Modal Fixed Point Logics

Gerhard JÄGER [1]
IAM, University of Bern

Abstract. The following notes are centered around multi-modal logics extended by the possibility to introduce least and greatest fixed points. We begin with discussing a range of traditional results and turn to more recent approaches dealing with finite and infinite derivations and explicit representations of proofs afterwards. Our focus is on foundational questions and a proof-theoretic perspective rather than practical applications.

Keywords. Proof theory, inductive definitions, modal mu-calculus, multi-modal logics for belief, knowledge and common knowledge, logic of proofs

Introduction

These notes are centered around multi-modal logics extended by the possibility to introduce least and greatest fixed points of suitable formulas. We begin with discussing the general framework which is provided by transition systems, monotone inductive definitions over those and the propositional modal μ-calculus. After dealing with some basic results such as the fundamental semantic theorem of the modal μ-calculus, we take a proof-theoretic perspective and look at infinitary and finite deductive systems.

Special emphasis is put on the multi-modal approach to knowledge and common knowledge. We follow the traditional path in designing Hilbert-style system for common knowledge and its appendant semantics before looking at common knowledge from a proof-theoretic perspective. The final section is about the recently established relationship between evidence and knowledge, starting off from the so-called logic of proofs.

The focus of these notes is on foundational questions rather than practical applications.

1. The general framework

Inductive definitions play an important rôle in many parts of mathematics and computer science; of particular interest in this context are the least and greatest fixed points generated by monotone operators. In mathematical logic there are prominent formalizations of inductive definitions, for example the celebrated theories $\mathbf{ID}_n$ (see, e.g., Buchholz, Feferman, Pohlers, and Sieg [9] for an overview). And in the context of modal logics we have the modal μ-calculus which turned out to be of seminal importance in at least two

[1] Institut für Informatik und angewandte Mathematik, Universität Bern, Neubrückstrasse 10, 3012 Bern, Switzerland; E-mail: jaeger@iam.unibe.ch.

ways: for theoretical foundational work and as a tool for setting up environments for the practical analysis of properties of systems and programs.

We begin this section with recalling some basic facts about monotone inductive definitions and introducing transition systems as the basic structures to be considered in the following. Afterwards we turn to the syntax and semantics of the propositional modal μ-calculus.

SOME PRELIMINARY REMARKS ABOUT FIXED POINTS

Suppose that we are given a set S and a mapping Φ from the power set of S to the power set of S, i.e.,

$$\Phi : \wp(S) \to \wp(S).$$

If $\Phi(M) \subseteq \Phi(N)$ whenever $M \subseteq N \subseteq S$, then Φ is called a *monotone operator on* S. If $\Phi(M) = M$ for $M \subseteq S$, then M is a *fixed point* of Φ. By a well-known theorem of Knaster and Tarski we know that any monotone operator Φ on a set S has a least fixed point $\mathit{lfp}(\Phi)$ and a greatest fixed point $\mathit{gfp}(\Phi)$ within the ordering $(\wp(S), \subseteq)$. Moreover, these two fixed points can be characterized as follows:

$$\begin{aligned}\mathit{lfp}(\Phi) &= \bigcap\{M \subseteq S : \Phi(M) = M\} = \bigcap\{M \subseteq S : \Phi(M) \subseteq M\},\\ \mathit{gfp}(\Phi) &= \bigcup\{M \subseteq S : \Phi(M) = M\} = \bigcup\{M \subseteq S : M \subseteq \Phi(M)\}.\end{aligned}$$

The fixed points $\mathit{lfp}(\Phi)$ and $\mathit{gfp}(\Phi)$ can also be generated by applying Φ repeatedly on the empty set and the set S, respectively. To make this precise, we define by recursion on the ordinals

$$\begin{aligned} I^0_\Phi &:= \emptyset, \quad I^{\sigma+1}_\Phi := \Phi(I^\sigma_\Phi), \quad I^\lambda_\Phi := \bigcup_{\xi<\lambda} I^\xi_\Phi \quad (\lambda \text{ limit}),\\ J^0_\Phi &:= S, \quad J^{\sigma+1}_\Phi := \Phi(J^\sigma_\Phi), \quad J^\lambda_\Phi := \bigcap_{\xi<\lambda} J^\xi_\Phi \quad (\lambda \text{ limit}).\end{aligned}$$

The monotonicity of Φ then yields $I^\sigma_\Phi \subseteq I^\tau_\Phi$ and $J^\tau_\Phi \subseteq J^\sigma_\Phi$ for any ordinals σ and τ such that $\sigma \leq \tau$. The least and greatest fixed points of Φ are approached by these stages I^σ_Φ and J^σ_Φ as

$$\mathit{lfp}(\Phi) = \bigcup_{\sigma \in On} I^\sigma_\Phi \quad \text{and} \quad \mathit{gfp}(\Phi) = \bigcap_{\sigma \in On} J^\sigma_\Phi.$$

In this sense, $\mathit{lfp}(\Phi)$ and $\mathit{gfp}(\Phi)$ are inductively defined sets. A simple cardinality argument even shows that there exist ordinals σ and τ of cardinalities less than or equal to the cardinality of S for which

$$\mathit{lfp}(\Phi) = I^\sigma_\Phi = I^{\sigma+1}_\Phi \quad \text{and} \quad \mathit{gfp}(\Phi) = J^\tau_\Phi = J^{\tau+1}_\Phi.$$

Proofs of all these observations can be found, for example, in the textbooks Barwise [5] and Moschovakis [16].

There exists an interesting duality between least and greatest fixed points. For a monotone operator Φ on a set S we introduce its dual operator Φ^d on S by defining, for any $M \subseteq S$,

$$\Phi^d(M) \;:=\; S \setminus \Phi(S \setminus M).$$

Obviously, Φ^d is monotone as well, and easy computations show that the least and greatest fixed points of one operator are the complements of the greatest and least fixed points of the dual operator, respectively:

$$\begin{aligned} \mathit{lfp}(\Phi^d) &= S \setminus \mathit{gfp}(\Phi), \qquad \mathit{gfp}(\Phi^d) = S \setminus \mathit{lfp}(\Phi), \\ \mathit{lfp}(\Phi) &= S \setminus \mathit{gfp}(\Phi^d), \qquad \mathit{gfp}(\Phi) = S \setminus \mathit{lfp}(\Phi^d). \end{aligned}$$

TRANSITION SYSTEMS

Transition systems provide a very abstract model of distributed systems and concurrent behavior. There are a set St of states and a set Ac of actions; actions act on states in the sense that for any action α there is a binary transition relation $Tr(\alpha)$ on St, and $(s,t) \in Tr(\alpha)$ is interpreted as "action α leads from state s to state t" or "state t can be reached from state s via action α".

Definition 1 *A* transition system *is a triple* $\mathfrak{T} = (St, Ac, Tr)$ *satisfying the following conditions:*

(TS.1) St *is a non-empty set of* states*,* Ac *a non-empty set of* actions*.*

(TS.2) Tr *is a mapping which assigns to any action* $\alpha \in Ac$ *a binary transition relation* $Tr(\alpha)$ *on* St*, i.e.,* $Tr(\alpha) \subseteq St \times St$*.*

If $\mathfrak{T}$ is the transition system (St, Ac, Tr), then we usually write $|\mathfrak{T}|$ for the set St of states of $\mathfrak{T}$. In addition, for $s, t \in St$ and $\alpha \in Ac$, the notations $s \overset{\mathfrak{T}(\alpha)}{\longmapsto} t$ or simply $s \overset{\alpha}{\longmapsto} t$, if $\mathfrak{T}$ is clear from the context, stand for $(s,t) \in Tr(\alpha)$.

Given a transition system $\mathfrak{T}$, a *run* from a state $s \in |\mathfrak{T}|$ is a finite or infinite sequence of the form

$$s \overset{\alpha_1}{\longmapsto} t_1 \overset{\alpha_2}{\longmapsto} t_2 \overset{\alpha_3}{\longmapsto} t_3 \overset{\alpha_4}{\longmapsto} \;\ldots\; .$$

Such a run from s is called an α*-run* from s if only the action α is involved, i.e., if it is of the form

$$s \overset{\alpha}{\longmapsto} t_1 \overset{\alpha}{\longmapsto} t_2 \overset{\alpha}{\longmapsto} t_3 \overset{\alpha}{\longmapsto} \;\ldots\; .$$

Clearly, each initial segment of a run from state s is a run from s as well. Also, given a run from state s, its end sequence that begins with state t is a run from t.

In building up our abstract framework we further assume that we are given a set BP of countably many *basic properties* $p_0, p_1, p_2, \ldots$ which may or may not hold at the individual states of a transition system $\mathfrak{T}$. Those states of $\mathfrak{T}$ which satisfy a basic property p are collected in the set $\mathfrak{T}(p)$.

Without going into details we just state that many elementary features of systems and runs can be expressed in this simple framework. For instance, the infinite run

$$t_0 \overset{\alpha_1}{\longmapsto} t_1 \overset{\alpha_2}{\longmapsto} t_2 \overset{\alpha_3}{\longmapsto} t_3 \overset{\alpha_4}{\longmapsto} \ \ldots$$

from t_0 has the property "eventually always p" if there exists a natural number i such that, for all $j \geq i$, $t_j \in \mathfrak{T}(p)$. The following examples, taken from Bradfield and Stirling [6], illustrate that more sophisticated properties can be expressed by making use of least and greatest fixed points of suitable operators.

