THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

On the Fusion of Coalgebraic Logics

Citation for published version:

Dahlqvist, F & Pattinson, D 2011, On the Fusion of Coalgebraic Logics. in *Algebra and Coalgebra in Computer Science: 4th International Conference, CALCO 2011, Winchester, UK, August 30 - September 2, 2011. Proceedings*. vol. 6859, Springer Berlin Heidelberg, pp. 161-175. https://doi.org/10.1007/978-3-642-22944-2_12

Digital Object Identifier (DOI):

[10.1007/978-3-642-22944-2_12](https://doi.org/10.1007/978-3-642-22944-2_12)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Algebra and Coalgebra in Computer Science

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

On the Fusion of Coalgebraic Logics

Fredrik Dahlqvist and Dirk Pattinson

Dept. of Computing, Imperial College London

Abstract. Fusion is arguably the simplest way to combine modal logics. For normal modal logics with Kripke semantics, many properties such as completeness and decidability are known to transfer from the component logics to their fusion. In this paper we investigate to what extent these results can be generalised to the case of arbitrary coalgebraic logics. Our main result generalises a construction of Kracht and Wolter and confirms that completeness transfers to fusion for a large class of logics over coalgebraic semantics. This result is independent of the rank of the logics and relies on generalising the notions of distance and box operator to coalgebraic models.

Keywords: modal logic, coalgebra, fusion, completeness

1 Introduction

The most common and simplest way to combine two modal logics L (for ‘left’) and R (for ‘right’) that we take as having disjoint sets of modal operators is their *fusion* $L \otimes R$, i.e. the smallest modal logic containing both L and R . In particular $L \otimes R$ does not contain any axioms combining operators from L with operators from R .

For modal logics with relational semantics, a large number of properties such as decidability and completeness are known to transfer from the component logics to their fusion [11, 17, 8, 13]. The situation for (non-normal) logics outside the realm of relational semantics is far less satisfactory, despite the fact that there is an ever-growing class of logics that fall into this category such as probabilistic modal logic [6] or the logic of (monotone) neighbourhood frames [3]. While decidability can be established, also in the non-normal case, by purely algebraic methods [1], transfer of completeness remains largely open.

In fact, the only result we are aware of is negative: it is shown in [7] that the construction known as ‘modalising’, also discussed in [13] cannot be used to transfer completeness in the case of non-normal modal logics.

The main result of this paper generalises a model-building construction of Kracht and Wolter [11] and confirms that completeness transfers to fusion for a large class of logics over coalgebraic semantics. This technique, also known as *iterated dovetailing* [13, 9], uses many familiar concepts such as successor states, distances and necessitation which are readily available in the context of Kripke semantics but need to be suitably adapted to be put to work in the coalgebraic setting. While all concepts above can be expressed in the coalgebraic

framework, we have to assume that the component logics allow to express a form of necessitation.

In contrast to existing work that addresses the combination of modal logics in the framework of coalgebraic semantics [4, 15] our result only assumes completeness of the component logics with respect to a subclass of (coalgebraic) frames, whereas *op.cit* assumes completeness with respect to the class of *all* coalgebras. In particular, we do not restrict to logics whose axiomatisation only uses so-called rank-1 axioms and our result can be applied to logics that incorporate arbitrary frame conditions, as long as they come equipped with a complete semantics and possess a necessitation operator.

2 Preliminaries

We fix a countable set V of atomic propositions throughout. A *modal signature* Λ is a set of (modal) operators with associated arities. A modal signature Λ' *extends* a modal signature Λ if $\Lambda' \supseteq \Lambda$. Given two modal signatures Λ_L and Λ_R , their disjoint union is denoted by $\Lambda_L + \Lambda_R$. The set $\mathcal{L}(\Lambda, V_0)$ of Λ -formulae over the set $V_0 \subseteq V$ of propositional variables is given by the grammar

$$\phi ::= p \mid \perp \mid \neg\phi \mid \phi \wedge \psi \mid \heartsuit(\phi_1, \dots, \phi_n)$$

where $p \in V_0$ ranges over atomic propositions and $\heartsuit \in \Lambda$ is n -ary. If $V_0 = V$ is the set of all propositional variables, we write $\mathcal{L}(\Lambda) = \mathcal{L}(\Lambda, V)$. The set of propositional variables that occur in a formula $\phi \in \mathcal{L}(\Lambda, V_0)$ is denoted by $\text{var}(\phi)$, $\text{sf}(\phi)$ is the set of subformulae of ϕ and $\text{md}(\phi)$ denotes the modal depth, i.e. the maximal nesting depth of modal operators in ϕ .

A Λ -logic L is a set of $\mathcal{L}(\Lambda)$ -formulae containing all propositional tautologies, and closed under modus ponens, uniform substitution and the congruence rules

$$\frac{p_1 \leftrightarrow q_1 \wedge \dots \wedge p_n \leftrightarrow q_n}{\heartsuit(p_1, \dots, p_n) \leftrightarrow \heartsuit(q_1, \dots, q_n)}$$

for each n -ary operator $\heartsuit \in \Lambda$. An L -*theorem* is a formula $\phi \in L$ and we write $\vdash_L \phi$ in this case, and a formula $\phi \in \mathcal{L}(\Lambda)$ is L -*consistent* if $\neg\phi \notin L$. The smallest congruential Λ -logic is denoted by $\mathbf{E}_\Lambda$.

Given two modal signatures Λ and Λ' such that Λ' extends Λ , we will say that a Λ' -logic M is an *extension* of a Λ -logic L if $L \subseteq M$. If, additionally, for every Λ -formula ϕ , $\phi \in L$ iff $\phi \in M$, then M is called a *conservative extension* of L . The *lattice of extensions* $\mathcal{E}(L)$ of a coalgebraic logic L is the set of all extensions $L \subseteq M$ of L with the meet and join operations given by the set intersection $\cap$ and union $\cup$ operations respectively.

On the semantical side, formulae are interpreted over T -coalgebras, where $T : \mathbf{Set} \rightarrow \mathbf{Set}$ is an endofunctor. A T -*coalgebra* is a pair $F = (W, \gamma)$ where W is a set (of worlds) and $\gamma : W \rightarrow TW$ is a (transition) function. Here, T -coalgebras play the role of frames, and we frequently refer to T -coalgebras as T -*frames*. A T -*model* is a triple $M = (W, \gamma, \sigma)$ where (W, γ) is a T -frame and $\sigma : V \rightarrow \mathcal{P}(W)$ is a valuation (and $\mathcal{P}(W)$ is the powerset of W). A T -model (W, γ, σ) is *based* on the T -frame (W, γ) .

