The Focused Constraint Inverse Method for Intuitionistic Modal Logics

Sean McLaughlin and Frank Pfenning

Carnegie Mellon University

Abstract. We present a focused inverse method for proof search in a variety of intuitionistic modal logics such as $\mathbf{K}$, $\mathbf{D}$, $\mathbf{T}$, $\mathbf{S_4}$ and $\mathbf{S_5}$. Unification of world-paths for such logics is non-unitary and therefore handled by adding constraints to sequents. We sketch proofs of soundness and completeness with respect to intuitionistic modal natural deduction and describe our implementation. Potential applications lie in multi-modal intuitionistic logics that have recently been proposed to reason about authorization and information flow security.

1 Introduction

Intuitionistic modal logics (IMLs) are extensions of intuitionistic logic that incorporate *modalities* for reasoning about judgments other than categorical truth. There are numerous applications of IMLs. They can, for instance, be used to reason about distributed computing environments where the modalities express which resources, such as data and processors, are accessible from which other resources [25]. IMLs are used to reason about authentication and security policies [10, 8]. For example, they can formalize questions such as: "Given a policy and Alice's current permissions, does she have permission to open Bob's file?".

Intuitionistic logics are preferable to classical logics when proofs have computational content or are otherwise are of primary importance. In the distributed computing example, an intuitionistic proof that a distributed program can be executed on a given network corresponds to a plan of which processors evaluate which data. In authentication logic, if a security policy maintains a log of the proof terms used during access and an unintended permissions violation occurs, the logged proofs can be used to audit the policy.

In this work we are interested in the theorem proving problem for intuitionistic modal logics. Adding a modality to a logic can make theorem proving considerably more difficult than in the underlying logic. An additional challenge for IMLs is one of software engineering. Since many modal logics are only slight variants of one another, we wanted to design the theorem prover in such a way that we could handle some different modal logics with only small changes to the system. This methodology has been used successfully for classical modal logics, e.g. [5]. This paper is a small step in the direction of building efficient and general theorem provers for IMLs. The main contributions of this paper are

- The design and implementation of a sound and complete focused inverse method theorem prover for the intuitionistic modal logics $\mathbf{K}$, $\mathbf{D}$, $\mathbf{K_4}$, $\mathbf{D_4}$, $\mathbf{T}$, $\mathbf{S_4}$ and $\mathbf{S_5}$.

The system is designed in a uniform manner by exploiting a form of Kripke semantics where the visibility relation can be determined by unification. The only operational difference between any two logics is the unification algorithm used during search.

- A novel use of *constraints* to delay computation (§4). Our calculus necessitates the use of a non-unitary unification. Applying multiple unifiers eagerly during proof search would quickly exhaust the available store. To mitigate this problem we add constraints to the sequent calculus in order to delay the application of substitutions resulting from unification.

2 Intuitionistic Modal Logic

We will consider the following syntax of propositions:

$$\text{Propositions A} ::= p \mid \top \mid \bot \mid A \wedge A \mid A \vee A \mid A \supset A \mid \Box A \mid \Diamond A$$

Intuitionistic modal logics differs from their classical analogs in that the underlying predicate logic is intuitionistic rather than classical. While the semantics of classical modal logic is typically understood in terms of its Kripke models (see e.g. [12]), there is considerable debate as to the proper interpretation of intuitionistic modal logic [24]. IMLs can be given a Kripke semantics where there are two different accessibility relations, one for the underlying intuitionistic logic and one for the modalities. However, it seems more intuitive to us to take the natural deduction calculus $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$ as the definition of IML.

2.1 Natural deduction

Figure 1 shows Simpson's natural deduction calculus for IML [24]. In a manner similar to labeled deduction, each proposition is relativized to an explicit world. We assume an infinite supply of *world variables* w. A labeled proposition has the form $A \mathbin{@} w$ and represents the proposition that A holds (intuitionistically) at world w. For example, the $\wedge$-elimination rules declares that if $A \wedge B$ holds at world w than A and B also hold at w. The rules with hypotheses in brackets are examples of hypothetical judgments, as in intuitionistic natural deduction.

2.2 The visibility relation

$\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$ is parametrized over the the visibility relation $\mathcal{R}$. For example, the $\Box$-elimination rule declares that if $\Box A$ holds at world w and $w \mathcal{R} w'$, then A holds at world w'. The properties of $\mathcal{R}$ differ from logic to logic. For instance, in $\mathbf{T}$, $\mathcal{R}$ is reflexive while in $\mathbf{S_4}$ it is both reflexive and transitive. In this paper we will consider the intuitionistic modal logics $\mathbf{K}$, $\mathbf{D}$, $\mathbf{K_4}$, $\mathbf{D_4}$, $\mathbf{T}$, $\mathbf{S_4}$ and $\mathbf{S_5}$. The visibility relations satisfy the properties given in Figure 5.

2.3 Sequent calculus

While Simpson's natural deduction system forms the most intuitive basis for a proof-theoretic semantics of IML, we prefer a sequent calculus called $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ for proof search. In $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ the assumptions $w \mathcal{R} w'$ are reified in a *world-graph* $\mathcal{G}$. The world graph is a compact description of the visibility relation. It has world variables for nodes and is rooted at the fixed initial world w_0. A (directed) edge between worlds w and w' indicates

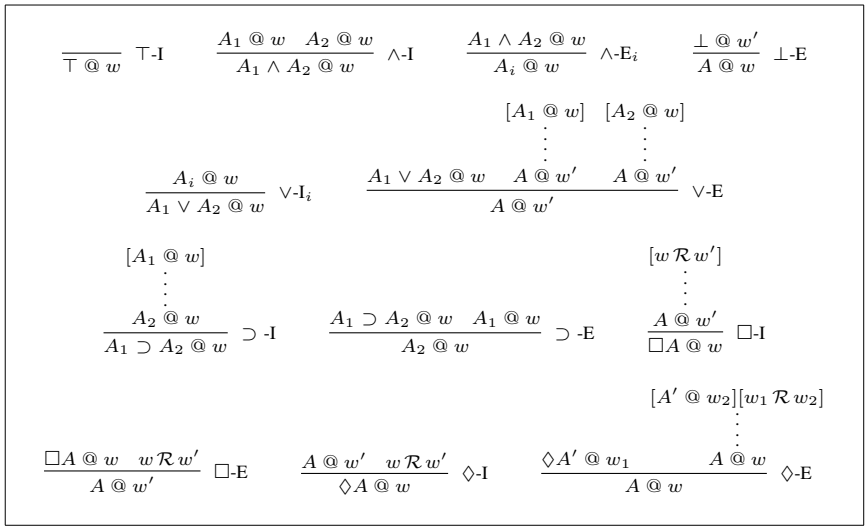

Fig. 1: $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$, natural deduction for IML

that w' is visible from w. An example of a world-graph is shown in Figure 4 (a). The edges of a world-graph do not (generally) completely describe the visibility relation. In the logics $\mathbf{K_4}$ for instance, since the relation is transitive an edge from w_1 to w_2 and another from w_2 to w_3 indirectly implies that w_3 is visible from w_1, regardless of whether there exists an edge between w_1 and w_3. We write the judgment regarding visibility in a world-graph as $\mathcal{G} \models w \mathcal{R} w'$.