LEAST FIXED POINTS AND LIVENESS

As above, $\mathfrak{T}$ is supposed to be a transition system. For an action α from $\mathfrak{T}$ and a basic property p we consider the operator Φ_0 on $|\mathfrak{T}|$ defined by, for any $M \subseteq |\mathfrak{T}|$,

$$\Phi_0(M) \ := \ \mathfrak{T}(p) \cup \{s \in |\mathfrak{T}| : \forall t(\text{if } s \overset{\alpha}{\longmapsto} t \text{ then } t \in M)\}.$$

Φ_0 is monotone, and for its least fixed point $\mathit{lfp}(\Phi)$ we have:

$$s \in \mathit{lfp}(\Phi_0) \quad \Longleftrightarrow \quad \text{Every infinite } \alpha\text{-run from } s \text{ contains a state } t \in \mathfrak{T}(p).$$

According to Bradfield and Stirling [6,7], this may be considered as a liveness property ("Something good eventually happens") since to contain a state in $\mathfrak{T}(p)$ is required to happen. To formulate that something has to happen on some path, pick a further basic property q and define an operator Φ_1 on $|\mathfrak{T}|$ by setting, for any $M \subseteq |\mathfrak{T}|$,

$$\Phi_1(M) \ := \ \mathfrak{T}(q) \cup (\mathfrak{T}(p) \cap \{s \in |\mathfrak{T}| : \exists t(s \overset{\alpha}{\longmapsto} t \text{ and } t \in M)\}).$$

This operator is also monotone; it satisfies the following property:

$$s \in \mathit{lfp}(\Phi_1) \quad \Longleftrightarrow \quad \begin{cases} \text{There exists an } \alpha\text{-run from } s \text{ containing a state} \\ t \in \mathfrak{T}(q) \text{ such that } r \in \mathfrak{T}(p) \text{ everywhere before } t. \end{cases}$$

GREATEST FIXED POINTS AND SAFETY

As least fixed points can be used to describe liveness, so some greatest fixed points reflect safety. For a transition system $\mathfrak{T}$, an action α and a basic property p we now introduce the operator Φ_2 on $|\mathfrak{T}|$ which s given by, for any $M \subseteq |\mathfrak{T}|$,

$$\Phi_2(M) \ := \ \mathfrak{T}(p) \cap \{s \in |\mathfrak{T}| : \forall t(\text{if } s \overset{\alpha}{\longmapsto} t \text{ then } t \in M)\}.$$

Clearly, Φ_2 is a monotone operator, but now we are interested in its greatest fixed point, for which we have:

$$s \in \mathit{gfp}(\Phi_2) \quad \Longleftrightarrow \quad \begin{cases} \text{For every } \alpha\text{-run from } s \text{ and every state} \\ t \text{ on this run we have } t \in \mathfrak{T}(p). \end{cases}$$

This is a safety property in the sense that it states for an s from $gfp(\Phi_2)$ that it can never happen that there is an α-run from s which contains a state t violating p. We conclude these examples by making the operator on $|\mathfrak{T}|$ a bit more complex; for any $M \subseteq |\mathfrak{T}|$,

$$\Phi_3(M) \;:=\; \mathfrak{T}(q) \cup (\mathfrak{T}(p) \cap \{s \in |\mathfrak{T}| : \forall t(\text{if } s \stackrel{\alpha}{\longmapsto} t \text{ then } t \in M)\}).$$

The greatest fixed point of this monotone operator provides for a further safety property, namely:

$$s \in gfp(\Phi_3) \quad \Longleftrightarrow \quad \begin{cases} \text{For every } \alpha\text{-run from } s, \text{ until } r \in \mathfrak{T}(q) \text{ is reached all} \\ \text{states } t \text{ prior to this } r \text{ belong to } \mathfrak{T}(p) \;\; (\text{“}p \text{ until } q\text{”}). \end{cases}$$

From the previous characterization we obtain, in particular, that state s belongs to $gfp(\Phi_3)$ if p holds everywhere.

SYNTAX OF MODAL μ-CALCULUS

Let Ac be an arbitrary but fixed set of actions. We formulate the propositional modal μ-calculus in a language $\mathcal{L}(\mu)$ (depending on Ac) which comprises the following syntactically different basic symbols: (i) countably many atomic propositions P, Q, R and countably many variables U, V, W, X, Y, Z (both possibly with subscripts) plus the connective $\sim$ for forming the complements of atomic propositions and variables; (ii) the propositional constants $\bot$ and $\top$ and the propositional connectives $\vee$ and $\wedge$; (iii) for any action α from Ac, the modal operators $\langle\alpha\rangle$ and $[\alpha]$; the fixed point operators μ and ν. As auxiliary symbols we allow parentheses, brackets and commas.

Definition 2 *The* formulas $A, B, C, \ldots$ *(possibly with subscripts) of $\mathcal{L}(\mu)$ are inductively defined as follows:*

1. *All atomic propositions P and variables X as well as their complements $\widetilde{P}$ and $\widetilde{X}$ are formulas of $\mathcal{L}(\mu)$.*
2. *The propositional constants $\bot$ and $\top$ are formulas of $\mathcal{L}(\mu)$.*
3. *If A and B are formulas of $\mathcal{L}(\mu)$, then $(A \vee B)$ and $(A \wedge B)$ are formulas of $\mathcal{L}(\mu)$.*
4. *If α is an action and A a formula of $\mathcal{L}(\mu)$, then $\langle\alpha\rangle A$ and $[\alpha]A$ are formulas of $\mathcal{L}(\mu)$.*
5. *If A is a formula of $\mathcal{L}(\mu)$ which does not contain occurrences of $\widetilde{X}$, then $(\mu X)A$ and $(\nu X)A$ are formulas of $\mathcal{L}(\mu)$.*

The syntactic requirement in the last clause ensures that we can later associate a monotone operator to A and use it for defining the semantic meaning of $(\mu X)A$ and $(\nu X)A$.

In general, we will only speak of formulas if it is clear that we refer to formulas of $\mathcal{L}(\mu)$ and often omit parentheses whenever there is no danger of confusion. The fixed point operators μ and ν may be understood as a sort of quantifiers. Therefore we can speak about free and bound occurrences of a variable X within a formula A as usual.

To introduce substitution of variables by formulas, we proceed in two steps: First, we confine ourselves to substituting occurrences of a variable X within formulas which do not contain occurrences of $\widetilde{X}$. Then, after having defined the negations of formulas, we deal with the general case.

If X is a variable, A a formula which does not contain free occurrences of $\widetilde{X}$, and B an arbitrary formula, then $A[B/X]$ denotes the formula obtained from A by simultaneously replacing all free occurrences of X by B. In order to avoid collision of variables, a renaming of bound variables of A may be necessary.

The *negation* $\neg A$ of an $\mathcal{L}(\mu)$ formula A is inductively defined by the usual laws of double negation, the standard dualities for the propositional connectives and modal operators and the dualities with respect to least and greatest fixed points:

$$\begin{array}{rlrl}
\neg P &:= \widetilde{P}, & \neg\widetilde{P} &:= P,\\
\neg X &:= \widetilde{X}, & \neg\widetilde{X} &:= X,\\
\neg\bot &:= \top, & \neg\top &:= \bot,\\
\neg(A \vee B) &:= (\neg A \wedge \neg B), & \neg(A \wedge B) &:= (\neg A \vee \neg B),\\
\neg\langle\alpha\rangle A &:= [\alpha]\neg A, & \neg[\alpha]A &:= \langle\alpha\rangle\neg A,\\
\neg(\mu X)A &:= (\nu X)\neg D & \neg(\nu X)A &:= (\mu X)\neg D,
\end{array}$$

where $D := A[\widetilde{X}/X]$; we observe that $\neg D$ then does not contain occurrences of $\widetilde{X}$, hence the definitions of $\neg(\mu X)A$ and $\neg(\nu X)A$ make sense.

Now suppose that we are given two formulas A, B and a variable X. Then $A[B/X]$ is the formula which is obtained from A by simultaneously replacing all free occurrences of the variable X by B and all occurrences of $\widetilde{X}$ by $\neg B$; in order to avoid collision of variables, a renaming of bound variables may be necessary. If the formula A is written as $C[X]$, then we often simply write $C[B]$ instead of $C[B/X]$. Further variants of this notation will be obvious.

Further logical connectives are now introduced as abbreviations, for example,

$$(A \to B) \;:=\; (\neg A \vee B) \quad \text{and} \quad (A \leftrightarrow B) \;:=\; ((A \to B) \wedge (B \to A)).$$

A formula is said to be *normal* if all bound variables are distinct and different from the free variables. In all systems we consider in these notes any formula can be transformed into an equivalent normal formula (by renaming bound variables).

Following Kozen [15], we now recall a Hilbert-style axiomatization $\mathbf{K}(\mu)$ of the propositional modal μ-calculus: the multi-modal version of normal modal logic is extended by closure axioms and induction rules for the least fixed point formulas $(\mu X)A$.

I. Logical axioms of $\mathbf{K}(\mu)$. All propositional tautologies and the distribution axioms, i.e., for all propositional tautologies A, all formulas B, C, and all actions α:

(TAU) $\qquad A,$

(DIS) $\qquad [\alpha](B \to C) \;\to\; ([\alpha]B \to [\alpha]C).$

II. Logical rules of $\mathbf{K}(\mu)$. Modus ponens and necessitation, i.e., for all formulas A, B and all actions α:

(MP) $$\frac{A \qquad A \to B}{B},$$

(NEC) $$\frac{A}{[\alpha]A}.$$

III. Closure axioms of $\mathbf{K}(\mu)$. For all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$:

(μ-CL) $$A[(\mu X)A[X]] \;\to\; (\mu X)A[X].$$

IV. Induction rules of $\mathbf{K}(\mu)$. For all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$ and all formulas B:

(μ-IND) $$\frac{A[B] \;\to\; B}{(\mu X)A[X] \;\to\; B}.$$

Provability of a formula A in the Hilbert system $\mathbf{K}(\mu)$ is defined as usual and written as $\mathbf{K}(\mu) \vdash A$.

Although the closure axioms and induction rules of $\mathbf{K}(\mu)$ are only formulated for formulas $(\mu X)A[X]$, it is an easy exercise to show that the duals of (μ-CL) and (μ-IND) can be derived in $\mathbf{K}(\mu)$ for formulas of the form $(\nu X)A[X]$. All we have to do in order to prove the following lemma is to take the respective contrapositions and to recall the definition of the negations of formulas.

Lemma 3 *For all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$ and for all formulas B we have:*

1. $\mathbf{K}(\mu) \vdash (\nu X)A[X] \;\to\; A[(\nu X)A[X]]$.
2. $\mathbf{K}(\mu) \vdash B \;\to\; A[B] \quad \Longrightarrow \quad \mathbf{K}(\mu) \vdash B \;\to\; (\nu X)A[X]$.

Note that the closure axiom (μ-CL) and the first part of the previous lemma only state that $(\mu X)A[X]$ and $(\nu X)A[X]$ are a pre-fixed point and a post-fixed point of $A[U]$, respectively. Now we show in $\mathbf{K}(\mu)$ that both have the fixed point property. Before proving this, we turn to two useful properties of $\mathbf{K}(\mu)$.

Lemma 4 (Substitution) *For all formulas $A[U]$ and B we have*

$$\mathbf{K}(\mu) \vdash A[U] \quad \Longrightarrow \quad \mathbf{K}(\mu) \vdash A[B].$$

To prove this result, one simply verifies that all axioms and rules of inference of $\mathbf{K}(\mu)$ are closed under substitution and then proceeds by induction on the derivation of $A[U]$.

Lemma 5 (Monotonicity) *For all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$ and all formulas B, C we have*

$$\mathbf{K}(\mu) \vdash B \to C \quad \Longrightarrow \quad \mathbf{K}(\mu) \vdash A[B] \to A[C].$$

PROOF. We assume $\mathbf{K}(\mu) \vdash B \to C$ and proceed by induction on the build-up of $A[U]$. If U does not occur in $A[U]$ or if $A[U]$ is the variable U, then our assertion is trivially satisfied. If $A[U]$ is of the form $(D_0[U] \vee D_1[U])$, $(D_0[U] \wedge D_1[U])$, $\langle\alpha\rangle D[U]$, or $[\alpha]D[U]$, the assertion follows from the induction hypothesis by straightforward reasoning in $\mathbf{K}(\mu)$.

