

On Constructing Fibred Tableaux for BDI Logics

Vineet Padmanabhan & Guido Governatori

School of ITEE, The University of Queensland, Australia
[vnair, guido]@itee.uq.edu.au

Abstract. In [11,13] we showed how to combine propositional BDI logics using Gabbay's *fibring* methodology. In this paper we extend the above mentioned works by providing a tableau-based decision procedure for the combined/fibred logics. We show how to uniformly construct a tableau calculus for the combined logic using Governatori's labelled tableau system **KEM**.

1 Introduction

The BDI model [15] is a rich and powerful logical framework developed in the early 90's focusing on three components of an agent: beliefs, desires and intentions. Its language is a propositional modal language with three families of modal operators $B_i, D_i, I_i, i \in Agents$. The logic **KD45** of modal logic is used for the belief operator B and the logic **KD** for D and I respectively. The semantics is given through standard Kripke frames by defining three families of accessibility relations, (b_i, d_i, i_i) for belief, desire and intention. The binary relations b_i are Euclidean, transitive and serial whereas the relations for intention i_i and desires d_i are serial. In addition to the above representation, the BDI framework impose constraints on beliefs, desires and intentions in the form of *interaction axioms* like, $I(\varphi) \rightarrow D(\varphi)$, $D(\varphi) \rightarrow B(\varphi)$: intentions are stronger than desires and desires than beliefs. Hence the basic BDI logic \mathbb{L} is the combination of different component logics plus the two interaction axioms as given below

$$\mathbb{L} \equiv (\otimes_{i=1}^n \mathbf{KD45}_{B_i}) \otimes (\otimes_{i=1}^n \mathbf{KD}_{D_i}) \otimes (\otimes_{i=1}^n \mathbf{KD}_{I_i}) + \{I_i\varphi \rightarrow D_i\varphi\} + \{D_i\varphi \rightarrow B_i\varphi\} \quad (1)$$

Any BDI theory, or for that matter any fully-fledged Multi-Agent-System (MAS) theory, modelling rational agents consists of a combined system of logic of beliefs, desires, goals and intentions as mentioned above. They are basically well understood standard modal logics *combined together* to model different facets of the agents. A number of researchers have provided such combined systems for different reasons and different applications. However, investigations into a general methodology for combining the different logics involved has been mainly neglected to a large extent. Recently [11,13] *fibring/dovetailing* [7] has been advanced as a semantic methodology to characterise BDI logics. Here we extend our previous work to provide a tableau decision procedure for the fibred logic based on the labelled tableau system **KEM** [9,8,1].

The key feature of our tableau system is that it is neither based on resolution nor on standard sequent/tableau techniques. It combines linear tableau expansion rules with natural deduction rules and an analytic version of the cut rule. The tableau rules are supplemented with a powerful and flexible label algebra that allows the system to deal

with a large class of intensional logics admitting possible world semantics (non-normal modal logic [10], multi-modal logics [9] and conditional logics [2]). The algebra simulates the possible world semantics and it is very strongly related with fibring [8].

As far as the field of *combining logics* is concerned, it has been an active research area for some time now and powerful results about the preservation of important properties of the logics being combined has been obtained [12,4,16,17]. Also, investigations related to using fibring as a combining technique in various domains has produced a wealth of results [7,18,5]. The novelty of combining logics is the aim to develop *general techniques* that allow us to produce combinations of *existing* and well understood logics. Such general techniques are needed for formalising complex systems in a systematic way. Such a methodology can help decompose the problem of designing a complex system into developing components (logics) and combining them.

The next section provides a brief introduction to the technique of fibring. Section 3 is divided into several sections describing the label formalism, unification, inference rules and proof search of the **KEM** tableau system, all within the context of a Multi-Agent-Scenario. The paper concludes with some final remarks.