Sequents have the form $\mathcal{G} \mid \Gamma \vdash \gamma$ where Γ is a set of labeled propositions, γ is a labeled proposition, and $\mathcal{G}$ is a world-graph. A world-graph can be extended by a new node using the notation $\mathcal{G} \cup \langle w, w' \rangle$, whereby we assume w is a node of $\mathcal{G}$ and w' is a world variable not in $\mathcal{G}$. The resulting graph has w' as a node with a directed edge from w to w'. Figure 2 shows the inference rules[1] of Simpson's sequent calculus $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$. Figure 3 shows an example proof in $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$. Using cut admissibility, Simpson proves that $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ is sound and complete with respect to $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$.

Theorem 1 (Simpson [24]). *$\mathcal{G}_0 \mid \cdot \vdash A @ w_0$ in $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ if and only if there exists a derivation of A in $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$.*

3 The World-Path Calculus

$\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ provides a basis for top-down[2] proof search. Given a method for determining the visibility relation with respect to a given world graph $\mathcal{G}$, one can in principal search

[1] Note that in the left rules of the backward calculi in this paper we assume the principal formula is copied to the premises. In some rules copying is necessary for completeness but overly verbose for presentation.

[2] *Top-down* and *backward* both refer to backward-chaining tableaux style search. *Bottom-up* and *forward* refer to forward-chaining resolution style proof search.

$$\frac{}{\mathcal{G} \mid \Gamma, A \mathbin{@} w \vdash A \mathbin{@} w}\ \text{Init} \qquad \frac{}{\mathcal{G} \mid \Gamma \vdash \top \mathbin{@} w}\ \top\text{R} \qquad \frac{}{\mathcal{G} \mid \Gamma, \bot \mathbin{@} w \vdash A \mathbin{@} w'}\ \bot\text{L}$$

$$\frac{\mathcal{G} \mid \Gamma \vdash A_1 \mathbin{@} w \quad \mathcal{G} \mid \Gamma \vdash A_2 \mathbin{@} w}{\mathcal{G} \mid \Gamma \vdash A_1 \wedge A_2 \mathbin{@} w}\ \wedge\text{R} \qquad \frac{\mathcal{G} \mid \Gamma, A_i \mathbin{@} w \vdash A \mathbin{@} w'}{\mathcal{G} \mid \Gamma, A_1 \wedge A_2 \mathbin{@} w \vdash A \mathbin{@} w'}\ \wedge\text{L}_i$$

$$\frac{\mathcal{G} \mid \Gamma, A_1 \mathbin{@} w \vdash A \mathbin{@} w' \quad \mathcal{G} \mid \Gamma, A_2 \mathbin{@} w \vdash A \mathbin{@} w'}{\mathcal{G} \mid \Gamma, A_1 \vee A_2 \mathbin{@} w \vdash A \mathbin{@} w'}\ \vee\text{L}$$

$$\frac{\mathcal{G} \mid \Gamma \vdash A_i \mathbin{@} w}{\mathcal{G} \mid \Gamma \vdash A_1 \vee A_2 \mathbin{@} w}\ \vee\text{R}_i \qquad \frac{\mathcal{G} \mid \Gamma, A_1 \mathbin{@} w \vdash A_2 \mathbin{@} w}{\mathcal{G} \mid \Gamma \vdash A_1 \supset A_2 \mathbin{@} w}\ \supset\text{R}$$

$$\frac{\mathcal{G} \mid \Gamma, A_2 \mathbin{@} w \vdash A \mathbin{@} w' \quad \mathcal{G} \mid \Gamma \vdash A_1 \mathbin{@} w}{\mathcal{G} \mid \Gamma, A_1 \supset A_2 \mathbin{@} w \vdash A \mathbin{@} w'}\ \supset\text{L}$$

$$\frac{\mathcal{G} \cup \langle w, w' \rangle \mid \Gamma \vdash A \mathbin{@} w'}{\mathcal{G} \mid \Gamma \vdash \Box A \mathbin{@} w}\ \Box\text{R}^{w'} \qquad \frac{\mathcal{G} \models w \mathcal{R} w_1 \quad \mathcal{G} \mid \Gamma, A \mathbin{@} w_1 \vdash A' \mathbin{@} w_2}{\mathcal{G} \mid \Gamma, \Box A \mathbin{@} w \vdash A' \mathbin{@} w_2}\ \Box\text{L}$$

$$\frac{\mathcal{G} \models w \mathcal{R} w' \quad \mathcal{G} \mid \Gamma \vdash A \mathbin{@} w'}{\mathcal{G} \mid \Gamma \vdash \Diamond A \mathbin{@} w}\ \Diamond\text{R} \qquad \frac{\mathcal{G} \cup \langle w_1, w \rangle \mid \Gamma, A \mathbin{@} w \vdash A' \mathbin{@} w_2}{\mathcal{G} \mid \Gamma, \Diamond A \mathbin{@} w_1 \vdash A' \mathbin{@} w_2}\ \Diamond\text{L}^{w}$$

Fig. 2: The Sequent Calculus $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\mathcal{G}_2 \mid P \mathbin{@} w_2 \vdash P \mathbin{@} w_2} \quad \mathcal{G}_2 \models w_0 \mathcal{R} w_2}{\mathcal{G}_2 \mid \Box P \mathbin{@} w_0 \vdash P \mathbin{@} w_2}}{\mathcal{G}_1 \mid \Box P \mathbin{@} w_0 \vdash \Box P \mathbin{@} w_1}}{\mathcal{G}_0 \mid \Box P \mathbin{@} w_0 \vdash \Box\Box P \mathbin{@} w_0}}{\mathcal{G}_0 \mid \vdash \Box P \supset \Box\Box P \mathbin{@} w_0}$$

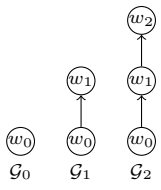

Observe that the judgment $G_2 \models w_0 \mathcal{R} w_2$ only holds if $\mathcal{R}$ is transitive. Thus, the formula is a theorem of $\mathbf{S_4}$ but not of $\mathbf{K}$.

Fig. 3: $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ proof

backward for all proofs of a given goal. There are a number of reasons, however, to prefer bottom-up search. In addition to the usual difficulties of backward search (e.g. meta-variables are global), constraint solving for the visibility relation in such a prover is complex. Once a proof skeleton is found (i.e., a proof whose leaves are axioms, but constraints remain to be checked), all the constraints need to be checked simultaneously. A failed check leads to backtracking.