Now let $A[U]$ be of the form $(\mu X)D[X, U]$. Then the induction hypothesis implies

$$\mathbf{K}(\mu) \vdash D[X, B] \to D[X, C]$$

from which we obtain in view of the previous lemma that

$$\mathbf{K}(\mu) \vdash D[(\mu X)D[X, C], B] \to D[(\mu X)D[X, C], C].$$

From this implication and the following axiom (μ-CL)

$$D[(\mu X)D[X, C], C] \to (\mu X)D[X, C]$$

we infer

$$\mathbf{K}(\mu) \vdash D[(\mu X)D[X, C], B] \to (\mu X)D[X, C].$$

Taking this implication as premise of an induction rule (μ-IND) for $(\mu X)D[X, B]$ permits us to conclude

$$\mathbf{K}(\mu) \vdash (\mu X)D[X, B] \to (\mu X)D[X, C],$$

and this is what we have to show. It only remains the case that $A[U]$ is of the form $(\nu X)D[X, U]$ which by Lemma 3 can be treated accordingly. □

Having this lemma at hand, it is now an easy matter to prove that $(\mu X)A[X]$ and $(\nu X)A[X]$ are fixed points of $A[U]$.

Lemma 6 (Fixed points) *For all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$ we have:*

1. $\mathbf{K}(\mu) \vdash (\mu X)A[X] \leftrightarrow A[(\mu X)A[X]]$.
2. $\mathbf{K}(\mu) \vdash (\nu X)A[X] \leftrightarrow A[(\nu X)A[X]]$.

PROOF. Axiom (μ-CL) states the direction from right to left of the first assertion. For the converse direction consider the formula $B := A[(\mu X)A[X]]$. By axiom (μ-CL) we therefore have

$$\mathbf{K}(\mu) \vdash B \to (\mu X)A[X],$$

so that the previous lemma implies

$$\mathbf{K}(\mu) \vdash A[B] \to B.$$

It only remains to apply (μ-IND), and we obtain what we need. The second assertion of this lemma follows by duality. □

Thus we know that, provably in $\mathbf{K}(\mu)$, the formulas $(\mu X)A[X]$ and $(\nu X)A[X]$ really stand for fixed points of the formula $A[U]$. That they are the least and greatest such fixed point immediately follows from (μ-IND) and the second part of Lemma 3.

SEMANTICS OF MODAL μ-CALCULUS

Semantically, the modal μ-calculus can be elegantly approached from transition systems. We simply have to tell at which states the atomic propositions and variables are satisfied and then extend the valuation to arbitrary formulas according to the usual rules of modal logic, with an extra proviso for the formulas $(\mu X)A$ and $(\nu X)A$.

Definition 7 *A* μ-structure *is a transition system* $\mathfrak{T} = (St, Ac, Tr)$*, where* Ac *is the set of actions of* $\mathcal{L}(\mu)$*, associating sets of states* $\mathfrak{T}(P)$ *and* $\mathfrak{T}(X)$ *to all atomic propositions* P *and variables* X*.*

If M is a subset of $|\mathfrak{T}|$, then we write $\mathfrak{T}[M{:}X]$ for the μ-structure which maps the variable X to M and otherwise agrees with $\mathfrak{T}$.

Definition 8 *Given a* μ*-structure* $\mathfrak{T}$*, the* truth set $\|A\|_{\mathfrak{T}}$ *of a formula* A *is inductively defined as follows:*

1. *For atomic propositions, variables, and propositional constants:*

$$\begin{aligned} \|P\|_{\mathfrak{T}} &:= \mathfrak{T}(P), & \|\widetilde{P}\|_{\mathfrak{T}} &:= |\mathfrak{T}| \setminus \mathfrak{T}(P), \\ \|X\|_{\mathfrak{T}} &:= \mathfrak{T}(X), & \|\widetilde{X}\|_{\mathfrak{T}} &:= |\mathfrak{T}| \setminus \mathfrak{T}(X), \\ \|\top\|_{\mathfrak{T}} &:= |\mathfrak{T}|, & \|\bot\|_{\mathfrak{T}} &:= \emptyset. \end{aligned}$$

2. *For disjunctions and conjunctions:*

$$\|A \vee B\|_{\mathfrak{T}} := \|A\|_{\mathfrak{T}} \cup \|B\|_{\mathfrak{T}}, \qquad \|A \wedge B\|_{\mathfrak{T}} := \|A\|_{\mathfrak{T}} \cap \|B\|_{\mathfrak{T}}.$$

3. *For formulas prefixed by a modal operator:*

$$\begin{aligned} \|\langle\alpha\rangle B\|_{\mathfrak{T}} &:= \{s \in |\mathfrak{T}| : \exists t(s \stackrel{\alpha}{\longmapsto} t \text{ and } t \in \|B\|_{\mathfrak{T}})\}, \\ \|[\alpha] B\|_{\mathfrak{T}} &:= \{s \in |\mathfrak{T}| : \forall t(\text{if } s \stackrel{\alpha}{\longmapsto} t \text{ then } t \in \|B\|_{\mathfrak{T}})\}. \end{aligned}$$

4. *For fixed point formulas: Given a formula* $A[X]$ *which does not contain occurrences of* $\widetilde{X}$*, we first introduce the monotone operator*

$$\Phi_{A[X]} : \wp(|\mathfrak{T}|) \to \wp(|\mathfrak{T}|), \quad \Phi_{A[X]}(M) \;:=\; \|A[X]\|_{\mathfrak{T}[M:X]}.$$

Based on this Φ*, we now set*

$$\|(\mu X)A[X]\|_{\mathfrak{T}} \;:=\; \mathit{lfp}(\Phi_{A[X]}) \quad \text{and} \quad \|(\nu X)A[X]\|_{\mathfrak{T}} \;:=\; \mathit{gfp}(\Phi_{A[X]}).$$

We say that a formula A is *valid* in the μ-structure $\mathfrak{T}$, written $\mathfrak{T} \models A$, if $|\mathfrak{T}| = \|A\|_{\mathfrak{T}}$. A formula is defined to be *μ-valid* if it is valid in all μ-structures; in this case we write $\mu \models A$. Finally, a formula A is called *μ-satisfiable* if there exists a μ-structure $\mathfrak{T}$ such that $\|A\|_{\mathfrak{T}} \neq \emptyset$.

It is quite easy to see that all axioms of $\mathbf{K}(\mu)$ are μ-valid and that the inference rules of $\mathbf{K}(\mu)$ preserve μ-validity. Hence $\mathbf{K}(\mu)$ is sound. The completeness of $\mathbf{K}(\mu)$ turned out to be rather complicated and was finally solved in Walukiewicz [20].

Theorem 9 (Soundness and completeness of $\mathbf{K}(\mu)$) *For all formulas A we have*

$$\mathbf{K}(\mu) \vdash A \quad \Longleftrightarrow \quad \mu \models A.$$

The system $\mathbf{K}(\mu)$ is unsuitable for proof search – most notably because of (MP) – and defiant against proper proof-theoretic analysis. Therefore we will later introduce sound and complete finite and infinite sequent systems which are better tailored for proof-theoretic research.

Coming back to the four operators which we considered in the previous section in connection with least and greatest fixed points, we can now write down the corresponding formulas of $\mathcal{L}(\mu)$:

$$\begin{array}{lcllcl} \mathit{lfp}(\Phi_0) & \approx & (\mu X)(P \vee [\alpha]X), & \mathit{lfp}(\Phi_1) & \approx & (\mu X)(Q \vee (P \wedge \langle\alpha\rangle X)) \\ \mathit{gfp}(\Phi_2) & \approx & (\nu X)(P \wedge [\alpha]X), & \mathit{gfp}(\Phi_3) & \approx & (\nu X)(Q \vee (P \wedge [\alpha]X)). \end{array}$$

2. Basic results

The central semantic result about the propositional modal μ-calculus is the the so-called *fundamental semantic theorem* due to Streett and Emerson [19]; the subsequent presentation, however, follows Bradfield and Stirling [7].

Definition 10 *A* pre-model *is a pair $(\mathfrak{T}, \Vdash)$ such that $\mathfrak{T}$ is a μ-structure and $\Vdash$ is a binary relation satisfying*

$$\begin{array}{lrcl} \textit{for arbitrary } A\textit{:} & (\mathfrak{T}, s) \Vdash \neg A & \Longleftrightarrow & (\mathfrak{T}, s) \nVdash A, \\ \textit{for literals } D\textit{:} & (\mathfrak{T}, s) \Vdash D & \Longleftrightarrow & s \in \|D\|_{\mathfrak{T}}, \\ \textit{for non-literals:} & (\mathfrak{T}, s) \Vdash A \vee B & \Longleftrightarrow & (\mathfrak{T}, s) \Vdash A \ \textit{or} \ (\mathfrak{T}, s) \Vdash B, \\ & (\mathfrak{T}, s) \Vdash A \wedge B & \Longleftrightarrow & (\mathfrak{T}, s) \Vdash A \ \textit{and} \ (\mathfrak{T}, s) \Vdash B, \\ & (\mathfrak{T}, s) \Vdash \langle\alpha\rangle B & \Longleftrightarrow & \exists t (s \stackrel{\alpha}{\longmapsto} t \ \textit{and} \ (\mathfrak{T}, t) \Vdash B), \\ & (\mathfrak{T}, s) \Vdash [\alpha] B & \Longleftrightarrow & \forall t (\textit{if } s \stackrel{\alpha}{\longmapsto} t \ \textit{then} \ (\mathfrak{T}, t) \Vdash B), \\ & (\mathfrak{T}, s) \Vdash (\mu X) A[X] & \Longleftrightarrow & (\mathfrak{T}, s) \Vdash A[(\mu X) A[X]], \\ & (\mathfrak{T}, s) \Vdash (\nu X) A[X] & \Longleftrightarrow & (\mathfrak{T}, s) \Vdash A[(\nu X) A[X]]. \end{array}$$

Thus in a pre-model formulas $(\mu X)A[X]$ and $(\nu X)A[X]$ are interpreted as arbitrary fixed points, not necessarily as least and greatest fixed points, respectively. As a con-

sequence, there are pre-models $(\mathfrak{T}, \Vdash)$, states s, and formulas A with $(\mathfrak{T}, s) \Vdash A$ and $s \notin \|A\|_{\mathfrak{T}}$. Thus validity with respect to all pre-models does not coincide with μ-validity. However, it is possible to characterize the those pre-models $(\mathfrak{T}, \Vdash)$ which behave "adequately" in this respect (see below).

