RWTH
LTCS–Report

Aachen University of Technology
Research group for
Theoretical Computer Science

The Inverse Method Implements the Automata Approach for Modal Satisfiability

Franz Baader and Stephan Tobies

LTCS-Report 01-03

RWTH Aachen
LuFg Theoretische Informatik
http://www-lti.informatik.rwth-aachen.de

Ahornstr. 55
52074 Aachen
Germany

The Inverse Method Implements the Automata Approach for Modal Satisfiability

Franz Baader Stephan Tobies

January 15, 2001

Abstract

Tableaux-based decision procedures for satisfiability of modal and description logics behave quite well in practice, but it is sometimes hard to obtain exact worst-case complexity results using these approaches, especially for ExpTime-complete logics. In contrast, automata-based approaches often yield algorithms for which optimal worst-case complexity can easily be proved. However, the algorithms obtained this way are usually not only worst-case, but also best-case exponential: they first construct an automaton that is always exponential in the size of the input, and then apply the (polynomial) emptiness test to this large automaton. To overcome this problem, one must try to construct the automaton "on-the-fly" while performing the emptiness test.

In this paper we will show that Voronkov's inverse method for the modal logic K can be seen as an on-the-fly realization of the emptiness test done by the automata approach for K. The benefits of this result are two-fold. First, it shows that Voronkov's implementation of the inverse method, which behaves quite well in practice, is an optimized on-the-fly implementation of the automata-based satisfiability procedure for K. Second, it can be used to give a simpler proof of the fact that Voronkov's optimizations do not destroy completeness of the procedure. We will also show that the inverse method can easily be extended to handle global axioms, and that the correspondence to the automata approach still holds in this setting. In particular, the inverse method yields an ExpTime-algorithm for satisfiability in K w.r.t. global axioms.

1 Introduction

Decision procedures for (propositional) modal logics and description logics play an important rôle in knowledge representation and verification. When developing such procedures, one is both interested in their worst-case complexity and in their behavior in practical applications. From the theoretical point of view, it is desirable to obtain an algorithm whose worst-case complexity matches the complexity of the problem. From the practical point of view it is more important to have an algorithm that is easy to implement and amenable to optimizations, such that it behaves well on practical instances of the decision problem. The most popular approaches for constructing decision procedures for modal logics are i) semantic tableaux and related methods [10, 2]; ii) translations into classical first-order logics [15, 1]; and iii) reductions to the emptiness problem for certain (tree) automata [17, 14].

Whereas highly optimized tableaux and translation approaches behave quite well in practice [11, 12], it is sometimes hard to obtain exact worst-case complexity results using these approaches. For example, satisfiability in the basic modal logic K w.r.t. global axioms is known to be ExpTime-complete [16]. However, the "natural" tableaux algorithm for this problem is a NExpTime-algorithm [2], and it is rather hard to construct a tableaux algorithm that runs in deterministic exponential time [6]. In contrast, it is folklore that the automata approach yields a very simple proof that satisfiability in K w.r.t. global axioms is in ExpTime. However, the algorithm obtained this way is not only worst-case, but also best-case exponential: it first constructs an automaton that is always exponential in the size of the input formulae (its set of states is the powerset of the set of subformulae of the input formulae), and then applies the (polynomial) emptiness test to this large automaton. To overcome this problem, one must try to construct the automaton "on-the-fly" while performing the emptiness test. Whereas this idea has successfully been used for automata that perform model checking [9, 5], to the best of our knowledge it has not yet been applied to satisfiability checking.

The original motivation of this work was to compare the automata and the tableaux approaches, with the ultimate goal of obtaining an approach that combines the advantages of both, without possessing any of the disadvantages. As a starting point, we wanted to see whether the tableaux approach could be viewed as an on-the-fly realization of the emptiness test done by the automata approach. At first sight, this idea was persuasive since a run of the automaton constructed by the automata approach (which is a so-called looping automaton working on infinite trees) looks very much like a run of the tableaux procedure, and the tableaux procedure does generate sets of formulae on-the-fly. However, the polynomial emptiness test for looping automata does *not* try to construct a run starting with the root of the tree, as done by the tableaux approach. Instead, it computes inactive states, i.e., states that can never occur on a successful run of the automaton, and tests whether all initial states are inactive. This computation starts "from the bottom" by locating obviously inactive states (i.e., states without successor states), and then "propagates" inactiveness along the transition relation. Thus, the emptiness test works in the opposite direction of the tableaux procedure. This observation suggested to consider an approach that inverts the tableaux approach: this is just the so-called inverse method. Recently, Voronkov [19] has applied this method to obtain a bottom-up decision procedure for satisfiability in K, and has optimized and implemented this procedure.

In this paper we will show that the inverse method for K can indeed be seen as an on-the-fly realization of the emptiness test done by the automata approach for K. The benefits of this result are two-fold. First, it shows that Voronkov's implementation, which behaves quite well in practice, is an optimized on-the-fly implementation of the automata-based satisfiability procedure for K. Second, it can be used to give a simpler proof of the fact that Voronkov's optimizations do not destroy completeness of the procedure. We will also show how the inverse method can be extended to handle global axioms, and that the correspondence to the automata approach still holds in this setting. In particular, the inverse method yields an ExpTime-algorithm for satisfiability in K w.r.t. global axioms.

2 Preliminaries

First, we briefly introduce the modal logic K and some technical definitions related to K-formulae, which are used later on to formulate the inverse calculus and the automata approach for K. Then, we define the type of automata used to decide satisfiability (w.r.t. global axioms) in K. These so-called looping automata [18] are a specialization of Büchi tree automata.

Modal Formulae

We assume the reader to be familiar with the basic notions of modal logic. For a thorough introduction to modal logics, refer to, e.g., [4].

K-formulae are built inductively from a countably infinite set $\mathcal{P} = \{p_1, p_2, \dots\}$ of propositional atoms using the Boolean connectives $\wedge$, $\vee$, and $\neg$ and the unary modal operators $\Box$ and $\Diamond$. The semantics of K-formulae is define as usual, based on Kripke models $\mathcal{M} = (W, R, V)$ where W is a non-empty set, $R \subseteq W \times W$ is an accessibility relation, and $V : \mathcal{P} \to 2^W$ is a valuation mapping propositional atoms to the set of worlds they hold in. The relation $\models$ between models, worlds, and formulae is defined in the usual way. Let G, H be K-formulae. Then G is *satisfiable* iff there exists a Kripke model $\mathcal{M} = (W, R, V)$ and a world $w \in W$ with $\mathcal{M}, w \models G$. The formula G is *satisfiable w.r.t. the global axiom* H iff there exists a Kripke model $\mathcal{M} = (W, R, V)$ and a world $w \in W$ such $\mathcal{M}, w \models G$ and $\mathcal{M}, w' \models H$ for all $w' \in W$. K-satisfiability is PSPACE-complete [13], and K-satisfiability w.r.t. global axioms is EXPTIME-complete [16].

A K-formula is in *negation normal form* (NNF) if $\neg$ occurs only in front of propositional atoms. Every K-formula can be transformed (in linear time) into an equivalent formula in NNF using de Morgan's laws and the duality of the modal operators.

For the automata and calculi considered here, sub-formulae of G play an important role and we will often need operations going from a formula to its super- or sub-formulae. As observed in [19], these operations become easier when dealing with "addresses" of sub-formulae in G rather than with the sub-formulae themselves.

Definition 1 (G-Paths) *For a K-formula G in NNF, the* set of G-paths Π_G *is a set of words over the alphabet* $\{\vee_l, \vee_r, \wedge_l, \wedge_r, \Box, \Diamond\}$. *The set Π_G and the sub-formula $G|_\pi$ of G addressed by $\pi \in \Pi_G$ are defined inductively as follows:*

- $\epsilon \in \Pi_G$ *and* $G|_\epsilon = G$
- *if* $\pi \in \Pi_G$ *and*
  - $G|_\pi = F_1 \wedge F_2$ *then* $\pi\wedge_l, \pi\wedge_r \in \Pi_G$, $G|_{\pi\wedge_l} = F_1$, $G|_{\pi\wedge_r} = F_2$, *and* π *is called* $\wedge$-path
  - $G|_\pi = F_1 \vee F_2$ *then* $\pi\vee_l, \pi\vee_r \in \Pi_G$, $G|_{\pi\vee_l} = F_1$, $G|_{\pi\vee_r} = F_2$, *and* π *is called* $\vee$-path

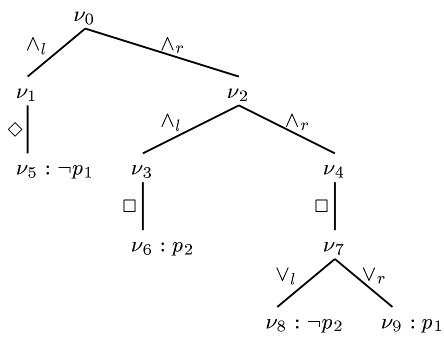

Figure 1: The set Π_G for $G = \Diamond\neg p_1 \wedge (\Box p_2 \wedge \Box(\neg p_2 \vee p_1))$

– $G|_\pi = \Box F$ *then* $\pi\Box \in \Pi_G$, $G|_{\pi\Box} = F$ *and* π *is called* $\Box$-path

– $G|_\pi = \Diamond F$ *then* $\pi\Diamond \in \Pi_G$, $G|_{\pi\Diamond} = F$ *and* π *is called* $\Diamond$-path

- Π_G *is the smallest set that satisfies the previous conditions.*

We use of $\wedge_*$ and $\vee_*$ as placeholders for $\wedge_l, \wedge_r$ and $\vee_l, \vee_r$, respectively. Also, we use $\wedge\!\!\!\!\vee$ and $\Box\!\!\!\!\Diamond$ as placeholders for $\wedge, \vee$ and $\Box, \Diamond$, respectively. If π is an $\wedge$- or and $\vee$-path then π is called $\wedge\!\!\!\!\vee$*-path*. If π is a $\Box$- or a $\Diamond$-path then π is called $\Box\!\!\!\!\Diamond$*-path*. Figure 1 shows an example of a K-formula G and the corresponding set Π_G, which can be read off the edge labels. For example, $\wedge_r\wedge_r$ is a G-path and $G|_{\wedge_r\wedge_r} = \Box(\neg p_2 \vee p_1)$

Looping Automata

For a natural number n, let $[n]$ denote the set $\{1, \ldots, n\}$. An *n-ary infinite tree over the alphabet* Σ is a mapping $t : [n]^* \to \Sigma$. An *n-ary looping tree automaton* is a tuple $\mathfrak{A} = (Q, \Sigma, I, \Delta)$, where Q is a finite set of states, Σ is a finite alphabet, $I \subseteq Q$ is the set of initial states, and $\Delta \subseteq Q \times \Sigma \times Q^n$ is the transition relation. Sometimes, we will view Δ as a function from $Q \times \Sigma$ to 2^{Q^n} and write $\Delta(q, \sigma)$ for the set $\{\mathbf{q} \mid (q, \sigma, \mathbf{q}) \in \Delta\}$. A *run* of $\mathfrak{A}$ on a tree t is a n-ary infinite tree r over Q such that

$$(r(p), t(p), (r(p1), \ldots, r(pn))) \in \Delta$$

for every $p \in [n]^*$. The automaton $\mathfrak{A}$ *accepts* t iff there is a run r of $\mathfrak{A}$ on t such that $r(\epsilon) \in I$. The set $L(\mathfrak{A}) := \{t \mid \mathfrak{A} \text{ accepts } t\}$ is the *language* accepted by $\mathfrak{A}$.