An inverse method proof search can be more efficient for non-classical logics [15, 16], and avoids the problem of global constraint solving by checking constraints locally at each sequent. However, there are difficulties with $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ (e.g. combining world-graphs in the forward direction is awkward) that make it imperfect for forward proof search. We instead consider an alternate presentation of IML using a variant of Ohlbach's world-paths [19, 20] rather than a world-graph.

3.1 World-paths

A different view of the visibility relation may be obtained by considering as primary the edges of the world-graph rather than the nodes (worlds) themselves. Consider Figure 4. Diagram (a) represents a world-graph. Diagram (b) names the edges between worlds rather than the worlds themselves. We call (b) the *world-path* representation. A *world-path* π consists of the empty world-path π_0, a single edge, or the (left-associative) concatenation of two world-paths $\pi_1 \cdot \pi_2$. There is an obvious bijection between the two representations of the world-graph. We can therefore define a translation $[\![w]\!]$ of the world-graph nodes as the path starting from π_0 and following a directed path in $\mathcal{G}$ to w. For example, $[\![w_4]\!] = \pi_0 \cdot e_{01} \cdot e_{14}$. We say a world-path π_2 *extends* a world-path π if there exists a π_1 such that $\pi_2 \equiv \pi \cdot \pi_1$. Figure 5 gives the relevant judgements for world-paths. Given an IML whose visibility relation has some subset of the properties of §2.2, we select inference rules for $\equiv$ such that the following theorem holds[3]. The theorem demonstrates the strong correspondence between visibility in world-graphs and extensions of world-paths.

Theorem 2. *Given a visibility relation $\mathcal{R}$, the corresponding rules for $\equiv$, a world-graph $\mathcal{G}$, and worlds $w, w' \in \mathcal{G}$, $\mathcal{G} \models w \mathcal{R} w'$ if and only if $[\![w']\!]$ extends $[\![w]\!]$.*

Theorem 2 actually consists of seven different theorems, one for each visibility relation and set of path axioms corresponding to $\mathbf{K}$, $\mathbf{D}$, $\mathbf{K_4}$, $\mathbf{D_4}$, $\mathbf{T}$, $\mathbf{S_4}$ and $\mathbf{S_5}$. For a given $\mathcal{R}$ and set of rules for $\equiv$ such that Theorem 2 holds, we write $\mathcal{R} \sim \equiv$.

Now that we have reduced the world-graph visibility problem to equivalence between paths, we are in a position to define the world-path sequent calculus (Figure 6). Sequents have the form $\Delta\,;\Gamma \implies \gamma$ where Δ is simply a context of edge parameters, Γ is a set of labeled propositions and γ is a labeled proposition. The soundness and completeness with respect to $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ (fixing a visibility relation and path algebra) is given in the following theorem:

Theorem 3. *If $\mathcal{R} \sim \equiv$ then there is a derivation $\mathcal{G} \mid \Gamma \vdash \gamma$ in $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ if and only if and there is a derivation $\cdot\,;[\![\Gamma]\!] \implies [\![\gamma]\!]$ in $\mathbf{P}^{\equiv}_{\Box\Diamond}$. (Here the notation $[\![A \mathbin{@} w]\!]$ means $A \mathbin{@} [\![w]\!]$, with the obvious extension to Γ.)*

[3] For now we ignore the seriality property. We return to it in §4.3.

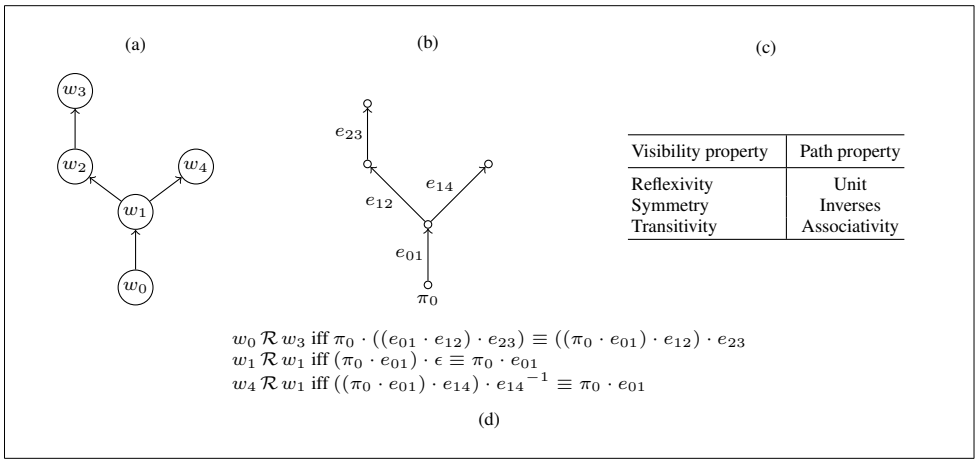

Fig. 4: Relationship between graphs and world-paths

Proof. Straightforward induction on the given derivation, using Theorem 2 to discharge the premises regarding the visibility relation.

Corollary 4 *If $\mathcal{R} \sim \equiv$ then there is a derivation $w_0 \mid \cdot \vdash A \mathbin{@} w_0$ in $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ if and only if and there is a derivation $\cdot ; \cdot \Longrightarrow A \mathbin{@} \pi_0$ in $\mathbf{P}^{\equiv}_{\Box\Diamond}$.*

4 The Inverse Method

The inverse method [14] is a generic bottom-up method for proof search. It is particularly useful for non-classical logics, where resolution is not available[4]. Following the inverse method "recipe" [6] for an inverse method theorem prover, the next step is to define a bottom-up version of the top-down calculus for forward proof search. This task is complicated by the presence of the hypotheses of the $\Box$ and $\Diamond$ rules regarding path equivalence and well-formedness. We solve this problem locally by solving these path equations when we can, and postponing them when we can not. We manage the postponement by adding a zone of *constraints* to the forward sequents. A forward sequent has the form $\Psi \mid \Gamma \longrightarrow \gamma$ where Ψ is a constraint, Γ is a set of labeled propositions, and γ is a set of a labeled propositions with at most one element. (This formulation of the consequent is necessary for incorporating falsehood and negation.) The constraint Ψ is constructed from the following grammar:

$$\text{Constraints } \Psi ::= \pi \text{ path} \mid \pi_1 \equiv \pi_2 \mid \top \mid \Psi \wedge \Psi \mid \forall e.\, e \text{ edge} \supset \Psi \mid \bot$$

An entailment relation $\Psi_1 \models \Psi_2$ on constraints is inherited from the equivalence axioms and path well-formedness rules. The forward ground world-path calculus is defined in Figure 7.

[4] This work can be seen as a reformulation of Ohlbach's work on resolution for classical modal logic [20] that applies to intuitionistic modal logics.