Definition 11 *Let $(\mathfrak{T}, \Vdash)$ be a pre-model. A function f which assigns to any state $s \in |\mathfrak{T}|$ and formula $(A \vee B)$ a formula $f(s, A \vee B)$ and to any state $s \in |\mathfrak{T}|$ and formula $\langle\alpha\rangle A$ a state $f(s, \langle\alpha\rangle A) \in |\mathfrak{T}|$ is called a* choice function *for $(\mathfrak{T}, \Vdash)$ if it satisfies the following two conditions:*

(C.1) For every $s \in |\mathfrak{T}|$ and $(A \vee B)$ we have

$$(\mathfrak{T}, s) \Vdash A \vee B \quad \Longrightarrow \quad f(s, A \vee B) \in \{A, B\} \ \textit{and} \ (\mathfrak{T}, s) \Vdash f(s, A \vee B).$$

(C.2) For every $s \in |\mathfrak{T}|$ and $\langle\alpha\rangle A$ we have

$$(\mathfrak{T}, s) \Vdash \langle\alpha\rangle A \quad \Longrightarrow \quad s \stackrel{\alpha}{\longmapsto} f(s, \langle\alpha\rangle A) \ \textit{and} \ (\mathfrak{T}, f(s, \langle\alpha\rangle A)) \Vdash A.$$

Given a pre-model $(\mathfrak{T}, \Vdash)$ and a choice function f for $(\mathfrak{T}, \Vdash)$ we introduce a *dependency relation* $\succ$ on expressions of the form $(\mathfrak{T}, f, s) \Vdash A$, where A is assumed to be normal, by requiring that

$$\begin{aligned}
(\mathfrak{T}, f, s) \Vdash A_1 \vee A_2 \ &\succ \ (\mathfrak{T}, f, s) \Vdash f(s, A_1 \vee A_2), \\
(\mathfrak{T}, f, s) \Vdash A_1 \wedge A_2 \ &\succ \ (\mathfrak{T}, f, s) \Vdash A_i \ \text{ for } \ i = 1, 2, \\
(\mathfrak{T}, f, s) \Vdash \langle\alpha\rangle B \ &\succ \ (\mathfrak{T}, f, f(s, \langle\alpha\rangle B)) \Vdash B, \\
(\mathfrak{T}, f, s) \Vdash [\alpha] B \ &\succ \ (\mathfrak{T}, f, t) \Vdash B \ \text{ for all } t \text{ such that } s \stackrel{\alpha}{\longmapsto} t, \\
(\mathfrak{T}, f, s) \Vdash (\mu X) A[X] \ &\succ \ (\mathfrak{T}, f, s) \Vdash A[(\mu X) A[X]], \\
(\mathfrak{T}, f, s) \Vdash (\nu X) A[X] \ &\succ \ (\mathfrak{T}, f, s) \Vdash A[(\nu X) A[X]].
\end{aligned}$$

A *trail* for $(\mathfrak{T}, f, s) \Vdash A$ is a maximal chain of dependencies

$$(\mathfrak{T}, f, s_0) \Vdash A_0 \ \succ \ (\mathfrak{T}, f, s_1) \Vdash A_1 \ \succ \ (\mathfrak{T}, f, s_2) \Vdash A_2 \ \succ \ \cdots$$

with s_0 being the state s and A_0 being the formula A. A choice function f for $(\mathfrak{T}, \Vdash)$ is called *well-founded* if for every state $s \in |\mathfrak{T}|$ and every formula A the following requirement is fulfilled: the outermost bounded variable occurring infinitely often in any trail for $(\mathfrak{T}, f, s) \Vdash A$ is bounded by ν. The pre-model $(\mathfrak{T}, \Vdash)$ is called *well-founded* if there exists a well-founded choice function for $(\mathfrak{T}, \Vdash)$.

Theorem 12 (Fundamental semantic theorem)

1. *Let $\mathfrak{T}$ be a μ-structure. Then there exists a well-founded pre-model $(\mathfrak{T}, \Vdash)$ such that for any $s \in |\mathfrak{T}|$ and any normal formula A*

$$s \in \|A\|_{\mathfrak{T}} \quad \Longleftrightarrow \quad (\mathfrak{T}, s) \Vdash A.$$

2. *Let $(\mathfrak{T}, \Vdash)$ be a well-founded pre-model. Then we have for all $s \in |\mathfrak{T}|$ and all normal formulas A that*

$$(\mathfrak{T}, s) \Vdash A \quad \Longrightarrow \quad s \in \|A\|_{\mathfrak{T}}.$$

For a proof of this theorem one may consult the original publication Streett and Emerson [19] or Bradfield and Stirling [7], where the central ideas are described. There are also more recent presentations of this result and its proof in an automata-theoretic environment; see, for example, Wilke [21].

The fundamental semantic theorem and the techniques developed for its proof enabled Streett and Emerson to obtain the decidability and small model property of the modal μ-calculus. Again we omit proofs and refer to Streett and Emerson [19].

Theorem 13 (Decidability and small model property)

1. *Given a formula A, it is decidable whether A is μ-satisfiable.*
2. *If the formula A is μ-satisfiable, then there exists a finite μ-structure $\mathfrak{T}$ of size exponential in the size of A such that $\|A\|_{\mathfrak{T}} \neq \emptyset$.*

Open problem 14 For a μ-structure $\mathfrak{T}$ with state $s \in |\mathfrak{T}|$ and a formula A a typical question of model-checking is: Do we have $s \in \|A\|_{\mathfrak{T}}$? For finite μ-structures $\mathfrak{T}$ this question is decidable and known to be in $\mathrm{NP} \cap \mathrm{coNP}$ with respect to the size of $\mathfrak{T}$ plus the size of A. But is it polynomial?

Now we leave the semantics of the modal μ-calculus and turn to some of its proof-theoretic aspects. As mentioned at the end of Section 1, Hilbert systems are inappropriate for proof-theoretic investigations. There exist a lot of proof-theoretically relevant work about tableau systems for the modal μ-calculus and game-theoretic approaches (e.g., Stirling and Walker [18], Niwiński and Walukiewicz [17]). Here we follow a different track and focus on a traditional sequent-style approach.

We present two Tait-style systems $\mathbf{K}^{\omega}(\mu)$ and $\mathbf{K}^{<\omega}(\mu)$ for the modal μ-calculus, which both are sound, complete and cut-free. $\mathbf{K}^{\omega}(\mu)$ is an infinitary deduction system, introducing greatest fixed points $(\nu X)A[X]$ by a sort of ω-rule (ω-ν); $\mathbf{K}^{<\omega}(\mu)$ is the finitization of $\mathbf{K}^{\omega}(\mu)$. In the formulation of the rule (ω-ν) we use the finite approximations of $(\nu X)A[X]$ which are inductively defined, for each natural number $n > 0$, as follows:

$$(\nu X)^1 A[X] \;:=\; A[\top] \quad \text{and} \quad (\nu X)^{n+1} A[X] \;:=\; A[(\nu X)^n A[X]].$$

Both, $\mathbf{K}^{\omega}(\mu)$ and $\mathbf{K}^{<\omega}(\mu)$, derive finite sets $\Gamma, \Delta, \Pi, \Sigma, \ldots$ (possibly with subscripts) of formulas rather than individual formulas. These finite sets of formulas are interpreted disjunctively, and in general we write Γ, A for $\Gamma \cup \{A\}$; similarly for expressions like Γ, Δ, A, B. In addition, if Γ is the set $\{A_1, \ldots, A_m\}$ and α some action, then we set

$$\langle\alpha\rangle\Gamma \;:=\; \{\langle\alpha\rangle A_1, \ldots, \langle\alpha\rangle A_m\} \quad \text{and} \quad \Gamma^{\vee} \;:=\; A_1 \vee \ldots \vee A_m.$$

$\mathbf{K}^{\omega}(\mu)$ contains the standard axioms and logical rules of multi-modal logic, the Tait-style analogues of the μ-closure-axioms plus the above mentioned infinitary rule for introducing $(\nu X)A[X]$.

I. Axioms of $\mathbf{K}^\omega(\mu)$. For all finite formula sets Γ, all atomic propositions P, and all variables X:

(Ax1) $$\Gamma,\ \top,$$

(Ax2) $$\Gamma,\ P,\ \widetilde{P},$$

(Ax3) $$\Gamma,\ X,\ \widetilde{X}.$$

II. Logical rules of $\mathbf{K}^\omega(\mu)$. For all finite formula sets Γ, Δ, all actions α, and all formulas A, B:

($\vee$) $$\frac{\Gamma,\ A,\ B}{\Gamma,\ A \vee B},$$

($\wedge$) $$\frac{\Gamma,\ A \qquad \Gamma,\ B}{\Gamma,\ A \wedge B},$$

(DIS) $$\frac{\Gamma,\ A}{\langle\alpha\rangle\Gamma,\ [\alpha]A,\ \Delta}.$$

III. μ-rules of $\mathbf{K}^\omega(\mu)$. For all finite formula sets Γ and all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$:

(μ) $$\frac{\Gamma,\ A[(\mu X)A[X]]}{\Gamma,\ (\mu X)A[X]}.$$

IV. ν-rules of $\mathbf{K}^\omega(\mu)$. For all finite formula sets Γ and all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$:

(ω-ν) $$\frac{\ldots\ \Gamma,\ (\nu X)^n A[X]\ \ldots \quad \text{(for all } 0 < n < \omega)}{\Gamma,\ (\nu X)A[X]}.$$

Provability of Γ in $\mathbf{K}^\omega(\mu)$ is defined as usual and denoted by $\mathbf{K}^\omega(\mu) \vdash \Gamma$. On account of the rule (ω-ν) there are derivations in $\mathbf{K}^\omega(\mu)$ which are infinitely branching trees of infinite depths.

In Jäger, Kretz, and Studer [14] the completeness of $\mathbf{K}^\omega(\mu)$ is proved by adapting the canonical saturated sets construction. Problems only arise because of the impredicativity of the rule (μ): the logical complexity of $A[(\mu X)A[X]]$ is greater than that of $(\mu X)A[X]$. As a consequence, proofs by induction on the lengths of formulas cannot be carried through directly. But by carefully assigning finite sequences of ordinals (rather than ordinals) to formulas and taking up ideas from Streett and Emerson [19], we achieve our goal.

Theorem 15 (Completeness of $\mathbf{K}^\omega(\mu)$) *For all sentences A we have*

$$\mu \models A \quad \Longrightarrow \quad \mathbf{K}^\omega(\mu) \vdash A.$$

Rather than showing the soundness of $\mathbf{K}^{\omega}(\mu)$ directly, we move on to its finitization $\mathbf{K}^{<\omega}(\mu)$. Looking at the rules of $\mathbf{K}^{\omega}(\mu)$, we immediately notice that only the rule (ω-ν) is responsible for possibly infinite derivations. Hence all proofs will be finite if we succeed in restricting the infinitely many premises of each application of (ω-ν) to a finite subset. Fortunately, this can be achieved by exploiting the small model property of the modal μ-calculus, see Theorem 13.