Assumption 1. Modulo a modification of its action on the empty set and empty mappings, any **Set**-endofunctor T can be assumed to preserve finite intersections, all monos as well as the inverse image of injective maps [10]. Since modifying T on the empty set only gives an isomorphic category of coalgebras we will assume throughout that all functors do possess these preservation properties.

The interpretation of Λ -formulae over T -models requires that T extends to a Λ -structure, i.e. T comes equipped with an interpretation of the operators in Λ . Concretely, a Λ -structure consists of an endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ together with an assignment of an n -ary *predicate lifting*, i.e. a set-indexed family of maps

$$(\llbracket \heartsuit \rrbracket_X : \mathcal{P}(X)^n \rightarrow \mathcal{P}(TX)_{X \in \mathbf{Set}})$$

to every n -ary operator $\heartsuit \in \Lambda$ that satisfies the *naturality condition*

$$\llbracket \heartsuit \rrbracket_X \circ (f^{-1})^n = (Tf)^{-1} \circ \heartsuit_Y$$

for all maps $f : X \rightarrow Y$. Categorically speaking, $\llbracket \heartsuit \rrbracket$ is a natural transformation $2^- \rightarrow 2^{T^-}$ where $2^- : \mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$ is the contravariant powerset functor. We will usually keep the assignment of predicate liftings implicit and just refer to Λ -structures by the underlying endofunctor.

Given a modal signature Λ and a Λ -structure T , the satisfaction relation between worlds of T -models and Λ -formulae is given inductively by

$$\begin{aligned} \mathcal{M}, w \models p &\text{ iff } w \in \sigma(p) & \mathcal{M}, w \models \neg\phi &\text{ iff not } \mathcal{M}, w \models \phi \\ \mathcal{M}, w \models \perp &\text{ never} & \mathcal{M}, w \models \phi \wedge \psi &\text{ iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \end{aligned}$$

$$\mathcal{M}, w \models \heartsuit(\phi_1, \dots, \phi_n) \text{ iff } \gamma(w) \in \llbracket \heartsuit \rrbracket_W(\llbracket \phi_1 \rrbracket, \dots, \llbracket \phi_n \rrbracket)$$

where $\mathcal{M} = (W, \gamma, \sigma)$ is a T -model. We write $\llbracket \phi \rrbracket_{\mathcal{M}} = \{w \in W \mid \mathcal{M}, w \models \phi\}$ for the *truth-set* of ϕ relative to $\mathcal{M}$.

A formula $\phi \in \mathcal{L}(\Lambda)$ is *valid* in a T -frame $F = (W, \gamma)$ if $\mathcal{M}, w \models \phi$ for all T -models $\mathcal{M}$ based on F , this is denoted by $F \models \phi$. If $\mathcal{C}$ is a class of T -frames, we say that ϕ is *valid* on $\mathcal{C}$ if $F \models \phi$ for all $F \in \mathcal{C}$, this is denoted $\mathcal{C} \models \phi$. The *logic of* $\mathcal{C}$ is the set of all formulae valid on $\mathcal{C}$, i.e.

$$\text{Log}(\mathcal{C}) = \{\phi \in \mathcal{L}(\Lambda) \mid \mathcal{C} \models \phi\}.$$

It is easy to check that $\text{Log}(\mathcal{C})$ is a Λ -logic.

3 Fusion and Transfer of Soundness and Consistency

Given two modal signatures Λ_L and Λ_R (where the subscripts stand for ‘left’ and ‘right’, respectively), the *fusion* of a Λ_L -logic with a Λ_R -logic is the smallest $\Lambda_L + \Lambda_R$ -logic that extends both. In other words, we may mix Λ_L and Λ_R operators freely in the fusion, but not stipulate any interaction between them.

Definition 2 (Fusion of Logics). Given a Λ_L -logic L and a Λ_R -logic R where Λ_L and Λ_R are two arbitrary modal signatures, the *fusion* $F = L \otimes R$ is the smallest $\Lambda_L + \Lambda_R$ -logic containing both L and R . The fusion is therefore a binary operation $- \otimes - : \mathcal{E}(\mathbf{E}_{\Lambda_L}) \times \mathcal{E}(\mathbf{E}_{\Lambda_R}) \rightarrow \mathcal{E}(\mathbf{E}_{\Lambda_L + \Lambda_R})$.

While we consider the modal operators of L and R to be disjoint by assumption, propositional connectives are shared. Given two structures for Λ_L and Λ_R , we can interpret the fusion over a $\Lambda_L + \Lambda_R$ -structure as follows:

Definition 3 (Fusion of Structures). Given a Λ_L -structure S and a Λ_R -structure T over modal signatures Λ_L and Λ_R , the *fusion* of S and T is the $\Lambda_L + \Lambda_R$ -structure over the endofunctor $S \times T$ where the assignment of predicate liftings is given by

$$\begin{aligned} \llbracket \heartsuit \rrbracket_X^1 &= \pi_1^{-1} \circ \llbracket \heartsuit \rrbracket_X & \text{for } \heartsuit \in \Lambda_L \\ \llbracket \spadesuit \rrbracket_X^2 &= \pi_2^{-1} \circ \llbracket \spadesuit \rrbracket_X & \text{for } \spadesuit \in \Lambda_R \end{aligned}$$

where π_1 and π_2 are the projections $SX \xleftarrow{\pi_1} SX \times TX \xrightarrow{\pi_2} TX$. Note that both $\llbracket \heartsuit \rrbracket^1$ and $\llbracket \spadesuit \rrbracket^2$ are predicate liftings of type $2^{(-)^n} \rightarrow 2^{S \times T}$ for n -ary operators $\heartsuit$ and $\spadesuit$.

In other words, the fusion of two structures S and T produces a structure for the disjoint union of operators. In particular $S \times T$ -frames carry both the structure of an S -frame and a T -frame, and the interpretation of modalities over the product of S and T is obtained by their interpretation over S and T by just projecting to the respective component. This induces a fusion operation on frame classes.

Definition 4 (Fusion of Frame Classes). Let S and T be two **Set**-endofunctors. If $\mathcal{C}_L$ and $\mathcal{C}_R$ are classes of S and T -frames, respectively, the *fusion* $\mathcal{C}_L \otimes \mathcal{C}_R$ of $\mathcal{C}_L$ and $\mathcal{C}_R$ is the class of $S \times T$ -frames given by

$$\mathcal{C}_L \otimes \mathcal{C}_R = \{(X \xrightarrow{\langle \gamma, \delta \rangle} SX \times TX) \mid (X, \gamma) \in \mathcal{C}_L \text{ and } (X, \delta) \in \mathcal{C}_R\}$$

where $\langle \gamma, \delta \rangle(x) = (\gamma(x), \delta(x))$.