Path formation

$$\begin{array}{ll} \text{Edges} & e \\ \text{Paths} & \pi \ ::= e \mid \pi_0 \mid \pi_1 \cdot \pi_2 \mid \epsilon \mid e^{-1} \\ \text{Contexts} & \Delta ::= \cdot \mid \Delta, e \end{array}$$

Path well-formedness

$$\frac{}{\Delta \models \pi_0 \text{ path}} \qquad \frac{e \in \Delta}{\Delta \models e \text{ path}} \qquad \frac{\Delta \models \pi_1 \text{ path} \quad \Delta \models \pi_2 \text{ path}}{\Delta \models \pi_1 \cdot \pi_2 \text{ path}}$$

$$\frac{}{\Delta \models \epsilon \text{ path}}\ * \qquad \frac{e \in \Delta}{\Delta \models e^{-1} \text{ path}}\ \dagger \qquad \frac{\Delta \models \pi_1 \text{ path} \quad \ldots \quad \Delta \models \pi_n \text{ path}}{\Delta \models \{\pi_1, \ldots, \pi_n\} \text{ paths}}$$

Path equivalence

$$\frac{}{\Delta \models \pi \equiv \pi} \qquad \frac{}{\Delta \models \pi \cdot \epsilon \equiv \pi}\ * \qquad \frac{}{\Delta \models \epsilon \cdot \pi \equiv \pi}\ *$$

$$\frac{}{\Delta \models e \cdot e^{-1} \equiv \epsilon}\ \dagger \qquad \frac{}{\Delta \models e^{-1} \cdot e \equiv \epsilon}\ \dagger \qquad \frac{}{\Delta \models (e^{-1})^{-1} \equiv e}\ \dagger$$

$$\frac{\Delta \models \pi_1 \equiv \pi_1' \quad \Delta \models \pi_2 \equiv \pi_2'}{\Delta \models \pi_1 \cdot \pi_2 \equiv \pi_1' \cdot \pi_2'} \qquad \frac{}{\Delta \models (\pi_1 \cdot \pi_2) \cdot \pi_3 \equiv \pi_1 \cdot (\pi_2 \cdot \pi_3)}\ \ddagger$$

(*) when $\equiv$ admits unit ($\mathcal{R}$ is reflexive)
(†) when $\equiv$ admits inverses ($\mathcal{R}$ is symmetric)
(‡) when $\equiv$ is associative ($\mathcal{R}$ is transitive)

Relation properties of modal logics

$\mathbf{K}$	no special properties
$\mathbf{K_4}$	transitive
$\mathbf{T}$	reflexive
$\mathbf{S_4}$	reflexive and transitive
$\mathbf{S_5}$	reflexive, symmetric and transitive

Fig. 5: Judgments regarding world-paths

$$\frac{}{\Delta\,;\,\Gamma, p \mathbin{@} \pi \Longrightarrow p \mathbin{@} \pi}\ \text{Init} \qquad \frac{}{\Delta\,;\,\Gamma \Longrightarrow \top \mathbin{@} \pi}\ \top\text{R} \qquad \frac{}{\Delta\,;\,\Gamma, \bot \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}\ \bot\text{L}$$

$$\frac{\Delta\,;\,\Gamma \Longrightarrow A_1 \mathbin{@} \pi \quad \Delta\,;\,\Gamma \Longrightarrow A_2 \mathbin{@} \pi}{\Delta\,;\,\Gamma \Longrightarrow A_1 \wedge A_2 \mathbin{@} \pi}\ \wedge\text{R} \qquad \frac{\Delta\,;\,\Gamma, A_i \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}{\Delta\,;\,\Gamma, A_1 \wedge A_2 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}\ \wedge\text{L}_i$$

$$\frac{\Delta\,;\,\Gamma, A_1 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi' \quad \Delta\,;\,\Gamma, A_2 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}{\Delta\,;\,\Gamma, A_1 \vee A_2 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}\ \vee\text{L} \qquad \frac{\Delta\,;\,\Gamma \Longrightarrow A_i \mathbin{@} \pi}{\Delta\,;\,\Gamma \Longrightarrow A_1 \vee A_2 \mathbin{@} \pi}\ \vee\text{R}_i$$

$$\frac{\Delta\,;\,\Gamma, A_1 \mathbin{@} \pi \Longrightarrow A_2 \mathbin{@} \pi}{\Delta\,;\,\Gamma \Longrightarrow A_1 \supset A_2 \mathbin{@} \pi}\ \supset\text{R} \qquad \frac{\Delta\,;\,\Gamma, A_2 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi' \quad \Delta\,;\,\Gamma \Longrightarrow A_1 \mathbin{@} \pi}{\Delta\,;\,\Gamma, A_1 \supset A_2 \mathbin{@} \pi \Longrightarrow A \mathbin{@} \pi'}\ \supset\text{L}$$

$$\frac{\Delta, e\,;\,\Gamma \Longrightarrow A \mathbin{@} \pi \cdot e}{\Delta\,;\,\Gamma \Longrightarrow \Box A \mathbin{@} \pi}\ \Box\text{R}^e \qquad \frac{\Delta, e\,;\,\Gamma, A \mathbin{@} \pi \cdot e \Longrightarrow A' \mathbin{@} \pi_2}{\Delta\,;\,\Gamma, \Diamond A \mathbin{@} \pi \Longrightarrow A' \mathbin{@} \pi_2}\ \Diamond\text{L}^e$$

$$\frac{\pi_1 \equiv \pi \cdot \pi' \quad \Delta \models \pi' \text{ path} \quad \Delta\,;\,\Gamma, A \mathbin{@} \pi_1 \Longrightarrow A' \mathbin{@} \pi_2}{\Delta\,;\,\Gamma, \Box A \mathbin{@} \pi \Longrightarrow A' \mathbin{@} \pi_2}\ \Box\text{L}$$

$$\frac{\pi_1 \equiv \pi \cdot \pi' \quad \Delta \models \pi' \text{ path} \quad \Delta\,;\,\Gamma \Longrightarrow A \mathbin{@} \pi_1}{\Delta\,;\,\Gamma \Longrightarrow \Diamond A \mathbin{@} \pi}\ \Diamond\text{R}$$

In the Init rule, p is an atomic formula. In $\Box$-R and $\Diamond$-L e is an eigenvariable not occurring in the remainder of the sequent.

Fig. 6: The World Path Calculus $\mathbf{P}^{\equiv}_{\Box\Diamond}$

4.1 Subsumption

Continuing with the inverse method recipe, we need to update the definition of subsumption for constraint sequents. Recall that an intuitionistic sequent $\Gamma_1 \longrightarrow \gamma_1$ *subsumes* $\Gamma_2 \longrightarrow \gamma_2$ if there exists a substitution θ such that $\Gamma_1\theta \subseteq \Gamma_2$ and $\gamma_1\theta \subseteq \gamma_2$. It is known that the inverse method can not directly prove any valid sequent, but in general can only prove a stronger one (i.e., one that can be weakened to the goal sequent.) In the modal case the constraints limit the validity of the remainder of the sequent. For example, the sequent $e_1 \equiv e_2 \mid \Gamma \longrightarrow \gamma$ is trivial when e_1 and e_2 are distinct edges Since *stronger* constraints limit the valid substitution instances of a sequent, the subsuming sequent must have a *weaker* constraint than the subsumed sequent in the following sense.