2 Fibring BDI Logics

The basic BDI logic \mathbb{L} given in (1) is defined from three component logics, **KD45_n** for belief, and **KD_n** for desires and intentions. For sake of clarity, consider just two of the component logics, $\nabla_1(\mathbf{KD45})$ and $\nabla_2(\mathbf{KD})$ and their corresponding languages $\mathcal{L}_{\nabla_1}, \mathcal{L}_{\nabla_2}$ built from the respective sets \mathfrak{A}_1 and \mathfrak{A}_2 of atoms having classes of models $\mathfrak{M}_{\nabla_1}, \mathfrak{M}_{\nabla_2}$ and satisfaction relations \models_1 and \models_2 . Hence we are dealing with two different systems S_1 and S_2 characterised, respectively, by the class of Kripke models \mathcal{K}_1 and \mathcal{K}_2 . For instance, we know how to evaluate $\Box_1\varphi$ ($B(\varphi)$) in \mathcal{K}_1 (**KD45**) and $\Box_2\varphi$ ($D(\varphi)$) in \mathcal{K}_2 (**KD**). We need a method for evaluating \Box_1 (resp. \Box_2) with respect to \mathcal{K}_2 (resp. \mathcal{K}_1). To do so, we link (fibre), via a *fibring* function the model for ∇_1 with a model for ∇_2 and build a fibred model of the combination. The fibring function evaluates (give yes/no) answers with respect to a modality in S_2 , being in S_1 and vice versa. The interpretation of a formula φ of the combined language in the fibred model at a state w can be given as $w \models \varphi$ if and only if $\mathfrak{f}(w) \models^* \varphi$, where \mathfrak{f} is a fibring function mapping a world to a model *suitable for interpreting φ* and \models^* is the corresponding satisfaction relation (\models_1 for ∇_1 or \models_2 for ∇_2). For example, let ∇_1, ∇_2 be two modal logics as given above and $\varphi = \Box_1\Diamond_2p_0$ be a formula on a world w_0 of the fibred semantics. φ belongs to the language $\mathcal{L}_{(1,2)}$ as the outer connective (\Box_1) belongs to the language \mathcal{L}_1 and the inner connective (\Diamond_2) belongs to the language \mathcal{L}_2 . By the standard definition we start evaluating \Box_1 of $\Box_1\Diamond_2$ at w_0 . According to the standard definition we have to check whether \Diamond_2p_0 is true at every w_1 accessible from w_0 since from the point of view of \mathcal{L}_1 this formula has the form \Box_1p (where $p = \Diamond_2p_0$ is atomic). But at w_1 we cannot interpret the operator \Diamond_2 , because we are in a model of ∇_1 , not of ∇_2 . To evaluate this we need the fibring function \mathfrak{f} which at w_1 points to a world v_0 , a world in a model suitable to interpret formulae from ∇_2 . Now all we have to check is whether \Diamond_2p_0 , is true at v_0 in this last model and this can be done in the usual way. Hence the fibred semantics for the combined language $\mathcal{L}_{(1,2)}$ has models of the form $(\mathcal{F}_1, w_1, v_1, \mathfrak{f}_1)$,

where $\mathcal{F}_1 = (W_1, R_1)$ is a frame, and \mathfrak{F}_1 is the fibring function which associates a model \mathfrak{M}_w^2 from \mathcal{L}_2 with w in \mathcal{L}_1 i.e. $\mathfrak{F}_1(w) = \mathfrak{M}_w^2$.

Let \mathbf{I} be a set of labels representing the modal operators for the intentional states (belief, desire, intention) for a set of agents, and $\nabla_i, i \in \mathbf{I}$ be modal logics whose respective modalities are $\Box_i, i \in \mathbf{I}$.

Definition 1 [7] A fibred model is a structure $(W, S, R, \mathbf{a}, \mathbf{v}, \tau, \mathbf{F})$ where

- W is a set of possible worlds;
- S is a function giving for each w a set of possible worlds, $S^w \subseteq W$;
- R is a function giving for each w , a relation $R^w \subseteq S^w \times S^w$;
- \mathbf{a} is a function giving the actual world \mathbf{a}^w of the model labelled by w ;
- \mathbf{v} is an assignment function $\mathbf{v}^w(q_0) \subseteq S^w$, for each atomic q_0 ;
- τ is the semantical identifying function $\tau : W \rightarrow \mathbf{I}$. $\tau(w) = i$ means that the model $(S^w, R^w, \mathbf{a}^w, \mathbf{v}^w)$ is a model in \mathcal{K}_i , we use W_i to denote the set of worlds of type i ;
- \mathbf{F} , is the set of fibring functions $\mathfrak{F} : \mathbf{I} \times W \mapsto W$. A fibring function \mathfrak{F} is a function giving for each i and each $w \in W$ another point (actual world) in W as follows:

$$\mathfrak{F}_i(w) = \begin{cases} w & \text{if } w \in S^{\mathfrak{M}} \text{ and } \mathfrak{M} \in \mathcal{K}_i \\ \text{a value in } W_i, & \text{otherwise} \end{cases}$$

such that if $w \neq w'$ then $\mathfrak{F}_i(w) \neq \mathfrak{F}_i(w')$. It should be noted that fibring happens when $\tau(w) \neq i$. Satisfaction is defined as follows with the usual boolean connections:

$$\begin{aligned} w \models q_0 & \text{ iff } \mathbf{v}(w, q_0) = 1, \text{ where } q_0 \text{ is an atom} \\ w \models \Box_i \varphi & \text{ iff } \begin{cases} w \in \mathfrak{M} \text{ and } \mathfrak{M} \in \mathcal{K}_i \text{ and } \forall w' (wRw' \rightarrow w' \models \varphi), \text{ or} \\ w \in \mathfrak{M}, \text{ and } \mathfrak{M} \notin \mathcal{K}_i \text{ and } \forall \mathfrak{F} \in \mathbf{F}, \mathfrak{F}_i(w) \models \Box_i \varphi. \end{cases} \end{aligned}$$

We say the model satisfies φ iff $w_0 \models \varphi$.