Since looping tree automata are special Büchi tree automata, emptiness of their accepted language can effectively be tested using the well-known (quadratic) emptiness test for Büchi automata [17]. However, for looping tree automata this algorithm can be specialized into a simpler (linear) one. Though this is well-known in the automata theory community, there appears to be no reference for the result.

Intuitively, the algorithm works by computing inactive states. A state $q \in Q$ is *active* iff there exists a tree t and a run of $\mathfrak{A}$ on t in which q occurs; otherwise, q is *inactive*. It is easy to see that a looping tree automaton accepts at least one tree iff it has an active initial state. How can the set of inactive states be computed? Obviously, a state from which no successor states are reachable is inactive. Moreover, a state is inactive if every transition possible from that state involves an inactive state. Thus, one can start with the set

$$Q_0 := \{q \in Q \mid \forall \sigma \in \Sigma . \Delta(q, \sigma) = \emptyset\}$$

of obviously inactive states, and then propagate inactiveness through the transition relation. We formalize this propagation process in a way that allows for an easy formulation of our main results.

A *derivation* of the emptiness test is a sequence $Q_0 \rhd Q_1 \rhd \ldots \rhd Q_k$ such that $Q_i \subseteq Q$ and $Q_i \rhd Q_{i+1}$ iff $Q_{i+1} = Q_i \cup \{q\}$ with

$$q \in \{q' \in Q \mid \forall \sigma \in \Sigma . \forall (q_1, \ldots, q_n) \in \Delta(q, \sigma) . \exists j . q_j \in Q_i\}.$$

We write $Q_0 \rhd^* P$ iff there is a $k \in \mathbb{N}$ and a derivation $Q_0 \rhd \ldots \rhd Q_k$ with $P = Q_k$. The emptiness test answers "$L(\mathfrak{A}) = \emptyset$" iff there exists a set of states P such that $Q_0 \rhd^* P$ and $I \subseteq P$.

Note that $Q \rhd P$ implies $Q \subseteq P$ and that $Q \subseteq Q'$ and $Q \rhd P$ imply $Q' \rhd^* P$. Consequently, the *closure* $Q_0^{\rhd}$ of Q_0 under $\rhd$, defined by $Q_0^{\rhd} =: \bigcup \{P \mid Q_0 \rhd P\}$, can be calculated starting with Q_0, and successively adding states q to the current set Q_i such that $Q_i \rhd Q_i \cup \{q\}$ and $q \notin Q_i$, until no more states can be added. It is easy to see that this closure consists of the set of inactive states, and thus $L(\mathfrak{A}) = \emptyset$ iff $I \subseteq Q_0^{\rhd}$. As described until now, this algorithm runs in time polynomial in the number of states. By using clever data structures and a propagation algorithm similar to the one for satisfiability of propositional Horn formulae [7], one can in fact obtain a linear emptiness test for looping tree automata.

3 Automata, Modal Formulae, and the Inverse Calculus

We first describe how to decide satisfiability in K using the automata approach and the inverse method, respectively. Then we show that both approaches are closely connected.

3.1 Automata and Modal Formulae

Given a K-formula G, we define an automaton $\mathfrak{A}_G$ such that $L(\mathfrak{A}_G) = \emptyset$ iff G is not satisfiable. In contrast to the "standard" automata approach, the states of our automaton $\mathfrak{A}_G$ will be subsets of Π_G rather than sets of subformulae of G. Using paths instead of subformulae is mostly a matter of notation. We also require the states to satisfy additional properties (i.e., we do not allow for arbitrary subsets of Π_G). This makes the proof of correctness of the automata approach only slightly more complicated, and it allows us to treat some important optimisations of the inverse calculus within our framework. The next definition introduces these properties.

Definition 2 (Propositionally expanded, clash) *Let G be a K-formula in NNF, Π_G the set of G-paths, and $\Phi \subseteq \Pi_G$. An $\wedge$-path $\pi \in \Phi$ is* propositionally expanded in Φ *iff $\{\pi\wedge_l, \pi\wedge_r\} \subseteq \Phi$. An $\vee$-path $\pi \in \Phi$ is* propositionally expanded in Φ *iff $\{\pi\vee_l, \pi\vee_r\} \cap \Phi \neq \emptyset$. The set Φ is* propositionally expanded *iff every $\wedge\!\!\!\vee$-path $\pi \in \Phi$ is propositionally expanded in Φ. We use "p.e." as an abbreviation for "propositionally expanded".*

The set Φ' is an expansion *of the set Φ if $\Phi \subseteq \Phi'$, Φ' is p.e. and Φ' is minimal w.r.t. set inclusion with these properties. For a set Φ, we define the set of its expansions as $\langle\!\langle \Phi \rangle\!\rangle := \{\Phi' \mid \Phi'$ is an expansion of $\Phi\}$.*

Φ contains a clash *iff there are two paths $\pi_1, \pi_2 \in \Phi$ such that $G|_{\pi_1} = p$ and $G|_{\pi_2} = \neg p$ for a propositional variable p. Otherwise, Φ is called* clash-free.

For a set of paths Ψ, the set $\langle\!\langle \Psi \rangle\!\rangle$ can effectively be constructed by successively adding paths required by the definition of p.e. A formal construction of the closure can be found in the proof of Lemma 4. Note that $\emptyset$ is p.e., clash-free, and $\langle\!\langle \emptyset \rangle\!\rangle = \{\emptyset\}$.

Definition 3 (Formula Automaton) *For a* K*-formula* G *in NNF, we fix an arbitrary enumeration* $\{\pi_1, \ldots, \pi_n\}$ *of the* $\Diamond$*-paths in* Π_G*. The* n*-ary looping automaton* $\mathfrak{A}_G$ *is defined by* $\mathfrak{A}_G := (Q_G, \Sigma_G, \langle\!\langle\{\epsilon\}\rangle\!\rangle, \Delta_G)$*, where* $Q_G := \Sigma_G := \{\Phi \subseteq \Pi_G \mid \Phi \text{ is p.e.}\}$ *and the transition relation* Δ_G *is defined as follows:*

- Δ_G *contains only tuples of the form* $(\Phi, \Phi, \ldots)$*.*
- *If* Φ *is clash-free, then we define* $\Delta_G(\Phi, \Phi) := \langle\!\langle\Psi_1\rangle\!\rangle \times \cdots \times \langle\!\langle\Psi_n\rangle\!\rangle$*, where*

$$\Psi_i = \begin{cases} \{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in \Phi \text{ is a } \Box\text{-path } \} & \text{if } \pi_i \in \Phi \\ \emptyset & \text{else} \end{cases}$$

- *If* Φ *contains a clash, then* $\Delta_G(\Phi, \Phi) = \emptyset$*, i.e., there is no transition from* Φ*.*

Note, that this definition implies $\Delta_G(\emptyset, \emptyset) = \{(\emptyset, \ldots, \emptyset)\}$ and only states with a clash have no successor states.

Theorem 1 *For a* K*-formula* G*,* G *is satisfiable iff* $L(\mathfrak{A}_G) \neq \emptyset$*.*

This theorem can be proved by showing that i) every tree accepted by $\mathfrak{A}_G$ induces a model of G; and ii) every model $\mathcal{M}$ of G can be turned into a tree accepted by $\mathfrak{A}_G$ by a) unraveling $\mathcal{M}$ into a tree model $\mathcal{T}$ for G; b) labeling every world of $\mathcal{T}$ with a suitable p.e. set depending on the formulae that hold in this world; and c) padding "holes" in $\mathcal{T}$ with $\emptyset$.

Proof. Let $\{\pi_1, \ldots, \pi_n\}$ be an enumeration of the $\Diamond$-paths in Π_G.

For the *if*-direction let $L(\mathfrak{A}_G) \neq \emptyset$, $t, r : [n]^* \to \{\Phi \subseteq \Pi_G \mid \Phi \text{ is p.e.}\}$ a tree that is accepted by $\mathfrak{A}_G$ and a corresponding run of $\mathfrak{A}_G$. By construction of $\mathfrak{A}_G$, $t(w) = r(w)$ for every $w \in [n]^*$. We construct a Kripke model $\mathcal{M} = (W, R, V)$ from t by setting

$$\begin{aligned} W &= \{w \in [n]^* \mid t(w) \neq \emptyset\} \\ R &= \{(w, wi) \in W \times W \mid i \in [n]\} \\ V &= \lambda P.\{p \in W \mid \exists \pi \in t(w).G|_\pi = P\} \quad \text{for all propositional atoms } P \end{aligned}$$

Claim. For all $w \in W$, if $\pi \in t(w)$ then $\mathcal{M}, w \models G|_\pi$.

Proof of the claim. The claim is proved by induction on the structure of K-formulae. Let $w \in W$ be a world and $\pi \in \Pi_G$ be a path such that $\pi \in t(w)$.

- if $G|_\pi = P$ is a propositional atom and $w \in W$, then $w \in V(P)$ and hence $\mathcal{M}, w \models G|_\pi$.

- if $G|_\pi = \neg P$ is a negated propositional atom, then, since $t(w)$ is clash free, there is no $\pi' \in t(w)$ such that $G|_{\pi'} = P$. Thus, $w \notin V(P)$ and hence $\mathcal{M}, w \models \neg P$.

- if $G|_\pi = F_1 \wedge F_2$ then π is an $\wedge$-path, and since $t(w)$ is p.e., $\{\pi\wedge_l, \pi\wedge_r\} \subseteq t(w)$. By induction, $\mathcal{M}, w \models G|_{\pi\wedge_*}$ and hence $\mathcal{M}, w \models G|_\pi$.

- if $G|_\pi = F_1 \vee F_2$ then π is an $\vee$-path, and since $t(w)$ is p.e., $\{\pi\vee_l, \pi\vee_r\} \cap t(w) \neq \emptyset$. By induction, $\mathcal{M}, w \models G|_{\pi\vee_l}$ or $\mathcal{M}, w \models G|_{\pi\vee_r}$ and hence $\mathcal{M}, w \models G|_\pi$.

- if $G|_\pi = \Diamond F$ then π is a $\Diamond$-path and, w.o.l.g., assume $\pi = \pi_i$. Since $\pi_i \in r(w)$, $\pi_i\Diamond \in r(wi) = t(wi)$ holds and hence $wi \in W$ and $(w, wi) \in R$. By induction, we have that $\mathcal{M}, wi \models G|_{\pi_i\Diamond}$ and hence $\mathcal{M}, w \models G|_{\pi_i}$.

- if $G|_\pi = \Box F$ and $(w, w') \in R$ then $w' = wi$ for some $i \in [n]$ and $t(wi) \neq \emptyset$ holds and by construction of $\mathfrak{A}_G$, this implies $\pi\Box \in r(wi) = t(wi)$. By induction, this implies $\mathcal{M}, wi \models G|_{\pi\Box}$ and since $wi = w'$ and w' has been chosen arbitrarily, $\mathcal{M}, w \models G|_\pi$.

This finishes the proof of the claim. Since $t(\epsilon) = r(\epsilon) \in \langle\langle\{\epsilon\}\rangle\rangle$ and hence $\epsilon \in t(\epsilon)$, $\mathcal{M}, \epsilon \models G|_\epsilon$ and $G = G|_\epsilon$ is satisfiable.

For the *only if*-direction, we first show an auxiliary claim: for a set $\Psi \subseteq \Pi_G$ we define $\mathcal{M}, w \models \Psi$ iff $\mathcal{M}, w \models G|_\pi$ for every $\pi \in \Psi$.

Claim. If $\Psi \subseteq \Pi_G$ and $w \in W$ such that $\mathcal{M}, w \models \Psi$, then there is a $\Phi \in \langle\langle\Psi\rangle\rangle$ such that $\mathcal{M}, w \models \Phi$.