The purpose of this paper is essentially to show that for a large class of functors Definitions 2 and 4 are counterparts of one another, i.e. $L \otimes R = \text{Log}(\mathcal{C}_L \otimes \mathcal{C}_R)$ iff $L = \text{Log}(\mathcal{C}_L)$ and $R = \text{Log}(\mathcal{C}_R)$. For now, let us introduce some examples of Λ -logics, their coalgebraic semantics as well as some examples of fusion.

Example 5. 1. For $\Lambda = \{\Box\}$, the logic $\mathbf{E}_{\{\Box\}}$ is the smallest congruential modal logic (classical modal logic in the terminology of [3]). It is sound and complete with respect to neighbourhood frames, i.e. $\mathcal{N}$ -coalgebras, where $\mathcal{N} = 2^{2^-}$ is the neighbourhood functor (see also [14]).

2. The smallest extension of $\mathbf{E}$ containing $\Box(p \rightarrow q) \rightarrow \Box p \rightarrow \Box q$ that is closed under the necessitation rule $(p/\Box p)$ is the logic $\mathbf{K}$ which is sound and complete with respect to the class of all $\mathcal{P}$ -coalgebras, where $\mathcal{P}$ is the covariant powerset functor and $\llbracket \Box \rrbracket_X(A) = \{B \in \mathcal{P}(X) \mid B \subseteq A\}$. A $\mathcal{P}$ -coalgebra $\gamma :$

$W \rightarrow \mathcal{P}W$ is equivalent to a relation $\mathcal{R}$ on W via $(x, y) \in \mathcal{R}$ iff $y \in \gamma(x)$. It follows that a $\mathcal{P} \times \mathcal{P}$ -coalgebra $\langle \gamma, \delta \rangle$ as defined by Def. 4 is equivalent to a pair of relations $\mathcal{R}_L$ and $\mathcal{R}_R$ via $(x, y) \in \mathcal{R}_L$ iff $y \in \pi_1 \circ \langle \gamma, \delta \rangle(x)$ and $(x, y) \in \mathcal{R}_R$ iff $y \in \pi_2 \circ \langle \gamma, \delta \rangle(x)$. We thus recover the standard definition of fusion of Kripke frames [13, 11].

3. If we restrict the previous example to the class $\mathcal{C}_{\mathbf{S5}}$ of $\mathcal{P}$ -coalgebras corresponding to Kripke frames where the relation is an equivalence, then $\text{Log}(\mathcal{C}_{\mathbf{S5}})$ is the epistemic logic $\mathbf{S5}$ [2]. Syntactically $\mathbf{S5}$, is the smallest extension of $\mathbf{K}$ containing the axioms for transitivity (4), reflexivity (T) and symmetry (B). It is customary write the operator of $\mathbf{S5}$ as K and to consider the n -fold fusion $\mathbf{S5}_n = \bigotimes_{i=1}^n \mathbf{S5}$. The modal signature of this logic is $\coprod_{i=1}^n \{K\} \simeq \{K_i\}_{1 \leq i \leq n}$. By Definition 2, $\mathbf{S5}_n$ is thus the smallest $\{K_i\}_{1 \leq i \leq n}$ -logic closed under a necessitation rule and a copy of $\Box(p \rightarrow q) \rightarrow \Box p \rightarrow \Box q$, T , B and 4 per operator.

4. The finitely supported probability distribution functor $\mathcal{D}(X) = \{\mu : X \rightarrow_f [0, 1] \mid \sum_{x \in X} \mu(x) = 1\}$ (where $\rightarrow_f$ denotes finite support) extends to a Λ -structure for $\Lambda = \{L_u \mid u \in [0, 1] \cap \mathbb{Q}\}$ consisting of unary operators L_u , read as ‘with probability at least u ’. The functor $\mathcal{D}$ extends to a Λ -structure by stipulating $\llbracket L_u \rrbracket_X(A) = \{\mu \in \mathcal{D}(X) \mid \mu(A) \geq u\}$. We therefore obtain *probabilistic epistemic logic* as the fusion of $\mathbf{S5}_n \otimes \text{Log}(\mathcal{D})$ where $\mathcal{D}$ is the class of all $\mathcal{D}$ -frames.

Our first goal is to show that consistency transfers under fusion, i.e. $\perp \notin L \otimes R$ whenever $\perp \notin L, R$. We will show this algebraically following [9] using terminology adapted from [5]. The transfer of consistency can be proved using purely coalgebraic arguments, but it requires the assumption that both constituents of the fusion be sound and complete w.r.t. some classes of coalgebras. Using algebraic arguments we can prove the transfer of consistency under fusion without any extra hypothesis. The proof is also shorter and more elegant.

Definition 6 (Algebraic Semantics). Given a modal signature Λ , a Λ -modal algebra $\mathfrak{A}$ is a boolean algebra $(A, 0, 1, \neg^{\mathfrak{A}}, \wedge^{\mathfrak{A}})$ together with an n -ary function $f_{\heartsuit} : A^n \rightarrow A$, for each $\heartsuit \in \Lambda$. The boolean algebra $(A, 0, 1, \neg^{\mathfrak{A}}, \wedge^{\mathfrak{A}})$ is called the *boolean reduct* of $\mathfrak{A}$. A *variable assignment* is a map $\sigma : V \rightarrow A$ into the carrier set of $\mathfrak{A}$. Every variable assignment σ extends to an *interpretation* (denoted by the same symbol) $\sigma : \mathcal{L}(\Lambda) \rightarrow \mathfrak{A}$ given by

$$\begin{aligned} \sigma(\phi \wedge \psi) &= \sigma(\phi) \wedge^{\mathfrak{A}} \sigma(\psi) & \sigma(\neg \phi) &= \neg^{\mathfrak{A}} \sigma(\phi) \\ \sigma(\heartsuit(\phi_0, \dots, \phi_n)) &= f_{\heartsuit}(\sigma(\phi_0), \dots, \sigma(\phi_n)) \end{aligned}$$

for $\heartsuit \in \Lambda$ n -ary. An *algebraic model* is a pair $\mathfrak{M} = (\mathfrak{A}, \sigma)$ where $\mathfrak{A}$ is a Λ -modal algebra and σ is an interpretation of $\mathcal{L}(\Lambda)$, $\mathfrak{M}$ is said to be *based on* $\mathfrak{A}$. We say that a formula ϕ is *satisfied* in $\mathfrak{M}$ if $\sigma(\phi) \neq 0$ and we say that ϕ is *true* in $\mathfrak{M}$ (notation $(\mathfrak{A}, \sigma) \models \phi$) when $\sigma(\phi) = 1$. Finally, we say that ϕ is *valid* in $\mathfrak{A}$ if ϕ is true for any model $(\mathfrak{A}, \sigma)$ based on $\mathfrak{A}$.