From the small model property of the modal μ-calculus we know that there exists a function ℓ, defined on all finite sets of formulas Γ and exponential in the number of symbols occurring in Γ, which has the following property: If $\Gamma^{\vee}$ is μ-satisfiable, then there exists a μ-structure $\mathfrak{T}$ such that the cardinality of $|\mathfrak{T}|$ is smaller than $\ell(\Gamma)$ and $\|\Gamma^{\vee}\|_{\mathfrak{T}} \neq \emptyset$.

With this bounding function ℓ at our disposal, the finite versions of the ν-rules are obtained. The rules (fin-ν) are the truncations of (ω-ν) at a bound provided by ℓ.

V. Finite ν-rules. For all finite formula sets Γ, Δ and all formulas $A[U]$ which do not contain occurrences of $\widetilde{U}$:

$$\text{(fin-}\nu\text{)} \qquad \frac{\ldots\ \Gamma,\ (\nu X)^n A[X]\ \ldots \quad \text{(for all } 0 < n < \ell(\Gamma, (\nu X)A[X]))}{\Gamma,\ (\nu X)A[X],\ \Delta}.$$

The system $\mathbf{K}^{<\omega}(\mu)$ is obtained from $\mathbf{K}^{\omega}(\mu)$ by replacing the ν-rules (ω-ν) by their finite variants (fin-ν); the notion $\mathbf{K}^{<\omega}(\mu) \vdash \Gamma$ is introduced in analogy to $\mathbf{K}^{\omega}(\mu) \vdash \Gamma$.

Observe that the number of premises of a finite ν-rule depends on the length of (the essential part of) its conclusion; the set Δ is added in the conclusions just to incorporate weakening.

Naturally, $\mathbf{K}^{<\omega}(\mu)$ is a finite system. Besides that, every derivation in $\mathbf{K}^{\omega}(\mu)$ collapses to a derivation in $\mathbf{K}^{<\omega}(\mu)$. The proof of this observation is by induction on the derivations in $\mathbf{K}^{\omega}(\mu)$, and one only has to observe that each application of a rule (ω-ν) in $\mathbf{K}^{\omega}(\mu)$ may be replaced by the appropriate rule (fin-ν) in $\mathbf{K}^{<\omega}(\mu)$.

Lemma 16 *For all finite sets Γ of formulas we have*

$$\mathbf{K}^{\omega}(\mu) \vdash \Gamma \quad \Longrightarrow \quad \mathbf{K}^{<\omega}(\mu) \vdash \Gamma.$$

Of course, this means that the completeness of $\mathbf{K}^{\omega}(\mu)$ transfers to $\mathbf{K}^{<\omega}(\mu)$; just combine the previous lemma with Theorem 15.

Corollary 17 (Completeness of $\mathbf{K}^{<\omega}(\mu)$) *For all sentences A we have*

$$\mu \models A \quad \Longrightarrow \quad \mathbf{K}^{<\omega}(\mu) \vdash A.$$

What remains is to show the soundness of $\mathbf{K}^{<\omega}(\mu)$. The following auxiliary consideration is an immediate consequence of the properties of approximations of greatest fixed points, which have been stated in Section 1.

Lemma 18 *Let $\mathfrak{T}$ be a μ-structure whose universe $|\mathfrak{T}|$ contains at most n elements (n a positive natural number). For all formulas $(\nu X)A[X]$ we then have*

$$\|(\nu X)A[X]\|_{\mathfrak{T}} \ = \ \|(\nu X)^n A[X]\|_{\mathfrak{T}}.$$

Combining this lemma with the small model property of the μ-calculus, we can now easily establish the soundness of $\mathbf{K}^{<\omega}(\mu)$.

Theorem 19 (Soundness of $\mathbf{K}^{<\omega}(\mu)$) *For all finite sets Γ of formulas we have*

$$\mathbf{K}^{<\omega}(\mu) \vdash \Gamma \quad \Longrightarrow \quad \mu \models \Gamma^{\vee}.$$

PROOF. The proof proceeds by induction on the derivation of Γ, and we distinguish the following cases:
1. Γ is an axiom or the conclusion of a logical rule of $\mathbf{K}^{<\omega}(\mu)$. Then our assertion is obvious or an immediate consequence of the induction hypothesis.
2. Γ is the conclusion of a μ-rule of $\mathbf{K}^{<\omega}(\mu)$. Then there exist a set Δ and a formula $(\mu X)A[X]$ so that Γ is the set $\Delta, (\mu X)A[X]$, and this rule has the form

$$\frac{\Delta,\ A[(\mu X)A[X]]}{\Delta,\ (\mu X)A[X]}\ .$$

Now the induction hypothesis yields

$$\mu \ \models\ \Delta^{\vee} \vee A[(\mu X)A[X]]. \tag{1}$$

But according to our semantics we also have

$$\mu \ \models\ A[(\mu X)A[X]] \rightarrow (\mu X)A[X], \tag{2}$$

and therefore the desired μ-validity of $\Gamma^{\vee}$ is a trivial from (1) and (2).
3. Γ is the conclusion of a finite ν-rule of $\mathbf{K}^{<\omega}(\mu)$. Then there exist sets Δ, Π and a formula $(\nu X)A[X]$ so that Γ is the set $\Delta, (\nu X)A[X], \Pi$, and this rule has the form

$$\frac{\ldots\ \Delta,\ (\nu X)^n A[X]\ \ldots \quad \text{(for all } 0 < n < \ell(\Delta, (\nu X)A[X]))}{\Delta,\ (\nu X)A[X],\ \Pi}\ .$$

In this case the induction hypothesis yields

$$\mu \ \models\ \Delta^{\vee} \vee (\nu X)^n A[X] \tag{3}$$

for all natural numbers n such that $0 < n < \ell(\Delta, (\nu X)A[X])$. Now assume that the formula $\Delta^{\vee} \vee (\nu X)A[X]$ is not μ-valid. Then $\neg\Delta^{\vee} \wedge \neg(\nu X)A[X]$ has to be μ-satisfiable, and we infer from the small model property that there exists a μ-structure $\mathfrak{T}$ such that the cardinality of $|\mathfrak{T}|$, we call it k, is smaller than $\ell(\Delta, (\nu X)A[X])$ and

$$\|\neg\Delta^{\vee} \wedge \neg(\nu X)A[X]\|_{\mathfrak{T}}\ \neq\ \emptyset. \tag{4}$$

In view of Lemma 18 this inequality can be rewritten as

$$\|\neg\Delta^{\vee} \wedge \neg(\nu X)^k A[X]\|_{\mathfrak{T}}\ \neq\ \emptyset, \tag{5}$$

implying that the formula $\Delta^\vee \vee (\nu X)^k A[X]$ is not μ-valid. However, this is in contradiction to (3), and therefore $\Delta^\vee \vee (\nu X)A[X]$ has to be μ-valid. This completes the proof of our theorem. $\Box$

Considering this theorem in the context of Lemma 16, it provides the soundness of the infinitary calculus $\mathbf{K}^\omega(\mu)$.

Corollary 20 *For all finite sets Γ of formulas we have*

$$\mathbf{K}^{<\omega}(\mu) \vdash \Gamma \quad \Longleftrightarrow \quad \mathbf{K}^{\omega}(\mu) \vdash \Gamma \quad \Longleftrightarrow \quad \mu \models \Gamma^\vee.$$

While the previous tells us that the finite Hilbert-style system $\mathbf{K}(\mu)$ and the infinitary Tait-style system $\mathbf{K}^\omega(\mu)$ prove the same sentences, we have no way (yet?) to take a proof of a sentence A in $\mathbf{K}(\mu)$ and transform it into a proof of A in $\mathbf{K}^\omega(\mu)$. This is due to the lack of an equivalent of Modus Ponens in $\mathbf{K}^\omega(\mu)$. To overcome this deficiency, we add a further rule.

The cut rule. For all finite formula sets Γ and all formulas A:

$$(\mathrm{cut}) \qquad \frac{\Gamma,\, A \qquad \Gamma,\, \neg A}{\Gamma}\,.$$

The formulas A and $\neg A$ are called the cut formulas of this cut.

It is fairly easy to see that every proof of a formula A of $\mathcal{L}(\mu)$ within the system $\mathbf{K}(\mu)$ can be translated – in a natural way – into a proof of A within the system $\mathbf{K}^\omega(\mu) + (\mathrm{cut})$, where (cut) takes over the rôle of (MP) in this translation.

Nevertheless, if only provability (and not the translation of proofs) is considered, the cut-rule is not needed. Since (cut) is obviously correct, semantic cut elimination follows from Corollary 20.

Corollary 21 (Semantic cut elimination) *For all finite sets Γ of formulas of $\mathcal{L}(\mu)$ we have:*

1. $\mathbf{K}^\omega(\mu) + (\mathrm{cut}) \vdash \Gamma \quad \Longrightarrow \quad \mathbf{K}^\omega(\mu) \vdash \Gamma.$
2. $\mathbf{K}^{<\omega}(\mu) + (\mathrm{cut}) \vdash \Gamma \quad \Longrightarrow \quad \mathbf{K}^{<\omega}(\mu) \vdash \Gamma.$

What we have achieved are a natural infinitary axiomatization of the propositional modal μ-calculus and its finitization $\mathbf{K}^{<\omega}(\mu)$, which are both sound and complete. They are cut-free, but because of their completeness, cut rules could be added without changing their strength.

$\mathbf{K}^{<\omega}(\mu)$ is the finite collapse of $\mathbf{K}^\omega(\mu)$, but one may argue how natural $\mathbf{K}^{<\omega}(\mu)$ is as a deductive system. However, the important purpose of this system is to provide an explicit proof that a cut-free adequate axiomatization of the propositional modal μ-calculus exists.

Open problems 22

1. Are there syntactic cut elimination procedures for $\mathbf{K}^\omega(\mu)$ and $\mathbf{K}^{<\omega}(\mu)$?
2. Is there a more natural finite derivation system for the modal μ-calculus which is cut-free, sound and complete?

3. Knowledge and common knowledge

In this section we consider subsystems of the full modal μ-calculus which play an important rôle in the context of epistemic logic and epistemic reasoning. We fix a natural number $n \geq 1$ and concentrate on transition systems whose set of actions is the set $\{1, \ldots, n\}$. Actions are now called *agents* and may stand for any nodes (e.g., persons, processors) in a complex distributed and possibly communicating environment.

Definition 23 *An* n-knowledge structure *is a* μ*-structure* $\mathfrak{T} = (St, Ac, Tr)$ *whose set of actions* Ac *is the set* $\{1, \ldots, n\}$.