Proof of the claim. Let $\Psi \subseteq \Pi_G$ and $w \in W$ such that $\mathcal{M}, w \models \Psi$. We will show how to construct an expansion of Ψ with the desired property. If Ψ is already p.e., then $\Psi \in \langle\langle\Psi\rangle\rangle$ and we are done. If Ψ is not p.e. then let $\pi \in \Psi$ be a $\wedge\!\!\!\vee$-path that is not p.e. in Ψ.

- If π is a $\wedge$-path then $G|_\pi = F_1 \wedge F_2$ and since $\mathcal{M}, w \models G|_\pi$, also $\mathcal{M}, w \models F_1 = G|_{\pi\wedge_l}$ and $\mathcal{M}, w \models F_2 = G|_{\pi\wedge_r}$. Hence $\mathcal{M}, w \models \Psi \cup \{\pi\wedge_l, \pi\wedge_r\}$

and $\Psi' = \Psi \cup \{\pi\wedge_l, \pi\wedge_r\}$ is a set with $\mathcal{M}, w \models \Psi'$ that is"one step closer" to being p.e. than Ψ.

- If π is a $\vee$-path then $G|_\pi = F_1 \vee F_2$ and since $\mathcal{M}, w \models G|_\pi$, also $\mathcal{M}, w \models F_1 = G|_{\pi\vee_l}$ or $\mathcal{M}, w \models F_2 = G|_{\pi\vee_r}$. Hence $\mathcal{M}, w \models \Psi \cup \{\pi\vee_l\}$ or $\mathcal{M}, w \models \Psi \cup \{\pi\vee_r\}$ and hence can obtain a set Ψ' with $\mathcal{M}, w \models \Psi'$ that is again "one step close" to being p.e. than Ψ.

Restarting this process with $\Psi = \Psi'$ eventually yields an expansion Φ of the initial set Ψ with $\mathcal{M}, w \models \Phi$, which proves the claim.

Let $\mathcal{M} = (W, R, V)$ be a model for G with $w \in W$ such that $\mathcal{M}, w \models G$. From $\mathcal{M}$ we construct a tree that is accepted by $\mathfrak{A}_G$. Using this claim, we inductively define a tree t accepted by $\mathfrak{A}_G$. To this purpose, we also inductively define a function $f : [n]^* \to W$ such that, if $\mathcal{M}, f(p) \models t(p)$ for all p.

We start by setting $f(\epsilon) = w$ for a $w \in W$ with $\mathcal{M}, w \models G$. and $t(\epsilon) = \Phi$ for a $\Phi \in \langle\langle\{\epsilon\}\rangle\rangle$ such that $\mathcal{M}, w \models \Phi$. From the claim we have that such a set Φ exists because $\mathcal{M}, w \models G = G|_\epsilon$.

If $f(p)$ and $t(p)$ are already defined, then, for $i \in [n]$, we define $f(pi)$ and $t(pi)$ as follows:

- if $\pi_i \in t(p)$ then $\mathcal{M}, f(p) \models G|_{\pi_i}$ and hence there is a $w' \in W$ such that $(f(p), w') \in R$ and $\mathcal{M}, w' \models G|_{\pi_i\Diamond}$. If $\pi \in t(p)$ is a $\Box$-path, then also $\mathcal{M}, w' \models G|_{\pi\Box}$ holds. Hence $\mathcal{M}, w' \models \{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in t(p) \text{ is a } \Box\text{-path }\}$. We set $f(pi) = w'$ and $t(pi) = \Phi$ for a $\Phi \in \langle\langle\{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in t(p) \text{ is a } \Box\text{-path }\}\rangle\rangle$ with $\mathcal{M}, w' \models \Phi$, which exist by the claim.

- if $\pi_i \notin t(p)$, then we set $f(pi) = w$ for an arbitrary $w \in W$ and $t(pi) = \emptyset$.

In both cases, we have define $f(pi)$ and $t(pi)$ such that $\mathcal{M}, f(pi) \models t(pi)$. It is easy to see that t is accepted by $\mathfrak{A}_G$ with the run $r = t$. Hence $L(\mathfrak{A}_G) \neq \emptyset$ which is what we needed to show. ■

Together with the emptiness test for looping tree automata, Theorem 1 yields a decision procedure for K-satisfiability. To test a K-formula G for unsatisfiability, construct $\mathfrak{A}_G$ and test whether $L(\mathfrak{A}_G) = \emptyset$ holds using the emptiness test for looping tree automata: $L(\mathfrak{A}_G) = \emptyset$ iff $\langle\langle\{\epsilon\}\rangle\rangle \subseteq Q_0^{\triangleright}$, where

$Q_0 \subseteq Q_G$ is the set of states containing a clash. The following is a derivation of a superset of $\langle\langle\{\epsilon\}\rangle\rangle$ from Q_0 for the example formula from Figure 1:

$$Q_0 = \{\underbrace{\{\nu_5,\nu_6,\nu_7,\nu_8\},\ \{\nu_5,\nu_6,\nu_7,\nu_9\}}_{=\langle\langle\nu_5,\nu_6,\nu_7\rangle\rangle},\ldots\} \rhd Q_0 \cup \underbrace{\{\{\nu_0,\nu_1,\nu_2,\nu_3,\nu_4\}\}}_{=\langle\langle\{\epsilon\}\rangle\rangle}$$

3.2 The Inverse Calculus

In the following, we introduce the inverse calculus for K. We stay close to the notation and terminology used in [19].

A *sequent* is a subset of Π_G. Sequents will be denoted by capital greek letters. The union of two sequents Γ and Λ is denote by Γ, Λ. If Γ is a sequent and $\pi \in \Pi_G$ then we denote $\Gamma \cup \{\pi\}$ by Γ, π. If Γ is a sequent that contains only $\Box$-paths then we write $\Gamma\Box$ to denote the sequent $\{\pi\Box \mid \pi \in \Gamma\}$. Since states of $\mathfrak{A}_G$ are also subsets of Π_G and hence sequents, we will later on use the same notational conventions for states as for sequents.

Definition 4 (The inverse path calculus) *Let G be a formula in NNF and Π_G the set of paths of G.* Axioms *of the inverse calculus are all sequents $\{\pi_1, \pi_2\}$ such that $G|_{\pi_1} = p$ and $G|_{\pi_2} = \neg p$ for some propositional variable p. The* rules *of the inverse calculus are given in Figure 2, where all paths occurring in a sequent are G-paths and, for every $\Diamond^+$ inference, π is a $\Diamond$-path. We refer to this calculus by IC_G.*[1]

We define $\mathcal{S}_0 := \{\Gamma \mid \Gamma$ is an axiom $\}$. A derivation *of IC_G is a sequence of sets of sequents $\mathcal{S}_0 \vdash \cdots \vdash \mathcal{S}_m$ where $\mathcal{S}_i \vdash \mathcal{S}_{i+1}$ iff $\mathcal{S}_{i+1} = \mathcal{S}_i \cup \{\Gamma\}$ such that there exists sequents $\Gamma_1, \ldots \Gamma_k \in \mathcal{S}_i$ and $\dfrac{\Gamma_1 \quad \ldots \quad \Gamma_k}{\Gamma}$ is an inference.*

We write $\mathcal{S}_0 \vdash^* \mathcal{S}$ iff there is a derivation $\mathcal{S}_0 \vdash \cdots \vdash \mathcal{S}_m$ with $\mathcal{S} = \mathcal{S}_m$. The *closure* $\mathcal{S}_0^{\vdash}$ of $\mathcal{S}_0$ under $\vdash$ is defined by $\mathcal{S}_0^{\vdash} = \bigcup\{\mathcal{S} \mid \mathcal{S}_0 \vdash^* \mathcal{S}\}$. Again, the closure can effectively be computed by starting with $\mathcal{S}_0$ and then adding sequents that can be obtained by an inference until no more new sequents can be added.

As shown in [19], the computation of the closure yields a decision procedure for K-satisfiability:

Fact 1 *G is unsatisfiable iff $\{\epsilon\} \in \mathcal{S}_0^{\vdash}$.*

[1] G appears in the subscript because the calculus is highly dependent of the input formula G: only G-paths can be generated by IC_G.

$$(\vee)\frac{\Gamma_l, \pi\vee_l \qquad \Gamma_r, \pi\vee_r}{\Gamma_l, \Gamma_r, \pi} \qquad (\wedge_l)\frac{\Gamma, \pi\wedge_l}{\Gamma, \pi} \qquad (\wedge_r)\frac{\Gamma, \pi\wedge_r}{\Gamma, \pi}$$

$$(\Diamond)\frac{\Gamma\Box, \pi\Diamond}{\Gamma, \pi} \qquad (\Diamond+)\frac{\Gamma\Box}{\Gamma, \pi}$$

Figure 2: Inference rules of IC_G

Figure 3 shows the inferences of IC_G that lead to $\nu_0 = \epsilon$ for the example formula from Figure 1.

3.3 Connecting the Two Approaches

The results shown in this subsection imply that IC_G can be viewed as an on-the-fly implementation of the emptiness for $\mathfrak{A}_G$. In addition to generating states on-the-fly, states are also represented in a compact manner: one sequent generated by IC_G represents several states of $\mathfrak{A}_G$.

Definition 5 *For the formula automaton $\mathfrak{A}_G$ with states Q_G and a sequent $\Gamma \subseteq \Pi_G$ we define $[\![\Gamma]\!] := \{\Phi \in Q_G \mid \Gamma \subseteq \Phi\}$, and for a set $\mathcal{S}$ of sequents we define $[\![\mathcal{S}]\!] := \bigcup_{\Gamma\in\mathcal{S}} [\![\Gamma]\!]$.*

The following theorem, which is one of the main contributions of this paper, establishes the correspondence between the emptiness test and IC_G.

Theorem 2 (IC_G and the emptiness test mutually simulate each other) *Let Q_0, $\mathcal{S}_0$, $\rhd$, and $\vdash$ be defined as above.*

1. *Let Q be a set of states such that $Q_0 \rhd^* Q$. Then there exists a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash^* \mathcal{S}$ and $Q \subseteq [\![\mathcal{S}]\!]$.*
2. *Let $\mathcal{S}$ be a set of sequents such that $\mathcal{S}_0 \vdash^* \mathcal{S}$. Then there exists a set of states $Q \subseteq Q_G$ with $Q_0 \rhd^* Q$ and $[\![\mathcal{S}]\!] \subseteq Q$.*

The first part of the theorem shows that IC_G can simulate each computation of the emptiness test for $\mathfrak{A}_G$. The set of states represented by the set of sequents computed by IC_G may be larger than the one computed by a particular derivation of the emptiness test. However, the second part of the theorem implies that all these states are in fact inactive since a possibly larger set of states can also be computed by a derivation of the emptiness test. In particular, the theorem implies that IC_G can be used to calculate a compact

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\wedge_l\Diamond,\ \wedge_r\wedge_r\Box\vee_r \quad | \quad \wedge_r\wedge_l\Box,\ \wedge_r\wedge_r\Box\vee_l}{\wedge_l\Diamond,\ \wedge_r\wedge_l\Box,\ ,\wedge_r\wedge_r\Box}(\vee)}{\wedge_l,\ \wedge_r\wedge_l,\ \wedge_r\wedge_r}(\Diamond)}{\wedge_l,\ \wedge_r,\ \wedge_r\wedge_l}(\wedge_r)}{\wedge_l,\ \wedge_r}(\wedge_l)}{\epsilon,\ \wedge_l}(\wedge_r)}{\epsilon}(\wedge_l)
$$

Figure 3: An example of inferences in IC_G

representation of $Q_0^{\triangleright}$. This is an on-the-fly computation since $\mathfrak{A}_G$ is never constructed explicitly.

