ’. For the special case of O_1 , where 1 is a singleton output set, we write 1 for O_1 .

Figure 1 shows the associated proof rules, already in a special format needed in Section 5. The rules for non-determinism and uncertainty are taken from (Schröder and Pattinson 2009), where we use an obviously equivalent simplification of the rule for uncertainty; the others are obtained by the same principles. The sum expression in the uncertainty rule refers to the arithmetic of

Nondeterminism :	$\frac{\bigwedge_{j=1}^n \alpha_j \rightarrow \beta}{\bigwedge_{j=1}^n \Box_a \alpha_j \rightarrow \Box_a \beta} \quad (n \geq 0, a \in A)$
Uncertainty :	$\frac{\sum_{i=0}^n r_i a_i \sqsupset \sum_{i=0}^n r_i p_i}{\bigvee_{0 \leq i \leq n} \text{sgn}(r_i) L_{p_i} a_i} \left(\begin{array}{l} n \geq 0; r_0, \dots, r_n \in \mathbb{Z} - \{0\}; \sqsupset \in \{>, \geq\}; \\ \sqsupset \text{ is } > \text{ if } r_j < 0 \text{ for all } j, \text{ and } \geq \text{ otherwise.} \end{array} \right)$
Coalition :	$\frac{\bigvee_{i=1}^n \neg \alpha_i}{\bigvee_{i=1}^n \neg [C_i] \alpha_i} \quad \frac{\bigwedge_{i=1}^n \alpha_i \rightarrow (\beta \vee \bigvee_{j=1}^m c_j)}{\bigwedge_{i=1}^n [C_i] \alpha_i \rightarrow ([D] \beta \vee \bigvee_{j=1}^m [N] c_j)}$ $(m, n \geq 0; C_j \subseteq D \text{ for } j = 1, \dots, m; C_i \cap C_j = \emptyset \text{ for } i \neq j)$
Choice :	$\frac{(\bigwedge_{j=1}^m \alpha_j \rightarrow \bigvee_{k=1}^n \beta_k) : 1 \quad (\bigwedge_{j=1}^m \gamma_j \rightarrow \bigvee_{k=1}^n \delta_k) : 2}{\bigwedge_{j=1}^m (\alpha_j + \gamma_j) \rightarrow \bigvee_{k=1}^n (\beta_k + \delta_k)} \quad (m, n \geq 0)$
Fusion :	$\frac{(\bigwedge_{j=1}^m \alpha_j \rightarrow \bigvee_{k=1}^n \beta_k) : i}{\bigwedge_{j=1}^m [\pi_i] \alpha_j \rightarrow \bigvee_{k=1}^n [\pi_i] \beta_k} \quad (i = 1, 2; m, n \geq 0)$ $(\gamma_i \leftrightarrow \gamma_j) : 1 \quad (i, j \in \{1, \dots, m+n\})$
Conditionality :	$\frac{(\bigwedge_{j=1}^m \alpha_j \rightarrow \bigvee_{k=1}^n \beta_k) : 2}{\bigwedge_{j=1}^m (\gamma_i \overset{\bullet}{\Rightarrow} \alpha_j) \rightarrow \bigvee_{k=1}^n (\gamma_{k+m} \overset{\bullet}{\Rightarrow} \beta_k)} \quad (m, n \geq 0)$
Input :	$\frac{(\bigwedge_{j=1}^m \alpha_j \rightarrow \bigvee_{k=1}^n \beta_k) : i}{\bigwedge_{j=1}^m (i) \alpha_j \rightarrow \bigvee_{k=1}^n (i) \beta_k} \quad (i \in I; m, n \geq 0)$
Output :	$\frac{}{o_1 \wedge o_2 \rightarrow \perp} \quad (o_1, o_2 \in O; o_1 \neq o_2) \quad \frac{}{\bigvee_{o \in O} o}$

Figure 1. Proof rules for the features of Example 3.2

characteristic functions (Schröder and Pattinson 2009); explicitly,

$$\sum_{i \in I} r_i \phi_i \geq k \equiv \bigwedge_{\substack{J \subseteq I \\ r(J) < k}} \left(\bigwedge_{j \in J} \phi_j \rightarrow \bigvee_{j \notin J} \phi_j \right),$$

where $r(J) = \sum_{j \in J} r_j$. Moreover, $\text{sgn}(r)\phi \equiv \phi$ when $r > 0$, and $\text{sgn}(r)\phi \equiv \neg\phi$ otherwise.

We note that the above list of features is by no means exhaustive. In particular, we can see every modal logic that is axiomatized by a set of one-step rules (Schröder and Pattinson 2008b, Definition 2) as a unary feature. Other natural features that arise in this way are for instance monotonicity, neighbourhoods, or integer weights, where the associated proof rules can be found in (Schröder and Pattinson 2009, Example 3.19 and Example 6.2).

Remark 3.3. In the nondeterminism feature, we have opted not to decompose the feature further into K and the input feature; this is both to illustrate that we are not forced to use the finest possible granularity in our choice of basic features, and to keep the running examples in tune with the introduction in Section 2.

The examples from Section 2 demonstrate that features can be combined in different ways. This is formalized by the notion of gluing.

Definition 3.4. Let Φ be a set of features, and let \mathcal{S} be a set of sorts. *Feature expressions* t (over \mathcal{S}) are terms over the set \mathcal{S} of variables where the features appear as function symbols, i.e.

$$t ::= a \mid F(t_1, \dots, t_n) \quad a \in \mathcal{S}, F \in \Phi \text{ } n\text{-ary.}$$

A *gluing* of Φ over \mathcal{S} is a family $G = (t_a)_{a \in \mathcal{S}}$ of feature expressions, denoted by $(a_1 \rightarrow t_{a_1}, \dots, a_n \rightarrow t_{a_n})$ for $\mathcal{S} = \{a_1, \dots, a_n\}$; in this case we also write $a_i \rightarrow t_{a_i} \in G$.

A gluing $G = (t_a)_{a \in \mathcal{S}}$ induces a **multi-sorted modal logic**, as follows. The set $\text{Types}(G)$ of *types* consists of the proper subterms of the t_a , where the sorts $a \in \mathcal{S}$ are called *base types* and the expressions $t \in \text{Types}(G) \setminus \mathcal{S}$ are the *composite types*. The reason the t_a themselves are not in general regarded as types is that we can instead represent them by their index sort a ; this is related to the semantic interpretation of the association $a \rightarrow t_a$ as a component of a multi-sorted coalgebra map, to be defined in Section 4. The types are related to the *ingredients* of (Jacobs 2001), and will serve to type modal formulas in our logic. In some sense, they also correspond to types of states, an intuition that will be made explicit in our flattening construction. We call a gluing *flat* if $\mathcal{S} = \text{Types}(G)$, i.e. there are no composite types, which is the case if every term t_a is of the form $F(a_1, \dots, a_n)$; intuitively, a gluing is flat if all types of states are already explicit in the sort set \mathcal{S} .

Typed G -formulas $\phi : s$, where $s \in \text{Types}(G)$, are inductively generated by closure under Boolean operators \perp, \neg, \wedge at each type (with further Boolean operators $\vee, \top, \rightarrow, \leftrightarrow$ defined in the standard way) and by the following typing rules for composite types (left) and base types (right)

$$\frac{\phi_1 : s_1 \quad \dots \quad \phi_n : s_n}{\heartsuit(\phi_{i_1}, \dots, \phi_{i_n}) : F(s_1, \dots, s_n)} \quad \frac{\phi_1 : s_1 \quad \dots \quad \phi_n : s_n}{\heartsuit(\phi_{i_1}, \dots, \phi_{i_n}) : a},$$

where the left hand rule has side condition $F(s_1, \dots, s_n) \in \text{Types}(G)$ and the right hand rule has side condition $a \rightarrow F(s_1, \dots, s_n) \in G$, and in both cases $\heartsuit : i_1, \dots, i_n \rightarrow *$ in F . The special status of base types is related to the fact that for $a \rightarrow F(s_1, \dots, s_n) \in G$, the feature expression $F(s_1, \dots, s_n)$ is not in general a type, being instead thought of as represented by the base type a . We write $\mathcal{F}_s(G)$ for the set of G -formulas of type s and denote the family $(\mathcal{F}_s(G))_{s \in \text{Types}(G)}$ by $\mathcal{F}(G)$.

Similarly, a gluing induces a typed **proof system**, described in terms of a $\text{Types}(G)$ -indexed family of derivability predicates $\vdash_s \subseteq \mathcal{F}_s(G)$. These predicates are defined inductively by closure under propositional reasoning at each type and the deduction rules for composite types (left) and base types (right), distinguished only by the type discipline in the same way as the above typing rules for formulas,

$$\frac{\vdash_{s_1} \phi_1 \sigma \quad \dots \quad \vdash_{s_n} \phi_n \sigma}{\vdash_{F(s_1, \dots, s_n)} \psi \sigma} \quad \frac{\vdash_{s_1} \phi_1 \sigma \quad \dots \quad \vdash_{s_n} \phi_n \sigma}{\vdash_a \psi \sigma}$$

where $F(s_1, \dots, s_n) \in \text{Types}(G)$ in the left hand rule, $a \rightarrow F(s_1, \dots, s_n) \in G$ in the right hand rule, and in both cases, $(\phi_1; \dots; \phi_n)/\psi$ is a rule of F and σ is a substitution mapping variables $a \in V_i$ to formulas $\sigma(a) : s_i$.

Remark 3.5. Formulas and proofs induced by a gluing $G = (t_a)_{a \in \mathcal{S}}$ need not have unique types, namely in case t_a is a subterm of t_b for some $a, b \in \mathcal{S}$. We therefore disambiguate formulas explicitly by their type when we define the semantics in Section 4.

A given logic can be syntactically generated by different gluings, typically including both flat and non-flat ones, determining different classes of semantic structures (see Section 4). One of the main contributions of this paper is the proof of logical equivalence for the respective semantics. Flat gluings are technically more tractable, while logics occurring in the literature, including the ones described in Section 2, are typically non-flat.

Example 3.6. From the features S , U , and N_A (Example 3.2), we can form gluings

$$G_1 \equiv (o \rightarrow S(U(o), N_A(o))) \text{ and } G_2 \equiv (o \rightarrow S(u, n), u \rightarrow U(o), n \rightarrow N_A(o)).$$

Here, G_1 has types o , $N_A(o)$, $U(o)$, whereas G_2 is flat with types o , n , u . Modulo identifications $N_A(o) = n$ and $U(o) = u$, both gluings give rise to the (typed) formulas describing alternating systems (Section 2.1). To illustrate the typing system, we briefly show how the formula $L_p\alpha + \Box_a\beta$ is typed in the non-flat gluing G_1 , assuming $\alpha : o$, $\beta : o$:

$$\frac{\frac{\alpha : o}{L_p\alpha : U(o)} \quad \frac{\beta : o}{\Box_a\beta : N_A(o)}}{L_p\alpha + \Box_a\beta : o}$$

Here, the last derivation step uses the typing rule for base types and the other derivation steps use the typing rule for composite types.