We will sketch the proof that if L and R are consistent (i.e. $\perp \notin L, R$) then $L \otimes R$ is a conservative extension of both L and R . For more details we refer the reader to [9] where this theorem is proved for normal unary multi-modal

logics. However, since the proof does not involve normality or the arity of the operators it extends trivially to our setup. To show that $L \otimes R$ is a conservative extension of both L and R the first step is to build the Lindenbaum-Tarski algebra of L which provides us with a Λ_L -modal algebra $\mathfrak{A}_L$ which validates all formulae of L . Secondly, by contraposition consider a Λ_L -formula ϕ such that $\not\vdash_L \neg\phi$. Then $\sigma_{\mathfrak{A}_L}(\neg\phi) \neq 1$, where $\sigma_{\mathfrak{A}_L}$ is the interpretation associated with the Lindenbaum-Tarski algebra of L . Thirdly, we build the Lindenbaum-Tarski algebra $\mathfrak{A}_R$ of R which validates all formulae in R . It is easy to show that both $\mathfrak{A}_L$ and $\mathfrak{A}_R$ are countably infinite atomless provided both similarity types are at most countable [9]. As any countably infinite boolean algebras are isomorphic, their boolean reducts (and thus carrier sets) can therefore be assumed to be equal. The $\Lambda_L + \Lambda_R$ -modal algebra consisting of this common boolean algebra plus all operators from $\mathfrak{A}_L$ and $\mathfrak{A}_R$ provides us with a $\Lambda_L + \Lambda_R$ -modal algebra $\mathfrak{A}_{L \otimes R}$ which validates L and R and in which $\sigma(\neg\phi) \neq 1$. By the completeness of algebraic semantics we can thus conclude that $\not\vdash_{L \otimes R} \neg\phi$.

Theorem 7. *If L and R are consistent logics over at most countable similarity types Λ_L and Λ_R (i.e. $\perp \notin L, R$) then $L \otimes R$ is a conservative extension of both L and R .*

A direct consequence of this proposition is that fusion preserves consistency.

Theorem 8 (Consistency transfers). *The fusion of two consistent Λ -logics is consistent, i.e. $\perp \notin L \otimes R$ whenever $\perp \notin L, R$.*

While the transfer of consistency can be based on a purely syntactic argument, the transfer of soundness involves the (coalgebraic) semantics, and in particular the fact that a $S \times T$ -model can be seen as an S (or T) model by simply forgetting the T (or S) structure.

Theorem 9 (Soundness Transfers). *Let L and R be Λ_L and Λ_R -logics, respectively. If $\mathcal{C}_L$ and $\mathcal{C}_R$ are classes of S -frames and T -frames, respectively, then $L \otimes R \subseteq \text{Log}(\mathcal{C}_L \otimes \mathcal{C}_R)$.*

4 Transfer of Completeness

The remainder of the paper is devoted to establishing the converse of Theorem 9: we establish that completeness transfers to the fusion of two logics over coalgebraic semantics. While an algebraic approach yields transfer of consistency without any further assumptions (in particular without assuming a complete semantics), the situation is different when it comes to transfer of completeness. A naive approach via categorical duality only yields completeness with respect to (the coalgebraic analogue of) descriptive general frames, or coalgebras over Stone spaces [12]. In particular, this form of algebraic completeness does not appear to yield completeness with respect to the fusion of frame classes, mainly because constructions like canonical extensions do not have a coalgebraic counterpart. We therefore adapt a classical construction that witnesses satisfiability

of (fusion-) consistent formulas [11] to the coalgebraic setting, and directly build (set-theoretic) models.

For the whole section, we fix two modal signatures Λ_L and Λ_R , two Λ_L - and Λ_R -structures over functors S and T respectively and two logics L and R that are sound and complete with respect to classes $\mathcal{C}_L$ and $\mathcal{C}_R$ of S and T -frames, respectively, that is, $L = \text{Log}(\mathcal{C}_L)$ and $R = \text{Log}(\mathcal{C}_R)$. Our goal is to show that the fusion $L \otimes R$ is complete with respect to the fusion of the corresponding frame classes, i.e. $L \otimes R = \text{Log}(\mathcal{C}_L \otimes \mathcal{C}_R)$. As usual, we consider the contrapositive and show that every $L \otimes R$ -consistent formula can be satisfied in a model that is based on a $\mathcal{C}_L \otimes \mathcal{C}_R$ -frame. We use the fact that we know how to build coalgebraic models for L - and R -consistent formulae. So the trick is to turn $L \otimes R$ -consistent formulae into L - and R -consistent ones and glue the satisfying models together in a suitable way. The passage from $\mathcal{L}(\Lambda_L + \Lambda_R)$ formulae to formulae of the component languages $\mathcal{L}(\Lambda_L)$ and $\mathcal{L}(\Lambda_R)$ is achieved by the following constructions introduced in [11].

4.1 Ersatz and Reconstruction

Definition 10 (Ersatz). Let Λ_L and Λ_R be two similarity types. If ϕ is a formula in the language of the fusion $\mathcal{L}(\Lambda_L + \Lambda_R)$, then its *L-ersatz* (or left-ersatz) ϕ^L is defined by putting $p^L = p$, $(\phi \wedge \psi)^L = \phi^L \wedge \psi^L$, $(\neg \phi)^L = \neg \phi^L$ and

$$\begin{aligned} (\heartsuit(\phi_1, \dots, \phi_n))^L &= \heartsuit(\phi_1^L, \dots, \phi_n^L) && \text{for all } \heartsuit \in \Lambda_L \\ (\spadesuit(\phi_1, \dots, \phi_m))^L &= q_{\spadesuit(\phi_1, \dots, \phi_m)} && \text{for all } \spadesuit \in \Lambda_R \end{aligned}$$

where $p \in V$ and $q_{\spadesuit(\phi_1, \dots, \phi_m)} \in V_R$ is a fresh variable from a set V_R of variables disjoint from V , called an *R-surrogate* (or right-surrogate). The *R-ersatz* (or right-ersatz) ϕ^R of ϕ is defined dually by switching the role of $\heartsuit$ and $\spadesuit$.

The ersatz operations $(\cdot)^L$ and $(\cdot)^R$ are thus an operations of type

$$(\cdot)^L : \mathcal{L}(\Lambda_L + \Lambda_R, V) \rightarrow \mathcal{L}(\Lambda_L, V \cup V_R) \quad (\cdot)^R : \mathcal{L}(\Lambda_L + \Lambda_R, V) \rightarrow \mathcal{L}(\Lambda_R, V \cup V_L).$$

Although not mentioned in [11], the construction relies on the fact that the ersatz operations preserve consistency.