Corollary 1 $Q_0^{\triangleright} = [\![\mathcal{S}_0^{\vdash}]\!]$.

Proof. If $\Phi \in Q_0^{\triangleright}$ then there exists a set of states Q such that $Q_0 \triangleright^* Q$ and $\Phi \in Q$. By Theorem 2.1, there exists a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash^* \mathcal{S}$ and $Q \subseteq [\![\mathcal{S}]\!]$. Hence $\Phi \in [\![\mathcal{S}_0^{\vdash}]\!]$. For the converse direction, if $\Phi \in [\![\mathcal{S}_0^{\vdash}]\!]$ then there exists a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash^* \mathcal{S}$ and $\Phi \in [\![\mathcal{S}]\!]$. By Theorem 2.2, there exists a set of states Q with $Q_0 \triangleright^* Q$ and $[\![\mathcal{S}]\!] \subseteq Q$ and hence $\Phi \in Q_0^{\triangleright}$. ∎

The proof of the second part of Theorem 2 is the easier one. It is a consequence of the next three lemmata. First, observe that the two calculi have the same starting points.

Lemma 1 *If $\mathcal{S}_0$ is the set of axioms of IC_G, and Q_0 is the set of states of $\mathfrak{A}_G$ that have no successor states, then $[\![\mathcal{S}_0]\!] = Q_0$.*

Proof. The set $\mathcal{S}_0$ is the set of all axioms i.e., the set of all clashes. Hence $[\![\mathcal{S}_0]\!] = \{\Phi \mid \Phi \text{ contains a clash}\} = Q_0$. ∎

Second, since states are assumed to be p.e., propositional inferences of IC_G do not change the set of states represented by the sequents.

Lemma 2 *Let $\mathcal{S} \vdash \mathcal{T}$ be a derivation of IC_G that employs a $\wedge_l$-, $\wedge_r$-, or a $\vee$-inference. Then $[\![\mathcal{S}]\!] = [\![\mathcal{T}]\!]$.*

Proof. Since $\mathcal{S} \subseteq \mathcal{T}$, $[\![\mathcal{S}]\!] \subseteq [\![\mathcal{T}]\!]$ holds immediately. To show $[\![\mathcal{T}]\!] \subseteq [\![\mathcal{S}]\!]$, we distinguish the different inferences used to obtain $\mathcal{T}$ from $\mathcal{S}$:

- If the employed inference is $(\wedge_*)\dfrac{\Gamma, \pi\wedge_*}{\Gamma, \pi}$ and $\mathcal{T} = \mathcal{S} \cup \{\Gamma, \pi\}$ with $\Gamma, \pi\wedge_* \in \mathcal{S}$. Then $[\![\mathcal{T}]\!] = [\![\mathcal{S}]\!] \cup [\![\Gamma, \pi]\!]$. Let $\Phi \in [\![\Gamma, \pi]\!]$. Φ is p.e. and hence $\pi \in \Phi$ implies $\pi\wedge_* \in \Phi$. Thus, $\Gamma, \pi\wedge_* \subseteq \Phi$ and $\Phi \in [\![\Gamma, \pi\wedge_*]\!] \subseteq [\![\mathcal{S}]\!]$.
- Assume that the employed inference is $(\vee)\dfrac{\Gamma_l, \pi\vee_l \qquad \Gamma_r, \vee_r}{\Gamma_l, \Gamma_r, \pi}$ and $\mathcal{T} = \mathcal{S} \cup \{\Gamma_l, \Gamma_r, \pi\}$ with $\Gamma_l, \pi\vee_l \in \mathcal{S}$, $\Gamma_r, \vee_r \in \mathcal{S}$. Then $[\![\mathcal{T}]\!] = [\![\mathcal{S}]\!] \cup [\![\Gamma_l, \Gamma_r, \pi]\!]$. Let $\Phi \in [\![\Gamma_l, \Gamma_r, \pi]\!]$. Φ is p.e. and hence, w.o.l.g., $\pi\vee_l \in \Phi$. Thus, $\Gamma_l, \pi\vee_l \subseteq \Phi$ and $\Phi \in [\![\Gamma_l, \pi\vee_l]\!] \subseteq [\![\mathcal{S}]\!]$. ∎

Third, modal inferences of IC_G can be simulated by derivations of the emptiness test.

Lemma 3 *Let $\mathcal{S} \vdash \mathcal{T}$ be derivation of IC_G that employs a $\Diamond$- or $\Diamond^+$-inference. If Q is a set of states with $[\![\mathcal{S}]\!] \cup Q_0 \subseteq Q$ then there exists a set of states P with $Q \rhd^* P$ and $[\![\mathcal{T}]\!] \subseteq P$.*

Proof. We only consider the $\Diamond$-inference, the case of a $\Diamond^+$-inference is analogous. If $\mathcal{S} \vdash \mathcal{T}$ by an application of a $\Diamond$-inference, then $\mathcal{T} = \mathcal{S} \cup \{\Gamma, \pi\}$ where Γ consists only of $\Box$-paths, π is a $\Diamond$-path (w.o.l.g., we assume $\pi = \pi_i$, the i-th path in the enumeration of $\Diamond$-paths in Π_G), $\Gamma\Box, \pi_i\Diamond \in \mathcal{S}$ and $(\Diamond)\dfrac{\Gamma\Box, \pi_i\Diamond}{\Gamma, \pi_i}$. Also, $[\![\mathcal{T}]\!] = [\![\mathcal{S}]\!] \cup [\![\Gamma, \pi_i]\!]$ holds.

Claim. Let $\Phi \in [\![\Gamma, \pi_i]\!]$ and R a set of states with $[\![\Gamma\Box, \pi_i\Diamond]\!] \cup Q_0 \subseteq R$. Then there exists a derivation $R \rhd^* R'$ with $\Phi \in R'$ and $[\![\Gamma\Box, (\pi_i\Diamond)]\!] \cup Q_0 \subseteq R'$

Proof of the Claim. If Φ contains a clash then $\Phi \in Q_0 \subseteq R$ and nothing has to be done. If Φ does not contain a clash, then $\Delta_G(\Phi, \Phi) = \langle\!\langle \Psi_i \rangle\!\rangle \times \cdots \times \langle\!\langle \Psi_n \rangle\!\rangle$ where the Ψ_i are defined as in Definition 3 and especially, since $\pi_i \in \Phi$,

$$\langle\!\langle \Psi_i \rangle\!\rangle = \langle\!\langle \underbrace{\{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in \Phi \text{ is a } \Box\text{-path }\}}_{\supseteq \Gamma\Box, \pi_i\Diamond} \rangle\!\rangle \subseteq [\![\Gamma\Box, \pi_i\Diamond]\!] \subseteq R$$

Since all states in $\langle\!\langle \Psi_i \rangle\!\rangle$ have been marked inactive, the emptiness test can also mark Φ inactive and derive $R \rhd R \cup \{\Phi\} = R'$, which proves the claim.

Using this claim, we prove the lemma as follows. Let $\Phi_i, \dots \Phi_k$ be an enumeration of $[\![\Gamma, \pi_i]\!]$. The set $P_0 = Q$ satisfies the requirements of the claim for R. Thus, we repeatedly use the claim and chain the derivations to obtain a derivation $Q = P_0 \rhd P_1 \rhd \dots \rhd P_k = P$ such that $\Phi_i \in P_i$. Since the sets grow monotonically, in the end $[\![\Gamma, \pi]\!] \subseteq P$ holds, which implies $[\![\mathcal{T}]\!] \subseteq P$. ∎

Given these lemmata, proving Theorem 2.2 is quite simple.

Proof of Theorem 2.2. The proof is by induction on the length m of the derivation $\mathcal{S}_0 \vdash \mathcal{S}_1 \cdots \vdash \mathcal{S}_m = \mathcal{S}$ of IC_G. The base case $m = 0$ is Lemma 1. For the induction step, $\mathcal{S}_{i+1}$ is either inferred from $\mathcal{S}_i$ using a propositional inference, which is dealt with by Lemma 2, or by a modal inference, which is dealt with by Lemma 3. Lemma 3 is applicable since, for every set of states Q with $Q_0 \rhd^* Q$, $Q_0 \subseteq Q$. ∎

Proving the first part of Theorem 2 is more involed because of the calculation of the propositional expansions implicit in the definition of $\mathfrak{A}_G$.

Lemma 4 *Let $\Phi \subseteq \Pi_G$ be a set of paths and $\mathcal{S}$ a set of sequents such that $\langle\!\langle \Phi \rangle\!\rangle \subseteq [\![\mathcal{S}]\!]$. Then there exists a set of sequents $\mathcal{T}$ with $\mathcal{S} \vdash^* \mathcal{T}$ such that there exists a sequent $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \Phi$.*

Proof. If Φ is p.e., then this is immediate, as in this case $\langle\!\langle \Phi \rangle\!\rangle = \{\Phi\} \subseteq [\![\mathcal{S}]\!]$.

If Φ is not p.e., then let select be an arbitrary *selection function*, i.e., a function that maps every set Ψ that is not p.e. to a $\mathbb{X}$-path $\pi \in \Psi$ that is not p.e. in Ψ. Let T_Φ be the following, inductively defined tree:

- The root of T_Φ is Φ.
- If a node Ψ of T_Φ is not p.e., then
  - if $\mathsf{select}(\Psi) = \pi$ is an $\wedge$-path, then Ψ has the successor node $\Psi, \pi\wedge_l, \pi\wedge_r$ and Ψ is called an $\wedge$-node.
  - if $\mathsf{select}(\Psi) = \pi$ is an $\vee$-path, then Ψ has the successor nodes $\Psi, \pi\vee_l$ and $\Psi, \pi\vee_l$ and Ψ is called an $\vee$-node.
- If a node Ψ of T_Φ is p.e., then it is a leaf of the tree.

Obviously, the construction is such that the set of leaves of T_Φ is $\langle\!\langle\Phi\rangle\!\rangle$.

Let $\Upsilon_1, \ldots \Upsilon_\ell$ be a post-order traversal of this tree, so the sons of a node occur before the node itself and $\Upsilon_\ell = \Phi$. Along this traversal we will construct a derivation $\mathcal{S} = \mathcal{T}_0 \vdash^* \cdots \vdash^* \mathcal{T}_\ell = \mathcal{T}$ such that, for every $1 \leq i \leq j \leq \ell$, $\mathcal{T}_j$ contains a sequent Λ_i with $\Lambda_i \subseteq \Upsilon_i$. Since the sets $\mathcal{T}_j$ grow monotonically, it su ces to show that, for every $1 \leq i \leq \ell$, $\mathcal{T}_i$ contains a sequent Λ_i with $\Lambda_i \subseteq \Upsilon_i$.

Whenever Υ_i is a leaf of T_Φ, then $\Upsilon_i \in \langle\!\langle\Phi\rangle\!\rangle \subseteq [\![\mathcal{S}]\!]$. Hence there is already a sequent $\Lambda_i \in \mathcal{T}_0$ with $\Lambda_i \subseteq \Upsilon_i$ and no derivation step is necessary. Particularly, in a post-order traversal, Υ_1 is a leaf.

We now assume that the derivation has been constructed up to $\mathcal{T}_i$.