A fibred model for $\nabla_{\mathbf{I}}^{\mathfrak{F}}$ can be generated from fibring the semantics for the modal logics $\nabla_i, i \in \mathbf{I}$. The detailed construction is given in [13]. Also, to accommodate the interaction axioms specific constraints need to be given on the fibring function. In [11] we outline the specific conditions required on the fibring function to accommodate axiom schemas of the type $G^{a,b,c,d} (\Diamond_a \Box_b \varphi \rightarrow \Box_c \Diamond_d \varphi)$. We do not want to get into the details here as the main theme of this paper is with regard to tableau decision procedures for fibred logics. Notice, however, that the fibring construction given in [11,13] works for normal (multi-)modal logics as well as non-normal modal logics.

3 Labelled Tableau for Fibred BDI Logic

In this section we show how to adapt **KEM**, a labelled modal tableau system, to deal with the fibred combination of BDI logics. A tableau system is a semantic based refutation method that systematically tries to build a (counter-) model for a set of formulas. A failed attempt to refute (invalidate) a set of formulas generates a model where the set of formulas is true. To show that a property P follows from a theory (set of formulas/axioms) A_1, \dots, A_n we verify whether a model for $\{A_1, \dots, A_n, \neg P\}$ exists. If it does not then P is a consequence of the theory.

In labelled tableau systems, the object language is supplemented by labels meant to represent semantic structures (possible worlds in the case of modal logics). Thus the formulas of a labelled tableau system are expressions of the form $A : i$, where A is a formula of the logic and i is a label. The intuitive interpretation of $A : i$ is that A is true at (the possible world(s) denoted by) i . **KEM**'s inferential engine is based on a combination of standard tableau linear expansion rules and natural deduction rules supplemented by an analytic version of the cut rule. In addition it utilises a sophisticated but powerful label formalism that enables the logic to deal with a large class of modal and non-classical logics. Furthermore the label mechanism corresponds to fibring and thus it is possible to define tableau systems for multi-modal logic by a seamless combination of the (sub)tableaux systems for the component logics of the combination. It is not possible in this paper to give a full presentation of **KEM** for fully fledged BDI logic supplemented with interaction axioms as given in (1) (for a comprehensive presentation see [8]). Accordingly we will limit ourselves to a single modal operator for each agent and show how to characterise the axioms corresponding to each individual agent as well as the interactions between different agents with the help of an example.

3.1 Label Formalism

KEM uses *Labelled Formulas* (L -formulas for short), where an L -formula is an expression of the form $A : i$, where A is a wff of the logic, and i is a label. For fibred BDI logic we need to have labels for various modalities (belief, desire, intention) for each agent. However, as we have just explained we will consider only one modality and thus will have only labels for the agents. The set of atomic labels, \mathfrak{S}_1 , is then given as $\mathfrak{S}_1 = \bigcup_{i \in \text{Agt}} \Phi^i$, where Agt is the set of agents. Every Φ^i is partitioned into (non-empty) sets of variables and constants: $\Phi^i = \Phi_V^i \cup \Phi_C^i$ where $\Phi_V^i = \{W_1^i, W_2^i, \dots\}$ and $\Phi_C^i = \{w_1^i, w_2^i, \dots\}$. Φ_C and Φ_V denote the set of constants and the set of variables. We also add a set of auxiliary un-indexed atomic labels, $\Phi^A = \Phi_V^A \cup \Phi_C^A$ where $\Phi_V^A = \{W_1, W_2, \dots\}$ and $\Phi_C^A = \{w_1, w_2, \dots\}$, that will be used in unifications and proofs. A label $u \in \mathfrak{S}_1$ is either (i) an element of the set Φ_C , or (ii) an element of the set Φ_V , or (iii) a path term (u', u) where (iiia) $u' \in \Phi_C \cup \Phi_V$ and (iiib) $u \in \Phi_C$ or $u = (v', v)$ where (v', v) is a label. As an intuitive explanation, we may think of a label $u \in \Phi_C$ as denoting a world (a *given* one), and a label $u \in \Phi_V$ as denoting a set of worlds (*any* world) in some Kripke model. A label $u = (v', v)$ may be viewed as representing a path from v to a (set of) world(s) v' accessible from v (the world(s) denoted by v). For any label $u = (v', v)$ we shall call v' the *head* of u , v the *body* of u , and denote them by $h(u)$ and $b(u)$ respectively. If $b(u)$ denotes the body of u , then $b(b(u))$ will denote the body of $b(u)$, and so on. We call each of $b(u)$, $b(b(u))$, etc., a *segment* of u . The length of a label u , $\ell(u)$, is the number of atomic labels in it. $s^n(u)$ will denote the segment of u of length n and we shall use $h^n(u)$ as an abbreviation for $h(s^n(u))$. Notice that $h(u) = h^{\ell(u)}(u)$. Let u be a label and u' an atomic label. For any label u , $\ell(u) > n$, we define the *counter-segment- n* of u , as follows (for $n < k < \ell(u)$):