Lemma 11 (Ersatz perserves Consistency). *If ϕ is an $L \otimes R$ -consistent formula then ϕ^L (respectively ϕ^R) is L -consistent (respectively R -consistent).*

Example 12. Consider the formula $\phi = L_u K_i L_v p \wedge L_u q$ in the language of $\mathbf{S}_5^n \otimes \mathbf{Prob}$ introduced in Example 5. We obtain

$$\phi^L = q_{L_u K_i L_v p} \wedge q_{L_u q} \text{ and } \phi^R = L_u q_{K_i L_v p} \wedge L_u q.$$

As can be seen from this example and from the definition, the left-ersatz construction transforms a subformula into a surrogate variable as soon as it sees

an R -operator. The nesting of operators following the outermost R -operator is therefore lost. One of the key ideas of the completeness transfer theorem is to alternate left and right ersatz constructions, building a model at each stage until we have drilled down to the level of propositional variables. In order to do this we need to be able to build ersatz formulae deeper inside the nesting of operators.

Definition 13 (Reconstruction). Let ϕ be free of left (respectively right) surrogate variables. We define the reconstruction operators $\uparrow_L$ and $\uparrow_R$ by

$$\begin{aligned}\uparrow_R \phi &= \phi \left(\spadesuit((\psi_1^i)^R, \dots, (\psi_m^i)^R) / q_{\spadesuit(\psi_1^i, \dots, \psi_m^i), p_j} \right) \text{ for all } \spadesuit \in \Lambda_R \\ \uparrow_L \phi &= \phi \left(\heartsuit((\psi_1^i)^L, \dots, (\psi_n^i)^L) / q_{\heartsuit(\psi_1^i, \dots, \psi_n^i), p_j} \right) \text{ for all } \heartsuit \in \Lambda_L\end{aligned}$$

where i ranges over the set of right (resp. left) surrogate variables in ϕ and j ranges over the set of non-surrogate variables in ϕ . These operations have type

$$\begin{aligned}\uparrow_L : \mathcal{L}(\Lambda_L + \Lambda_R, V \cup V_L) &\rightarrow \mathcal{L}(\Lambda_L + \Lambda_R, V \cup V_R) \\ \uparrow_R : \mathcal{L}(\Lambda_L + \Lambda_R, V \cup V_R) &\rightarrow \mathcal{L}(\Lambda_L + \Lambda_R, V \cup V_L),\end{aligned}$$

that is $\uparrow_L$ maps formulae free of R -surrogates to formulae free of L -surrogates, dually for $\uparrow_R$. As a consequence, left and right reconstructions can be alternated. To simplify notation we write $\uparrow$ for both and $\uparrow^n$ for the n -fold iteration of $\uparrow$.

Example 14. Let ϕ be as in Example 12. Then we have the following:

$$\begin{aligned}\uparrow \phi^L &= \uparrow(q_{L_u K_i L_v p} \wedge q_{L_u q}) = L_u(K_i L_v p)^R \wedge (L_u q)^R = L_u q_{K_i L_v p} \wedge L_u q \\ \uparrow^2 \phi^L &= \uparrow(L_u q_{K_i L_v p} \wedge L_u q) = L_u K_i (L_v p)^L \wedge L_u q = L_u K_i q_{L_v p} \wedge L_u q \\ \uparrow^3 \phi^L &= \uparrow(L_u K_i q_{L_v p} \wedge L_u q) = L_u K_i L_v (p)^R \wedge L_u q = L_u K_i L_v p \wedge L_u q = \phi\end{aligned}$$

It is intuitively clear that repeated reconstruction reconstructs the original formula step-by-step. We note this as:

Lemma and Definition 15. *Let ϕ be a formula on which the reconstruction operator $\uparrow$ is defined. Then there exists $n \in \omega$ such that $\uparrow^n \phi = \uparrow^{n+1} \phi$. This is the case exactly when the $\uparrow^n \phi$ has no surrogate variables (i.e. has the same variables as ϕ). We call $\uparrow^n \phi$ the total reconstruction of ϕ , denoted by $\phi^\uparrow$.*

4.2 Consistency Sets and Consistency Formulae

We are now equipped with two constructions that allow us to ‘project’ an $L \otimes R$ -consistent formula ϕ onto L - and R -consistent formulae (by Lemma 11) for which we know how to build coalgebraic models. Let us for instance start by building ϕ^L . Since it is L -consistent and L is complete with respect to a class $\mathcal{C}_L$ of S -coalgebras, we can build a coalgebraic model for ϕ^L . However, we quickly run into trouble: any model for ϕ^L cannot take into account the actual meaning of its surrogate variables, i.e. of their indices. In particular, it can assign to two surrogate variables q_ϕ, q_ψ a truth value of *true* at the same point even if $\phi \wedge \psi$ is $L \otimes R$ -inconsistent. To avoid this problem we need some way of enforcing that

a model of ϕ^L gives $L \otimes R$ -consistent valuations to surrogate variables. This is achieved by using the following constructions. By convention we will always consider reconstructions based on the left-ersatz (this can be done without loss of generality, since $L \otimes R \simeq R \otimes L$). To simplify the notation we will write the i^{th} reconstruction of ϕ^L as $\phi_i = \uparrow^i(\phi^L)$, and we define for each i the *index set* $\mathbb{S}_i$ of indices of surrogate variables of the i^{th} reconstruction and their subformulae. Formally:

$$\mathbb{S}_i = \text{sf}\{\psi \mid q_\psi \in \text{var}(\phi_i)\} \cup \text{var}(\phi)$$

$\mathbb{S}_0$ for example, regroups the indices of all surrogate variables of ϕ^L as well as their subformulae. To enforce $L \otimes R$ -consistent valuations of the surrogate variables of ϕ^L , the idea is to first list all the possible $L \otimes R$ -consistent combinations of formulae in $\mathbb{S}_0$ into a ‘consistency set’.

Definition 16 (Consistency Sets). Let L be a logic and let Δ be a finite set of formulae in the language of L . The *L-consistency set* $\Sigma(\Delta)$ is defined by

$$\Sigma(\Delta) = \{\bigwedge M \mid M \subseteq \Delta \cup \neg\Delta \mid M \text{ maximally consistent}\}$$

where $\neg\Delta = \{\neg\phi \mid \phi \in \Delta\}$ and $s_\Delta^L = \bigvee \Sigma(\Delta)$ is the *L-consistency formula* of Δ .