Definition 1 (Subsumption). *Sequent* $\Psi_1 \mid \Gamma_1 \longrightarrow \gamma_1$ subsumes $\Psi_2 \mid \Gamma_2 \longrightarrow \gamma_2$ *if there exists a substitution* θ *such that* $\Gamma_1\theta \subseteq \Gamma_2$, $\gamma_1\theta \subseteq \gamma_2$ *and* $\Psi_2 \models \Psi_1\theta$.

Soundness and completeness theorems then establish the connection between the forward calculus $\mathbf{P}^{\equiv Inv}_{\Box\Diamond}$ and $\mathbf{P}^{\equiv}_{\Box\Diamond}$.

Theorem 5. $\Delta\,;\,\Gamma \Longrightarrow A \mathbin{@} \pi_0$ *if and only if there exists* Ψ, Γ', γ' *such that* $\Psi \mid \Gamma' \longrightarrow \gamma'$ *and* $\Psi \mid \Gamma' \longrightarrow \gamma'$ *subsumes* $\top \mid \Gamma \longrightarrow A \mathbin{@} \pi_0$.

Proof. This proof is complicated by the constraints. Care must be taken because the constraints have different forms in the two calculi (cf. $\Box$-R). To prove it we actually first define a backward calculus $\mathbf{P}'^{\equiv}_{\Box\Diamond}$ that is closer in spirit to the forward calculus than

$\mathbf{P}^{\equiv}_{\Box\Diamond}$. Then we prove that $\mathbf{P}^{\equiv}_{\Box\Diamond}$ is sound and complete with respect to $\mathbf{P'}^{\equiv}_{\Box\Diamond}$ and that $\mathbf{P'}^{\equiv}_{\Box\Diamond}$ is sound and complete with respect to $\mathbf{P}^{\equiv Inv}_{\Box\Diamond}$.

Corollary 6 *$\cdot\,;\cdot \Longrightarrow A \mathbin{@} \pi_0$ if and only if $\top \mid \cdot \longrightarrow A \mathbin{@} \pi_0$.*

$$\frac{}{\pi \text{ path} \mid p \mathbin{@} \pi \longrightarrow p \mathbin{@} \pi}\ \text{Init} \qquad \frac{}{\pi \text{ path} \mid \cdot \longrightarrow \top \mathbin{@} \pi}\ \top\text{R}$$

$$\frac{\Psi \mid \Gamma, A_i \mathbin{@} \pi \longrightarrow \gamma}{\Psi \mid \Gamma, A_1 \wedge A_2 \mathbin{@} \pi \longrightarrow \gamma}\ \wedge\text{L}_i \qquad \frac{\Psi \mid \Gamma \longrightarrow A_i \mathbin{@} \pi}{\Psi \mid \Gamma \longrightarrow A_1 \vee A_2 \mathbin{@} \pi}\ \vee\text{R}_i$$

$$\frac{\Psi_1 \mid \Gamma_1 \longrightarrow A_1 \mathbin{@} \pi_1 \quad \Psi_2 \mid \Gamma_2 \longrightarrow A_2 \mathbin{@} \pi_2}{\pi_1 \equiv \pi_2 \wedge \Psi_1 \wedge \Psi_2 \mid \Gamma_1, \Gamma_2 \longrightarrow A_1 \wedge A_2 \mathbin{@} \pi_1}\ \wedge\text{R}$$

$$\frac{\Psi_1 \mid \Gamma_1, A_1 \mathbin{@} \pi_1 \longrightarrow \gamma_1 \quad \Psi_2 \mid \Gamma_2, A_2 \mathbin{@} \pi_2 \longrightarrow \gamma_2}{\pi_1 \equiv \pi_2 \wedge \Psi_1 \wedge \Psi_2 \mid \Gamma_1, \Gamma_2, A_1 \vee A_2 \mathbin{@} \pi_1 \longrightarrow \gamma_1 \cup \gamma_2}\ \vee\text{L}$$

$$\frac{\Psi \mid \Gamma, A_1 \mathbin{@} \pi_1 \longrightarrow A_2 \mathbin{@} \pi_2}{\pi_1 \equiv \pi_2 \wedge \Psi \mid \Gamma \longrightarrow A_1 \supset A_2 \mathbin{@} \pi_1}\ \supset\text{R}_1 \qquad \frac{\Psi \mid \Gamma, A_1 \mathbin{@} \pi \longrightarrow \cdot}{\Psi \mid \Gamma \longrightarrow A_1 \supset A_2 \mathbin{@} \pi}\ \supset\text{R}_2$$

$$\frac{\Psi \mid \Gamma \longrightarrow A_2 \mathbin{@} \pi}{\Psi \mid \Gamma \longrightarrow A_1 \supset A_2 \mathbin{@} \pi}\ \supset\text{R}_3 \qquad \frac{}{\pi \text{ path} \mid \bot \mathbin{@} \pi \longrightarrow \cdot}\ \bot\text{L}$$

$$\frac{\Psi_1 \mid \Gamma_1, A_2 \mathbin{@} \pi_1 \longrightarrow \gamma \quad \Psi_2 \mid \Gamma_2 \longrightarrow A_1 \mathbin{@} \pi_2}{\pi_1 \equiv \pi_2 \wedge \Psi_1 \wedge \Psi_2 \mid \Gamma_1, \Gamma_2, A_1 \supset A_2 \mathbin{@} \pi_1 \longrightarrow \gamma}\ \supset\text{L} \qquad \frac{\Psi \mid \Gamma, A \mathbin{@} \pi_1, A \mathbin{@} \pi_2 \longrightarrow \gamma}{\pi_1 \equiv \pi_2 \wedge \Psi \mid A \mathbin{@} \pi_1, \Gamma \longrightarrow \gamma}\ \text{Contr}$$

$$\frac{\Psi \mid \Gamma \longrightarrow A \mathbin{@} \pi'}{\forall e.\ e \text{ edge} \supset (\pi' \equiv \pi \cdot e \wedge \Psi) \mid \Gamma \longrightarrow \Box A \mathbin{@} \pi}\ \Box\text{R}^e$$

$$\frac{\Psi \mid \Gamma, A \mathbin{@} \pi_2 \longrightarrow \gamma}{\pi_2 \equiv \pi \cdot \pi_1 \wedge \pi_1 \text{ path} \wedge \Psi \mid \Gamma, \Box A \mathbin{@} \pi \longrightarrow \gamma}\ \Box\text{L}$$

$$\frac{\Psi \mid \Gamma \longrightarrow A \mathbin{@} \pi_2}{\pi_2 \equiv \pi \cdot \pi_1 \wedge \pi_1 \text{ path} \wedge \Psi \mid \Gamma \longrightarrow \Diamond A \mathbin{@} \pi}\ \Diamond\text{R}$$

$$\frac{\Psi \mid \Gamma, A \mathbin{@} \pi' \longrightarrow \gamma}{\forall e.\ e \text{ edge} \supset (\pi' \equiv \pi \cdot e \wedge \Psi) \mid \Gamma, \Diamond A \mathbin{@} \pi \longrightarrow \gamma}\ \Diamond\text{L}^e$$