In the context of n-knowledge structures $\mathfrak{T}$, given a natural number α such that $1 \leq \alpha \leq n$ and states $s, t \in |\mathfrak{T}|$, we propose to read

$$s \stackrel{\alpha}{\longmapsto} t \quad \text{as} \quad \text{agent } \alpha \text{ at state } s \text{ considers state } t \text{ as possible.}$$

The modal-logic approach to modeling the knowledge of an agent α is to identify α's knowledge with what is the case in all states that α considers possible at the present state:

$$\alpha \text{ knows } A \text{ at } s \quad \Longleftrightarrow \quad A \text{ holds at all states } t \text{ that } \alpha \text{ considers possible at } s.$$

As this coincides with the semantics of $[\alpha]A$, the (informal) interpretation of $[\alpha]A$ as "agent α knows that A" is justified.

A word of caution: Often knowing a statement A is supposed to imply the truth of A. If we want this to be the case here as well, then, for all agents α and all formulas A, the standard truth axioms

$$\text{(T)} \qquad [\alpha]A \;\rightarrow\; A$$

have to be added and only n-knowledge structures with reflexive transition relations must be considered. Without (T), the formula $[\alpha]A$ is then the formalization of "agent α believes that A". Further possible strengthenings of knowledge add the axioms about positive introspection

$$\text{(PI)} \qquad [\alpha]A \;\rightarrow\; [\alpha][\alpha]A$$

or even positive introspection plus negative introspection

$$\text{(NI)} \qquad \neg[\alpha]A \;\rightarrow\; [\alpha]\neg[\alpha]A,$$

which on the semantic side corresponds to the restriction to n-knowledge structures whose transition relations are reflexive-transitive relations or equivalence relations, respectively. This all is more or less a matter of taste or context and not relevant for us in the following. Therefore we confine us here to the most elementary case without (T), (PI) or (NI).

With n agents around, "everybody knows that A" is written $\mathbb{E}[A]$ and defined by

$$\mathbb{E}[A] \;:=\; [1]A \wedge \ldots \wedge [n]A.$$

It must not be confused with the common knowledge of A. To see why, we recall the famous muddy children puzzle taken, in this formulation, from Fagin, Moses, Halpern, and Vardi [10]:

> There are n children playing together. During their play some of the children, say k of them, get mud on their foreheads. Each can see the mud on others but not on his own forehead. Along comes a father, who says, "At least one of you has mud on your forehead". He then asks the following question, over and over: "Can any of you prove that you have mud on your forehead?" Assuming that all the children are perceptive, intelligent, truthful, and that they answer simultaneously, what will happen?
>
> There is a proof that the first $k-1$ times the father asks the question, the children will all say "no" but that the k-th time the children that are dirty will answer "yes".

The rôle of the father's announcement is that all children know that at least one of them has mud on his/her forehead *and that all know that the others also know that this is the case*; actually, this fact becomes common knowledge. Moreover, whenever that father repeats his question, all children can deduce that so far his question could not be answered. Try to find out what happens without the father's announcement.

The iterations $\mathbb{E}^m[A]$ of "everybody knows" are inductively defined, for any natural number m, by

$$\mathbb{E}^0[A] \;:=\; A \qquad \text{and} \qquad \mathbb{E}^{m+1}[A] \;:=\; \mathbb{E}[\mathbb{E}^m[A]],$$

and the infinite conjunction $\bigwedge_{m\geq 1} \mathbb{E}^m[A]$ reflects the intuitive idea that A is common knowledge.

SYNTAX OF $\mathbf{K}_n(\mathbb{C})$

The language $\mathcal{L}_n(\mathbb{C})$ for n agents and common knowledge is the modification of the language $\mathcal{L}(\mu)$ obtained by specifying the set Ac to be the set $\{1,\ldots,n\}$, dropping the fixed point operators μ and ν and adding instead two new operators $\mathbb{C}$ and $\widetilde{\mathbb{C}}$. The formulas $A, B, C, \ldots$ (possibly with subscripts) of $\mathcal{L}_n(\mathbb{C})$ are defined by the following grammar:

$$A \;::=\; P \mid \widetilde{P} \mid X \mid \widetilde{X} \mid \bot \mid \top \mid (A \vee A) \mid (A \wedge A) \mid \langle\alpha\rangle A \mid [\alpha]A \mid \mathbb{C}(A) \mid \widetilde{\mathbb{C}}(A),$$

where P and X range over atomic propositions and variables, respectively, and α is a natural number, $1 \leq \alpha \leq n$. The negation $\neg A$ of an $\mathcal{L}_n(\mathbb{C})$ formula A is defined as before with the clauses

$$\neg\mathbb{C}(A) \;:=\; \widetilde{\mathbb{C}}(\neg A) \qquad \text{and} \qquad \neg\widetilde{\mathbb{C}}(A) \;:=\; \mathbb{C}(\neg A)$$

for the operators $\mathbb{C}$ and $\widetilde{\mathbb{C}}$. Picking some variable X which does not occur in A and replacing $\mathbb{C}(A)$ by $(\nu X)\mathbb{E}[A \wedge X]$ and $\widetilde{\mathbb{C}}(A)$ by $(\mu X)\mathbb{D}[A \vee X]$ with

$$\mathbb{D}[U] \;:=\; \langle 1\rangle U \;\vee\; \ldots \;\vee\; \langle n\rangle U.$$

yields an embedding of $\mathcal{L}_n(\mathbb{C})$ into $\mathcal{L}(\mu)$. Later we will see that this translation does exactly what is intended.

Now we recall a Hilbert-style axiomatization $\mathbf{K}_n(\mathbb{C})$ for n agents and common knowledge as presented, for example, in Fagin, Moses, Halpern, and Vardi [10]. Its logical axioms and logical rules are the same as for $\mathbf{K}(\mu)$, formulated for the language $\mathcal{L}_n(\mathbb{C})$. In addition, for all $\mathcal{L}_n(\mathbb{C})$ formulas A, the system $\mathbf{K}_n(\mathbb{C})$ comprises a co-closure axiom for the operator $\mathbb{C}$ and the corresponding induction principle.

Co-closure axioms of $\mathbf{K}_n(\mathbb{C})$. For all $\mathcal{L}_n(\mathbb{C})$ formulas A:

($\mathbb{C}$-CCL) $$\mathbb{C}(A) \;\rightarrow\; \mathbb{E}[A \wedge \mathbb{C}(A)].$$

Induction rules of $\mathbf{K}_n(\mathbb{C})$. For all $\mathcal{L}_n(\mathbb{C})$ formulas A, B:

($\mathbb{C}$-IND) $$\frac{B \;\rightarrow\; \mathbb{E}[A \wedge B]}{B \rightarrow \mathbb{C}(A)}.$$

Of course, provability of an $\mathcal{L}_n(\mathbb{C})$ formula A in the Hilbert system $\mathbf{K}_n(\mathbb{C})$ is denoted by $\mathbf{K}_n(\mathbb{C}) \vdash A$. The co-closure axioms and induction rules are the syntactic form of expressing that $\mathbb{C}(A)$ is the greatest fixed point of the formula $\mathbb{E}[A \wedge U]$.

SEMANTICS OF $\mathbf{K}_n(\mathbb{C})$

To set up the semantics of $\mathbf{K}_n(\mathbb{C})$, we take an n-knowledge structure $\mathfrak{T}$ and proceed in defining the truth set $\|A\|_{\mathfrak{T}}$ of an $\mathcal{L}_n(\mathbb{C})$ formula A as in Definition 8 provided that A does not begin with $\mathbb{C}$ or $\widetilde{\mathbb{C}}$ and set otherwise:

$$\|\mathbb{C}(A)\|_{\mathfrak{T}} \;:=\; \bigcap_{m \geq 1} \|\mathbb{E}^m[A]\|_{\mathfrak{T}} \qquad \text{and} \qquad \|\widetilde{\mathbb{C}}(A)\|_{\mathfrak{T}} \;:=\; |\mathfrak{T}| \setminus \|\mathbb{C}(\neg A)\|_{\mathfrak{T}}.$$

Then an $\mathcal{L}_n(\mathbb{C})$ formula A is called $(n, \mathbb{C})$-*valid*, denoted by $(n, \mathbb{C}) \models A$, if $|\mathfrak{T}| = \|A\|_{\mathfrak{T}}$ for all n-knowledge structures $\mathfrak{T}$.

It is easily verified that $\|\mathbb{C}(A)\|_{\mathfrak{T}}$ is the greatest fixed point of the monotone operator Φ on $|\mathfrak{T}|$, satisfying, for any $M \subseteq |\mathfrak{T}|$,

$$\Phi(M) \;=\; \|\mathbb{E}[A \wedge X]\|_{\mathfrak{T}[M:X]},$$

and $\|\widetilde{\mathbb{C}}(A)\|_{\mathfrak{T}}$ is the least fixed point of the monotone operator Ψ on $|\mathfrak{T}|$, satisfying, for any $M \subseteq |\mathfrak{T}|$,

$$\Psi(M) \;=\; \|\mathbb{D}[A \vee X]\|_{\mathfrak{T}[M:X]},$$

where X is not to occur in A. $\mathbf{K}_n(\mathbb{C})$ can be shown to be sound an complete; see, for example, Fagin, Moses, Halpern, and Vardi [10].

Theorem 24 (Soundness and completeness of $\mathbf{K}_n(\mathbb{C})$) *For all $\mathcal{L}_n(\mathbb{C})$ formulas A we have*

$$\mathbf{K}_n(\mathbb{C}) \vdash A \quad \Longleftrightarrow \quad (n, \mathbb{C}) \models A.$$

From what we have mentioned in Section 2 it follows that $(n, \mathbb{C})$-validity of an $\mathcal{L}_n(\mathbb{C})$ formula is decidable. Without going into details we also mention that checking for validity is EXPTIME-complete in the size of the input formula and refer to Halpern and Moses [12] for further details.

A TAIT-STYLE REFORMULATION OF $\mathbf{K}_n(\mathbb{C})$

Our interest is in the proof theory of common knowledge. Since $\mathbf{K}_n(\mathbb{C})$ can be regarded as a subsystem of $\mathbf{K}(\mu)$, we can, of course, proceed as in the previous section and go over from $\mathbf{K}_n(\mathbb{C})$ to an infinitary Tait-style version $\mathbf{K}_n^{\omega}(\mathbb{C})$ and its finitization $\mathbf{K}_n^{<\omega}(\mathbb{C})$. Both systems are cut-free and provide sound and complete axiomatizations of common knowledge. $\mathbf{K}_n^{<\omega}(\mathbb{C})$ has, more or less, the same positive properties as $\mathbf{K}^{<\omega}(\mu)$ and should only be considered as a basis for more research about cut-free common knowledge. See Jäger, Kretz, and Studer [13] and Brünnler and Studer [8] for more work in this direction.