$$c^n(u) = h(u) \times (\dots \times (h^k(u) \times (\dots \times (h^{n+1}(u), w_0))))$$

where w_0 is a dummy label, i.e., a label not appearing in u (the context in which such a notion occurs will tell us what w_0 stands for). The counter-segment- n defines what

remains of a given label after having identified the segment of length n with a ‘dummy’ label w_0 . In the context of fibring w_0 can be thought of as denoting the actual world obtained via the fibring function from the world denoted by $s^n(u)$.

So far we have provided definitions about the structure of the labels without regard to the elements they are made of. The following definitions will be concerned with the type of world symbols occurring in a label. We say that a label u is *i-preferred* iff $h(u) \in \Phi^i$; a label u is *i-pure* iff each segment of u of length $n > 1$ is *i-preferred*.

3.2 Label Unifications

The basic mechanism of **KEM** is its logic dependent label unification. In the same way as each modal logic is characterised by a combination of modal axioms (or semantic conditions on the model), **KEM** defines a unification for each modality and axiom/semantic condition and then combines them in a recursive and modular way. In particular we use what we call unification to determine whether two labels can be mapped to the same possible world in the possible worlds semantics. The second key issue is the ability to split labels and to work with parts of labels. The mechanism permits the encapsulation of operations on sub-labels. This is an important feature that, in the present context, allows us to correlate unifications and fibring functions. Given the modularity of the approach the first step of the construction is to define unifications (pattern matching for labels) corresponding to the single modality in the logic we want to study. Every unification is built from a basic unification defined in terms of a substitution $\rho : \mathfrak{S}_1 \mapsto \mathfrak{S}$ such that: $\Phi_V^i \mapsto \mathfrak{S}^i$ for every $i \in \text{Agt}$, $\Phi_V^A \mapsto \mathfrak{S}$. Accordingly we have that two atomic labels u and v σ -unify iff there is a substitution ρ such that $\rho(u) = \rho(v)$. We shall use $[u; v]\sigma$ both to indicate that there is a substitution ρ for u and v , and the result of the substitution. The σ -unification is then extended to composite labels:

$$[i; j]\sigma = k \text{ iff } \exists \rho : h(k) = \rho(h(i)) = \rho(h(j)) \text{ and } b(k) = [b(i); b(j)]\sigma$$

Clearly σ is symmetric, i.e., $[u; v]\sigma$ iff $[v; u]\sigma$. Moreover this definition offers a flexible and powerful mechanism: it allows for an independent computation of the elements of the result of the unification.

3.3 An Example (Friends Puzzle [3])

Consider the agents Peter, John and Wendy with modalities \Box_p, \Box_j , and \Box_w . John and Peter have an *appointment*. Suppose that Peter knows the *time* of appointment. Peter knows that John knows the *place* of their appointment. Wendy knows that if Peter knows the *time* of appointment, then John knows that too (since John and Peter are friends). Peter knows that if John knows the *place* and the *time* of their appointment, then John knows that he has an *appointment*. Peter and John satisfy the axioms T and 4. Also, if Wendy knows something then Peter knows the same thing (suppose Wendy is Peter’s wife) and if Peter knows that John knows something then John knows that Peter knows the same thing. The Knowledge/belief base of the example is given in Fig.1. So we have a modal language consisting of three modalities \Box_p, \Box_j and \Box_w denoting respectively the agents Peter, John and Wendy and characterised by the set $A = \{A_i \mid i = 1, \dots, 6\}$ of interaction axioms. Suppose now that one wants to show that *each of the friends knows that the other one knows that he has an appointment*, i.e.,

$$\Box_j \Box_p \text{appointment} \wedge \Box_p \Box_j \text{appointment} \quad (2)$$

1. $\Box_p time$	$A_1 T_p : \Box_p \varphi \rightarrow \varphi$
2. $\Box_p \Box_j place$	$A_2 4_p : \Box_p \varphi \rightarrow \Box_p \Box_p \varphi$
3. $\Box_w (\Box_p time \rightarrow \Box_j time)$	$A_3 T_j : \Box_j \varphi \rightarrow \varphi$
4. $\Box_p \Box_j (place \wedge time \rightarrow appt)$	$A_4 4_j : \Box_j \varphi \rightarrow \Box_j \Box_j \varphi$
	$A_5 I_{wp} : \Box_w \varphi \rightarrow \Box_p \varphi$
	$A_6 S_{pj} : \Box_p \Box_j \varphi \rightarrow \Box_j \Box_p \varphi$

Fig. 1. Knowledge base related to the Friend's puzzle.