- If Υ_{i+1} is a leaf of T_Φ, then nothing has to be done as there exists a $\Lambda_{i+1} \in \mathcal{T}_0 \subseteq \mathcal{T}_i$ with $\Lambda_{i+1} \subseteq \Upsilon_{i+1}$

- If Υ_{i+1} is an $\wedge$-node with selected $\wedge$-path $\pi \in \Upsilon_{i+1}$. Then, the successor of Υ_{i+1} in T_Φ is $\Upsilon_{i+1}\pi\wedge_l, \pi\wedge_r$ and appears before Υ_{i+1} in the traversal. By construction there exists a sequent $\Lambda \in \mathcal{T}_i$ with $\Lambda \subseteq \Upsilon_{i+1}, \pi\wedge_l, \pi\wedge_r$. If $\Lambda \cap \{\pi\wedge_l, \pi\wedge_r\} = \emptyset$ then we are done because then also $\Lambda \subseteq \Upsilon_{i+1}$. If one or both of $\pi\wedge_l, \pi\wedge_r$ occur in Λ, then

  - if $\Lambda = \Gamma, \pi\wedge_l$ for some Γ with $\pi\wedge_r \notin \Gamma$ then this implies that the inference
  $$(\wedge_l)\frac{\Gamma, \pi\wedge_l}{\Gamma, \pi} \tag{1}$$
  can be used to derive $\mathcal{T}_i \vdash \mathcal{T}_i \cup \{\Gamma, \pi\} = \mathcal{T}_{i+1}$ and $\Gamma, \pi \subseteq \Upsilon_{i+1}$ holds.

  - the case $\Lambda = \Gamma, \pi\wedge_r$ for some Γ with $\pi\wedge_l \notin \Gamma$ if analogous.

  - if $\Lambda = \Gamma, \pi\wedge_l, \pi\wedge_r$ for some Γ with $\{\pi\wedge_l, \pi\wedge_r\} \cap \Gamma = \emptyset$ then the inferences
  $$(\wedge_l)\frac{\Gamma, \pi\wedge_l, \pi\wedge_r}{(\wedge_r)\dfrac{\Gamma, \pi, \pi\wedge_r}{\Gamma, \pi, \pi}} \tag{2}$$
  can be used in the derivation $\mathcal{T}_i \vdash \mathcal{T}_i \cup \{\Gamma, \pi, \pi\wedge_r\} \vdash \mathcal{T}_i \cup \{\Gamma, \pi, \pi\wedge_r\} \cup \{\Gamma, \pi\} = \mathcal{T}_{i+1}$ and by construction $\Gamma, \pi \subseteq \Upsilon_{i+1}$ holds.

- If Υ_{i+1} is an $\vee$-node with selected $\vee$-path $\pi \in \Upsilon_{i+1}$. Then, the successors of Υ_{i+1} in T_Φ are $\Upsilon_{i+1}, \pi\vee_l$ and $\Upsilon_{i+1}, \pi\vee_r$, and by construction there exist sequences $\Lambda_l, \Lambda_r \in \mathcal{T}_i$ with $\Lambda_* \subseteq \Upsilon_{i+1}, \pi\vee_*$.

If $\pi\vee_l \notin \Lambda_l$ or $\pi\vee_r \notin \Lambda_r$, then $\Lambda_l \subseteq \Upsilon_{i+1}$ or $\Lambda_r \subseteq \Upsilon_{i+1}$ holds and hence already $\mathcal{T}_i$ contains a sequent Λ with $\Lambda \subseteq \Upsilon_{i+1}$.

If $\Lambda_l = \Gamma_l, \pi\vee_l$ and $\Lambda_r = \Gamma_r, \pi\vee_r$ with $\pi\vee_* \notin \Gamma_*$ then IC_G can use the inference

$$(\vee)\frac{\Gamma_l, \pi\vee_l \qquad \Gamma_r, \pi\vee_r}{\Gamma_l, \Gamma_r, \pi} \tag{3}$$

to derive $\mathcal{T}_i \vdash \mathcal{T}_i \cup \{\Gamma_l, \Gamma_r, \pi\} = \mathcal{T}_{i+1}$, and and $\Gamma_l, \Gamma_r, \pi \subseteq \Upsilon_{i+1}$ holds as follows: assume there is a $\pi' \in \Gamma_l, \Gamma_r, \pi$ with $\pi' \notin \Upsilon_{i+1}$. Since $\pi \in \Upsilon_{i+1}$, w.o.l.g., $\pi' \in \Gamma_l$. But then also $\Gamma_l \nsubseteq \Upsilon_{i+1}, \pi\vee_l$ would hold, since $\pi' \neq \pi\vee_l$ because $\pi\vee_l \notin \Gamma_l$.

Proceeding in this manner, starting from $\mathcal{T}_0 = \mathcal{S}$, we can construct a derivation that yields a set $\mathcal{T} = \mathcal{T}_k$ of states containing a sequent Λ such that $\Lambda \subseteq \Upsilon_\ell = \Phi$. ∎

Proof of Theorem 2.1. We show this by induction on the number k of steps in the derivation $Q_0 \triangleright \ldots \triangleright Q_k = Q$. Again, Lemma 1 yields the base case.

For the induction step, let $Q_0 \triangleright \ldots \triangleright Q_i \triangleright Q_{i+1} = Q_i \cup \{\Phi\}$ be a derivation of the emptiness test and $\mathcal{S}_i$ a set of sequents such that $\mathcal{S} \vdash^* \mathcal{S}_i$ and $Q_i \subseteq [\![\mathcal{S}_i]\!]$. Such a set exists by the induction hypothesis because the derivation $Q_0 \triangleright \ldots \triangleright Q_i$ is of length i. Now let $Q_i \triangleright Q_i \cup \{\Phi\} = Q_{i+1}$ be the derivation of the emptiness test. If already $\Phi \in Q_i$ then $Q_{i+1} \subseteq [\![\mathcal{S}_i]\!]$ and we are done.

If $\Phi \notin Q_i$, then $Q_0 \subseteq Q_i$ implies that $\Delta_G(\Phi, \Phi) \neq \emptyset$. Since $\emptyset$ is an active state, we know that $\emptyset \notin Q_i$, and for $Q_i \triangleright Q_{i+1}$ to be a possible derivation of the emptiness test, $\Delta_G(\Phi, \Phi) = \langle\!\langle \Psi_1 \rangle\!\rangle \times \cdots \times \langle\!\langle \Psi_n \rangle\!\rangle \neq \{(\emptyset, \ldots, \emptyset)\}$ must hold, i.e., there must be a $\Psi_i \neq \emptyset$ such that $\langle\!\langle \Psi_i \rangle\!\rangle \subseteq Q_i \subseteq [\![\mathcal{S}_i]\!]$. Hence $\pi_i \in \Phi$ and $\Psi_i = \{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in \Phi \text{ is a } \Box\text{-path}\}$.

Lemma 4 yields the existence of a set of sequents $\mathcal{T}_i$ with $\mathcal{S}_i \vdash^* \mathcal{T}$ containing a sequent Λ with $\Lambda \subseteq \Psi_i$. This sequent is either of the form $\Lambda = \Gamma\Box, \pi_i\Diamond$ or $\Lambda = \Gamma\Box$ for some $\Gamma \subseteq \Phi$. In the former case, IC_G can use a $\Diamond$-inference

$$(\Diamond)\frac{\Gamma\Box, \pi_i\Diamond}{\Gamma, \pi_i}$$

and in the latter case a $\Diamond^+$-inference

$$(\Diamond^+)\frac{\Gamma\Box}{\Gamma, \pi_i}$$

to derive $S_0 \vdash^* \mathcal{S}_i \vdash^* \mathcal{T} \vdash \mathcal{T} \cup \{\Gamma, \pi_i\} = \mathcal{S}$ and $\Phi \subseteq [\![\Gamma, \pi_i]\!]$ holds. ■

4 Optimizations

Since the inverse calculus can be seen as an on-the-fly implementation of the emptiness test, optimizations of the inverse calculus also yield optimizations of the emptiness test. We use the connection between the two approaches to provide an easier proof of the fact that the optimizations of IC_G introduced by Voronkov [19] do not destroy completeness of the calculus.

4.1 Unreachable states / redundant sequents

States that cannot occur on any run starting with an initial state have no effect on the language accepted by the automaton. We call such states *unreachable.* In the following, we will determine certain types of unreachable states.

Definition 6 *Let $\pi, \pi_1, \pi_2 \in \Pi_G$.*

- *The* modal length *of π is the number of occurrences of $\Box$ and $\Diamond$ in π.*
- *$\pi_1, \pi_2 \in \Pi_G$ form a* $\vee$-fork *if $\pi_1 = \pi\vee_l\pi_1'$ and $\pi_2 = \pi\vee_r\pi_2'$ for some π, π_1', π_2'.*
- *π_1, π_2 are* $\Diamond$-separated *if $\pi_1 = \pi_1'\Diamond\pi_1''$ and $\pi_2 = \pi_2'\Diamond\pi_2''$ such that π_1', π_2' have the same modal length and $\pi_1' \neq \pi_2'$.*

Lemma 5 *Let $\mathfrak{A}_G$ be the formula automaton for a K-formula G in NNF and $\Phi \in Q$. If Φ contains a $\vee$-fork, two $\Diamond$-separated paths, or two paths of different modal length, then Φ is unreachable.*

The lemma shows that we can remove such states from $\mathfrak{A}_G$ without changing the accepted language. Sequents containing a $\vee$-fork, two $\Diamond$-separated paths, or two paths of different modal length represent only unreachable states, and are thus redunant, i.e., inferences involving such sequents need not be considered.

Definition 7 (Reduced automaton) *Let $\bar{Q}$ be the set of states of $\mathfrak{A}_G$ that contain a $\vee$-fork, two $\Diamond$-separated paths, or two paths of different modal length. The* reduced *automaton $\mathfrak{A}'_G = (Q'_G, \Sigma_G, \langle\langle\{\epsilon\}\rangle\rangle, \Delta'_G)$ is defined by*

$$Q'_G := Q_G \setminus \bar{Q} \quad \text{and} \quad \Delta'_G := \Delta_G \cap (Q'_G \times \Sigma_G \times Q'_G \times \cdots \times Q'_G).$$

Since the states in $\bar{Q}$ are unreachable, $L(\mathfrak{A}_G) = L(\mathfrak{A}'_G)$. From now on, we consider $\mathfrak{A}'_G$ and define $[\![\cdot]\!]$ relative to the states on $\mathfrak{A}'_G$: $[\![\Gamma]\!] = \{\Phi \in Q'_G \mid \Gamma \subseteq \Phi\}$.

4.2 G-orderings / redundant inferences

In the following, the applicability of the propositional inferences of the inverse calculus will be restricted to those where the affected paths are maximal w.r.t. a total ordering of Π_G. In order to maintain completeness, one cannot consider arbitrary orderings in this context.

Two paths π_1, π_3 are *brothers* iff there exists a $\between$-path π such that $\pi_1 = \pi\between_l$ and $\pi_3 = \pi\between_r$ or $\pi_1 = \pi\between_r$ and $\pi_3 = \pi\between_l$.

Definition 8 (G-ordering) *Let G be a K-formula in NNF. A total ordering $\succ$ of Π_G is called a G-*ordering *iff*

1. *$\pi_1 \succ \pi_2$ whenever*
   - (a) *the modal length of π_1 is strictly greater than the modal length of π_2; or*
   - (b) *π_1, π_2 have the same modal length, the last symbol of π_1 is $\between_*$, and the last symbol of π_2 is $\davidsstar$; or*
   - (c) *π_1, π_2 have the same modal length and π_2 is a prefix of π_1*
2. *There is no path between brothers, i.e., there exist no G-paths π_1, π_2, π_3 such that $\pi_1 \succ \pi_2 \succ \pi_3$ and π_1, π_3 are brothers.*

For the example formula G of Figure 1, a G-ordering $\succ$ can be defined by setting $\nu_9 \succ \nu_8 \succ \cdots \succ \nu_1 \succ \nu_0$. Voronkov [19] shows that G-orderings exist for every K-formula G in NNF. Using an arbitrary, but fixed G-ordering $\succ$, the applicability of the propositional inferences is restricted as follows.