In the rule Init, p is an atomic formula. In the rules $\Box$-R and $\Diamond$-L, e is a new eigenvariable, not occurring elsewhere in the sequent. In the rule $\vee$-L, the consequents are combined with the $\cup$ operator. By this we mean that if either of the consequents are empty, the result is the other consequent. If both are nonempty, then the consequents must have the form $A \mathbin{@} \pi_3$, $A \mathbin{@} \pi_4$ and we add the constraint $\pi_3 \equiv \pi_4$ to the constraint zone.

Fig. 7: The Forward World Path Calculus

4.2 Unification

As in first-order logic, the next step is to *lift* the ground calculus described in the last section to allow free path variables, thus making finite the number of initial sequents. A

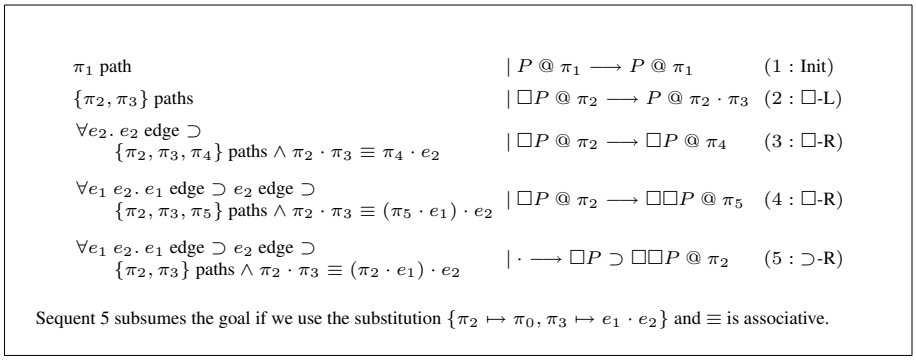

Fig. 8: Example inverse method proof

sequent with free variables then stands for all of its substitution instances. This is typically done using unification and most general unifiers [6]. Unfortunately, the world-path unification problem does not always admit most general unifiers (though the set of unifiers is always finite in the cases we are considering). For example, when the visibility relation is transitive, the equivalence $e_1 \cdot e_2 \cdot e_3 \equiv x_1 \cdot x_2$ has (at least) the following unifiers, none of which is more general than another: $\{x_1 \mapsto e_1, x_2 \mapsto e_2 \cdot e_3\}, \{x_1 \mapsto e_1 \cdot e_2, x_2 \mapsto e_3\}$. While a number of authors have developed algorithms [20, 27] for such equivalences, we consider here the *T-string unification* algorithms of Otten and Kreitz [22]. For each modal logic in this paper they give a list of transformation (rewrite) rules that applies to a set of T-string unification equations. They prove that the rule application terminate with a minimal set of most general unifiers. Since our world-paths satisfy the T-string property[5], we can use their algorithms directly on systems of world-path equations. The problem then is to transform a constraint Ψ into a system of equivalences that can be solved by T-string unification. This is achieved by transforming Ψ into a normal form where the equivalences are immediate. Call a constraint inconsistent if $\Psi \models \bot$. A constraint that is not inconsistent is consistent.

Definition 2 (Constraint normal form). *A constraint Ψ is in* normal form *if it has the form* $\forall e_1 \ldots e_n.\ (e_1$ *edge* $\wedge \ldots \wedge\ e_n$ *edge*$) \supset (\bigwedge_i \pi_i \equiv \pi'_i \wedge \{\pi_1, \ldots, \pi_n\}$ *paths*$)$

Theorem 7. *Every consistent constraint Ψ is equivalent to a constraint Ψ' in constraint normal form.*

Proof. Because of the restriction on quantifier structure given by the grammar for Φ, by alpha-renaming we can prenex all quantifiers and rearrange the conjunctions into the desired form.

[5] Note that the unification problem as we described it does not seem to precisely fit the T-string framework. In T-string unification the concatenation operator is always associative, and has no inverses or units. The different properties of the visibility relation are obtained there by the selection of transformation rules and restricting what can be instantiated for a variable. It is nevertheless a straightforward matter to transform our presentation to satisfy the T-string property.

Given the equational part of the constraint normal form, $\forall e_1 \ldots e_n.\ \bigwedge_i \pi_i \equiv \pi'_i$ the universally quantified variables serve as constants in the unification equations, while all free variables represent unification variables that can be instantiated. This transformed problem is then passed to the T-string unification algorithm. Note that rather than enumerating the unifiers, for completeness we need only check for unifiability of the constraints.

A lifted calculus (omitted for brevity) can be defined and shown to have all the properties necessary for a complete inverse method: 1) a finite number of axioms 2) starting with a finite set of sequents, and since we only generate subformulas of a given goal sequent, there are only a finite number of new sequents derivable using the inference rules. Thus the method outlined above is a sound and complete method for proof search in the IMLs for which we have a unification procedure.

4.3 Serial worlds

The existence of logics with non-serial visibility relations is problematic in every presentation of theorem provers for modal logic. Since our approach differs from those we could find in the literature, we will describe an example. In the backward calculus, seriality becomes significant in the rules $\Box$-L and $\Diamond$-R with the π path predicate. Figure 9 shows the skeleton of a proof of $\Diamond\top$. Since when dealing with a non-serial visibility relation, it is not guaranteed that we can find an edge by which to traverse from π to π_0. Thus the proof should fail. The inverse method attempt is given in the same figure. The first step unifies the path variable π with the path concatenation $\pi' \cdot \pi''$ (essentially moving downward on the world-graph). Then the second sequent is unified with the goal to test subsumption, and π' becomes π_0. Since there is nothing with which to prove π'' path the subsumption check, and thus the proof, fails. This mechanism allows us to treat $\mathbf{D}$, which is $\mathbf{K}$ plus serality, and $\mathbf{D_4}$ which is seriality and transitivity.