In other words one want to show that (2) is a theorem of the knowledge-base. The tableaux proof of (2) using the **KEM** tableau procedure is given in Fig.3. But before getting into the proof details we should understand how the label unification introduced in the previous section works for the modal operators, \Box_w , \Box_j and \Box_p . We can capture the relationship between \Box_w and \Box_p by extending the substitution ρ by allowing a variable of type w to be mapped to labels of the same type and of type p .

$$\rho^w(W^w) \in \mathfrak{S}^w \cup \mathfrak{S}^p$$

Then the unification σ^w is obtained from the basic unification σ by replacing ρ with the extended substitution ρ^w . This procedure must be applied to all pairs of modalities \Box_1, \Box_2 related by the interaction axiom $\Box_1 \varphi \rightarrow \Box_2 \varphi$. For the unifications for \Box_p and \Box_j (σ^p and σ^j) we assume that the labels involved are i -pure. First we notice that these two modal operators are **S4** modalities and thus have to use the unification for this logic.

$$[u; v] \sigma^{S4} = \begin{cases} [u; v] \sigma^D & \text{if } \ell(u) = \ell(v) \\ [u; v] \sigma^T & \text{if } \ell(u) < \ell(v), h(u) \in \Phi_C \\ [u; v] \sigma^4 & \text{if } \ell(u) < \ell(v), h(u) \in \Phi_V \end{cases} \quad (3)$$

It is worth noting that the conditions on axiom unifications are needed in order to provide a deterministic unification procedure. The σ^T and σ^4 are defined as follows: (It should be noted that for the rest of the unifications, given two labels u and v we will assume that $\ell(u) > \ell(v)$. The conditions for $\ell(v) > \ell(u)$ are symmetric). Thus,

$$\begin{aligned} [u; v] \sigma^T &= [s^{\ell(v)}(u); v] \sigma \text{ if } \forall n \geq \ell(v), [h^n(u); h(v)] \sigma = [h(u); h(v)] \sigma \\ [u; v] \sigma^4 &= c^{\ell(v)}(u) \text{ if } h(v) \in \Phi_V \text{ and } w_0 = [s^{\ell(v)}(u); v] \sigma \end{aligned}$$

σ^T allows us to unify two labels such that the segment of the longest with the length of the other label and the other label unify, provided that all remaining elements of the longest have a common unification with the head of the shortest. For example let $u = (w_3, (W_2, (w_2, w_1)))$ and $v = (w_3, (W_1, w_1))$. Here $[W_2; w_3] \sigma = [w_3; w_3] \sigma$ and the two labels σ^T -unify to $(w_3, (w_2, w_1))$. This means that after a given point the head of the shortest is always included in its extension, and thus it is accessible from itself, and consequently we have reflexivity. For σ^4 we have that the shortest label unifies with the segment with the same length of the longest and that the head of the shortest is variable. A variable stands for all worlds accessible from its predecessor. Thus, given transitivity, every element extending the segment with length of the shortest is accessible from this

point. Then a unification corresponding to axiom A_6 from Figure 1 is

$$[u; v]\sigma^{S_{p,j}} = \begin{cases} c^{m+n}(v) \text{ if } h(u) \in \Phi_V^j \text{ and } c^n(v) \text{ is } p\text{-pure, and} \\ h^{\ell(u)-1}(u) \in \Phi_V^p \text{ and } c^n(v) \text{ is } j\text{-pure, and} \\ w_0 = [s^{\ell(u)-2}(u); s^m(v)]\sigma \end{cases}$$

This unification allows us to unify two labels when in one label we have a sequence of a variable of type p followed by a variable of type j and a label where we have a sequence of labels of type j followed by a sequence of labels of type p . The unification for \Box_p and \Box_j are just the combination of the three unifications given above. Finally the unification for the logic \mathbf{L} defined by the axioms $A_1 - A_6$ is obtained from the following recursive unification

$$[u; v]\sigma_L = \begin{cases} [u; v]\sigma^{w,p,j} \\ [c^m(u); c^n(v)]\sigma^{w,p,j} \text{ where } w_0 = [s^m(u); s^n(v)]\sigma_L \end{cases}$$

$\sigma^{w,p,j}$ is the simple combination of the unifications for the three modal operators. Having accounted for the label unification we now give the inference rules used in **KEM** proofs.