Definition 9 (Optimized Inverse Calculus) *For a sequent Γ and a path π we write $\pi \succ \Gamma$ iff $\pi \succ \pi'$ for every $\pi' \in \Gamma$.*

- *An inference* $(\wedge_*)\dfrac{\Gamma, \pi\wedge_*}{\Gamma, \pi}$ respects $\succ$ *iff* $\pi\wedge_* \succ \Gamma$.

- *An inference* $(\vee)\dfrac{\Gamma_l, \pi\vee_l \qquad \Gamma_r, \pi\vee_r}{\Gamma_l, \Gamma_r, \pi}$ respects $\succ$ *iff* $\pi\vee_l \succ \Gamma_l$ *and* $\pi\vee_r \succ \Gamma_r$.

- *The* $\Diamond$*- and* $\Diamond^+$*-inferences always respect* $\succ$*.*

The optimized inverse calculus $\mathsf{IC}_G^{\succ}$ *works as* IC_G*, but for each derivation* $\mathcal{S}_0 \vdash \cdots \vdash \mathcal{S}_k$ *the following* restrictions *must hold:*

- *For every step* $\mathcal{S}_i \vdash \mathcal{S}_{i+1}$*, the employed inference respects* $\succ$*, and*

- $\mathcal{S}_i$ *must not contain* $\vee$*-forks,* $\Diamond$*-separated paths, or paths of different modal length.*

To distinguish derivations of IC_G and $\mathsf{IC}_G^{\succ}$, we will use the symbol $\vdash_{\succ}$ in derivations of $\mathsf{IC}_G^{\succ}$. In [19], correctness of $\mathsf{IC}_G^{\succ}$ is shown.

Fact 2 ([19]) *Let* G *be a* K*-formula in NNF and* $\succ$ *a* G*-ordering. Then* G *is unsatisfiable iff* $\{\epsilon\} \in \mathcal{S}_0^{\vdash_{\succ}}$*.*

Using the correspondence between the inverse method and the emptiness test of $\mathfrak{A}'_G$, we will now give an alternative, and in our opinion simpler, proof of this fact. Since $\mathsf{IC}_G^{\succ}$ is merely a restriction of IC_G, soundness (i.e., the if-direction of the fact) is immediate.

Completeness requires more work. In particular, the proof of Lemma 4 needs to be reconsidered since the propositional inferences are now restricted: we must show that the $\mathbb{X}$-inferences employed in that proof respect (or can be made to respect) $\succ$. To this purpose, we will follow [19] and introduce the notion of $\succ$-compactness. For $\succ$-compact sets, we can be sure that all applicable $\mathbb{X}$-inferences respect $\succ$. To ensure that all the sets Υ_i constructed in the proof of Lemma 4 are $\succ$-compact, we again follow Voronkov and employ a special selection strategy.

Definition 10 ($\succ$-compact, $\mathsf{select}_{\succ}$) *Let* G *be a* K*-formula in NNF and* $\succ$ *a* G*-ordering. An arbitrary set* $\Phi \subseteq \Pi_G$ *is* $\succ$-compact *iff, for every* $\mathbb{X}$*-path* $\pi \in \Phi$ *that is not p.e. in* Φ*,* $\pi\mathbb{X}_* \succ \Phi$*.*

The selection function $\mathsf{select}_{\succ}$ *is defined as follows: if* Φ *is not p.e., then let* $\{\pi_1, \ldots, \pi_m\}$ *be the set of* $\mathbb{X}$*-paths that are not p.e. in* Φ*. From this* *set,* $\mathsf{select}_\succ$ *selects the path* π_i *such that the paths* $\pi_i\smashtimes_*$ *are the two smallest elements in* $\{\pi_j\smashtimes_* \mid 1 \leq j \leq m\}$.

The function $\mathsf{select}_\succ$ is well-defined because of Condition (2) of G-orderings. The definition of compact ensures that $\smashtimes$-inferences applicable to not propositionally expanded sequents respect $\succ$.

Lemma 6 *Let* G *be a* K*-formula in NNF,* $\succ$ *a* G*-ordering, and* $\mathsf{select}_\succ$ *the selection function as defined above. Let* $\Phi = \{\epsilon\}$ *or* $\Phi = \Gamma\Box, \pi_i\Diamond$ *with* $\Box$*-paths* Γ *and a* $\Diamond$*-path* π*, all of equal modal length. If* T_Φ*, as defined in the proof of Lemma 4, is generated using* $\mathsf{select}_\succ$ *as selection function, then every node* Ψ *of* T_Φ *is* $\succ$*-compact.*

Proof. The proof is similar to the proof of Lemma 5.8.3 in [19]. It is given by induction on the depths of the node Ψ in the tree T_Φ. For the root Φ there are two possibilities. If $\Phi = \{\epsilon\}$ and ϵ is a $\smashtimes$-path, then $\smashtimes_l$ and $\smashtimes_r$ have the same modal length as ϵ and $\smashtimes_* \succ \epsilon$ by Condition (1c) of G-orderings. If $\Phi = \Gamma\Box, \pi_i\Diamond$ and $\pi \in \Phi$ is a $\smashtimes$-path, then $\pi\smashtimes_* \succ \Phi$ holds by Condition (1b) of G-orderings because the last symbol of every path in Φ is $\Box\!\!\!\!\Diamond$.

For the induction step, let Ψ be a node in T_Φ which we have already shown to be $\succ$-compact. We show that then also its successor nodes (if any) are $\succ$-compact.

- If Ψ is an $\wedge$-node with selected $\wedge$-path $\pi \in \Psi$, then the successor node of Ψ is $\Psi' = \Psi, \pi\wedge_l, \pi\wedge_r$. Let $\pi' \in \Phi'$ be a $\smashtimes$-path that is not p.e. in Φ'. There are two possibilities:
  - $\pi' = \pi\wedge_*$. In this case, since $\pi\wedge_*\smashtimes_* \succ \pi\wedge_*$ by Condition (1c) of G-orderings and $\pi\wedge_* \succ \Psi$, $\pi'\smashtimes_* \succ \Psi'$ holds.
  - $\pi' \neq \pi\wedge_*$. Then, $\pi' \in \Psi$ and $\pi' \neq \pi$ holds because π is p.e. in Ψ'. Since Ψ is $\succ$-compact, $\psi'\smashtimes_* \succ \nu$ for every $\nu \in \Psi$. It remains to show that $\pi'\smashtimes_* \succ \pi\smashtimes_*$, which follows from the fact that π was selected by $\mathsf{select}_\succ$.
- If Ψ is an $\vee$-path and the selected $\vee$-path is $\pi \in \Psi$, then, w.o.l.g., $\Phi = \Psi, \pi\vee_l$. The same arguments as before apply. ∎

Given this lemma, it is easy to show that the construction employed in the proof of Lemma 4 also works for $\mathsf{IC}_G^{\succ}$, provided that we restrict the set Φ as in Lemma 6:

Lemma 7 *Let $\Phi = \{\epsilon\}$ or $\Phi = \Gamma\Box, \pi_i\Diamond$ with $\Box$-paths Γ and a $\Diamond$-path π all of equal modal length and $\mathcal{S}$ a set of sequents such that $\langle\!\langle\Phi\rangle\!\rangle \subseteq [\![\mathcal{S}]\!]$. Then there exists a set of sequents $\mathcal{T}$ with $\mathcal{S} \vdash_{\succ}^{*} \mathcal{T}$ such that there exists $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \Phi$.*

Proof. We use the same construction as in the proof of Lemma 4, but the special selection function $\mathsf{select}_{\succ}$ as above. From Lemma 6 we have that all nodes Υ_i in T_Φ are $\succ$-compact. All we have to do is to make sure that the employed inferences respect $\succ$. We refer to the inferences by number assigned to them in the proof of Lemma 4.

(1) Since Υ_{i+1} is compact and $\pi \in \Upsilon_{i+1}$ is not p.e. in Υ_{i+1}, $\pi\wedge_l \succ \Upsilon_{i+1}$ and hence $\pi\wedge_l \succ \Gamma$ because $\Gamma \subseteq \Upsilon_{i+1}$.

(2) W.l.o.g., assume $\pi\wedge_l \succ \pi\wedge_r$. (If this is not the case, then reverse the order of the two inferences.) Since Υ_{i+1} is compact, $\Gamma \subseteq \Upsilon_{i+1}$ and $\pi \in \Upsilon_{i+1}$ is not p.e., $\pi\wedge_l \succ \Gamma$ holds as well as $\pi\wedge_l \succ \pi\wedge_r$. Also $\pi\wedge_r \succ \Gamma$ holds, which means that both inferences respect $\succ$.

(3) Since Υ_{i+1} is compact and $\pi \in \Upsilon_{i+1}$ is not p.e. we have $\pi\vee_* \succ \Upsilon_{i+1}$ and since both Γ_l and Γ_r are subsets of Υ_{i+1}, also $\pi\vee_l \succ \Gamma_l$ and $\pi\vee_r \succ \Gamma_r$ holds. ∎

Alternative Proof of Fact 2. As mentioned before, soundness (the if-direction) is immediate. For the only-if-direction, if G is not satisfiable, then $L(\mathfrak{A}'_G) = \emptyset$ and there is a set of states Q with $Q_0 \rhd^* Q$ and $\langle\!\langle\{\epsilon\}\rangle\!\rangle \subseteq Q$. Using Lemma 7 we show that there is a derivation of $\mathsf{IC}_G^{\succ}$ that simulates this derivation, i.e., there is a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash_{\succ}^{*} \mathcal{S}$ and $Q \subseteq [\![\mathcal{S}]\!]$.

The proof is by induction on the length m of the derivation $Q_0 \rhd \ldots \rhd Q_m = Q$ and is totally analogous to the proof of Theorem 2. The base case is Lemma 1, which also holds for $\mathsf{IC}_G^{\succ}$ and the reduced automaton. The induction step uses Lemma 7 instead of Lemma 4, but this is the only difference.

Hence, $Q_0 \rhd^* Q$ and $\langle\!\langle\{\epsilon\}\rangle\!\rangle \subseteq Q$ implies that there exist a derivation $\mathcal{S}_0 \vdash_{\succ}^{*} \mathcal{S}$ such that $\langle\!\langle\{\epsilon\}\rangle\!\rangle \subseteq [\![\mathcal{S}]\!]$. Lemma 7 yields a derivation $\mathcal{S} \vdash_{\succ}^{*} \mathcal{T}$ with $\{\epsilon\} \in \mathcal{T} \subseteq \mathcal{S}_0^{\vdash_{\succ}}$. ∎

5 Global axioms

When considering satisfiability of G w.r.t. the global axiom H, we must take subformulae of G and H into account. We address subformulae using paths in G and H.