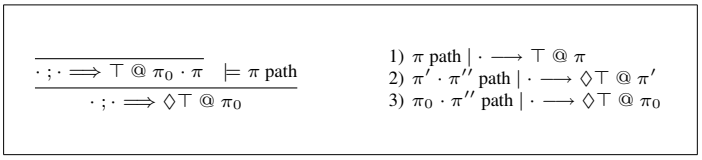

Fig. 9: Serial world example

5 Implementation

We implemented an experimental prototype of the constraint sequent calculus described above[6]. The implementation extends our implementation of a theorem prover for intuitionistic propositional and first-order logic called Imogen [15, 16]. The overall implementation is about 17K lines of Haskell. The amount that needed to be added to handle the modal operators and constraints was about 3K lines. The bulk of the work went

[6] The implementation can be found on the first author's website [1].

into implementing the T-string unification and managing the constraint entailment relation. In the remainder of this section we describe a few significant properties of our implementation that differ from the formal presentation.

5.1 Focusing

An important optimization for sequent calculus proof search is *focusing* [4]. In focusing we distinguish between connectives that are invertible on the right (negative) and left (positive). New connectives called *shifts* convert between positive and negative formulas. $\Box$ and $\Diamond$ are positive and negative respectively.

$$\begin{array}{ll} \text{Positive formulas} & A^+ ::= p^+ \mid A^+ \stackrel{*}{\wedge} A^+ \mid \stackrel{*}{\top} \mid A^+ \vee A^+ \mid \bot \mid \downarrow A^- \mid \exists x.\, A^+ \mid \Diamond A^+ \\ \text{Negative formulas} & A^- ::= p^- \mid A^- \wedge A^- \mid \top \mid A^+ \supset A^- \mid \uparrow A^+ \mid \forall x. A^- \mid \Box A^- \end{array}$$

Among other benefits, focusing allows for a dramatic reduction in the size of the search space [15, 16]. An important detail of this particular formulation is that unlike in other modal logics such as lax logic, linear logic and the judgmental formulation of modal logic [23], the modal operators share the polarity of their immediate subformula. This extends the focusing phases which makes for a smaller search space. The focused version of $\mathbf{P}^{\equiv Inv}_{\Box\Diamond}$ is implemented in Imogen. The completeness proof for the focused calculus is analogous to the numerous other focusing proofs for non-classical logics, e.g. [13].

5.2 Quantification

Though we did not describe it in our presentation thus far, our implementation allows first order quantification. We chose the fixed-domain semantics because it posed the fewest conceptual difficulties. For example, Figure 10 gives a proof of the Barcan formula.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{e \text{ edge} \,;\, p(c) \mathbin{@} \pi_0 \cdot e \Longrightarrow p(c) \mathbin{@} \pi_0 \cdot e} \quad e \text{ edge} \models e \text{ path}}{e \text{ edge} \,;\, \Box p(c) \mathbin{@} \pi_0 \Longrightarrow p(c) \mathbin{@} \pi_0 \cdot e}}{e \text{ edge} \,;\, \forall x.\, \Box p(x) \mathbin{@} \pi_0 \Longrightarrow p(c) \mathbin{@} \pi_0 \cdot e}}{e \text{ edge} \,;\, \forall x.\, \Box p(x) \mathbin{@} \pi_0 \Longrightarrow \forall x.\, p(x) \mathbin{@} \pi_0 \cdot e}}{\cdot \,;\, \forall x.\, \Box p(x) \mathbin{@} \pi_0 \Longrightarrow \Box \forall x.\, p(x) \mathbin{@} \pi_0}}{\cdot \,;\, (\forall x.\, \Box p(x)) \supset (\Box \forall x.\, p(x)) \mathbin{@} \pi_0 \Longrightarrow}$$

Fig. 10: A $\mathbf{P}^{\equiv}_{\Box\Diamond}$ proof of the Barcan formula

5.3 Constraints

Constraints are fundamental to the efficiency of the theorem prover. The constraint of every new sequent whose antecedents and consequent match the goal needs to be checked for unifiability. In addition, when we add a sequent to the database, we check

to see if the unification problem is unitary. If so, we apply the unifier throughout the sequent and simplify the constraint. In this sense we delay splitting a sequent due to non-unitary unification. While they are an clear benefit in allowing us to delay the computation of non-unitary unifiers, they can be difficult to manage. At present, we solve constraint entailments $\Psi_1 \models \Psi_2$ only in the special cases where $\Psi_1 = \top$ or $\Psi_2 = \bot$. This is sufficient for completeness, because to subsume the final goal we only need to verify that constraints are valid ($\Psi_1 = \top$). To eliminate many inconsistent sequents we only need the case where $\Psi_2 = \bot$. In future work we plan to develop practical algorithms solving further constraint entailments in order to further reduce redundancy in the search space.

6 Related Work

The work on automated deduction for modal logics can be roughly partitioned into the following areas.

- **Resolution methods.** Classical resolution for modal logic is the work nearest to ours. Ohlbach [19, 20] shows how to use resolution with a more sophisticated unification algorithm to prove theorems in classical modal logic in a top-down manner. We use a slight variant of his path calculus. The primary difference is that our underlying logic is intuitionistic. Voronkov presents an inverse method for a number of non-classical logics [6] and describes an implementation of an inverse method theorem prover for classical $\mathbf{K}$ [26]. It would be interesting to compare our focusing prover to his prover that is unfocused but uses optimizations he describes in the above papers.
- **Tableaux methods.** Wallen [27] describes a generalization of the classical connection method for modal logics using paths and path-unification. Otten [21] uses Wallen's approach to design efficient theorem provers for intuitionistic logic and some classical modal logics. Catach [5] uses a general tableaux strategy for a larger family of modal logics. Howe [11] implements a tableaux style prover for intuitionistic $\mathbf{S_4}$ and Lax logic. Amati and Perri [3] show a tableaux method for a large family of intuitionistic modal logics, though it does not seem to have been implemented. Garg [9] describes both a goal directed tableaux search and a saturation method similar to Datalog in his authorization logic $\mathbf{BL_0}$.
- **Translations.** A popular way to reason about modalities is via a translation to a non-modal logic. Abadi and Manna [2] translate the modalities into first-order classical logic with equality. Nonnengart [18] extends Ohlbach's work by developing a semi-functional translation from temporal logic and some modal logics into classical first-order logic. The equational theory then determines the modal logic. Egly [7] translates the Lax modality directly into first-order intuitionistic logic.