3.4 Inference Rules

For the inference rules we use the Smullyan-Fitting unifying notation [6] (Figure.2). The α -rules are just the familiar linear branch-expansion rules of the tableau method. The β -rules are nothing but natural inference patterns such as Modus Ponens, Modus Tollens and Disjunctive syllogism generalised to the modal case. To apply such rules the labels of the premises must unify and the label of the conclusion is the result of the unification. The ν and π rules are the normal expansion rules for modal operators of

$\frac{\alpha : u}{\alpha_1 : u} (\alpha)$	$\frac{\beta : u \quad \beta_i^c : v \quad (i = 1, 2)}{\beta_{3-i} : [u; v]\sigma} (\beta)$	$\frac{\nu^i : u}{\nu_0^i : (W_n^i, u) \clubsuit} (\nu) \quad \frac{\pi^i : u}{\pi_0^i : (w_n^i, u) \clubsuit} (\pi)$
$\alpha_2 : u$	$A : u$	
$\frac{}{A : u \mid \neg A : u} (\text{PB})$	$\frac{\neg A : v}{\times} [\text{if } [u; v]\sigma] (\text{PNC})$	
		$(\clubsuit) W_n^i, w_n^i \text{ are new labels.}$

Fig. 2. Inference Rules of **KEM** using the Smullyan-Fitting Notation

labelled tableaux with free variable. The intuition for the ν rule is that if $\Box_i A$ is true at u , then A is true at all worlds accessible via R_i from u , and this is the interpretation of the label (W_n^i, u) ; similarly if $\Box_i A$ is false at u (i.e., $\neg \Box_i A$ is true), then there must be a world, let us say w_n^i accessible from u , where $\neg A$ is true. A similar intuition holds when u is not i -preferred, but the only difference is that we have to make use of the fibring function instead of the accessibility relation. PB represents the semantic counterpart of the cut rule of the sequent calculus (intuitive meaning: a formula A is either true or false in

any given world). Accordingly it is a zero-premise inference rule, so in its unrestricted version can be applied whenever we like. However, we impose a restriction on its application. PB can be only applied w.r.t. immediate sub-formulas of unanalysed β -formulas, that is β formulas for which we have no immediate sub-formulas with the appropriate labels in the tree. PNC states that two labelled formulas are σ_L -complementary when the two formulas are complementary and their labels σ_L -unify.

3.5 Proof Search

Let $\Gamma = \{X_1, \dots, X_m\}$ be a set of formulas. Then \mathcal{T} is a **KEM-tree** for Γ if there exists a finite sequence $(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n)$ such that (i) \mathcal{T}_1 is a 1-branch tree consisting of $\{X_1 : t_1, \dots, X_m : t_m\}$; (ii) $\mathcal{T}_n = \mathcal{T}$, and (iii) for each $i < n$, \mathcal{T}_{i+1} results from \mathcal{T}_i by an application of a rule of **KEM**. A branch θ of a **KEM-tree** \mathcal{T} of L -formulas is said to be σ_L -closed if it ends with an application of *PNC*, open otherwise. As usual with tableau methods, a set Γ of formulas is checked for consistency by constructing a **KEM-tree** for Γ . Moreover we say that a formula A is a **KEM-consequence** of Γ ($\Gamma \vdash_{\mathbf{KEM}(L)} A$) if a **KEM-tree** for $\{X_1 : u_1, \dots, X_n : u_n, \neg A : v\}$ is closed using the unification for the logic L , where $v \in \Phi_C^A$, and $u_i \in \Phi_V^A$. The intuition behind this definition is that A is a consequence of Γ when we take Γ as a set of global assumptions [6], i.e., true in every world in a Kripke model.

We now describe a systematic procedure for **KEM**. First we define the following notions. Given a branch θ of a **KEM-tree**, we shall call an L -formula $X : u$ *E-analysed* in θ if either (i) X is of type α and both $\alpha_1 : t$ and $\alpha_2 : u$ occur in θ ; or (ii) X is of type β and one of the following conditions is satisfied: (a) if $\beta_1^C : v$ occurs in θ and $[u; v]\sigma$, then also $\beta_2 : [u; v]\sigma$ occurs in θ , (b) if $\beta_2^C : v$ occurs in θ and $[u; v]\sigma$, then also $\beta_1 : [u; v]\sigma$ occurs in θ ; or (iii) X is of type μ and $\mu_0 : (u', u)$ occurs in θ for some appropriate u' of the right type, not previously occurring in θ . We shall call a branch θ of a **KEM-tree** *E-completed* if every L -formula in it is *E-analysed* and it contains no complementary formulas which are not σ_L -complementary. We shall say a branch θ of a **KEM-tree** *completed* if it is *E-completed* and all the L -formulas of type β in it are either analysed or cannot be analysed. We shall call a **KEM-tree** *completed* if every branch is completed.