Definition 11 ((*G, H*)-Paths) *For* K*-formulae* G, H *in NNF, the set of* (G, H)*-paths* $\Pi_{G,H}$ *is a subset of* $\{\epsilon_G, \epsilon_H\} \cdot \{\vee_l, \vee_r, \wedge_l, \wedge_r, \Box, \Diamond\}^*$*. The set* $\Pi_{G,H}$ *and the subformula* $(G, H)|_\pi$ *of* G, H *addressed by a path* $\pi \in \Pi_{G,H}$ *are defined inductively as follows:*

- $\epsilon_G \in \Pi_{G,H}$ *and* $(G, H)|_{\epsilon_G} = G$*, and* $\epsilon_H \in \Pi_{G,H}$ *and* $(G, H)|_{\epsilon_H} = H$
- *if* $\pi \in \Pi_{G,H}$ *and* $(G, H)|_\pi = F_1 \wedge F_2$ *then* $\pi\wedge_l, \pi\wedge_r \in \Pi_{G,H}$, $(G, H)|_{\pi\wedge_l} = F_1$, $(G, H)|_{\pi\wedge_r} = F_2$*, and* π *is called* $\wedge$-path.
- *The other cases are defined analogously (see also Definition 1).*
- $\Pi_{G,H}$ *is the smallest set that satisfies the previous conditions.*

The definitions of *p.e.* and *clash* are extended to subsets of $\Pi_{G,H}$ in the obvious way, with the *additional requirement* that, for $\Phi \neq \emptyset$ to be p.e., $\epsilon_H \in \Phi$ must hold. This additional requirement enforces the global axiom.

Definition 12 (Formula Automaton with Global Axioms) *For* K*-formulae* G, H *in NNF, let* $\{\pi_1, \ldots, \pi_n\}$ *be an enumeration of the* $\Diamond$*-paths in* $\Pi_{G,H}$*. The* n*-ary looping automaton* $\mathfrak{A}_{G,H}$ *is defined by*

$$\mathfrak{A}_G := (Q_{G,H}, \Sigma_{G,H}, \langle\langle\{\epsilon_G\}\rangle\rangle, \Delta_{G,H}),$$

where $Q_{G,H} := \Sigma_{G,H} := \{\Phi \in \Pi_{G,H} \mid \Phi$ *is p.e.*$\}$ *and the transition relation* $\Delta_{G,H}$ *is defined as for the automaton* $\mathfrak{A}_G$ *in Definition 3.*

Theorem 3 *G is satisfiable w.r.t. the global axiom H iff* $L(\mathfrak{A}_{G,H}) \neq \emptyset$*.*

Proof. The proof is totally analogous to the proof of Theorem 1. We use the same constructions for both directions.

Let $\{\pi_1, \ldots, \pi_n\}$ be an enumeration of the $\Diamond$-paths in $\Pi_{G,H}$.

For the *if*-direction let $L(\mathfrak{A}_{G,H}) \neq \emptyset$, $t, r : [n]^* \to \{\Phi \subseteq \Pi_{G,H} \mid \Phi \text{ is p.e.}\}$ a tree that is accepted by $\mathfrak{A}_{G,H}$ and a corresponding run of $\mathfrak{A}_{G,H}$. By construction of $\mathfrak{A}_{G,H}$, $t(w) = r(w)$ for every $w \in [n]^*$. We construct a Kripke model $\mathcal{M} = (W, R, V)$ from t by setting

$$
\begin{aligned}
W &= \{w \in [n]^* \mid t(w) \neq \emptyset\} \\
R &= \{(w, wi) \in W \times W \mid i \in [n]\} \\
V &= \lambda P.\{p \in W \mid \exists \pi \in t(w).(G,H)|_\pi = P\} \quad \text{for all propositional atoms } P
\end{aligned}
$$

Claim. For all $w \in W$, if $\pi \in t(w)$ then $\mathcal{M}, w \models (G,H)|_\pi$.

Proof of the claim. The claim is proved by induction on the structure of K-formulae. Let $w \in W$ be a world and $\pi \in \Pi_G$ be a path such that $\pi \in t(w)$.

- if $(G,H)|_\pi = P$ is a propositional atom and $w \in W$, then $w \in V(P)$ and hence $\mathcal{M}, w \models (G,H)|_\pi$.
- if $(G,H)|_\pi = \neg P$ is a negated propositional atom, then, since $t(w)$ is clash free, there is no $\pi' \in \Pi_{G,H}$ such that $(G,H)|_{\pi'} = P$. Thus, $w \notin V(P)$ and hence $\mathcal{M}, w \models \neg P$.
- if $(G,H)|_\pi = F_1 \wedge F_2$ then π is an $\wedge$-paths, and since $t(w)$ is p.e., $\{\pi\wedge_l, \pi\wedge_r\} \subseteq t(w)$. By induction, $\mathcal{M}, w \models (G,H)|_{\pi\wedge_*}$ and hence $\mathcal{M}, w \models (G,H)|_\pi$.
- if $(G,H)|_\pi = F_1 \vee F_2$ then π is an $\vee$-paths, and since $t(w)$ is p.e., $\{\pi\vee_l, \pi\vee_r\} \cap t(w) \neq \emptyset$. By induction, $\mathcal{M}, w \models (G,H)|_{\pi\vee_l}$ or $\mathcal{M}, w \models (G,H)|_{\pi\vee_r}$ and hence $\mathcal{M}, w \models (G,H)|_\pi$.
- if $(G,H)|_\pi = \Diamond F$ then π is a $\Diamond$-path and, w.o.l.g., assume $\pi = \pi_i$. Since $\pi_i \in r(w)$, $\pi_i\Diamond \in r(wi) = t(wi)$ holds and hence $wi \in W$ and $(w, wi) \in R$. By induction, we have that $\mathcal{M}, wi \models (G,H)|_{\pi_i\Diamond}$ and hence $\mathcal{M}, w \models (G,H)|_{\pi_i}$.
- if $(G,H)|_\pi = \Box F$ and $(w, w') \in R$ then $w' = wi$ for some $i \in [n]$ and $t(wi) \neq \emptyset$ holds and by construction of $\mathfrak{A}_{G,H}$, this implies $\pi\Box \in r(wi) = t(wi)$. By induction, this implies $\mathcal{M}, wi \models (G,H)|_{\pi\Box}$ and since $wi = w'$ and w' has been chosen arbitrarily, $\mathcal{M}, w \models (G,H)|_\pi$.

This finishes the proof of the claim. Since $t(\epsilon) = r(\epsilon) \in \langle\!\langle \{\epsilon_G\} \rangle\!\rangle$ and hence $\epsilon_G \in t(\epsilon)$, $\mathcal{M}, \epsilon \models (G, H)|_{\epsilon_G}$ and $G = (G, H)|_{\epsilon_G}$ is satisfiable.

Also, since $t(w)$ is p.e., $\epsilon_H \in t(w)$ for every $w \in W$ and, by the claim, $\mathcal{M}, w \models H = (G, H)|_{\epsilon_H}$ holds for every $w \in W$. Hence G is satisfiable w.r.t. the global axiom H.

For the *only if*-direction, we first show an auxiliary claim: for a set $\Psi \subseteq \Pi_{G,H}$ we define $\mathcal{M}, w \models \Psi$ iff $\mathcal{M}, w \models (G, H)|_\pi$ for every $\pi \in \Psi$.

Claim. If $\Psi \subseteq \Pi_{G,H}$ and $w \in W$ such that $\mathcal{M}, w \models \Psi$, then there is a $\Phi \in \langle\!\langle \Psi \rangle\!\rangle$ such that $\mathcal{M}, w \models \Phi$.

Proof of the claim. Let $\Psi \subseteq \Pi_{G,H}$ and $w \in W$ such that $\mathcal{M}, w \models \Psi$. We will show how to construct an expansion of Ψ with the desired property. If Ψ is already p.e., then $\Psi \in \langle\!\langle \Psi \rangle\!\rangle$ and we are done.

- If Ψ is not p.e. because $\epsilon_H \notin \Psi$ then, because $\mathcal{M}, w \models H$, $\Psi' = \Psi \cup \{\epsilon_H\}$ is a set with $\mathcal{M}, w \models \Psi$ that is "one step closer" to being p.e. than Ψ.
- If Ψ is not p.e. and $\epsilon_H \in \Psi$ then let $\pi \in \Psi$ be a $\between$-path that is not p.e. in Ψ.
  - If π is a $\wedge$-path then $(G, H)|_\pi = F_1 \wedge F_2$ and since $\mathcal{M}, w \models (G, H)|_\pi$, also $\mathcal{M}, w \models F_1 = (G, H)|_{\pi \wedge_l}$ and $\mathcal{M}, w \models F_2 = (G, H)|_{\pi \wedge_r}$. Hence $\mathcal{M}, w \models \Psi \cup \{\pi \wedge_l, \pi \wedge_r\}$ and $\Psi' = \Psi \cup \{\pi \wedge_l, \pi \wedge_r\}$ is a set with $\mathcal{M}, w \models \Psi'$ that is "one step closer" to being p.e. than Ψ.
  - If π is a $\vee$-path then $(G, H)|_\pi = F_1 \vee F_2$ and since $\mathcal{M}, w \models (G, H)|_\pi$, also $\mathcal{M}, w \models F_1 = (G, H)|_{\pi \vee_l}$ or $\mathcal{M}, w \models F_2 = (G, H)|_{\pi \vee_r}$. Hence $\mathcal{M}, w \models \Psi \cup \{\pi \vee_l\}$ or $\mathcal{M}, w \models \Psi \cup \{\pi \vee_r\}$ and hence can obtain a set Ψ' with $\mathcal{M}, w \models \Psi'$ that is again "one step close" to being p.e. than Ψ.

Restarting this process with $\Psi = \Psi'$ eventually yields an expansion Φ of the initial set Ψ with $\mathcal{M}, w \models \Phi$, which proves the claim.

Let $\mathcal{M} = (W, R, V)$ be a model for G with $w \in W$ such that $\mathcal{M}, w \models G$. From $\mathcal{M}$ we construct a tree that is accepted by $\mathfrak{A}_{G,H}$. Using this claim, we inductively define a tree t accepted by $\mathfrak{A}_{G,H}$. To this purpose, we also inductively define a function $f : [n]^* \to W$ such that, if $\mathcal{M}, f(p) \models t(p)$ for all p.

We start by setting $f(\epsilon) = w$ for a $w \in W$ with $\mathcal{M}, w \models G$. and $t(\epsilon) = \Phi$ for a $\Phi \in \langle\!\langle \{\epsilon\} \rangle\!\rangle$ such that $\mathcal{M}, w \models \Phi$. From the claim we have that such a set Φ exists because $\mathcal{M}, w \models G = (G, H)|_\epsilon$.

If $f(p)$ and $t(p)$ are already defined, then, for $i \in [n]$, we define $f(pi)$ and $t(pi)$ as follows:

- if $\pi_i \in t(p)$ then $\mathcal{M}, f(p) \models (G, H)|_{\pi_i}$ and hence there is a $w' \in W$ such that $(f(p), w') \in R$ and $\mathcal{M}, w' \models (G, H)|_{\pi_i \Diamond}$. If $\pi \in t(p)$ is a $\Box$-path, then also $\mathcal{M}, w' \models (G, H)|_{\pi\Box}$ holds. Hence $\mathcal{M}, w' \models \{\pi_i \Diamond\} \cup \{\pi\Box \mid \pi \in t(p) \text{ is a } \Box\text{-path }\}$. We set $f(pi) = w'$ and $t(pi) = \Phi$ for a $\Phi \in \langle\!\langle \{\pi_i\Diamond\} \cup \{\pi\Box \mid \pi \in t(p) \text{ is a } \Box\text{-path }\} \rangle\!\rangle$ with $\mathcal{M}, w' \models \Phi$, which exist by the claim.