The remaining example logics from Section 2 are captured by the following gluings:

Probabilistic modal logic of Segala Systems: This is the single-sorted gluing $n \rightarrow N_A(U(n))$.

It has types n and $U(n)$, the latter corresponding to the sort u in the discussion in Section 2.

We can generate the same logic from the flat gluing $(n \rightarrow N_A(p), p \rightarrow U(n))$.

Fusion: The fusion of logics \mathcal{L}_a and \mathcal{L}_b as in Section 2.2, regarded as features, is $f \rightarrow P(\mathcal{L}_a(f), \mathcal{L}_b(f))$.

Extended conditional logic: This may be captured by the single-sorted gluing $c \rightarrow C(c, K(c))$.

Note in particular that in the induced proof system, we can derive the standard rule

$$(RCK) \frac{\bigwedge_{i=1}^n \alpha_i \rightarrow \beta}{\bigwedge_{i=1}^n (\gamma \Rightarrow \alpha_i) \rightarrow (\gamma \Rightarrow \beta)} \quad (n \geq 0)$$

of the conditional logic CK (Chellas 1980) (presented for the sake of readability in non-injective form, i.e. with repeated occurrences of γ), where $\gamma \Rightarrow \alpha$ abbreviates $\gamma \overset{\bullet}{\Rightarrow} \Box\alpha$, as follows:

$$\frac{\frac{\vdash_c \gamma \leftrightarrow \gamma \quad \frac{\vdash_c \bigwedge_{i=1}^n \alpha_i \rightarrow \beta}{\vdash_{K(c)} \bigwedge_{i=1}^n \Box\alpha_i \rightarrow \Box\beta}}{\vdash_c \bigwedge_{i=1}^n \gamma \overset{\bullet}{\Rightarrow} \Box\alpha_i \rightarrow \gamma \overset{\bullet}{\Rightarrow} \Box\beta}}$$

Here, the first step in the right-hand branch applies the rule for K following the typing discipline for composite types, and the second step applies the rule for C following the typing discipline for base types.

The conditional logic $CKCEM$ (see, e.g., (Olivetti, Pozzato, and Schwind 2007)), which additionally obeys the axiom of *conditional excluded middle*

$$(\alpha \Rightarrow \beta) \vee (\alpha \Rightarrow \neg\beta),$$

can be captured (in extended form) as the gluing

$$c \rightarrow \mathsf{C}(c, \mathsf{S}(c, 1)).$$

Here, we encode the standard conditional arrow $\phi \Rightarrow \psi$ as $\phi \xrightarrow{\bullet} (\psi + \top)$.

Probabilistic Conditional Logic A probabilistic variant of conditional logic is given by the gluing

$$c \rightarrow \mathsf{C}(c, \mathsf{U}(c)),$$

in which formulas such as $\phi \xrightarrow{\bullet} L_p \psi$ express ‘if ϕ , then we can normally expect ψ with confidence at least p ’.

Probabilistic Coalition Logic A combination of quantitative uncertainty and coalition logic is captured by the gluing

$$c \rightarrow \mathsf{G}_N(\mathsf{U}(c))$$

that involves the features Uncertainty and Coalition described in Example 3.2. It allows us to speak about games with elements of chance; e.g. the formula $[C]L_p \phi$ expresses that coalition C has a joint strategy to bring about ϕ with probability at least p .

4. Multi-Sorted vs. Single-Sorted Coalgebraic Semantics

We now generalize the coalgebraic interpretation of modal logic, introduced in (Pattinson 2004), to the multi-sorted case. Crucially, we interpret multi-sorted logics over multi-sorted coalgebras. The parametricity over signature functors for coalgebras is the key feature of our framework that allows for uniform results that can be instantiated to a large number of structurally different systems and logics. We recall some basic notions of multi-sorted coalgebra (see, e.g., (Mossakowski et al. 2006)), generalising the single-sorted setting (Rutten 2000):

Definition 4.1. We write \mathbf{Set} for the category of sets and functions. Let $\mathbf{Set}^{\mathcal{S}}$ denote the category of \mathcal{S} -sorted sets and \mathcal{S} -sorted functions, with objects being families $X = (X_a)_{a \in \mathcal{S}}$ (or just (X_a)) of sets X_a , and morphisms $f : (X_a) \rightarrow (Y_a)$ being families $f = (f_a)_{a \in \mathcal{S}}$ of maps $f_a : X_a \rightarrow Y_a$. We shall often introduce \mathcal{S} -sorted sets and maps using just one letter, say X , and then implicitly understand X_a as a notation for the a -th component of X . We write \mathbf{Set}^n for $\mathbf{Set}^{\{1, \dots, n\}}$. A functor $T : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}^{\mathcal{S}}$ may be regarded a family $T = (T_a)_{a \in \mathcal{S}}$ of functors $T_a : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}$. A T -coalgebra $A = (X, \xi)$ is a pair (X, ξ) where X is an \mathcal{S} -sorted set and $\xi = (\xi_a) : X \rightarrow TX$ is an \mathcal{S} -sorted function (i.e. $\xi_a : X_a \rightarrow T_a X$) called the *transition function*. A *morphism* between T -coalgebras (X, ξ) and (Y, ζ) is an \mathcal{S} -sorted function $f : X \rightarrow Y$ such that $(Tf)\xi = \zeta f$ in $\mathbf{Set}^{\mathcal{S}}$, i.e. at every sort $a \in \mathcal{S}$, the diagram

$$\begin{array}{ccc} X_a & \xrightarrow{f_a} & Y_a \\ \xi_a \downarrow & & \downarrow \zeta_a \\ T_a X & \xrightarrow{T_a f} & T_a Y \end{array}$$

commutes.

We view coalgebras as generalized transition systems: the transition function maps states to structured sets of observations and successor states, the latter taken from the available sorts as specified by T .

Assumption 4.2. We can assume w.l.o.g. that T preserves tuples of injective maps, by an easy extension of a construction from (Barr 1993) from single-sorted to the multi-sorted maps. For convenience of notation, we will in fact assume that (multi-sorted) subset inclusions $X \hookrightarrow Y$ are mapped to subset inclusions $TX \hookrightarrow TY$.

The interpretation of modal operators is based on predicate liftings (Pattinson 2004; Schröder 2008); in the multi-sorted setting, this takes the following shape.

Definition 4.3. A *predicate lifting* λ of profile $\lambda : i_1, \dots, i_k \xrightarrow{\bullet} *$ for a functor $T : \mathbf{Set}^n \rightarrow \mathbf{Set}$, where $i_1, \dots, i_k \leq n$, is a natural transformation

$$\lambda : (\mathcal{Q} \circ P_{i_1}^{op}) \times \dots \times (\mathcal{Q} \circ P_{i_k}^{op}) \rightarrow \mathcal{Q} \circ T^{op}$$

between functors $(\mathbf{Set}^n)^{op} \rightarrow \mathbf{Set}$, where \mathcal{Q} denotes the contravariant powerset functor $\mathbf{Set}^{op} \rightarrow \mathbf{Set}$ (i.e. $\mathcal{Q}X = \mathcal{P}X$, and $\mathcal{Q}(f)(A) = f^{-1}[A]$) and $P_i : \mathbf{Set}^n \rightarrow \mathbf{Set}$ is the i -th projection.

We now construct a compositional coalgebraic semantics of the logic $\mathcal{F}(\mathbb{G})$ induced by a gluing \mathbb{G} from structures associated with the features combined by \mathbb{G} . We first describe the notion of structure associated with a single feature, and then the combination of such structures along a gluing.

Definition 4.4. Let $\mathbb{F} = (\Lambda, \mathcal{R})$ be an n -ary feature. A *structure* for \mathbb{F} consists of a functor $\llbracket \mathbb{F} \rrbracket : \mathbf{Set}^n \rightarrow \mathbf{Set}$ and an assignment of a predicate lifting $\llbracket \heartsuit \rrbracket : i_1, \dots, i_k \xrightarrow{\bullet} *$ for T to every modal operator $\heartsuit : i_1, \dots, i_k \rightarrow *$ in Λ , subject to the condition that every rule $R = \phi_1; \dots; \phi_n / \psi$ over V in \mathcal{R} is *one-step sound*: for every n -sorted set X and every assignment τ of subsets $\tau(a) \subseteq X_i$ to the variables $a \in V_i$, if $\llbracket \phi_i \rrbracket_{X, \tau} = X_i$ for all i , then $\llbracket \psi \rrbracket_{TX, \tau} = TX$, where $\llbracket \phi_i \rrbracket_{X, \tau} \subseteq X_i$ and $\llbracket \psi \rrbracket_{TX, \tau} \subseteq TX$ are defined by the usual clauses for Boolean operators and $\llbracket \heartsuit(a_1, \dots, a_k) \rrbracket_{TX, \tau} = \llbracket \heartsuit \rrbracket_X(\tau(a_1), \dots, \tau(a_k))$.

When features are equipped with structures, every feature expression t over the set \mathcal{S} of sorts defines a functor $\llbracket t \rrbracket : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}$ by

$$\llbracket a \rrbracket = P_a : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set} \quad (a \in \mathcal{S}) \quad \text{and} \quad \llbracket \mathbb{F}(t_1, \dots, t_n) \rrbracket = \llbracket \mathbb{F} \rrbracket \circ \langle \llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket \rangle,$$

where P_a is projection to the a -th component and $\langle \cdot \rangle$ represents tupling. Thus, a gluing $\mathbb{G} = (t_a)_{a \in \mathcal{S}}$ induces a functor $\llbracket \mathbb{G} \rrbracket = \langle \llbracket t_a \rrbracket \rangle_{a \in \mathcal{S}} : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}^{\mathcal{S}}$. We briefly refer to $\llbracket \mathbb{G} \rrbracket$ -coalgebras as \mathbb{G} -coalgebras. Strictly speaking, the semantics $\llbracket t \rrbracket : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}$ of a feature expression t depends on the sort context \mathcal{S} ; when necessary, we make this explicit by writing $\llbracket \mathcal{S} \triangleright t \rrbracket$ in place of $\llbracket t \rrbracket$. We have an obvious *substitution lemma* stating that substitution of feature expressions corresponds to composition of functors.