What we want to do now is to present a natural Tait-style reformulation $\overline{\mathbf{K}}_n(\mathbb{C})$ of $\mathbf{K}_n(\mathbb{C})$ which allows us to control all cuts involved, but unfortunately, does not permit full cut elimination. $\overline{\mathbf{K}}_n(\mathbb{C})$ derives finite sets of formulas, comprises the usual axioms and rules of Tait-calculi for multi-modal logic plus additional rules for the epistemic operators. If Γ is the set $\{A_1, \ldots, A_m\}$ of $\mathcal{L}_n(\mathbb{C})$ formulas, we set

$$\widetilde{\mathbb{C}}(\Gamma) \;:=\; \{\widetilde{\mathbb{C}}(A_1), \ldots, \widetilde{\mathbb{C}}(A_m)\}.$$

I. Axioms of $\overline{\mathbf{K}}_n(\mathbb{C})$. For all finite formula sets Γ of $\mathcal{L}_n(\mathbb{C})$ formulas, all atomic propositions P and all variables X:

(Ax1) $$\Gamma,\ \top,$$

(Ax2) $$\Gamma,\ P,\ \widetilde{P},$$

(Ax3) $$\Gamma,\ X,\ \widetilde{X}.$$

II. Logical rules of $\overline{\mathbf{K}}_n(\mathbb{C})$. For all finite formula sets Γ, Δ, Π of $\mathcal{L}_n(\mathbb{C})$ formulas, all agents α $(1 \leq \alpha \leq n)$, and all $\mathcal{L}_n(\mathbb{C})$ formulas A, B:

($\vee$) $$\frac{\Gamma,\ A,\ B}{\Gamma,\ A \vee B},$$

($\wedge$) $$\frac{\Gamma,\ A \qquad \Gamma,\ B}{\Gamma,\ A \wedge B},$$

(DIS) $$\frac{\Gamma,\ A,\ \widetilde{\mathbb{C}}(\Delta)}{\langle\alpha\rangle\Gamma,\ [\alpha]A,\ \widetilde{\mathbb{C}}(\Delta),\ \Pi}.$$

III. $\mathbb{C}$-rules of $\overline{\mathbf{K}}_n(\mathbb{C})$. For all finite formula sets Γ, Δ, Π of $\mathcal{L}_n(\mathbb{C})$ formulas and all $\mathcal{L}_n(\mathbb{C})$ formulas A:

$$(\widetilde{\mathbb{C}}) \qquad \frac{\Gamma,\ \neg\mathbb{E}[\neg A]}{\Gamma,\ \widetilde{\mathbb{C}}(A)},$$

$$(\mathbb{C}) \qquad \frac{\mathbb{E}[A],\ \widetilde{\mathbb{C}}(\Delta)}{\mathbb{C}(A),\ \widetilde{\mathbb{C}}(\Delta),\ \Pi}.$$

IV. Induction rules of $\overline{\mathbf{K}}_n(\mathbb{C})$. For all finite formula sets Δ, Π of $\mathcal{L}_n(\mathbb{C})$ formulas and all $\mathcal{L}_n(\mathbb{C})$ formulas A, B:

$$\text{(Ind)} \qquad \frac{B,\ \mathbb{E}[A],\ \widetilde{\mathbb{C}}(\Delta) \qquad\qquad B,\ \mathbb{E}[\neg B],\ \widetilde{\mathbb{C}}(\Delta)}{B,\ \mathbb{C}(A),\ \widetilde{\mathbb{C}}(\Delta),\ \Pi}.$$

The axioms and rules of our Tait-style reformulation of $\mathbf{K}_n(\mathbb{C})$ do not comprise cuts. We will mention them explicitly in order to emphasize which cuts are being used. Let Ω be a collection of $\mathcal{L}_n(\mathbb{C})$ formulas closed under negations. Then the Ω-cuts are all cuts whose cut formulas belong to Ω.

V. Ω-cuts. For all finite formula sets Γ of $\mathcal{L}_n(\mathbb{C})$ formulas and all formulas $A \in \Omega$ (the designated formulas A and $\neg A$ are the cut formulas of this cut):

$$(\Omega\text{-cut}) \qquad \frac{\Gamma,\ A \qquad\qquad \Gamma,\ \neg A}{\Gamma}.$$

Derivability of a finite set Γ of $\mathcal{L}_n(\mathbb{C})$ formulas within $\overline{\mathbf{K}}_n(\mathbb{C})$ with possible additional Ω-cuts is defined as usual and written as $\overline{\mathbf{K}}_n(\mathbb{C}) + (\Omega\text{-cut}) \vdash \Gamma$.

It is relatively easy to show that, if arbitrary cuts are permitted, this Tait-style system proves the same formulas as $\mathbf{K}_n(\mathbb{C})$. Some care is only needed to check that the co-closure axioms for $\mathbb{C}$ are provable in $\overline{\mathbf{K}}_n(\mathbb{C})$ and that $\overline{\mathbf{K}}_n(\mathbb{C}) + (\mathcal{L}_n(\mathbb{C})\text{-cut})$ is closed under ($\mathbb{C}$-IND); see Alberucci and Jäger [1] for all details.

Theorem 25 *For all finite sets Γ of $\mathcal{L}_n(\mathbb{C})$ formulas we have that*

$$\overline{\mathbf{K}}_n(\mathbb{C}) + (\mathcal{L}_n(\mathbb{C})\text{-cut}) \vdash \Gamma \quad\Longleftrightarrow\quad \mathbf{K}_n(\mathbb{C}) \vdash \Gamma^{\vee}.$$

The rule $(\mathcal{L}_n(\mathbb{C})\text{-cut})$ is the stumbling block to a decent proof-theoretic analysis of common knowledge within $\overline{\mathbf{K}}_n(\mathbb{C}) + (\mathcal{L}_n(\mathbb{C})\text{-cut})$. Moreover, on the basis of $\overline{\mathbf{K}}_n(\mathbb{C})$ cuts cannot be avoided completely. To see why, pick two different atomic propositions P and Q and consider the formula A defined by

$$A \ :=\ \langle 1\rangle(\widetilde{P} \vee \widetilde{\mathbb{C}}(\widetilde{Q})) \ \vee\ \langle 2\rangle(\widetilde{Q} \vee \widetilde{\mathbb{C}}(\widetilde{P})) \ \vee\ \mathbb{C}(P \vee Q).$$

Then it is easily checked that $(2, \mathbb{C}) \models A$, implying $\overline{\mathbf{K}}_2(\mathbb{C}) + (\mathcal{L}_2(\mathbb{C})\text{-cut}) \vdash A$ because of Theorem 24 and Theorem 25. On the other hand, it is also not difficult to show that A cannot be derived in $\overline{\mathbf{K}}_2(\mathbb{C})$.

What we can achieve, however, is a formalism, in which all necessary cuts can be controlled by means of the Γ which is to be derived. To do so we first introduce the so-called Fischer-Ladner closure $\mathrm{FL}(A)$ of an $\mathcal{L}_n(\mathbb{C})$ formula A.

Definition 26 *The* Fischer-Ladner closure $\mathrm{FL}(A)$ *of an* $\mathcal{L}_n(\mathbb{C})$ *formula* A *is the set of* $\mathcal{L}_n(\mathbb{C})$ *formulas which is inductively defined as follows:*

(FL1) A *belongs to* $\mathrm{FL}(A)$.

(FL2) If B *belongs to* $\mathrm{FL}(A)$*, then* $\neg B$ *belongs to* $\mathrm{FL}(A)$.

(FL3) If $(B \vee C)$ *belongs to* $\mathrm{FL}(A)$*, then* B *and* C *belong to* $\mathrm{FL}(A)$.

(FL4) If $\langle\alpha\rangle B$ *belongs to* $\mathrm{FL}(A)$*, then* B *belongs to* $\mathrm{FL}(A)$.

(FL5) If $\mathbb{C}(B)$ *belongs to* $\mathrm{FL}(A)$*, then* B*,* $\mathbb{E}[B]$*, and* $\mathbb{E}[\mathbb{C}(B)]$ *belong to* $\mathrm{FL}(A)$.

The Fischer-Ladner closure of any $\mathcal{L}_n(\mathbb{C})$ formula is finite and, according to Fischer and Ladner [11], the number of elements of $\mathrm{FL}(A)$ is of order $\mathcal{O}(|A|)$, where $|A|$ denotes the length of the formula A.

For a finite set Γ of $\mathcal{L}_n(\mathbb{C})$ formulas we set $\mathrm{FL}(\Gamma) := \mathrm{FL}(\Gamma^\vee)$. Furthermore, $\mathrm{DC}_1(\Gamma)$ is defined to be the closure of $\mathrm{FL}(\Gamma)$ under conjunctions (without repetitions) and $\mathrm{DC}_2(\Gamma)$ the closure of $\mathrm{DC}_1(\Gamma)$ under disjunctions (without repetitions). Then the *disjunctive-conjunctive closure* of Γ is given by

$$\mathrm{DC}(\Gamma) := \mathrm{DC}_2(\Gamma) \cup \{\neg A : A \in \mathrm{DC}_2(\Gamma)\}.$$

In Alberucci and Jäger [1] we showed that cuts with cut formulas from $\mathrm{DC}(\Gamma)$ are sufficient in order to derive a valid finite set Γ of $\mathcal{L}_n(\mathbb{C})$ formulas. The proof is by constructing a canonical n-knowledge structure whose worlds are the maximal $\mathrm{DC}(\Gamma)$-consistent sets.

Theorem 27 *For all finite sets* Γ *of* $\mathcal{L}_n(\mathbb{C})$ *formulas we have that*

$$\overline{\mathbf{K}}_n(\mathbb{C}) + (\mathrm{DC}(\Gamma)\text{-cut}) \vdash \Gamma \quad \Longleftrightarrow \quad (n, \mathbb{C}) \models \Gamma^\vee.$$

This theorem says that for a proof of a valid formula A only cuts are needed which belong to the bounded set $\mathrm{DC}(\{A\})$ and thus permits a control of the cuts. From the point of view of computational complexity and proof search, the size of $\mathrm{DC}(\{A\})$ is still infeasible. We know that the restriction to cuts from $\mathrm{DC}(\{A\})$ is far from being optimal, but it is an interesting open question how far we can go.

4. Evidence and knowledge

We end this overview by presenting some connections between Artemov's so-called *logic of proofs* and the previously considered epistemic systems. Good comprehensive introductions into the logic of proofs or justification logic (as it is often called recently) are presented in Artemov [2] and Artemov and Beklemishev [4].

One of the basic ideas is to extend the framework of multi-modal logic with n agents by a system of terms for representing evidence and expressions of the form "$(a : A)$" expressing the idea that "a provides evidence for A". In the original logic of proofs these terms acted as explicit representations of proofs, but their interpretation as evidence witnesses makes sense as well. What we are going to sketch now is a first attempt to combine knowledge and evidence; it partly follows Artemov [3].