- if $\pi_i \notin t(p)$, then we set $f(pi) = w$ for an arbitrary $w \in W$ and $t(pi) = \emptyset$

In both cases, we have define $f(pi)$ and $t(pi)$ such that $\mathcal{M}, f(pi) \models t(pi)$. It is easy to see that t is accepted by $\mathfrak{A}_{G,H}$ with the run $r = t$. Hence $L(\mathfrak{A}_{G,H}) \neq \emptyset$ which is what we needed to show. ■

Definition 13 (The Inverse Calculus w. Global Axiom) *Let G, H be* K- *formula in NNF and $\Pi_{G,H}$ the set of paths of G, H. Sequents are subsets of $\Pi_{G,H}$, and operations on sequents are defined as before.*

In addition to the inferences from Figure 2, the inverse calculus for G w.r.t. the global axiom H, $\mathsf{IC}^{ax}_{G,H}$, employs the inference

$$(ax)\frac{\Gamma, \epsilon_H}{\Gamma}.$$

From now on, $[\![\cdot]\!]$ is defined w.r.t. the states of $\mathfrak{A}_{G,H}$, i.e., $[\![\Gamma]\!] := \{\Phi \in Q_{G,H} \mid \Gamma \subseteq \Phi\}$.

Theorem 4 ($\mathsf{IC}^{ax}_{G,H}$ and the emptiness test for $\mathfrak{A}_{G,H}$ simulate each other) *Let $\vdash_{ax}$ denote derivation steps of $\mathsf{IC}^{ax}_{G,H}$, and $\rhd$ derivation steps of the emptiness test for $\mathfrak{A}_{G,H}$.*

1. *Let $Q \subseteq Q_{G,H}$ be a set of states such that $Q_0 \rhd^* Q$. Then there exists a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash^*_{ax} \mathcal{S}$ and $Q \subseteq [\![\mathcal{S}]\!]$.*

2. *Let $\mathcal{S}$ be a set of sequents such that $\mathcal{S}_0 \vdash^*_{ax} \mathcal{S}$. Then there exists a set of states $Q \subseteq Q_G$ with $Q_0 \rhd^* Q$ and $[\![\mathcal{S}]\!] \subseteq Q$.*

Lemma 1, 2, and 3, restated for $\mathfrak{A}_{G,H}$ and $\mathsf{IC}^{ax}_{G,H}$, can be shown as before. The following lemma deals with the ax-inference of $\mathsf{IC}^{ax}_{G,H}$.

Lemma 8 *Let $\mathcal{S} \rhd \mathcal{T}$ be a derivation of $\mathsf{IC}^{ax}_{G,H}$ that employs an ax-inference. Then $[\![\mathcal{S}]\!] = [\![\mathcal{T}]\!]$.*

Proof. Let $\mathcal{T} = \mathcal{S} \cup \{\Gamma\}$ with $\{\Gamma, \epsilon_H\} \in \mathcal{S}$. Then we know that $(ax)\dfrac{\Gamma, \epsilon_H}{\Gamma}$. $[\![\mathcal{T}]\!] = [\![\mathcal{S}]\!] \cup [\![\Gamma]\!]$. Since $\mathcal{S} \subseteq \mathcal{T}$, $[\![\mathcal{S}]\!] \subseteq [\![\mathcal{T}]\!]$ holds immediately. If $\Phi \in [\![\Gamma]\!]$, then, since Φ is p.e., $\epsilon_H \in \Phi$ and $\Phi \in [\![\Gamma, \epsilon_H]\!] \subseteq [\![\mathcal{S}]\!]$. ∎

The proof of Theorem 4.2 is now analogous to the proof of Theorem 2.2. For the proof of Theorem 4.1, Lemma 4 needs to be re-proved because the change in the definition of p.e. now also implies that $\epsilon_H \in \Phi$ holds for every set $\Phi \in \langle\!\langle \Psi \rangle\!\rangle$ for any $\Psi \neq \emptyset$ (see Lemma 9). This is where the new inference ax comes into play. In all other respects, the proof of Theorem 4.1 is analogous to the proof of Theorem 2.1.

Lemma 9 *Let $\Phi \subseteq \Pi_G$ a set of paths and $\mathcal{S}$ a set of sequents such that $\langle\!\langle \Phi \rangle\!\rangle \subseteq [\![\mathcal{S}]\!]$. Then there exists a set of sequents $\mathcal{T}$ with $\mathcal{S} \vdash^*_{ax} \mathcal{T}$ such that there exists $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \Phi$.*

Proof. If $\epsilon_H \in \Phi$ than we can use the same construction used in the proof of Lemma 4 to construct the set $\mathcal{T}$ such that $\mathcal{S} \vdash^*_{ax} \mathcal{T}$ and there is a $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \Phi$.

If $\epsilon_H \notin \Phi$, then set $\Psi = \Phi, \epsilon_H$ and again use the construction from the proof of Lemma 4 to construct a set $\mathcal{T}$ such that $\mathcal{S} \vdash^*_{ax} \mathcal{T}$ and there is a $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \Psi$. If $\epsilon_H \notin \Lambda$ then we are done since then also $\Lambda \subseteq \Phi$. If $\Lambda = \Gamma, \epsilon_H$ for some Γ with $\epsilon_H \notin \Gamma$, then $\Gamma \subseteq \Phi$ and $\mathcal{T} \vdash_{ax} \mathcal{T} \cup \{\Gamma\}$ can be derived by $\mathsf{IC}^{ax}_{G,H}$ using the inference $(ax)\dfrac{\Gamma, \epsilon_H}{\Gamma}$. ∎

Corollary 2 *$\mathsf{IC}^{ax}_{G,H}$ yields an* ExpTime *decision procedure for satisfiability w.r.t. global axioms in K.*

The following algorithm yields the desired procedure:

Algorithm 1 *Let G, H be* K*-formulae in NNF. To test satisfiability of G w.r.t. H, calculate $\mathcal{S}_0^{\models_{ax}}$. If $\{\emptyset, \{\epsilon_G\}\} \cap \mathcal{S}_0^{\models_{ax}} \neq \emptyset$, then answer "not satisfiable," and "satisfiable" otherwise.*

Correctness of this algorithm follows from Theorem 3 and 4. If G is not satisfiable w.r.t. H, then $L(\mathfrak{A}_{G,H}) = \emptyset$, and there exists a set of states Q with $Q_0 \rhd^* Q$ and $\langle\!\langle \{\epsilon_G\} \rangle\!\rangle \subseteq Q$. Thus, there exists a set of sequents $\mathcal{S}$ with $\mathcal{S}_0 \vdash_{ax}^* \mathcal{S}$ such that $Q \subseteq [\![\mathcal{S}]\!]$. With (the appropriately reformulated) Lemma 4 there exists a set of sequents $\mathcal{T}$ with $\mathcal{S} \vdash_{ax}^* \mathcal{T}$ such that there is a sequent $\Lambda \in \mathcal{T}$ with $\Lambda \subseteq \{\epsilon_G\}$. Consequently, $\Lambda = \emptyset$ or $\Lambda = \{\epsilon_G\}$.

Since $\mathcal{S}_0 \vdash_{ax}^* \mathcal{S}_0^{\models_{ax}}$, there exists a set of (inactive) states Q such that $Q_0 \rhd^* Q$ and $[\![\mathcal{S}_0^{\models_{ax}}]\!] \subseteq Q$. Since $\langle\!\langle \{\epsilon_G\} \rangle\!\rangle \subseteq [\![\{\epsilon_G\}]\!] \subseteq [\![\emptyset]\!]$, we know that $\{\emptyset, \{\epsilon_G\}\} \cap \mathcal{S}_0^{\models_{ax}} \neq \emptyset$ implies $\langle\!\langle \{\epsilon_G\} \rangle\!\rangle \subseteq Q$. Consequently, $L(\mathfrak{A}_{G,H}) = \emptyset$ and thus G is not satisfiable w.r.t. H.

For the complexity, note that there are only exponentially many sequents. Consequently, it is easy to see that the saturation process that leads to $\mathcal{S}_0^{\models_{ax}}$ can be realized in time exponential in the size of the input formulae.

6 Future Work

There are several interesting directions in which to continue this work. First, satisfiability in K (without global axioms) is PSPACE-complete whereas the inverse method yields only an EXPTIME-algorithm. Can suitable optimizations turn this into a PSPACE-procedure? Second, can the optimizations considered in Section 4 be extended to the inverse calculus with global axioms? Third, Voronkov considers additional optimizations. Can they also be handled within our framework? Finally, can the correspondence between the automata approach and the inverse method be used to obtain inverse calculi and correctness proofs for other modal or description logics?

References

[1] C. Areces, R. Gennari, J. Heguiabehere, and M. de Rijke. Tree-based heuristics in modal theorem proving. In W. Horn, editor, *Proc. of ECAI2000*, Berlin, Germany, 2000. IOS Press Amsterdam.

[2] F. Baader and U. Sattler. An overview of tableau algorithms for description logics. *Studia Logica*, 2000. To appear.

[3] F. Baader and S. Tobies. The inverse method implements the automata approach for modal satisfiability. LTCS-Report 01-03, LuFG Theoretical Computer Science, RWTH Aachen, Germany, 2001. See http://www-lti.informatik.rwth-aachen.de/Forschung/Reports.html.

[4] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001. Publishing date May 2001, preliminary version available online from `http://www.mlbook.org/`.

[5] C. Courcoubetis, M. Vardi, P. Wolper, and M. Yannakakis. Memory e cient algorithms for the verification of temporal properties. In E. M. Clarke and R. P. Kurshan, editors, *Proc. of CAV '90*, volume 531 of LNCS, pages 233–242. Springer Verlag, 1991.

[6] F. M. Donini and F. Massacci. EXPTIME tableaux for ALC. *Artificial Intelligence*, 124(1):87–138, 2000.

[7] W. F. Dowling and J. H. Gallier. Linear-time algorithms for testing the satisfiability of propositional horn formulae. *Journal of Logic Programming*, 1(3):267–28, 1984.

[8] R. Dyckhoff, editor. *Proc. of TABLEAUX 2000*, number 1847 in LNAI, St Andrews, Scotland, UK, 2000. Springer Verlag.

[9] R. Gerth, D. Peled, M. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *Proc. of the 15th International Symposium on Protocol Specification, Testing, and Verification*, pages 3–18, Warsaw, Poland, 1995. Chapman & Hall.

[10] R. Goré. Tableau methods for modal and temporal logics. In M. D'Agostino, D. M. Gabbay, R. Hähnle, and J. Posegga, editors, *Handbook of Tableau Methods*. Kluwer, Dordrecht, 1998.

[11] I. Horrocks. Benchmark analysis with FaCT. In Dyckhoff [8], pages 62–66.

[12] U. Hustadt and R. A. Schmidt. MSPASS: Modal reasoning by translation and first-order resolution. In Dyckhoff [8], pages 67–71.

[13] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

[14] C. Lutz and U. Sattler. The complexity of reasoning with boolean modal logic. In Wolter F., H. Wansing, M. de Rijke, and M. Zakharyaschev, editors, *Preliminary Proc. of AiML2000*, Leipzig, Germany, 2000.

[15] R. A. Schmidt. Resolution is a decision procedure for many propositional modal logics. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyaschev, editors, *Advances in Modal Logic, Volume 1*, volume 87 of *Lecture Notes*, pages 189–208. CSLI Publications, Stanford, 1998.

[16] E. Spaan. *Complexity of Modal Logics*. PhD thesis, Univ. van Amsterdam, 1993.

[17] Moshe Y. Vardi and Pierre Wolper. Automata-theoretic techniques for modal logics of programs. *Journal of Computer and System Sciences*, 32:183–221, 1986.

[18] M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.

[19] A. Voronkov. How to optimize proof-search in modal logics: new methods of proving redundancy criteria for sequent calculi. *ACM Transactions on Computational Logic*, 1(4), 2001. Available online from `http://www.cs.man.ac.uk/~voronkov/recent.html` .