We read ‘maximally consistent’ above as maximal among the subsets of $\Delta \cup \neg\Delta$, and $\Sigma(\Delta)$ contains all different possible realisations of combinations of formulae in Δ and the consistency formula s_Δ amounts to requiring that one such combination can be satisfied in a model. Returning to our problem and setting $\Delta = \mathbb{S}_0$, it is easy to see that if the consistency formula $s_{\mathbb{S}_0}^L$ is true at a certain point of an L -model, then we are guaranteed that the combination of surrogate variables which are valuated as ‘true’ at that point stand for an $L \otimes R$ -consistent combination of formulae of $\mathbb{S}_0$. We crucially have that:

Lemma 17. *L-consistency formulae are L-theorems.*

Since $s_{\mathbb{S}_0}$ is an $L \otimes R$ -theorem, $\phi \wedge s_{\mathbb{S}_0}$ is $L \otimes R$ -consistent, thus by Lemma 11, $\phi^L \wedge s_{\mathbb{S}_0}^L$ is L -consistent, and by completeness of L a model can thus be build for it. Such a model of ϕ^L will necessarily have an $L \otimes R$ -consistent valuation for surrogate variables.

4.3 Necessity operators and distances

We have just solved a problem in the construction of our model, but we are almost immediately confronted by another one. Indeed, we may have avoided $L \otimes R$ -inconsistencies at *one* point in the L -model of ϕ^L (namely the point w making ϕ^L true), but $L \otimes R$ -inconsistent valuations of surrogate variables could still happen elsewhere in the model. We therefore need to ‘propagate’ consistency. In Kripke frames we can simply use the necessitation rule and the box operator $\Box$ to propagate $s_{\mathbb{S}_0}^L$, but in a coalgebraic interpretation this is in general not possible. This problem is the biggest hurdle, but also the most interesting, in generalising Kracht and Wolter’s construction [11] to coalgebraic semantics.

Definition 18 (One-step successors). Given a T -coalgebra $W \xrightarrow{\gamma} TW$, and an element $w \in W$, we define the *1-step successors* $S_1(w)$ of w to be the set

$$S_1(w) = \bigcap \{U \subseteq W \mid \gamma(w) \in Ti[TU]\}$$

where $i : U \hookrightarrow W$ is inclusion and $Ti[TU]$ is the direct image of TU under Ti .

Intuitively, $S_1(w)$ is the smallest subset of W providing a ‘support’ for $\gamma(w)$. Note however, that Definition 18 is in general not well defined, as arbitrary intersections need not be preserved by **Set**-endofunctors (unlike finite intersections, see Assumption 1). The filter functor provides an easy example of such a behaviour. In fact, [10, Corollary 4.8] provides an elegant criterion for the existence of the set of 1-step successors of a point in a coalgebra.

Proposition 19. *T preserves infinite intersections iff for any $u \in TW$ there is a smallest $U \subseteq W$ with $u \in Ti[TU]$.*

Assumption 20. From now on we will therefore assume that we are dealing with **Set**-endofunctors that preserve arbitrary intersections. Note that for logics with the finite model property, since we can always assume that the carrier set of any coalgebraic model is finite and since all **Set**-endofunctor preserve finite intersections, we drop this assumption. In particular, all complete rank-1 coalgebraic logic have the finite model property.

The notion of 1-step successor allows us to define a notion of distance on the points of a coalgebraic model $W \xrightarrow{\gamma} TW$.

Definition 21 (Distance). We say that there is a *path of length n* between $x, y \in W$ if there is a sequence of elements $(x_i)_{1 \leq i \leq n}$ such that $x = x_1$, $x_n = y$ and $x_{i+1} \in S_1(x_i)$ for all $1 \leq i < n$. The *distance* $dist(x, y)$ between $x, y \in W$ is the length of the shortest path between x and y , or ∞ if no such path exists. Any T -coalgebra (W, γ) induces a distance function $dist : W \times W \rightarrow \mathbb{N} \cup \{\infty\}$. Based on the notion of distance we can generalise the set of one-step successors to the following sets of *n -step successors* of w (ball and sphere of radius n around w):

$$S_n(w) = \{x \in W \mid dist(w, x) = n\} \quad B_n(w) = \{x \in W \mid dist(w, x) \leq n\}$$

Finally, we will say that there is a path between two points w and x and write $w \rightsquigarrow x$ if there is a path of finite length between w and x .

Remark 22. 1. The notion of one-step successor is not symmetric and the notion of distance defined above is therefore not a true metric. However, it is an (extended) quasimetric, i.e. it is non-negative, satisfies the triangle inequality and is zero iff the two arguments are equal.

2. For an $S \times T$ -coalgebra $W \xrightarrow{\langle \gamma, \delta \rangle} SW \times TW$, there are three notions of distance: the S -distance based on the notion of S -successors $B_1^S(w) := \bigcap \{U \subseteq W \mid \gamma(w) \in Si[SU]\}$, the T -distance based on the notion of T -successors $B_1^T(w) := \bigcap \{U \subseteq W \mid \delta(w) \in Ti[TU]\}$ and the combined $S \times T$ -distance based on the notion of $S \times T$ -successors $B_1^{S \times T}(w) := \bigcap \{U \subseteq W \mid (\gamma(w), \delta(w)) \in Si[SU] \times Ti[TU]\}$.

A key feature of Kripke semantics is the local aspect of truth, i.e. that the truth of a formula of modal depth n at a world w depends only on points at most n -steps away from w . Similarly, in the coalgebraic semantics it is intuitively clear that if $\phi \in L$ is of modal depth $\text{md}(\phi) = n$ and $\gamma : W \rightarrow TW$ is a coalgebraic frame, then the truth value of ϕ at $w \in W$ will also only depend on the valuations of propositional variables at points $x \in B_n(w)$.

Theorem 23 (Coalgebraic semantics is local). *Let ϕ be a Λ -formula of modal depth $\text{md}(\phi) = n$. If $M = (W, \gamma, \sigma)$ and $M' = (W, \gamma, \sigma')$ are T -models based on the same frame and T preserves arbitrary intersections, then we have for all $w \in W$ that*

$$M, w \models \phi \iff M', w \models \phi$$

whenever $\sigma(p) \cap B_n(w) = \sigma'(p) \cap B_n(w)$ for all $p \in \text{var}(\phi)$.

Returning to our problem, this result tells us how far we need to propagate the truth of our consistency formula $s_{S_0}^L$: if the modal depth of ϕ^L is n , we only need to concentrate our efforts on enforcing $s_{S_0}^L$ in a ball of radius n around the point w where ϕ^L will be satisfied. But how can this be done in a coalgebraic model? Over relational semantics we can use the box operator $\Box$ to enforce the truth of a formula ψ at all 1-step away successors but this ability of Kripke semantics to enforce truth on successor states is in general not available in the coalgebraic framework. We will therefore need a coalgebraic generalisation of the necessitation operator.