7 Conclusion and Future Work

We have described a focused constraint inverse method for automated theorem proving in a number of intuitionistic modal logics. From our target semantics of $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$, a natural deduction parametrized over a visibility relation, we presented the sequent calculus $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$, the world path calculus $\mathbf{P}^{\equiv}_{\Box\Diamond}$ and its focused variant. Soundness and complete-

ness results, given by Theorems 1, 4, 6 respectively, achieve our goal of a sound and complete focused inverse method with respect to $\mathbf{N}^{\mathcal{R}}_{\Box\Diamond}$.

$$\mathbf{N}^{\mathcal{R}}_{\Box\Diamond} \longleftrightarrow \mathbf{L}^{\mathcal{R}}_{\Box\Diamond} \overset{\mathcal{R}\sim\ \equiv}{\longleftrightarrow} \mathbf{P}^{\equiv}_{\Box\Diamond} \longleftrightarrow \mathbf{P}^{\equiv Inv}_{\Box\Diamond}$$

While the world path calculus was convenient from an automated deduction standpoint, it is not altogether satisfactory. There are other features of the visibility graphs of $\mathbf{L}^{\mathcal{R}}_{\Box\Diamond}$ that we can not currently see how to model using the path calculus. For instance, it is not clear to us how to use the algebraic properties of paths to represent, e.g. *directedness* or *Euclideanness*.

$$\forall w_1\ w_2\ w_3.\ w_1 \mathcal{R} w_2 \wedge w_1 \mathcal{R} w_3 \supset \exists w_4.\ w_2 \mathcal{R} w_4 \wedge w_3 \mathcal{R} w_4$$
$$\forall w_1\ w_2\ w_3.\ w_1 \mathcal{R} w_2 \wedge w_2 \mathcal{R} w_3 \supset w_1 \mathcal{R} w_3$$

Indeed, it seems almost serendipitous that the most common and useful properties can be represented algebraically. Perhaps there are extensions of path unification that can capture such properties.

We hope to extend the ideas presented here to prove theorems in some non-traditional intuitionistic modal logics. One such logic, designed for use with security and authentication protocols, is DKAL [10]. DKAL extends the intuitionistic propositional calculus with two (indexed) modal operators **said** and **implied**. The two modalities behave differently than our $\Box$ and $\Diamond$. While **said** has the same behavior as $\Box$ in (a multi-modal version of) $\mathbf{K}$, the rule for **implied** is unusual. It allows evidence of, e.g., `alice` **said** A to be used in verifying a proposition of the form `alice` **implied** B.

$$\frac{\Gamma \vdash A}{\Delta, \texttt{alice}\ \mathbf{said}\ \Gamma \vdash \texttt{alice}\ \mathbf{said}\ A} \qquad \frac{\Gamma_1, \Gamma_2 \vdash A}{\Delta, \texttt{alice}\ \mathbf{said}\ \Gamma_1, \texttt{alice}\ \mathbf{implied}\ \Gamma_2 \vdash \texttt{alice}\ \mathbf{implied}\ A}$$

We are still at work designing an efficient unification algorithm that takes this interference of the modalities into account. Once the necessary unification algorithm is in place, we intend to extend Imogen to allow direct reasoning in these modalities. Recently Mera and Bjørner [17] show how to translate the DKAL modalities into first-order classical logic with equality and arithmetic constraints that they can solve using the *Z3* SMT solver. We hope to be able to compare these two methods of theorem proving in the near future.

Acknowledgments

This research was supported by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office.

References

1. Imogen website. http://seanmcl.com/projects/imogen/.
2. M. Abadi and Z. Manna. Modal theorem proving. In J. H. Siekmann, editor, *CADE*, volume 230 of *LNCS*, pages 172–189. Springer, 1986.
3. G. Amati and F. Pirri. A uniform tableau method for intuitionistic modal logics I. *Studia Logica*, 53(1):29–60, 1994.
4. J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347, 1992.

5. L. Catach. TABLEAUX: A general theorem prover for modal logics. *Journal of Automated Reasoning*, 7(4):489–510, 1991.
6. A. Degtyarev and A. Voronkov. The inverse method. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 4, pages 179–272. Elsevier Science, 2001.
7. U. Egly. Embedding lax logic into intuitionistic logic. In A. Voronkov, editor, *CADE*, volume 2392 of *LNCS*, pages 78–93. Springer, 2002.
8. D. Garg. *Proof Theory for Authorization Logic and Its Application to a Practical File System*. PhD thesis, Carnegie Mellon University, 2000.
9. D. Garg. Proof search in an authorization logic. Technical Report CMU-CS-09-121, Computer Science Department, Carnegie Mellon University, April 2009.
10. Y. Gurevich and I. Neeman. DKAL: Distributed-knowledge authorization language. In *Computer Security Foundations*, pages 149–162. IEEE Computer Society, 2008.
11. J. M. Howe. *Proof Search Issues in Some Non-Classical Logics*. PhD thesis, University of St. Andrews, Scotland, 1998.
12. G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.
13. C. Liang and D. Miller. Focusing and polarization in intuitionistic logic. In J. Duparc and T. A. Henzinger, editors, *Computer Science Logic*, pages 451–465. Springer, 2007.
14. S. Y. Maslov. An inverse method for establishing deducibility in classical predicate calculus. *Doklady Akademii nauk SSSR*, 159:17–20, 1964.
15. S. McLaughlin and F. Pfenning. Imogen: Focusing the polarized focused inverse method for intuitionistic propositional logic. In I. C. et al., editor, *LPAR*, volume 5330 of *LNCS*, pages 174–181, 2008.
16. S. McLaughlin and F. Pfenning. Efficient intuitionistic theorem proving with the polarized inverse method. In R. A. Schmidt, editor, *CADE*, volume 5663 of *LNCS*, pages 230–244. Springer, 2009.
17. S. Mera and N. Bjørner. DKAL and Z3: A logic embedding experiment. In *To appear in Essays in honor of Yuri Gurevich's 70th birthday*, LNCS. Springer, 2010.
18. A. Nonnengart. Resolution-based calculi for modal and temporal logics. In M. A. McRobbie and J. K. Slaney, editors, *CADE*, volume 1104 of *LNAI*, pages 598–612. Springer, 1996.
19. H.-J. Ohlbach. *A Resolution Calculus for Modal Logics*. PhD thesis, Kaiserslautern, 1988.
20. H. J. Ohlbach. A resolution calculus for modal logics. In E. L. Lusk and R. A. Overbeek, editors, *CADE*, volume 310 of *LNCS*, pages 500–516. Springer, 1988.
21. J. Otten and C. Kreitz. A connection based proof method for intuitionistic logic. In P. B. et al., editor, *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 918 of *LNAI*, pages 122–137. Springer, 1995.
22. J. Otten and C. Kreitz. T-string-unification: unifying prefixes in non-classical proof methods. In P. M. et al., editor, *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 1071 of *LNAI*, pages 244–260. Springer, 1996.
23. F. Pfenning and R. Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11(4):511–540, 2001.
24. A. K. Simpson. *The proof theory and semantics of intuitionistic modal logic*. PhD thesis, University of Edinburgh, 1994.
25. T. M. VII. *Mobile Types for Mobile Code*. PhD thesis, Carnegie Mellon University, 2008.
26. A. Voronkov. Я: A theorem prover for K. In H. Ganzinger, editor, *CADE*, volume 1632 of *LNAI*, pages 383–387. Springer, 1999.
27. L. A. Wallen. *Automated proof search in non-classical logics : efficient matrix proof methods for modal and intuitionistic logics*. M.I.T. Press, 1990.