The coalgebraic semantics of $\mathcal{F}(\mathbb{G})$ is now given w.r.t. \mathbb{G} -coalgebras $C = (X, \xi)$. For a type $s \in \text{Types}(\mathbb{G})$, an s -state of C is an element $x \in \llbracket s \rrbracket X$ (recall that $\llbracket s \rrbracket : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}$ denotes a functor). The semantics of a formula $\phi : s$ is a set $\llbracket \phi \rrbracket_C \subseteq \llbracket s \rrbracket X$ of s -states. We have the usual clauses for propositional connectives, and the semantics of modal operators is given by the

following clauses for composite types (top, assuming $F(s_1, \dots, s_n) \in \text{Types}(\mathbb{G})$) and base types (bottom, for $a \rightarrow F(s_1, \dots, s_n) \in \mathbb{G}$):

$$\begin{aligned} \llbracket \heartsuit(\phi_1, \dots, \phi_n) : F(s_1, \dots, s_n) \rrbracket_C &= \llbracket \heartsuit \rrbracket(\llbracket \phi_1 \rrbracket_C, \dots, \llbracket \phi_n \rrbracket_C) \\ \llbracket \heartsuit(\phi_1, \dots, \phi_n) : a \rrbracket_C &= \xi_a^{-1} \circ \llbracket \heartsuit \rrbracket(\llbracket \phi_1 \rrbracket_C, \dots, \llbracket \phi_n \rrbracket_C) \end{aligned}$$

where, in both cases, $\heartsuit : i_1, \dots, i_n \rightarrow *$ in F . We write $x \models_C^s \phi$ if $\phi : s$ and $x \in \llbracket \phi \rrbracket_C$, omitting sub- and superscripts whenever they are clear from the context.

Note that the requirement that rules are one-step sound immediately yields soundness of the logic w.r.t. the semantics described above; this is as in (Cîrstea and Pattinson 2007).

Example 4.5. The standard semantics for the features of Example 3.2 is induced by the following structures.

Non-Determinism: A structure for \mathbb{N}_A is given by $\llbracket \mathbb{N}_A \rrbracket = \mathcal{P}(A \times _)$ and

$$\llbracket \square_a \rrbracket_X(C) = \{B \in \mathcal{P}(A \times X) \mid \{y \in X \mid (a, y) \in B\} \subseteq C\}.$$

Note that (single-sorted) coalgebras for $\mathcal{P}(A \times _)$ are labelled transition systems, as they associate to each state x a set of pairs (a, y) consisting of an action a and a successor state y , which we can read as ‘there is an a -transition from x to y ’. The lifting associated with \square_a gives rise to the usual semantics of Hennessy-Milner logic (Pattinson 2004): a state x satisfies a formula $\square_a \phi$ iff all states y to which x has an a -transition satisfy ϕ .

Uncertainty: Put $\llbracket \mathbb{U} \rrbracket = \mathcal{D}$, where \mathcal{D} is the *discrete distribution functor* \mathcal{D} that maps a set X to the set of discrete probability distributions on X ; i.e. $P \in \mathcal{D}(X)$ is given as $P(A) = \sum_{x \in A} p(x)$ by a map $p : X \rightarrow [0, 1]$ such that $\sum_{x \in X} p(x) = 1$ (which implies that $p(x) > 0$ for at most countably many x). The modal operators L_p are interpreted by

$$\llbracket L_p \rrbracket_X(A) = \{P \in \mathcal{D}X \mid PA \geq p\}.$$

(Single-sorted) \mathcal{D} -coalgebras are probabilistic transition systems. For $\mathbb{G} = (n \rightarrow \mathbb{N}_A(\mathbb{U}(n)))$ (Example 3.6), we have $\llbracket \mathbb{G} \rrbracket = \mathcal{P} \circ \mathcal{D}$, so that \mathbb{G} -coalgebras are precisely Segala systems, while coalgebras for the corresponding flat gluing $(n \rightarrow \mathbb{N}_A(p), p \rightarrow \mathbb{U}(n))$ have an explicit separation between non-deterministic and probabilistic states, and thus correspond to strictly alternating systems in the sense of (Hansson 1994).

Coalition: We opt for a semantics with finite strategy sets, following the standard semantics of alternating-time temporal logic (Alur, Henzinger, and Kupferman 2002). Thus, we put

$$\llbracket \mathbb{G}_N \rrbracket X = \left\{ (d, f) \mid d : N \rightarrow \mathbb{N}, f : \left(\prod_{a \in N} [d(a)] \right) \rightarrow X \right\},$$

where $[d(a)] = \{0, \dots, d(a)\}$ is the set of *strategies* available to agent a , and interpret $[C]$ for a coalition $C \subseteq N$ by

$$\begin{aligned} \llbracket [C] \rrbracket_X A = \left\{ (d, f) \in \llbracket \mathbb{G}_N \rrbracket X \mid \exists \sigma_C \in \prod_{a \in C} [d(a)]. \right. \\ \left. \forall \sigma_{N-C} \in \prod_{a \in N-C} [d(a)]. f(\sigma_C, \sigma_{N-C}) \in A \right\} \end{aligned}$$

where (σ_C, σ_{N-C}) denotes the obvious element of $\prod_{a \in N} [d(a)]$. That is, the informal reading of $[C]\phi$, ‘ C can enforce ϕ ’, translates into the requirement that there exists a collaborative

strategy for C such that whatever collaborative strategy is chosen by the opponents $N - C$, the outcome satisfies ϕ .

Choice: Let $\llbracket S \rrbracket$ be the disjoint sum functor $\llbracket S \rrbracket(X, Y) = X + Y$, and interpret the modality $+$ by

$$\llbracket + \rrbracket_{X,Y}(A, B) = A + B \subseteq X + Y.$$

Thus the condition A applies in the left-hand alternative of the sum, and the condition B applies in the right-hand alternative, in accordance with the informal reading of the $+$ operator given in Example 3.2.

Fusion: Let $\llbracket P \rrbracket$ be the binary product functor $\llbracket P \rrbracket(X, Y) = X \times Y$, and put

$$\llbracket \pi_1 \rrbracket_{X,Y} A = \{(x, y) \mid x \in A\} \text{ and } \llbracket \pi_2 \rrbracket_{X,Y} B = \{(x, y) \mid y \in B\}.$$

Again, this definition is in agreement with the informal description of the operators $[\pi_1]$, $[\pi_2]$ given in Example 3.2.

Conditionality: Define the functor $\llbracket C \rrbracket$ by $\llbracket C \rrbracket(X, Y) = \mathcal{Q}X \rightarrow Y$, with \mathcal{Q} denoting contravariant powerset and \rightarrow denoting function space, and put

$$\llbracket \bullet \rrbracket_{X,Y}(A, B) = \{f : \mathcal{Q}X \rightarrow Y \mid f(A) \in B\}.$$

For $G = (c \rightarrow C(c, K(c)))$ (Example 3.6), we have $\llbracket G \rrbracket X = \mathcal{Q}X \rightarrow \mathcal{P}X$, and G -coalgebras are conditional frames (Chellas 1980), i.e. associate to each state x and each proposition A a set of A -successor states of x .

Input: We put $\llbracket I_I \rrbracket X = I \rightarrow X$, and $\llbracket (i) \rrbracket_X A = \{f : I \rightarrow X \mid f(i) \in A\}$. Thus, $\llbracket I_I \rrbracket$ -coalgebras associate to each state x and each input i a successor state y , and then x satisfies $(i)\phi$ iff y satisfies ϕ .

Output: We put $\llbracket O_O \rrbracket = O$, and $\llbracket o \rrbracket_X = \{o\}$. Thus, for $G = (a \rightarrow O_O)$, a $\llbracket G \rrbracket$ -coalgebra ξ just maps each state x to an output $\xi(x) \in O$, and then x satisfies o iff $\xi(x) = o$.

Modal logic talks only about the observable behaviour of states; this is formally expressed as invariance of the logic under morphisms:

Proposition 4.6. Let $f : C \rightarrow D$ be a morphism of G -coalgebras. Then for each G -formula $\phi : s$ and each s -state x in C , $x \models_C^s \phi$ iff $(\llbracket s \rrbracket f)(x) \models_D^s \phi$.

Proof. Induction over the formula structure, with the case of modal operators taken care of by naturality of predicate liftings. \square

We can now state the (local) *satisfiability problem* for multi-sorted modal logics.

Definition 4.7. A G -formula $\phi : s$ is *G-satisfiable* if there exist a G -coalgebra C and an s -state x in C such that $x \models_C^s \phi$.

The central contribution of this paper is to show that for every gluing, we can construct a flat gluing with an equivalent satisfiability problem. For flat gluings, one can easily generalize existing model constructions and complexity results for coalgebraic modal logic, e.g. (Schröder 2007; Schröder and Pattinson 2008c; 2009), essentially by just writing additional indices for the sorts, and the relevant criteria reduce to the component logics; by way of an example, this is discussed in more detail in Section 5 for the shallow-model-based PSPACE algorithm of (Schröder and

Pattinson 2009). We thus obtain compositional algorithmic methods also for the standard single-sorted semantics present in the literature.

The core of the flattening result is the following construction, which converts a feature expression s that occurs in a gluing into a new sort. Let $G = (t_a)_{a \in \mathcal{S}}$ be a gluing over the set \mathcal{S} in which the feature expression s occurs, and pick a fresh sort \bar{s} (the sort into which the feature expression s is converted). That is, assume that

$$t_a = t'_a[s/\bar{s}]$$

for all $a \in \mathcal{S}$, where $t[s/\bar{s}]$ denotes the result of replacing all occurrences of (the fresh sort) \bar{s} by (the feature expression) s , and the t'_a are feature expressions over $\mathcal{S}' := \mathcal{S} \cup \{\bar{s}\}$. We then construct a gluing $G' = (t'_a)_{a \in \mathcal{S}'}$ over \mathcal{S}' by extending the family $(t'_a)_{a \in \mathcal{S}}$ of feature expressions to a gluing $(t'_a)_{a \in \mathcal{S}'}$ by putting

$$t'_{\bar{s}} = s.$$

To simplify notation, we regard \bar{s} as the ‘first’ sort in \mathcal{S}' , and correspondingly denote, for any \mathcal{S} -indexed tuple $x = (x_a)_{a \in \mathcal{S}}$ (e.g. an \mathcal{S} -sorted set or a \mathcal{S} -sorted map), the \mathcal{S}' -sorted tuple whose a -component is x_a for $a \in \mathcal{S}$ and whose \bar{s} -component is y by (y, x) . Conversely, given an \mathcal{S}' -indexed tuple $x = (x_a)_{a \in \mathcal{S}'}$, we denote the \mathcal{S} -indexed tuple $(x_a)_{a \in \mathcal{S}}$ by $x|_{\mathcal{S}}$.

Example 4.8. Consider the gluing

$$G = (a \rightarrow F(G(b, c), H(a)), b \rightarrow G(b, c), c \rightarrow F(a, c))$$

and let $s = G(b, c)$. This allows us to write

$$G = (a \rightarrow F(\bar{s}, H(a))[s/\bar{s}], b \rightarrow \bar{s}[s/\bar{s}], c \rightarrow F(a, c)[s/\bar{s}])$$

and we obtain a new gluing G'

$$G' = (\bar{s} \rightarrow s, a \rightarrow F(\bar{s}, H(a)), b \rightarrow \bar{s}, c \rightarrow F(a, c))$$

over the set $\mathcal{S}' = \{a, b, c, \bar{s}\}$ where we have given the \bar{s} -component first according to the convention above.