The following procedure starts from the 1-branch, 1-node tree consisting of $\{X_1 : u, \dots, X_m : v\}$ and applies the inference rules until the resulting **KEM-tree** is either closed or completed. At each stage of proof search (i) we choose an open non completed branch θ . If θ is not *E-completed*, then (ii) we apply the 1-premise rules until θ becomes *E-completed*. If the resulting branch θ' is neither closed nor completed, then (iii) we apply the 2-premise rules until θ becomes *E-completed*. If the resulting branch θ' is neither closed nor completed, then (iv) we choose an L -formula of type β which is not yet analysed in the branch and apply *PB* so that the resulting *LS*-formulas are $\beta_1 : u'$ and $\beta_1^C : u'$ (or, equivalently $\beta_2 : u'$ and $\beta_2^C : u'$), where $u = u'$ if u is restricted (and already occurring when $h(u) \in \Phi_C$), otherwise u' is obtained from u by instantiating $h(u)$ to a constant not occurring in u ; (v) (“Modal *PB*”) if the branch is not *E-completed* nor closed, because of complementary formulas which are not σ_L -complementary, then we have to see whether a restricted label unifying with both the labels of the complementary formulas occurs previously in the branch; if such a label exists, or can be built using

already existing labels and the unification rules, then the branch is closed, (vi) we repeat the procedure in each branch generated by *PB*. Fig.3. shows a **KEM** tableaux proof using the inference rules in Fig.2. and following the proof search mentioned above to solve the first conjunct of (2).

1. $\mathbf{F}\Box_j\Box_p\text{appt}$	w_0	9. $\mathbf{T}(\text{place} \wedge \text{time} \rightarrow \text{appt})$	(W_1^j, W_1^p, w_0)
2. $\mathbf{T}\Box_p\Box_j(\text{place} \wedge \text{time} \rightarrow \text{appt})$	W_0	10. $\mathbf{F}\text{place} \wedge \text{time}$	(w_1^p, w_1^j, w_0)
3. $\mathbf{T}\Box_w(\Box_p\text{time} \rightarrow \Box_j\text{time})$	W_0	11. $\mathbf{T}\Box_p\text{time} \rightarrow \Box_j\text{time}$	(W_1^w, w_0)
4. $\mathbf{T}\Box_p\Box_j\text{place}$	W_0	12. $\mathbf{T}\Box_j\text{place}$	(W_2^p, w_0)
5. $\mathbf{T}\Box_p\text{time}$	W_0	13. $\mathbf{T}\text{place}$	(W_2^j, W_2^p, w_0)
6. $\mathbf{F}\Box_p\text{appt}$	(w_1^j, w_0)	14. $\mathbf{F}\text{time}$	(w_1^p, w_1^j, w_0)
7. $\mathbf{F}\text{appt}$	(w_1^p, w_1^j, w_0)	15. $\mathbf{T}\Box_p\text{time}$	(w_1^j, w_0)
8. $\mathbf{T}\Box_j(\text{place} \wedge \text{time} \rightarrow \text{appt})$	(W_1^p, w_0)	16. $\mathbf{T}\text{time}$	(W_3^p, w_1^j, w_0)
		\times	

Fig. 3. Proof of $\Box_j\Box_p$ using **KEM** representation.

The proof goes as follows; 1. is the negation of the formula to be proved. The formulas in 2–5 are the global assumptions of the scenario and accordingly they must hold in every world of every model for it. Hence we label them with a variable W_0 that can unify with every other label. This is used to derive 12. from 11. and 5. using a β -rule, and for introducing 15.; 6. is from 1., and 7. from 6. by applying π rule. Similarly we get 8. from 2., 9. from 8. using \vee rule. 10. comes from 9. and 7. through the use of modus tollens. Applying \vee rule twice we can derive 11. from 3. as well as 13. from 12. Through propositional reasoning we get 14. from 10. and by a further use of \vee rule on 15. we get 16. (14. and 16.) are complementary formulas indicating a contradiction and this results in a closed tableaux because the labels in 14. and 16. unify, denoting that the contradiction holds *in the same world*.

4 Conclusions and Related Work

In this paper we have argued that BDI logics can be explained in terms of fibring as a combination of simpler modal logics. We then outlined a decision procedure based on the **KEM** tableaux system showing the correlation between **KEM** unification and fibring. To evaluate the features of our tableau system we demonstrated how it can deal with a multi-agent scenario like the Friend's puzzle.