Definition 24 (Necessity and necessity operators). Let L be a Λ -logic. Then L has *weak necessity* over a Λ -structure T if, for every L -consistent formula ϕ there exists an L -consistent formula $\text{nec}(\phi)$ such that

$$S_1(w) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}} \text{ whenever } \mathcal{M}, w \models \text{nec}(\phi)$$

for all T -models $\mathcal{M} = (W, \gamma, \sigma)$ and all $w \in W$. We will say that L has *strong necessity* over T if

$$S_1(w) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}} \text{ iff } \mathcal{M}, w \models \text{nec}(\phi)$$

A unary operator $\heartsuit \in \Lambda$ is a *necessity operator* over T if

$$S_1(w) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}} \text{ whenever } \mathcal{M}, w \models \heartsuit \phi$$

for every T -model $\mathcal{M} = (W, \gamma, \sigma)$ and all $w \in W$. We usually use the symbol $\Box$ for necessity operators.

We will solely focus on notions of necessity arising from necessity operators. In most practical cases such an operator can either be found directly in the logic itself or can be *simulated* within the logic as a boolean combination of existing operators. The following result shows how frequent necessity operators are.

Proposition 25. *Suppose that $T : \mathbf{Set} \rightarrow \mathbf{Set}$ preserves weak pullbacks and Λ is a modal signature containing $\Box$. Then there exists a predicate lifting $\lambda : 2^- \rightarrow 2^{T^-}$ making $\Box$ a necessity operator over T .*

In other words, a necessity operator $\Box$ exists for any endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ that preserves weak-pullbacks. It is the predicate lifting associated with the subset $T1$ of $T2$ as described in [16].

Example 26. Many logics have a necessitation operator. This is evident for extensions of (multi-modal) K . It is easy to see that probabilistic modal logic also has a necessitation operator. In the terminology of Example 5, it can be seen easily that $\Box = L_0$ is a necessitation operator over all extensions of probabilistic modal logic, as long as the latter is interpreted over the structure presented in *loc.cit.* An easy calculation shows that this operator indeed arises from the set $\mathcal{D}(1)$ via the construction described in [16].

In the construction of a satisfiable model for $L \otimes R$ -consistent formulae, we will use necessitation operators to propagate consistency formulae.

Lemma 27 (Necessity operators satisfy necessitation). *Suppose that L is a Λ -logic and $\Box \in \Lambda$ is a necessity operator over a Λ -structure T , $\phi \in \mathcal{L}(\Lambda)$ and $\mathcal{M}$ is a T -model. Then $\mathcal{M} \models \Box\phi$ whenever $\mathcal{M} \models \phi$.*

We now return to the construction of satisfying $S \times T$ -models. We impose $L \otimes R$ -consistency at all relevant points in an L -model of ϕ^L as follows: since s_{s_0} is a $L \otimes R$ -theorem, so is $\Box^{\leq md(\phi^L)} s_{s_0}$ (by Lemma 27) and by using the same reasoning as earlier we can thus build an L -model

$$(W_0, \gamma_0, \sigma_0), w \models \phi^L \wedge \Box^{\leq md^L(\phi)} s_{s_0}^L \quad (1)$$

that satisfies the consistency formula at all points that influence the interpretation of ϕ (see Theorem 23).

4.4 Generated submodels

What do we do with the points of the model in Equation 1 that cannot affect the truth of ϕ at w ? They may not affect ϕ at w but we still need to build an $L \otimes R$ -consistent model. We deal with these points in two ways: first we ensure that our model has no truly excessive points by using generated submodels, secondly we will use non-standard valuations during the construction of the model for points that cannot influence ϕ at w . Only at the last step of the construction will we return to standard (boolean) valuations. Let us first deal with the first point.

Definition 28 (Generated Submodels). Given a coalgebra $W \xrightarrow{\gamma} TW$, we define the *subcoalgebra generated* by $w \in W$ as the set of points reachable via a $\rightsquigarrow$ -trace from w , i.e.

$$\text{Tr}(w) = \{x \in W \mid w \rightsquigarrow x\}$$

together with the map $\delta : \text{Tr}(w) \rightarrow T(\text{Tr}(w)), x \mapsto Ti^{-1}(\gamma(x))$ where i is the injection of $\text{Tr}(w)$ in W .

It follows from Proposition 19 that the above is well-defined. In fact, we can show slightly more:

Proposition 29. *Given a coalgebra $W \xrightarrow{\gamma} TW$ and $w \in W$, $\text{Tr}(w)$ is the smallest subcoalgebra containing w , i.e.*

$$\text{Tr}(w) = \bigcap \{S \subseteq W \mid (S, \delta) \subseteq (W, \gamma) \text{ is a subcoalgebra for some } \delta : S \rightarrow TS\}.$$

Clearly the passage from points in satisfying models to generated submodels does not change the validity of formulae at that point.

Proposition 30. *Let ϕ be a Λ -formula for some modal signature Λ for which we assume a Λ -structure has been fixed and let $\mathcal{M} = (W, \gamma, \sigma)$ be a coalgebraic model. Then for any $w \in W$*

$$\mathcal{M}, w \models \phi \Leftrightarrow \text{Tr}(w), w \models \phi.$$

Once we have pruned our model of irrelevant points by considering only generated submodels we deal with the remaining points that cannot influence ϕ^L by weakening the valuations at these points with a non-standard valuation [11].

4.5 Characteristic sets and formulae

Finally, how do we give the surrogate variables of ϕ^L an interpretation that reflects their index? The idea is to build at each $x \in W_0$ an R -model which will provide an R -interpretation of the index of all surrogate variables true at x . We must therefore ‘sum-up’ all that is true at a certain point.

Definition 31 (Characteristic Sets). Given an L -model $\mathcal{M} = (W, \sigma, \gamma)$ and a set of $L \otimes R$ -formulae Δ , the L -characteristic set $X_L^{\sigma, \Delta}(t)$ at $t \in W$ is

$$X_L^{\sigma, \Delta}(t) = \{\psi \mid \psi \in \Delta \text{ and } \mathcal{M}, t \models \psi^L\} \cup \{\neg\psi \mid \psi \in \Delta \text{ and } \mathcal{M}, t \not\models \psi^L\}$$

The L -characteristic formula $\chi_L^{\sigma, \Delta}(t)$ is defined as:

$$\chi_L^{\sigma, \Delta}(t) = \bigwedge X_L^{\sigma, \Delta}(t)$$

The R -characteristic sets and formulae are defined dually.