We have a functor Pad from G -coalgebras to G' -coalgebras, which extends a G -coalgebra (X, ξ) to the G' -coalgebra

$$\text{Pad}(X, \xi) = ((\llbracket \mathcal{S} \triangleright s \rrbracket X, X), (id_{\llbracket \mathcal{S} \triangleright s \rrbracket X}, \xi))$$

(where for $a \in \mathcal{S}$, the maps ξ_a have the right type for a G' -coalgebra due to the choice of the \bar{s} -component). Similarly, Pad extends a morphism f of G -coalgebras to the morphism $\text{Pad}(f) = (\llbracket \mathcal{S} \triangleright s \rrbracket f, f)$ of G' -coalgebras. In short, Pad adds to (X, ξ) an identity \bar{s} -component.

Example 4.9. We continue Example 4.8 and assume that each feature is equipped with a structure, denoted by $\llbracket \cdot \rrbracket$. The gluing G induces the functor

$$\llbracket G \rrbracket = (\llbracket F \rrbracket \circ \langle \llbracket G \rrbracket \circ \langle P_b, P_c \rangle, \llbracket H \rrbracket \circ P_a \rangle, \llbracket G \rrbracket \circ \langle P_b, P_c \rangle, \llbracket F \rrbracket \circ \langle P_a, P_c \rangle) : \mathbf{Set}^{\mathcal{S}} \rightarrow \mathbf{Set}^{\mathcal{S}}$$

whereas $\llbracket G' \rrbracket$ is the four-sorted functor

$$\llbracket G' \rrbracket = (\llbracket G \rrbracket \circ \langle P_b, P_c \rangle, \llbracket F \rrbracket \circ \langle P_{\bar{s}}, \llbracket H \rrbracket \circ P_a \rangle, P_{\bar{s}}, \llbracket F \rrbracket \circ \langle P_a, P_c \rangle) : \mathbf{Set}^{\mathcal{S}'} \rightarrow \mathbf{Set}^{\mathcal{S}'}$$

where again the \bar{s} -component appears in the first position. The induced functor Pad then acts on $\llbracket G \rrbracket$ -coalgebras by

$$((X_a, X_b, X_c), (\xi_a, \xi_b, \xi_c)) \mapsto (\llbracket s \rrbracket(X_a, X_b, X_c), X_a, X_b, X_c), (id, \xi_a, \xi_b, \xi_c),$$

i.e. by inserting the identity on $\llbracket s \rrbracket(X_a, X_b, X_c)$ in the first component.

Our next goal is to establish that the semantics of modal logics is invariant under Pad , and we collect some simple properties.

Lemma 4.10. The functor Pad is full, faithful, and injective on objects.

Proof. The second and third claim are clear. To see that Pad is full, let $f : \text{Pad}(X, \xi) \rightarrow \text{Pad}(Y, \zeta)$ be a morphism of G' -coalgebras. Then $f = \text{Pad}(f|_S)$, as by the morphism property in the \bar{s} -component, the diagram

$$\begin{array}{ccc} \llbracket s \rrbracket X & \xlongequal{\quad} & X_{\bar{s}} \xrightarrow{f_{\bar{s}}} Y_{\bar{s}} \xlongequal{\quad} \llbracket s \rrbracket Y \\ & & \downarrow id \qquad \qquad \downarrow id \\ & & \llbracket s \rrbracket X \xrightarrow[\llbracket s' \triangleright s \rrbracket f]{\llbracket S \triangleright s \rrbracket (f|_S)} \llbracket s \rrbracket Y \end{array}$$

commutes. □

In other words, the category of G -coalgebras can be identified with the full subcategory of G' -coalgebras with identity \bar{s} -component. We will show that this subcategory is reflective, which will give rise to a reflection functor Comp (*compose*) from G' -coalgebras to G -coalgebras. For a G' -coalgebra (X, ξ) , $\text{Comp}(X, \xi)$ is defined as the G -coalgebra $(X|_S, \bar{\xi})$ where

$$\bar{\xi}_a = \llbracket t'_a \rrbracket(\xi_{\bar{s}}, id_{X|_S}) \circ \xi_a \text{ for } a \in S,$$

with type information visualized as follows:

$$X_a \xrightarrow{\xi_a} \llbracket t'_a \rrbracket X \xrightarrow{\llbracket t'_a \rrbracket(\xi_{\bar{s}}, id_{X|_S})} \llbracket t'_a \rrbracket(\llbracket s \rrbracket(X|_S), X|_S) = \llbracket t_a \rrbracket(X|_S),$$

where the last equality is by the substitution lemma. We have a canonical map

$$\eta = (\xi_{\bar{s}}, id_{X|_S}) : (X, \xi) \rightarrow \text{Pad}(\text{Comp}(X, \xi)).$$

Lemma 4.11. The map η is a morphism of G' -coalgebras, and as such a reflective arrow; i.e. every morphism of the form $f : \xi \rightarrow \text{Pad}(Y, \zeta)$ factors uniquely as $f = \text{Pad}(f^\#) \circ \eta$ for some $f^\# : \text{Comp}(X, \xi) \rightarrow (Y, \zeta)$:

$$\begin{array}{ccc} (X, \xi) & \xrightarrow{\eta} & \text{Pad}(\text{Comp}(X, \xi)) & & \text{Comp}(X, \xi) \\ & \searrow f & \downarrow \text{Pad}(f^\#) & & \downarrow f^\# \\ & & \text{Pad}(Y, \zeta) & & (Y, \zeta). \end{array}$$

Specifically, $f^\# = f|_S$.

Proof. To begin, we check that η is a coalgebra morphism. In the \bar{s} -component, the morphism property amounts to the commutation of

$$\begin{array}{ccc} X_{\bar{s}} & \xrightarrow{\xi_{\bar{s}}} & \llbracket s \rrbracket(X|_{\mathcal{S}}) \\ \xi_{\bar{s}} \downarrow & & \downarrow id \\ \llbracket s \rrbracket(X|_{\mathcal{S}}) & \xrightarrow{\llbracket s \rrbracket id_{X|_{\mathcal{S}}}} & \llbracket s \rrbracket(X|_{\mathcal{S}}). \end{array}$$

Moreover, in the a -component for $a \in \mathcal{S}$, the morphism property is the evident commutation of

$$\begin{array}{ccc} X_a & \xrightarrow{id} & X_a \\ \xi_a \downarrow & & \downarrow \llbracket t'_a \rrbracket(\xi_{\bar{s}, id_{X|_{\mathcal{S}}}) \circ \xi_a \\ \llbracket t'_a \rrbracket(X) & \xrightarrow{\llbracket t'_a \rrbracket(\xi_{\bar{s}, id_{X|_{\mathcal{S}}})} & \llbracket t'_a \rrbracket(\llbracket s \rrbracket(X|_{\mathcal{S}}), X|_{\mathcal{S}}). \end{array}$$

Secondly, we have to verify the universal property. Let $f : (X, \xi) \rightarrow \text{Pad}(Y, \zeta)$ for a G-coalgebra (Y, ζ) . Then the only candidate for a morphism $f^\# : \text{Comp}(X, \xi) \rightarrow (Y, \zeta)$ such that $\text{Pad}(f^\#) \circ \eta = f$ is $f^\# = f|_{\mathcal{S}}$ as in the claim. To see that indeed $\text{Pad}(f^\#) \circ \eta = f$, it remains only to check commutation in the \bar{s} -component, i.e. $\llbracket s \rrbracket(f|_{\mathcal{S}}) \circ \xi_{\bar{s}} = f_{\bar{s}}$. This follows from the morphism property of f , which in the \bar{s} -component amounts to commutation of

$$\begin{array}{ccc} X_{\bar{s}} & \xrightarrow{f_{\bar{s}}} & \llbracket s \rrbracket Y \\ \xi_{\bar{s}} \downarrow & & \downarrow id \\ \llbracket s \rrbracket(X|_{\mathcal{S}}) & \xrightarrow{\llbracket s \rrbracket(f|_{\mathcal{S}})} & \llbracket s \rrbracket Y. \end{array}$$

Moreover, we have to verify that $f^\#$ is a morphism $\text{Comp}(X, \xi) \rightarrow (Y, \zeta)$. For $a \in \mathcal{S}$, we have to prove commutation of

$$\begin{array}{ccc} X_a & \xrightarrow{f_a} & Y_a \\ \llbracket t'_a \rrbracket(\xi_{\bar{s}, id_{X|_{\mathcal{S}}}) \circ \xi_a \downarrow & & \downarrow \zeta_a \\ \llbracket t'_a \rrbracket(\llbracket s \rrbracket(X|_{\mathcal{S}}), X|_{\mathcal{S}}) & \xrightarrow[\llbracket t'_a \rrbracket(f|_{\mathcal{S}})]{\llbracket t'_a \rrbracket(\llbracket s \rrbracket(f|_{\mathcal{S}}), f|_{\mathcal{S}})} & \llbracket t'_a \rrbracket(\llbracket s \rrbracket Y, Y) \end{array}$$

where the equality on the bottom arrow is by the substitution lemma. Now by the previous diagram, the lower-left path in this diagram equals $\llbracket t'_a \rrbracket(f) \circ \xi_a$, so that commutation follows from the morphism property of $f : (X, \xi) \rightarrow \text{Pad}(Y, \zeta)$ in the a -component. \square

The previous lemma implies that *Comp extends to a functor from G'-coalgebras to G-coalgebras, which is a left adjoint of Pad* (see e.g. (Adámek, Herrlich, and Strecker 1990)). Since *Pad* is full and faithful, the arising co-unit $\epsilon : \text{Comp} \circ \text{Pad} \rightarrow id$ is necessarily an isomorphism; in fact, we have $\text{Comp} \circ \text{Pad} = id$, and ϵ is identity.

Morally, the above means in particular that G-coalgebras are equivalent to G'-coalgebras w.r.t. the behaviour of states: G-coalgebras can, via *Pad*, be regarded as G'-coalgebras, and every G'-coalgebra has a behaviour-preserving map into a G-coalgebra. This will mean in particular that

G' -coalgebras and G -coalgebras are equivalent w.r.t. satisfiability of modal formulae; we proceed to make this more precise.

We note that the types of G' are essentially the same as the types of G , except that we have introduced a new base type \bar{s} as a synonym for s , and removed some or all occurrences of the composite type s . The following lemma, proved by straightforward induction over typing derivations, states that G and G' induce essentially the same modal syntax:

Lemma 4.12. The formulas of type \bar{s} in G' are precisely the formulas of type s in G ; moreover, for $r \in \text{Types}(G') - \{\bar{s}\}$, the formulas of type r in G and G' coincide.