Elsewhere [14] we have shown why other labelled tableaux approaches (both *path* & *graph*) are not suited for fibring. The path approach (where prefixes are sequences of integers representing a world as a path in the model connecting the initial world to the one at hand) requires the definition of new inference rules for each logic with a simple labelling mechanism. It is not clear how the path-based approach can be extended to more complex cases of fibring, for example when we consider non-normal modal operators for the mental attitudes of the agents. The graph approach (where accessibility relations are given explicitly) on the other hand does not require, in general, any new rule, since it uses the semantic structure to propagate formulas to the appropriate labels. It is then suitable for an approach based on fibring, since the relationships between

two labels can be given in terms of fibring. However, when the structure of the model is more complicated (for example when the models for the logics are given in terms of neighbourhood models) then the approach might not be applicable since it assumes relationships between labels/worlds in a model and not more complex structures. In addition, the system does not give a decision procedure unless the relationships among labels are restricted to decidable fragments of first-order logic. Thus it is not possible to represent logic that are not first-order definable and the designer of an agent logic has to verify that she is operating within a decidable fragment of first order logic.

KEM, similar to the graph approach, does not need logic dependent rules, however, similar to the path approach, it needs logic dependent label unifications. The label algebra can be seen as a form of fibring [8], thus simple fibring does not require special attention; therefore **KEM** allows for a seamless composition of (sub)tableaux for modal logics. The label algebra contrary to the graph reasoning mechanism is not based on first order logic and thus can deal with complex structure and is not limited to particular fragments. Indeed **KEM** has been used with complex label schema for non-normal modal logics in a uniform way [10] as well as other intensional logics such as conditional logics [2]. For these reasons we believe that **KEM** offers a suitable framework for decision procedure for multi-modal logic for MAS. As we only described the static fragment of BDI logics (no temporal evolution was considered) the future work is to extend the tableaux framework so as to accommodate temporal modalities.

Acknowledgements

This work was supported by the Australian Research Council (ARC) under Discovery Project No. DP0452628 on "Combining Modal Logics for Dynamic and Multi-Agent Systems".

References

1. A. Artosi, P. Benassi, G. Governatori, and A. Rotolo. Shakespearian modal logic: A labelled treatment of modal identity. In *Advances in Modal Logic*, volume 1. CSLI, 1998.
2. A. Artosi, G. Governatori, and A. Rotolo. Labelled tableaux for non-monotonic reasoning: Cumulative consequence relations. *Journal of Logic and Computation*, 12(6):1027–1060, 2002.
3. M. Baldoni. *Normal Multimodal Logics: Automatic Deduction and Logic Programming Extension*. PhD thesis, Universita degli Studi di, Italy, 1998.
4. P. Blackburn and M. Rijke. Zooming in, zooming out. *Journal of Logic, Language and Information*, 1996.
5. A. S. d' Avila Garcez and D. M. Gabbay. Fibring neural networks. In *American National Conference on Artificial Intelligence (AAAI-2004)*, pages 342–347. AAAI/MIT Press, 2004.
6. M. Fitting. *Proof Methods for Modal and Intuitionistic Logics*. Reidel, Dordrecht, 1983.
7. D. M. Gabbay. *Fibring Logics*. Oxford University Press, Oxford, 1999.
8. Dov M. Gabbay and Guido Governatori. *Fibred modal tableaux*. Kluwer Publishers, 2000.
9. G. Governatori. Labelled tableau for multi-modal logics. In *TABLEAUX*, volume 918, pages 79–94. Springer, 1995.
10. G. Governatori and A. Luppi. Labelled tableaux for non-normal modal logics. In *AI*IA 99*, number 1792 in LNAI, pages 119–130. Springer-Verlag, 2000.

11. G. Governatori, V. Padmanabhan, and A. Sattar. On Fibring Semantics for BDI Logics. In *Logics in Artificial Intelligence (JELIA-02), Italy*, LNAI-2424. Springer, 2002.
12. M. Kracht and F. Wolter. Properties of independently axiomatizable bimodal logics. *The Journal of Symbolic Logic*, 56(4):1469–1485, 1991.
13. V. Padmanabhan. *On Extending BDI Logics*. PhD thesis, School of Information Technology, Griffith University, Brisbane, Australia, 2003.
14. V. Padmanabhan and G. Governatori. A fibred tableaux calculi for BDI logics. To be published in the proceedings of DALT-06 (Declarative Agent Languages & Technologies).
15. A. S. Rao and M. P. Georgeff. Formal models and decision procedures for multi-agent systems. Technical note 61, Australian Artificial Intelligence Institute, 1995.
16. A. Sernadas, C. Sernadas, and M. Ryan. Combining logics: Synchronising and fibring. Technical report, Department of Mathematics, Instituto superior Tecnico, Lisbon, Portugal, 1996.
17. F. Wolter. The decision problem for combined (modal) logics. Technical report, Universitat Leipzig, Germany, www.informatik.uni-leipzig.de/~wolter/, September 9, 1999.
18. A. Zanardo, A. Sernadas, and C. Sernadas. Fibring: Completeness preservation. *Journal of Symbolic Logic*, 66(1):414–439, 2001.