Evidence terms $a, b, c, \ldots$ (possibly with subscripts) are built from *evidence constants* $u, v, w, \ldots$ and *evidence variables* $x, y, z, \ldots$ (all possibly with subscripts) by the following grammar:

$$a \ ::= \ u \mid \ x \mid (a \cdot a) \mid (a + a) \mid !a,$$

where $\cdot$ (application) and $+$ (union) are binary operations on terms while ! (inspection) is a unary operation on terms.

The language $\mathcal{L}_n^e$ for n agents and evidence is similar to the language $\mathcal{L}_n(\mathbb{C})$, but instead of the formulas $\mathbb{C}(A)$ and $\widetilde{\mathbb{C}}(A)$ we have $(a \ : \ A)$ and $\widetilde{(a : A)}$, respectively. Accordingly, the formulas $A, B, C, \ldots$ (possibly with subscripts) of $\mathcal{L}_n^e$ are defined by the following grammar:

$$A \ ::= \ P \mid \ \widetilde{P} \mid X \mid \widetilde{X} \mid \bot \mid \top \mid (A \vee A) \mid (A \wedge A) \mid \langle\alpha\rangle A \mid [\alpha]A \mid (a : A) \mid \widetilde{(a : A)},$$

where P and X range over atomic propositions and variables, respectively, and α is a natural number, $1 \leq \alpha \leq n$. The negation $\neg A$ of an $\mathcal{L}_n^e$ formula A is defined as before with the clauses

$$\neg(a : A) \ := \ \widetilde{(a : A)} \qquad \text{and} \qquad \neg\widetilde{(a : A)} \ := \ (a : A).$$

The system $\mathbf{T}_n^e$ provides a Hilbert-style formalization of knowledge with evidence. Its knowledge axioms are as in $\mathbf{K}_n(\mathbb{C})$ with the additional claim that knowledge implies truth. Then there are specific axioms and rules for evidence and a principle connecting evidence and knowledge.

I. Logical axioms of $\mathbf{T}_n^e$. All propositional tautologies, the distribution axioms, and the truth axioms, i.e., for all propositional tautologies A of $\mathcal{L}_n^e$, all $\mathcal{L}_n^e$ formulas B and C and all agents α $(1 \leq \alpha \leq n)$:

(TAU) $$A,$$

(DIS) $$[\alpha](B \to C) \ \to \ ([\alpha]B \to [\alpha]C),$$

(T) $$[\alpha]B \ \to \ B.$$

II. Logical rules of $\mathbf{T}_n^e$. Modus ponens and necessitation, i.e., for all $\mathcal{L}_n^e$ formulas A and B, and all agents α $(1 \leq \alpha \leq n)$:

(MP) $$\frac{A \qquad A \to B}{B},$$

(NEC) $$\frac{A}{[\alpha]A}.$$

III. Evidence axioms of $\mathbf{T}_n^e$. For all $\mathcal{L}_n^e$ formulas A, B and all evidence terms a, b:

(Application) $$a : A \ \wedge \ b : (A \to B) \ \to \ (b \cdot a) : B,$$

(Union) $\quad a : A \ \rightarrow \ (a + b) : A \qquad$ and $\qquad b : A \ \rightarrow \ (a + b) : A,$

(Inspection) $\quad a : A \ \rightarrow !a : (a : A),$

(Reflexivity) $\quad a : A \ \rightarrow \ A.$

IV. Evidence-to-knowledge axioms of $\mathbf{T}_n^e$. For all $\mathcal{L}_n^e$ formulas A, all evidence terms a, and all agents α $(1 \leq \alpha \leq n)$:

(EK) $\quad a : A \ \rightarrow \ [\alpha]A.$

V. Constant specifications of $\mathbf{T}_n^e$. For all axioms A of group I, group III, and group IV and all evidence constants u:

(CS) $\quad u : A.$

As always previously, provability of a formula A in the Hilbert system $\mathbf{T}_n^e$ is written as $\mathbf{T}_n^e \vdash A$.

The theory $\mathbf{T}_n^e$ shares many features of the logic of proofs. In particular, it has the internalization property which states that every derivation in $\mathbf{T}_n^e$ is witnessed by an evidence term. For its proof see Artemov [3].

Theorem 28 (Internalization) *If $\mathbf{T}_n^e \vdash A$ for some $\mathcal{L}_n^e$ formula A, then there exists an evidence term a such that $\mathbf{T}_n^e \vdash a : A$.*

Interesting in our context and establishing a connection to common knowledge is the observation that all formulas $a : A$ are fixed points of $\mathbb{E}[A \wedge U]$.

Theorem 29 *For all $\mathcal{L}_n^e$ formulas A and all evidence terms a we have that*

$$\mathbf{T}_n^e \ \vdash \ a : A \ \leftrightarrow \ \mathbb{E}[A \wedge a : A].$$

PROOF. From $\mathbb{E}[A \wedge a : A]$ we deduce $[1](A \wedge a : A)$, and from that $a : A$ follows in view of (T). This settles the right-to-left part of our theorem. For the converse direction, use (EK) to derive $a : A \ \rightarrow \ [\alpha]A$ for any agent α, hence

(1) $\quad a : A \ \rightarrow \ \mathbb{E}[A].$

By (Inspection) we also have $a : A \ \rightarrow !a : (a : A)$. The axioms (EK) then yield $a : A \ \rightarrow \ [\alpha](a : A)$ for all agents α, hence

(2) $\quad a : A \ \rightarrow \ \mathbb{E}[a : A].$

The assertions (1) and (2) and simple reasoning in modal logic conclude the proof of the direction from left to right. □

This theorem does not say, however, that the formulas $(a : A)$ are greatest fixed points of $\mathbb{E}[A \wedge U]$. If we add the operators $\mathbb{C}$ and $\widetilde{\mathbb{C}}$ plus the respective axioms, then

$a : A \rightarrow \mathbb{C}(A)$ becomes provable. Hence evidence of A is stronger than common knowledge of A.

Even though this interplay between evidence, knowledge and common knowledge sheds light on an interesting area of epistemic logic and brings in some new and interesting parameters, it should be far from being the final answer. There are still several shortcomings of this approach which deserve substantial further research. We conclude this article with mentioning some of them.

Open problems 30

1. It seems unnatural that there are only global evidence assertions $(a : A)$. More flexibility is gained by adding evidence with respect to agents, $(a :_\alpha A)$. Some first interesting steps in this direction are due to Yavorskaya [22], but more – most notably proof-theoretic – research about such systems is needed.
2. An evidence-based version of common knowledge in the proper sense (greatest fixed point of $\mathbb{E}[A \wedge U]$) does not exist yet. There is promising work by S. Bucheli, and it seems that only minor technicalities are left to be straightened out.
3. Artemov [2] presents a sequent-style reformulation of his logic of proofs $\mathbf{LP}$, and it should be easy to obtain the same for $\mathbf{T}_n^e$. But these systems are only free of *external* cuts; internal cuts in form of the application axiom cannot be eliminated. Is there a system of evidence terms, axioms and rules equivalent to $\mathbf{T}_n^e$ which permits the elimination of internal and external cuts? It may well be that we have to introduce a form of reduction of evidence terms.

Acknowledgement

I wish to thank Roman Kuznets for his careful reading of an earlier version of this paper and for his many constructive and helpful remarks.

References

[1] L. Alberucci and G. Jäger, *About cut elimination for logics of common knowledge*, Annals for Pure and Applied Logic **133** (2005), 73–99.
[2] S. Artemov, *Explicit provability and constructive semantics*, Bulletin of Symbolic Logic **7** (2001), 1–36.
[3] ———, *Justified common knowledge*, Theoretical Computer Science **357** (2006), 4–22.
[4] S. Artemov and L. Beklemishev, *Provability Logic*, Handbook of Philosophical Logic, 2nd ed. (D. Gabbay and F. Guenthner, eds.), vol. 13, Springer, 2005, pp. 189–360.
[5] J. Barwise, *Admissible Sets and Structures*, Perspectives in Mathematical Logic, Springer, 1975.
[6] J. Bradfield and C. Stirling, *Modal Logics and mu-Calculi: An Introduction*, Handbook of Process Algebra (J. Bergstra, A. Ponse, and S. Smolka, eds.), Elsevier, 2001, pp. 293–330.
[7] ———, *Modal mu-Calculi*, Handbook of Modal Logic (P. Blackburn, J. van Benthem, and F. Wolter, eds.), Studies in Logic and Practical Reasoning, vol. 3, Elsevier, 2007, pp. 721–756.
[8] K. Brünnler and T. Studer, *Syntactic cut-elimination for common knowledge*, Annals for Pure and Applied Logic **160** (2009), 82–95.
[9] W. Buchholz, S. Feferman, W. Pohlers, and W. Sieg, *Iterated Inductive Definitions and Subsystems of Analysis: Recent Proof-Theoretical Studies*, Lecture Notes in Mathematics, vol. 897, Springer, 1981.
[10] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi, *Reasoning about Knowledge*, MIT Press, 1995.
[11] M.J. Fischer and R.E. Ladner, *Propositional dynamic logic of regular programs*, Journal of Computer and System Sciences **18** (1979), 194–211.

[12] J.Y. Halpern and Y. Moses, *A guide to completeness and complexity for modal logics of knowledge and belief*, Artificial Intelligence **54** (1992), 319–379.

[13] G. Jäger, M. Kretz, and T. Studer, *Cut-free common knowledge*, Journal of Applied Logic **5** (2007), 681–689.

[14] ———, *Canonical completeness of infinitary μ*, Journal of Logic and Algebraic Programming **76** (2008), 270–292.

[15] D. Kozen, *Results on the propositional μ-calculus*, Theoretical Computer Science **27** (1983), 333–354.

[16] Y.N. Moschovakis, *Elementary Induction on Abstract Structures*, Studies in Logic and the Foundations of Mathematics, North-Holland, 1974.

[17] D. Niwiński and I. Walukiewicz, *Games for the μ-calculus*, Theoretical Computer Science **163** (1996), 99–116.

[18] C. Stirling and D. Walker, *Local model checking in the modal mu-calculus*, Theoretical Computer Science **89** (1991), 161–177.

[19] R.S. Streett and E.A. Emerson, *An automata theoretic decision procedure for the propositional mu-calculus*, Information and Computation **81** (1989), 249–264.

[20] I. Walukiewicz, *Completeness of Kozen's axiomatisation of the propositional μ-calculus*, Information and Computation **157** (2000), 142–182.

[21] T. Wilke, *Alternating tree automata, parity games, and modal μ-calculus*, Bulletin of the Belgian Mathematical Society – Simon Stevin **8** (2001), 359–391.

[22] T. Yavorskaya, *Interacting explicit evidence systems*, Theory of Computing Systems **43** (2008), 272–293.