In short, the formulas of G and G' are the same up to interchange of the types s and \bar{s} , and we shall silently identify the two sets of formulas altogether. The same holds, *mutatis mutandis*, for proofs in G and G' . By the adjoint situation between Comp and Pad described above, together with invariance of modal formulas under coalgebra morphisms (Proposition 4.6), it is immediate that the transition from G to G' leaves modal semantics invariant:

Corollary 4.13. A formula is G' -satisfiable iff it is G -satisfiable. □

For the time being, we now assume that \mathcal{S} is finite. This serves to simplify some of the arguments, and is not a restriction for purposes of the main application pursued in the present work, since a modal formula to be checked for satisfiability will always involve only finitely many sorts. All our results below hold also for infinite sets of sorts, as discussed in Remark 4.21.

Under the assumption of finiteness of \mathcal{S} , it is clear that we can turn *all* composite types of a gluing into sorts, by just iterating the above-described process of promoting a type of a gluing to a new sort. We thus obtain, given a gluing G , a flat gluing G^b that we call the *flattening* of G . Explicitly, G^b is described as follows. The set of sorts is $\mathcal{S}^b = \{\bar{s} \mid s \in \text{Types}(G)\}$. When convenient, we identify a with \bar{a} , and hence regard \mathcal{S} as a subset of \mathcal{S}^b . For a feature expression over \mathcal{S} of the form $t = F(s_1, \dots, s_n)$, we write $t^b = F(\bar{s}_1, \dots, \bar{s}_n)$, and for $a \in \mathcal{S}$, we write $a^b = \bar{a}$. Then

$$G^b = (u_{\bar{s}})_{\bar{s} \in \mathcal{S}^b}, \text{ where } u_{\bar{s}} = \begin{cases} (t_a)^b & \text{for } s = a \in \mathcal{S} \\ s^b & \text{otherwise.} \end{cases}$$

Example 4.14. Given the gluings $G_1 \equiv (o \rightarrow S(U(o), N_A(o)))$ and $G_2 \equiv (o \rightarrow S(u, n), u \rightarrow U(o), n \rightarrow N_A(o))$ from Example 3.6, G_2 is the flattening of G_1 , up to renaming the sorts \bar{o} , $\bar{U}(o)$, and $\bar{N}_A(o)$ of the flattening into o , u , and n , respectively.

By iterated application of Lemma 4.12 (and the following comment), *the flattening G^b syntactically induces the same logic as G* , i.e. the types, formulas, and proof systems coincide (up to renaming s into \bar{s} for $s \in \text{Types}(G)$). Moreover, iterated application of Corollary 4.13 yields our main result:

Theorem 4.15. A formula is G -satisfiable iff it is G^b -satisfiable. □

We observe that the order in which composite types are promoted to sorts in the formation of the flattening of a gluing G does not matter – not only is it clear that the resulting flat gluing G^b is independent of this order, but also the arising embedding functor from G -coalgebras into G^b -coalgebras is evidently insensitive to the order of steps: the embedding is composed of the rele-

vant Pad functors, and hence in any case just adds identity components for all new sorts. Uniqueness of left adjoints then implies that the reflection functor from G^b -coalgebras to G -coalgebras that arises from the relevant Comp functors is independent of the order of steps. We therefore reuse Pad and Comp to denote also the embedding and reflection functors, respectively, between G -coalgebras and G^b -coalgebras. By picking a definite order in which types are promoted to sorts, say bottom up and leftmost first, we can turn the iterated construction of the flattening into a recursive definition. This yields exactly the definition of Comp given in (Schröder and Pattinson 2007); specifically, the object part of Comp maps a given $\llbracket G^b \rrbracket$ -coalgebra (X, ξ) to the $\llbracket G \rrbracket$ -coalgebra $(X|_{\mathcal{S}}, \bar{\xi})$, where

$$\bar{\xi}_a = \llbracket F \rrbracket(\zeta_{s_1}, \dots, \zeta_{s_n}) \circ \xi_a \quad (1)$$

for $a \rightarrow F(s_1, \dots, s_n)$ in G , and the maps $\zeta_s : X_s \rightarrow \llbracket s \rrbracket(X_a)_{a \in \mathcal{S}}$ for $s \in \text{Types}(G)$ are defined recursively by

$$\zeta_a = id_{X_a} \quad (a \in \mathcal{S}) \quad \text{and} \quad \zeta_{F(s_1, \dots, s_n)} = \llbracket F \rrbracket(\zeta_{s_1}, \dots, \zeta_{s_n}) \circ \xi_{F(s_1, \dots, s_n)} \quad (2)$$

(with some simplification using functoriality of $\llbracket F \rrbracket$ already applied). The unit $\eta : (X, \xi) \rightarrow \text{Pad}(\text{Comp}(X, \xi))$ is then given by $\eta_{\bar{s}} = \zeta_s$ for $s \in \text{Types}(G)$.

We record explicitly:

Theorem 4.16. Via Pad, the G -coalgebras form a full reflective subcategory of the G^b -coalgebras, with Comp as reflector.

In our running example, the situation is as follows:

Example 4.17. Consider the gluings G_1 and G_2 over $\mathcal{S} = \{o, u, n\}$ from Example 3.6 and recall from Example 4.14 that $G_1^b = G_2$. Let $C = (Y, \zeta : X \rightarrow \mathcal{D}X + \mathcal{P}(A \times X))$ be a G_1 -coalgebra. Then C can be converted into a G_2 -coalgebra $\text{Pad}(C) = (X, \xi)$ where $X_o = Y, X_u = \mathcal{D}Y, X_n = \mathcal{P}(A \times Y), \xi_o = \zeta, \xi_u = id_{X_u},$ and $\xi_n = id_{X_n}$; i.e. Pad just adds identities at the new sorts u, n .

Conversely, given a G_2 -coalgebra $D = (X, \xi)$, we construct a G_1 -coalgebra $\text{Comp}(D) = (Y, \zeta)$ (which has a single sort, o) by putting $Y_o = X_o$ and $\zeta_o = (\xi_u + \xi_n) \circ \xi_o$. The triple (id_{X_o}, ξ_u, ξ_n) is a coalgebra morphism $D \rightarrow \text{Pad}(\text{Comp}(D))$, the unit of the adjunction between Pad and Comp.

Using general properties of reflective subcategories, we obtain the following from Theorem 4.16.

Corollary 4.18. If C is a final G -coalgebra, then $\text{Pad}(C)$ is a final G^b -coalgebra. Conversely, if D is a final G^b -algebra, then D is isomorphic to $\text{Pad}(\text{Comp}(D))$, and $\text{Comp}(D)$ is a final G -coalgebra.

Remark 4.19. In the case of polynomial functors, the dual of the previous corollary is semi-folklore: nested recursive datatypes can be flattened into mutually recursive datatypes, as done in the theorem prover Isabelle/HOL (see, e.g., (Berghofer and Wenzel 1999) and references therein). We do not know of a formal treatment of this fact at the level of generality we employ here, nor of a statement of the dual of Theorem 4.16 in the literature.

Remark 4.20. If a final G^b -coalgebra exists, one can conclude Theorem 4.15 from Corollary 4.18 alone; however, in many cases of interest final coalgebras do not exist (e.g. already in the basic case of unbounded non-determinism).

Remark 4.21. Our temporary assumption that \mathcal{S} is finite is inessential for our results presented above, specifically the construction of the flattening, Theorems 4.15 and 4.16, and Corollary 4.18. The quickest way to see this in the current approach is to use the recursive definition of flattening (Equations (1) and (2)) directly and then verify the requisite properties by induction, with the inductive step being a minor generalization of the proof of Lemma 4.11.

Remark 4.22. While we restrict the current development to purely modal logics, mostly in the interest of readability, we point out that Theorem 4.16 implies also a general invariance result for coalgebraic *hybrid* logics, i.e. logics that include additional expressive means for individual states (Myers, Pattinson, and Schröder 2009; Schröder, Pattinson, and Kupke 2009; Goré et al. 2010), as long as these are restricted to states that live in the base types of the original gluing G : while such hybrid logics are no longer invariant under arbitrary coalgebra morphisms, they are certainly invariant under coalgebra homomorphisms that are carried by bijective maps at the base types, as is the case for the unit and the counit of the adjunction of Theorem 4.16.

Remark 4.23. Our definition of flat gluings admits *unguarded* sorts, i.e. in a flat gluing $G = (t_a)_{a \in \mathcal{S}}$ we may have $t_b = c$ for some $b, c \in \mathcal{S}$. One can turn such gluings into *guarded* flat gluings (where all sorts are associated to composite feature expressions), essentially getting rid of the sort b , as follows. We have to distinguish two subcases: if b is distinct from c , we can form a new gluing $G^{-b} = (t_a[c/b])_{a \in \mathcal{S} - \{b\}}$; the equivalence of G and G^{-b} is then just a special case of Corollary 4.13, applied to $s = c$ and with \bar{c} renamed to b . The slightly more complicated case is where $b = c$. In this case, we can form the new gluing $G^{-b} = (t_a[1/b])_{a \in \mathcal{S} - \{b\}}$, where 1 is the (trivial) output feature for the singleton output set 1. One easily checks equivalence of G^{-b} with G . The gluing G^{-b} is not in general flat, but can of course be flattened into a guarded flat gluing, which in the end means that we replace all sorts b such that $b \rightarrow b \in G$ with a single sort $\bar{1}$ that we associate with the feature expression 1 (which by the definition given in Section 3 is composite). The intuition behind these constructions is that if $b \rightarrow c \in G$ for $b \neq c$, then the observations on b are precisely the ones that arise by transiting to c and then making observations on c , so that b may be identified with c ; if, on the other hand, $b \rightarrow b \in G$, then no observations at all can be made on b , so that b may be identified with 1.

Remark 4.24. It is evident that the preceding development is entirely independent of the choice of **Set** as the base category, except that the definition of predicates and predicate liftings as presented here is somewhat set-oriented; a generalization to arbitrary base categories would, e.g., work with fibrations of predicates. Indeed one does not even need to assume that all sorts live in the same base category. The core message of our flattening construction for a gluing G is that G -coalgebras can be regarded as particular G^b -coalgebras, and every behaviour that is realized in a G -coalgebra is, via the reflective arrow into the corresponding G^b -coalgebra, realized also in a G^b -coalgebra; in summary, *G -coalgebras and G^b -coalgebras realize the same behaviours*. From the perspective of coalgebraic modal logics, this will mean that over any base category

(or base categories), modal semantics will be invariant under the transition from gluings to their flattenings as long as the logic is *adequate*, i.e. invariant under coalgebra morphisms.

As indicated above, the main benefit of the flattening result is that functor composition vanishes from the picture, and is replaced by multi-sortedness. This means that results that have earlier been phrased as compositionality results, i.e. results that reduce properties of composite functors to properties of their components, now just become generalizations of single-sorted results to the multi-sorted setting, which typically require no more than the annotation of previous results and proofs with sort indices. Roughly speaking, any result on coalgebraic modal logic (and more generally any property of coalgebras that is invariant under behavioural equivalence) will generalize to multi-sorted coalgebras provided that it uses only such properties of the base category that are stable under products of categories.