It is easy to see that by construction characteristic formulae are consistent and sum-up all the formulae of Δ that are true at a point. By taking Δ to be the index set $\mathbb{S}_0$ we can now give a meaning to the surrogate variables in ϕ^L , for if $(\chi^{\sigma_0, \mathbb{S}_0}(t))^R$ is satisfied at a point x_1^t in an R -model $(W_1^t, \delta_1^t, \sigma_1^t)$, then a surrogate variable q_ψ is true at $t \in W_0$ iff ψ^R is true at $x_1^t \in W_1^t$.

4.6 Transfer of Completeness

We have just seen how we can unravel the R -meaning of surrogate variables of ϕ^L in our original model $(W_0, \gamma_0, \sigma_0)$. The next step in the construction of the model is to perform this operation at every point of W_0 , i.e. for each $t \in W_0$ we build an R -model (a ‘fibre’) of $(\chi^{\sigma_0, \mathbb{S}_0}(t))^R$. We then glue all our models together

by identifying t and x_1^t (which as mentioned above is legitimate from the point of view of truth values). This gives us a model for the first reconstruction of ϕ which contains a host of R -surrogate variables which much be given a proper L -interpretation and the process we have described is thus iterated. By alternating and gluing these $L \otimes R$ -consistent L and R -models in the way we described we eventually reach a model of the final reconstruction of ϕ , i.e. ϕ itself.

Theorem 32 (Completeness Transfer). *Let Λ_L and Λ_R be two modal signatures and let L and R be consistent Λ_L and Λ_R -logics respectively. If both L and R have a necessitation operator over their respective structures, then*

$$L = \text{Log}(\mathcal{C}_L) \text{ and } R = \text{Log}(\mathcal{C}_R) \text{ iff } L \otimes R = \text{Log}(\mathcal{C}_L \otimes \mathcal{C}_R)$$

that is, completeness transfers to the fusion of coalgebraic logics.

Example 33. Given that both **S5** and **Prob** have necessitation over their respective structures (Example 26), we may apply Theorem 32 to show that **S5** $_n \otimes \mathbf{Prob}$ is sound and complete with respect to the class $\bigotimes_{i=1}^n (\mathcal{C}_{\mathbf{S5}})_i \otimes \mathcal{D}$.

Note that if L and R have the finite model property, then at any stage of our construction, the L - and R -models can be chosen to be finite, and since the total number of steps in the construction is finite (bounded by the modal depth of the formula), the final model is also finite.

Theorem 34. *Under the assumptions of Theorem 32, the finite model property transfers.*

5 Discussion and future work

Several questions emerge from our generalisation of Kracht and Wolter's construction. Firstly, can we drop the assumption that functors need to preserve all intersections? As we mentioned above, the case of logics with the finite model property offers a partial solution to this problem which includes any complete rank-1 logic, and in particular the classical logic **E** for which the completeness transfer problem is still open.

Secondly, in order to deal with logics whose semantics is given by non weak-pullback preserving functors we cannot use notions of necessity that are given by simply applying a unary necessity operator to formulae. Instead we may need more complex formulae, i.e. our general notion of necessity (Definition 24). But can the formula $\text{nec}(\phi)$ be constructed or be proven to exist in general, and if not what are the restrictions that prevent it from happening? Less ambitiously, but perhaps more realistically, is there a generalization of our unary necessity operators to the n -ary case for weak-pullback preserving functors?

Finally, is the Kracht and Wolter construction necessary at all? Could some duality argument interpret the syntactic fusion as a binary operation on algebras/theories whose dual would be the semantic fusion operation on (classes of) coalgebras/models. Syntactically, there is a strong connection between the

fusion and co-product constructions. The signature of a fusion can be seen as the co-product $\Lambda_L + \Lambda_R$ of its constituent signatures. Moreover, if we view the algebraic models of a Λ_L -logic L as models of a Lawvere theory T_L (the theory of Λ_L modal algebras) and the algebraic models of a Λ_R -logic R as models of a Lawvere theory T_R , then the models of $L \otimes R$ are models of the pushout $T_L \xleftarrow{i_1} T_B \xrightarrow{i_2} T_R$ where T_B is the Lawvere theory of boolean algebras and i_1, i_2 are just the inclusion as boolean reducts. Dually, the fusion of models is based on products, but the correct categorical framework in which to view the operation of fusion on classes of coalgebras as a kind of product or pullback is not clear.

References

1. F. Baader, S. Ghilardi, and C. Tinelli. A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. *Information and Computation*, 204(10):1413–1452, 2006.
2. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.
3. B. Chellas. *Modal Logic*. Cambridge, 1980.
4. C. Cirstea and D. Pattinson. Modular proof systems for coalgebraic logics. *Theoret. Comput. Sci.*, 388:83–108, 2007.
5. K. Došen. Duality between modal algebras and neighbourhood frames. *Studia Logica*, 48(2):219–234, June 1989.
6. R. Fagin and J. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41:340–367, 1994.
7. R. Fajardo and M. Finger. How not to combine modal logics. In B. Prasad, editor, *IJCAI*, pages 1629–1647. IJCAI, 2005.
8. K. Fine and G. Schurz. Transfer theorems for stratified modal logics. In J. Copeland, editor, *Logic and Reality, Essays in Pure and Applied Logic. In Memory of Arthur Prior*, pages 169–213. Oxford University Press, 1996.
9. D. M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyashev. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2003.
10. H. P. Gumm. From t-coalgebras to filter structures and transition systems. In *Algebra and Coalgebra in Computer Science*, pages 194–212. Springer, 2005.
11. M. Kracht and F. Wolter. Properties of independently axiomatizable bimodal logics. *J. Symb. Logic*, 56(4):1469–1485, 1991.
12. C. Kupke, A. Kurz, and Y. Venema. Stone coalgebras. *Theor. Comput. Sci.*, 327(1-2):109–134, 2004.
13. A. Kurucz. Combining modal logics. In J. V. B. Patrick Blackburn and F. Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 869 – 924. Elsevier, 2007.
14. L. Schröder and D. Pattinson. Rank-1 logics are coalgebraic. *Journal of Logic and Computation*, 20(5):1113–1147, 2010.
15. L. Schröder and D. Pattinson. Modular algorithms for heterogeneous modal logics via multi-sorted coalgebra. *Math. Struct. Comput. Sci.*, 21, 2011.
16. L. Schröder. Expressivity of coalgebraic modal logic: The limits and beyond. *Theor. Comput. Sci.*, 390(2-3):230–247, 2008.
17. F. Wolter. Fusions of modal logics revisited. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyashev, editors, *Advances in Modal Logic 1096*, CSLI Lecture Notes, pages 361–379. CSLI Publications, 1998.