We conclude this section with a further brief comment on the relationship between Segala systems and strictly-alternating systems.

Example 4.25. The flattening of the gluing $n \rightarrow N_A(U(n))$ is the gluing $(n \rightarrow N_A(p), p \rightarrow U(n))$ (Example 3.6); as discussed in Example 4.5, the corresponding coalgebras are Segala systems and strictly alternating systems, respectively. The relationship between Segala systems and strictly alternating systems has received some recent attention (Segala and Turrini 2005; Segala 2006). Our results show that both types of systems admit an interpretation of the same variant of probabilistic modal logic (Section 2) and validate the same formulas of this logic. Moreover, the notion of behavioural equivalence on the sort of non-deterministic states (which unlike the sort of probabilistic states is explicitly present in both types of systems) is the same in both cases, and since the distribution functor preserves weak pullbacks (Moss 1999; de Vink and Rutten 1999; Sokolova, Vink, and Woracek 2009), this equivalence extends to coalgebraic bisimulation, which captures standard notions of strong bisimulation. In particular, the notions of strong bisimulation coincide for Segala systems and strictly alternating systems, a result not apparently claimed in (Segala and Turrini 2005).

5. Applications to Model Construction and Complexity

We have seen in Section 3 that the same multi-sorted logic can arise from different gluings of given features, where the difference manifests itself only on a semantic level. The different interpretations of the logic are related by Theorem 4.15 which shows that the satisfiability problem for a given gluing is equivalent to that of its *flattening*. We now show that the generic shallow model construction and the ensuing PSPACE decision procedure from (Schröder and Pattinson 2009) generalize to flat gluings; this enables us to derive upper PSPACE bounds for arbitrary gluings, in particular for heterogeneous logics equipped with their standard single-sorted semantics as in Section 2. The entire point of this exercise is to show that the generalization from the single-sorted case to the multi-sorted case amounts to no more than additional bookkeeping for the sorts; the minor differences that do exist between the proof presented below and the one given in (Schröder and Pattinson 2009) are entirely due to presentational choices made in the context of *op. cit.*

The shallow model construction requires the involved structures to be strictly one-step complete in the following sense, where the notation $\llbracket - \rrbracket_{X,\tau}$ and $\llbracket - \rrbracket_{TX,\tau}$ is as in Definition 4.4.

(Strict) one-step completeness implies weak completeness of the rule system (Pattinson 2004; Schröder 2007).

Definition 5.1. An n -ary feature F is *strictly one-step complete* for a structure $T = \llbracket F \rrbracket : \mathbf{Set}^n \rightarrow \mathbf{Set}$ if, whenever $\llbracket \chi \rrbracket_{TX, \tau} = T(X_1, \dots, X_n)$ for a sorted set $V = (V_1, \dots, V_n)$ of variables, an assignment τ of subsets $\tau(a) \subseteq X_i$ to variables $a \in V_i$, and a clause χ over atoms of the form $\heartsuit(a_{i_1}, \dots, a_{i_k})$, where $\heartsuit : i_1, \dots, i_k \rightarrow *$ in F and $a_{i_j} \in V_{i_j}$, then χ is provable from propositional formulas valid under τ using a single instance of a rule of F ; formally: χ is propositionally entailed by a clause $\psi\sigma$ (i.e. either χ is a propositional tautology, or χ contains all literals of $\psi\sigma$), where $(\phi_i)/\psi$ is a rule of F and σ is a (V_1, \dots, V_n) -substitution (i.e. $\sigma(a) \in V_i$ for $a \in V_i$) such that $\llbracket \phi_i \sigma \rrbracket_{X, \tau} = X_i$ for all i .

Remark 5.2. In the above definition, we can assume that all V_i are finite and that χ is not a propositional tautology.

One shows analogously to the single-sorted case (Schröder 2007) that the set of all one-step sound rules for a given F -structure is strictly one-step complete, so that strictly one-step complete axiomatizations always exist. Formally:

Proposition 5.3. Suppose $F = (\Lambda, \mathcal{R})$ is an n -ary feature with an associated structure over $\llbracket F \rrbracket : \mathbf{Set}^n \rightarrow \mathbf{Set}$. If \mathcal{R} consists of all one-step rules that are one-step sound for $\llbracket F \rrbracket$, then F is strictly one-step complete for $\llbracket F \rrbracket$.

Proof. Let $T = \llbracket F \rrbracket$ be the underlying functor of the given structure for F . Assume $\llbracket \chi \rrbracket_{X, \tau} = X$ with X, τ, χ, V as in Definition 4.4 and V assumed to be finite in each sort. For $1 \leq i \leq n$ let ϕ_i be the propositional theory of τ over i , i.e. the finite conjunction of all clauses ρ over V_i such that $\llbracket \rho \rrbracket_{X, \tau} = X_i$. We will now show that the one-step rule $R \equiv (\phi_1; \dots; \phi_n)/\chi$ over V is one-step sound; it then follows that χ is derivable as required.

Thus, let $Y \in \mathbf{Set}^n$ and let σ be a $\mathcal{P}Y$ -valuation such that $\llbracket \phi_i \rrbracket_{\sigma} = Y_i$ for all $i = 1, \dots, n$. We have to show $\llbracket \chi \rrbracket_{\sigma} = TY$. We claim that for each i and each $y \in Y_i$, there exists $x \in X_i$ such that

$$x \in \tau(v) \iff y \in \sigma(v)$$

for all $v \in V_i$. Assume, for a contradiction, that $i \in \{1, \dots, n\}$ violates this property. Take ρ to be the propositional theory of y under σ , i.e. the (finite) conjunctive clause

$$\rho = \bigwedge_{\{v \in V_i \mid y \in \sigma(v)\}} v \wedge \bigwedge_{\{v \in V_i \mid y \notin \sigma(v)\}} \neg v.$$

Then $\llbracket \neg \rho \rrbracket_{X, \tau} = X_i$ by assumption, and $y \in \llbracket \rho \rrbracket_{\sigma}$ by construction, in contradiction to $\llbracket \phi_i \rrbracket_{\sigma} = Y_i$.

Thus, we have $f_i : Y_i \rightarrow X_i$ such that

$$\sigma(v) = f_i^{-1}[\tau(v)] \quad \text{for all } v \text{ in } V_i.$$

By naturality of predicate liftings, and since preimages commute with intersections and complements, we now obtain

$$\llbracket \chi \rrbracket_{\sigma} = (Tf)^{-1}[\llbracket \chi \rrbracket_{\tau}],$$

where $f = (f_i)_{1 \leq i \leq n}$. Since $\llbracket \chi \rrbracket_{TX, \tau} = TX$, we conclude $\llbracket \chi \rrbracket_{\sigma} = TY$ as required. \square

In (Schröder and Pattinson 2009), *rule resolution*, a systematic procedure for obtaining strictly one-step complete rule sets, has been described, which straightforwardly generalizes to the multi-sorted setting.

Throughout this section, we fix a gluing G of a set Φ of features over a set \mathcal{S} of sorts; moreover we assume that every feature is equipped with a structure.

Definition 5.4. The set $MA(\phi)$ of *modal atoms* of an G -formula ϕ is defined recursively by $MA(\phi \wedge \psi) = MA(\phi) \cup MA(\psi)$, $MA(\neg\phi) = MA(\phi)$, and $MA(\heartsuit(\rho_1, \dots, \rho_n)) = \{\heartsuit(\rho_1, \dots, \rho_n)\}$. A *pseudovaluation* for ϕ is a subset H of $MA(\phi)$. We define satisfaction of propositional formulas χ over $MA(\phi)$ by H ($H \models \chi$) inductively in the obvious way, with $H \models \chi \iff \chi \in H$ for $\chi \in MA(\phi)$.

Assuming that $s = F(s_1, \dots, s_n) \in \text{Types}(G)$ if s is composite and $a \rightarrow F(s_1, \dots, s_n) \in G$ if $s = a$ is a base type, we say that a rule $R = (\phi_1; \dots; \phi_n)/\psi$ associated with the feature F *matches* a pseudovaluation H for $\phi : s$ if there is a substitution σ such that $\psi\sigma$ is a clause over $MA(\phi)$ with $H \not\models \psi\sigma$. In this case, the pair (R, σ) is called a *matching* of H .

The significance of rule matching arises in relation to satisfiability checking; essentially, a matching rule for H is one that threatens to prove *unsatisfiability* of H . In more detail, given a pseudovaluation H and a rule $R = (\phi_1; \dots; \phi_n)/\psi$ that matches H , via a substitution σ , H can only be satisfiable if $\neg\psi\sigma$ is satisfiable, which by soundness of the rule can only be the case if one the premises of R , instantiated by σ , is not valid, i.e. if one of the $\neg\phi_i\sigma$ is satisfiable. This observation is the basis of both the shallow model construction and of the ensuing decision procedure; completeness of this principle is guaranteed by strict one-step completeness.

Explicitly, our shallow model theorem now takes the following form.

Theorem 5.5. If every feature in G is strictly one-step complete, then a formula $\phi : s$ is satisfiable in a G -model iff $H \models \phi$ for some pseudovaluation H for ϕ such that for every matching $((\phi_1; \dots; \phi_n)/\psi, \sigma)$ of H , one of the formulas $\neg\phi_i\sigma$ is satisfiable.

Proof. By Theorem 4.15 and Remark 4.23, we can assume that $G = (t_a)_{a \in \mathcal{S}}$ is a guarded flat gluing, i.e. $\mathcal{S} = \text{Types}(G)$, and in particular $s = b$ for some sort b and $b \rightarrow F(a_1, \dots, a_n) \in G$ for some n -ary feature F and sorts a_1, \dots, a_n .

‘Only if’: If $x \models_C^b \phi$, let $H = \{\chi \in MA(\phi) \mid x \models_C^b \chi\}$; then use soundness of the rules.

‘If’: Let Π be a set of formulas consisting of one satisfiable formula $\neg\phi_i\sigma : a_i$ for each matching $((\phi_1; \dots; \phi_n)/\psi, \sigma)$ of H , as guaranteed to exist by assumption. For each $\pi \in \Pi$, there exists a G -coalgebra $C_\pi = (X^\pi, \xi^\pi) = ((X_a^\pi), (\xi_a^\pi)_{a \in \mathcal{S}})$ and $x^\pi \in X_{a_i}^\pi$ such that $x^\pi \models_{C_\pi}^{a_i} \pi$; we can assume that for all a , the X_a^π are pairwise disjoint. Put

$$X_a = \bigcup_{\pi \in \Pi} X_a^\pi \text{ for } a \neq b, \text{ and } X_b = \{x_0\} \cup \bigcup_{\pi \in \Pi} X_b^\pi,$$

where x_0 is a fresh element. Let Y be the \mathcal{S} -sorted set of the x^π , $\pi \in \Pi$. For $\chi : a$, put

$$\hat{\chi} = \llbracket \chi \rrbracket_{C_\pi} \cap Y_a.$$

We define a G -coalgebra $C = (X, \xi) = ((X_a), (\xi_a))$ as follows. For $x \in X_a^\pi$, we put $\xi_a(x) = \xi_a^\pi(x) \in \llbracket t_a \rrbracket X^\pi \subseteq \llbracket t_a \rrbracket X$. Then (independently of the pending definition of $\xi_b(x_0)$) for $\chi : a$

and $x \in X_a^\pi$,

$$x \models_C^a \chi \iff x \models_{C^\pi}^a \chi \quad (3)$$

as C^π is a subcoalgebra of C . In particular, $\hat{\chi} = \llbracket \chi \rrbracket_C \cap Y_a$ for all $\chi : a$.

Crucially, we show that there exists $r \in \llbracket \mathbb{F} \rrbracket Y$ such that for $\heartsuit(\rho_1, \dots, \rho_k)$ in $MA(\phi)$,

$$r \in \llbracket \heartsuit \rrbracket(\hat{\rho}_1, \dots, \hat{\rho}_k) \iff \heartsuit(\rho_1, \dots, \rho_k) \in H. \quad (4)$$

Putting $\xi_b(x_0) = r$, we then obtain

$$x_0 \models_C \rho \iff H \models \rho \quad (5)$$

for all propositional formulas ρ over $MA(\phi)$: this follows immediately from (3) and (4) for $\rho \in MA(\phi)$, and then extends easily to $\text{Prop}(MA(\phi))$. By (5), in particular $x_0 \models_C \phi$, and we are done.

It remains to prove that r satisfying (4) exists. Assume the contrary. Let V be the set of propositional variables b_ρ , indexed over all formulas ρ that occur as arguments of modal atoms in ϕ , i.e. $\rho = \rho_i$ for some $\heartsuit(\rho_1, \dots, \rho_k) \in MA(\phi)$ and some $i = 1, \dots, k$. Let the clause θ over $MA(\phi)$ consist of the literals $\neg \heartsuit(b_{\rho_1}, \dots, b_{\rho_k})$ for $\heartsuit(\rho_1, \dots, \rho_k) \in H$ and $\heartsuit(b_{\rho_1}, \dots, b_{\rho_k})$ for $\heartsuit(\rho_1, \dots, \rho_k) \in MA(\phi) - H$; i.e. letting σ denote the substitution given by

$$\sigma(b_\rho) = \rho \text{ for all } b_\rho \in V,$$

we have

$$H \models \neg \theta \sigma. \quad (6)$$

Now by non-existence of r satisfying (4),

$$\llbracket \theta \rrbracket_{\llbracket \mathbb{F} \rrbracket Y, \tau_Y} = \llbracket \mathbb{F} \rrbracket Y,$$

where τ_Y is the $\mathcal{P}Y$ -valuation given by

$$\tau_Y(b_\rho) = \hat{\rho}.$$

By strict one-step completeness of \mathbb{F} , it follows that θ is derivable using a single rule of \mathbb{F} from propositional formulas over V that are valid under τ_Y ; explicitly: there exists a rule $(\phi_1; \dots; \phi_n)/\psi$ in \mathbb{F} and a substitution η such that

$$\llbracket \phi_i \eta \rrbracket_{Y, \tau_Y} = Y_{a_i} \text{ for } i = 1, \dots, n \quad (7)$$

and $\psi \eta$ propositionally entails θ , i.e., assuming w.l.o.g. that θ is not a propositional tautology, θ contains $\psi \eta$. In particular, $\psi \eta$ is a clause over $MA(\phi)$.

We now prove that

$$H \models \psi \eta \sigma, \quad (8)$$

where σ is the substitution taking b_ρ to ρ . This will finish the proof as it implies $H \models \theta \sigma$, in contradiction to (6).

We proceed to prove (8) by contradiction. So assume that $H \not\models \psi \eta \sigma$. By construction of Π , it follows that $\pi := \neg \phi_i \eta \sigma \in \Pi$ for some i ; recall that $\pi : a_i$ and $x^\pi \models_{C^\pi}^{a_i} \pi$. Let τ^π denote the $\mathcal{P}(X^\pi)$ -valuation taking b_ρ to $\llbracket \rho \rrbracket_{C^\pi}$. By naturality of predicate liftings, (7) implies $x^\pi \in \llbracket \phi_i \eta \rrbracket_{X^\pi, \tau^\pi}$, because $\tau_Y(b_\rho) = \tau^\pi(b_\rho) \cap Y_a$ for $\rho : a$. Since $\llbracket \phi_i \eta \rrbracket_{X^\pi, \tau^\pi} = \llbracket \phi_i \eta \sigma \rrbracket_{C^\pi}$, we have arrived at a contradiction to $x^\pi \models_{C^\pi}^{a_i} \pi \equiv \neg \phi_i \eta \sigma$. \square

From Theorem 5.5, we obtain a multi-sorted version of the generic PSPACE decision procedure of (Schröder and Pattinson 2009). This requires to compute matchings of given pseudovaluations, and we require that the rules associated with features are *closed under contraction*, i.e. it suffices to consider matchings $((\phi_1, \dots, \phi_n)/\psi, \sigma)$ where $\psi\sigma$ does not contain duplicate literals, with the consequence that there are only finitely many matches to check in every recursion step. Formally, these notions are defined as follows.

Definition 5.6. We call a clause *contracted* if all its literals are distinct. An instance of a rule is *contracted* if its conclusion is a contracted clause, and a matching is contracted if the corresponding rule instance is contracted. Finally, a feature $F = (\Lambda, \mathcal{R})$ is *closed under contraction* if every application of a rule instance with duplicate literals in the conclusion can be replaced by a contracted instance of a (possibly different) rule. Formally, for every V -instance $R\sigma$ of a rule $R \equiv (\phi_1; \dots; \phi_n)/\psi$ over V in \mathcal{R} , there exists a contracted V -instance $R'\sigma'$ of a rule $R' \equiv (\phi'_1; \dots; \phi'_n)/\psi' \in \mathcal{R}$ such that $\psi\sigma$ and $\psi'\sigma'$ are propositionally equivalent and for all i , $\phi'_i\sigma'$ is propositionally entailed by $\phi_i\sigma$.

To close a feature $F = (\Lambda, \mathcal{R})$ under contraction, just add a rule $(\phi'_1; \dots; \phi'_n)/\psi'$ for every rule $(\phi_1; \dots; \phi_n)/\psi$ over V in \mathcal{R} and every V -substitution σ , where ϕ'_i is some suitably chosen propositional equivalent of $\phi_i\sigma$ and ψ' is obtained from $\psi\sigma$ by removing duplicate literals. It is clear that the arising set of rules inherits one-step soundness and strict one-step completeness from \mathcal{R} .

From Theorem 5.5, one obtains immediately

Corollary 5.7. If every feature in G is strictly one-step complete and closed under contraction, then a formula $\phi : s$ is satisfiable over a $[[G]]$ -structure iff $H \models \phi$ for some pseudovaluation H for ϕ such that for every contracted matching $((\phi_1; \dots; \phi_n)/\psi, \sigma)$ of H , one of the formulas $\neg\phi_i\sigma$ is satisfiable.

Since rules are generally too large to pass around directly, we assume that every rule is represented by a *code*. For the features discussed in Example 3.2, the codes can be taken as the parameters of the rules. Formal definitions are as follows.

Definition 5.8. Let $F = (\Lambda, \mathcal{R})$ be a feature. We assume that the rules in \mathcal{R} are represented by *codes*, i.e. strings over a given finite alphabet. Two rules are *equivalent* if their conclusions are identical and their premises are propositionally equivalent.

Under the assumptions of the above corollary, we have the following algorithm on an alternating Turing machine (Chandra and Stockmeyer 1981), generalising an algorithm from (Vardi 1989). Essentially, the algorithm just applies the satisfiability criterion of Corollary 5.7 recursively, decreasing the modal depth of the formula in every recursion step. To achieve polynomial space usage, the algorithm proceeds by depth-first search; technically, this is encapsulated in the use of alternation.

Algorithm 5.9. (Decide satisfiability of a G -formula $\phi : s$, where $s = F(s_1, \dots, s_n) \in \text{Types}(G)$ if s is composite and $a \rightarrow F(s_1, \dots, s_n) \in G$ if $s = a$ is a base type.)

1. (Initialize) Construct the set $MA(\phi)$.
2. (Existential) Guess a pseudovaluation H for ϕ with $H \models \phi$.

3. (Universal) Nondeterministically choose a contracted matching $((\phi_1; \dots; \phi_n)/\psi, \sigma)$ of H .
4. (Existential) Guess $i \in \{1, \dots, n\}$ and a clause χ from the conjunctive normal form (CNF) of ϕ_i and recursively check that $\neg\chi\sigma : s_i$ is satisfiable.

The algorithm succeeds if all possible choices at steps marked *universal* lead to successful termination, and for all steps marked *existential*, there exists a choice leading to successful termination. Note that the algorithm terminates successfully in Step 3 if there are no matching rules. In particular, the algorithm terminates either in Step 2 or in Step 3 if ϕ has rank 0. Step 3 breaks down into choosing a contracted clause ρ over $MA(\phi)$ such that $H \models \neg\rho$, and then choosing a code of a rule that matches ρ . Note that it suffices to guess one rule from each equivalence class of rules (Definition 5.8).

The crucial requirement for the effectivity of Algorithm 5.9 is that Steps 3 and 4 can be performed in polynomial time; the central condition here is that one may restrict oneself to guess only rule codes of polynomially bounded size in Step 3. Formally, we have the following set of conditions.

Definition 5.10. A feature $F = (\Lambda, \mathcal{R})$ is *PSPACE-tractable* if all contracted matchings have a code of polynomially bounded size, and it can be decided in *NP*

- whether a given code is the code of some rule in \mathcal{R} ;
- whether a rule matches a given pseudovaluation; and
- whether a clause belongs to the CNF of a given premise of a given rule.

Under these conditions, Algorithm 5.9 runs in polynomial time, which proves the following Theorem 5.11, as $\text{APTIME} = \text{PSPACE}$ (Chandra and Stockmeyer 1981); details are as in (Schröder and Pattinson 2009).

Theorem 5.11 (Space Complexity). If every feature in G is strictly one-step complete, closed under contraction, and PSPACE-tractable, then the satisfiability problem for $\mathcal{F}(G)$ -formulas over $\llbracket G \rrbracket$ -coalgebras is in PSPACE.

All rule sets presented in Fig. 1 are strictly one-step complete, closed under contraction, and PSPACE-tractable (this is either clear or shown in (Schröder and Pattinson 2009)). Thus, the above theorem guarantees that satisfiability for logics arising through arbitrary gluings of the features from Example 3.2 are in PSPACE.

Remark 5.12. The recursive structure of the algorithm allows for a modular implementation which interconnects separate matching routines for each feature. In particular, the same algorithmic structure may alternatively be applied to effective heuristic matching routines, leading to more efficient overall reasoning.

6. Conclusions

We have introduced a calculus of *gluings*, which describe ways of combining logical features like uncertainty, non-determinism, and choice. We have shown that the satisfaction problem of a gluing is in PSPACE if this is true for the involved features. This has been achieved by equipping the logics under consideration with a multi-sorted coalgebraic semantics. Crucially, we have

shown that the satisfiability problem of a gluing is equivalent to that of a corresponding flattened gluing; flat gluings are technically more tractable than general gluings, and essentially allow for straightforward generalization of arbitrary generic results from single-sorted coalgebraic logic by just adding multi-sorted notation. As an example, we have carried this out in detail for the generic PSPACE algorithm for satisfiability (Schröder and Pattinson 2009). However, the same principle applies without further ado to other results such as generic EXPTIME algorithms for TBox reasoning in coalgebraic modal and hybrid logics (Schröder, Pattinson, and Kupke 2009; Goré, Kupke, and Pattinson 2010; Goré et al. 2010) or Hennessy-Milner-type expressivity results (Pattinson 2004; Schröder 2007); in fact, Schröder, Pattinson, and Kupke (2009) make extensive use of this mechanism in the examples. Our results thus pave the way for modularized tool support for a large class of heterogeneous logics. The nature of the type of logic combination studied here is essentially that of a typed fusion; the study of further combination mechanisms such as products or \mathcal{E} -connections (Kutz et al. 2004) in the coalgebraic framework is the subject of future work. A further point of interest is to impose frame conditions on heterogeneous coalgebras, thus obtaining axiomatic interaction between the component logics; initial results in this direction may be found in (Pattinson and Schröder 2008). The modularity mechanisms described here are the basis for proof support for heterogeneous modal logics in the generic modal proof tool CoLoSS (Coalgebraic Logic Satisfiability Solver) (Calin et al. 2009).

Acknowledgements The authors wish to thank the anonymous referees for helpful comments, and Erwin R. Catesbeiana for suggestions on multisorted consistency.

References

- Adámek, J.; Herrlich, H.; and Strecker, G. E. 1990. *Abstract and Concrete Categories*. Wiley Interscience.
- Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *J. ACM* 49:672–713.
- Barr, M. 1993. Terminal coalgebras in well-founded set theory. *Theoret. Comput. Sci.* 114:299–315.
- Bartels, F.; Sokolova, A.; and de Vink, E. 2003. A hierarchy of probabilistic system types. In Gumm, H.-P., ed., *Coalgebraic Methods in Computer Science, CMCS 03*, volume 82 of *Electron. Notes Theoret. Comput. Sci.* Elsevier.
- Berghofer, S., and Wenzel, M. 1999. Inductive datatypes in HOL - lessons learned in formal logic engineering. In Bertot, Y.; Dowek, G.; Hirschowitz, A.; Paulin, C.; and Théry, L., eds., *Theorem Proving in Higher Order Logics, TPHOLs 1999*, volume 1690 of *Lect. Notes Comput. Sci.* Springer. 19–36.
- Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*. Cambridge.
- Calin, G.; Myers, R.; Pattinson, D.; and Schröder, L. 2009. CoLoSS: The Coalgebraic Logic Satisfiability Solver (system description). In Areces, C., and Demri, S., eds., *Methods for Modalities, M4M-5, 2007*, volume 231 of *Electron. Notes Theoret. Comput. Sci.*, 41–54. Elsevier.
- Chandra, A., and Stockmeyer, L. 1981. Alternation. *Journal of the ACM* 28(1):114–133.
- Chellas, B. 1980. *Modal Logic*. Cambridge.

- Cîrstea, C., and Pattinson, D. 2007. Modular construction of modal logics. *Theoret. Comput. Sci.* 388(1-3):83–108.
- D’Agostino, G., and Visser, A. 2002. Finality regained: A coalgebraic study of Scott-sets and multisets. *Arch. Math. Logic* 41:267–298.
- de Vink, E., and Rutten, J. 1999. Bisimulation for probabilistic transition systems: A coalgebraic approach. *Theoret. Comput. Sci.* 221:271–293.
- Fine, K. 1972. In so many possible worlds. *Notre Dame J. Formal Logic* 13:516–520.
- Goré, R.; Kupke, C.; Pattinson, D.; and Schröder, L. 2010. Global caching for coalgebraic description logics. In Giesl, J., and Haehnle, R., eds., *International Joint Conference on Automated Reasoning, IJCAR 2010*, volume 6173 of *Lect. Notes Comput. Sci.*, 46–60. Springer.
- Goré, R.; Kupke, C.; and Pattinson, D. 2010. Optimal tableau algorithms for coalgebraic logics. In Esparza, J., and Majumdar, R., eds., *Tools and Algorithms for the Construction and Analysis of Systems, TACAS 10*, volume 6015 of *Lect. Notes Comput. Sci.*, 114–128. Springer.
- Halpern, J. Y. 2003. *Reasoning About Uncertainty*. MIT Press.
- Hansson, H., and Jonsson, B. 1990. A calculus for communicating systems with time and probabilities. In *Real-Time Systems, RTSS 90*, 278–287. IEEE.
- Hansson, H. A. 1994. *Time and Probability in Formal Design of Distributed Systems*. Elsevier.
- Heifetz, A., and Mongin, P. 2001. Probabilistic logic for type spaces. *Games and Economic Behavior* 35:31–53.
- Hemaspaandra, E. 1994. Complexity transfer for modal logic. In Abramsky, S., ed., *Logic in Computer Science, LICS 94*, 164–173. IEEE.
- Jacobs, B. 2001. Many-sorted coalgebraic modal logic: a model-theoretic study. *Theor. Inform. Appl.* 35:31–59.
- Jonsson, B.; Yi, W.; and Larsen, K. G. 2001. Probabilistic extensions of process algebras. In Bergstra, J.; Ponse, A.; and Smolka, S., eds., *Handbook of Process Algebra*. Elsevier.
- Kurucz, A. 2006. Combining modal logics. In van Benthem, J.; Blackburn, P.; and Wolter, F., eds., *Handbook of Modal Logic*. Elsevier.
- Kutz, O.; Lutz, C.; Wolter, F.; and Zakharyashev, M. 2004. \mathcal{E} -connections of abstract description systems. *Artif. Intell.* 156:1–73.
- Larsen, K., and Skou, A. 1991. Bisimulation through probabilistic testing. *Inform. Comput.* 94:1–28.
- Lewis, D. 1975. Intensional logics without iterative axioms. *J. Philos. Logic* 3:457–466.
- Moss, L. 1999. Coalgebraic logic. *Ann. Pure Appl. Logic* 96:277–317.
- Mossakowski, T.; Schröder, L.; Roggenbach, M.; and Reichel, H. 2006. Algebraic-coalgebraic specification in CoCASL. *J. Log. Algebr. Prog.* 67:146–197.
- Myers, R.; Pattinson, D.; and Schröder, L. 2009. Coalgebraic hybrid logic. In de Alfaro, L., ed., *Foundations of Software Science and Computation Structures, FoSSaCS 2009*, volume 5504 of *Lect. Notes Comput. Sci.*, 137–151. Springer.
- Olivetti, N.; Pozzato, G. L.; and Schwind, C. 2007. A sequent calculus and a theorem prover for standard conditional logics. *ACM Trans. Comput. Log.* 8.
- Pattinson, D., and Schröder, L. 2008. Beyond rank 1: Algebraic semantics and finite models for coalgebraic logics. In Amadio, R., ed., *Foundations of Software Science and Computation*

- Structures, FoSSaCS 2008*, volume 4962 of *Lect. Notes Comput. Sci.*, 66–80. Springer.
- Pattinson, D. 2004. Expressive logics for coalgebras via terminal sequence induction. *Notre Dame J. Formal Logic* 45:19–33.
- Pauly, M. 2002. A modal logic for coalitional power in games. *J. Logic Comput.* 12:149–166.
- Rutten, J. 2000. Universal coalgebra: A theory of systems. *Theoret. Comput. Sci.* 249:3–80.
- Schröder, L., and Pattinson, D. 2007. Modular algorithms for heterogeneous modal logics. In Arge, L.; Tarlecki, A.; and Cachin, C., eds., *Automata, Languages and Programming, ICALP 2007*, volume 4596 of *Lect. Notes Comput. Sci.*, 459–471. Springer.
- Schröder, L., and Pattinson, D. 2008a. How many toes do I have? Parthood and number restrictions in description logics. In Brewka, G., and Lang, J., eds., *Principles of Knowledge Representation and Reasoning, KR 2008*, 307–218. AAAI Press.
- Schröder, L., and Pattinson, D. 2008b. Rank-1 modal logics are coalgebraic. *J. Logic Comput.* In print.
- Schröder, L., and Pattinson, D. 2008c. Shallow models for non-iterative modal logics. In Dengel, A.; Berns, K.; Breuel, T.; Bomarius, F.; and Roth-Berghofer, T., eds., *Advances in Artificial Intelligence, KI 2008*, volume 5243 of *Lect. Notes Artif. Intell.*, 324–331. Springer.
- Schröder, L., and Pattinson, D. 2009. PSPACE bounds for rank-1 modal logics. *ACM Trans. Comput. Logic* 10(2:13):1–33.
- Schröder, L.; Pattinson, D.; and Kupke, C. 2009. Nominals for everyone. In Boutilier, C., ed., *International Joint Conferences on Artificial Intelligence, IJCAI 09*, 917–922. AAAI Press.
- Schröder, L. 2007. A finite model construction for coalgebraic modal logic. *J. Log. Algebr. Prog.* 73:97–110.
- Schröder, L. 2008. Expressivity of coalgebraic modal logic: The limits and beyond. *Theoret. Comput. Sci.* 390:230–247.
- Segala, R., and Turrini, A. 2005. Comparative analysis of bisimulation relations on alternating and non-alternating probabilistic models. In *Quantitative Evaluation of Systems, QEST 2005*, 44–53. IEEE.
- Segala, R. 1995. *Modelling and Verification of Randomized Distributed Real-Time Systems*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Segala, R. 2006. Probability and nondeterminism in operational models of concurrency. In *Concurrency Theory, CONCUR 2006*, volume 4137 of *Lect. Notes Comput. Sci.*, 64–78. Springer.
- Sokolova, A.; Vink, E. d.; and Woracek, H. 2009. Coalgebraic weak bisimulation for action-type systems. *Scientif. Ann. Comput. Sci.* 19:93–144.
- Tobies, S. 2001. PSPACE reasoning for graded modal logics. *J. Logic Comput.* 11:85–106.
- Vardi, M. 1989. On the complexity of epistemic reasoning. In *Logic in Computer Science, LICS 89*, 243–251. IEEE.
- Wolter, F. 1998. Fusions of modal logics revisited. In Kracht, M.; de Rijke, M.; and Wansing, H., eds., *Advances in modal logic, AiML 98*, volume 1 of *CSLI Lect. Notes*. CSLI. 361–379.