Information and Computation **157**, 142–182 (2000)
doi:10.1006/inco.1999.2836, available online at http://www.idealibrary.com on IDEAL®

Completeness of Kozen's Axiomatisation of the Propositional μ-Calculus[1]

Igor Walukiewicz[2]

Basic Research in Computer Science, Centre of the Danish National Research Foundation and Department of Computer Science, University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, Denmark
E-mail: igw@mimuw.edu.pl

Propositional μ-calculus is an extension of the propositional modal logic with the least fixpoint operator. In the paper introducing the logic Kozen posed a question about completeness of the axiomatisation which is a small extension of the axiomatisation of the model system K. It is shown that this axiomatisation is complete. © 2000 Academic Press

1. INTRODUCTION AND SUMMARY

We consider μ-calculus as defined by Kozen [4]. This is the logic obtained from modal logic by adding the least fixpoint operator: $\mu X.\alpha(X)$. The intended models of the logic are Kripke structures. Kozen's axiomatisation consists of the axiomatisation of the modal system K together with one axiom and one rule characterising the least fixpoint operator:

$$\alpha(\mu X.\alpha(X)) \Rightarrow \mu X.\alpha(X) \qquad \frac{\alpha(\varphi) \Rightarrow \varphi}{\mu X.\alpha \Rightarrow \varphi}$$

The completeness theorem considered here is sometimes called weak completeness because it deals with validity relation; it says that every valid formula is provable. Strong completeness refers to an axiomatisation of the semantic consequence relation. It is not possible to have finitary strongly complete axiomatisation for the μ-calculus because the compactness theorem fails for the logic. In the following completeness means weak completeness and provability means provability in Kozen's system unless explicitly stated otherwise.

In [4] Kozen showed that the axiom system proves negations of all unsatisfiable formulas of a special kind called *aconjunctive formulas*. In [10] another finitary axiomatisation was proposed and proved to be complete for the whole μ-calculus.

[1] This work was partially supported by Polish KBN Grant 2 P301 009 06.

[2] On leave from the Institute of Informatics, Warsaw University, Banacha 2, 02-097 Warsaw, Poland.

0890-5401/00 $35.00
Copyright © 2000 by Academic Press
All rights of reproduction in any form reserved.

This solved one part of the problem posed in [4] but the question of the completeness of the original axiomatisation remained still open. We give an affirmative answer to this question.

There are other reasons, apart from curiosity, to investigate the problem of the completeness of Kozen's system. The axiomatisation proposed in [10] makes essential use of the small model theorem for the μ-calculus; this makes it impossible to use it for extensions of the logic not enjoying the finite model theorem. The other reason is that Kozen's system is very natural, one may say as natural as the notion of Kripke structures. Hence, it is good to know that the class of Kripke structures is a complete subclass of a quasi-variety defined by Kozen's system.

Let us review some methods used in previous approaches to the completeness problem. The first step is a tableau method of model construction of Streett and Emerson [9]. For a given formula one constructs a *tableau*; if the formula is satisfiable then one can construct a model from a part of this tableau. It was shown in [7] that if the initial formula is unsatisfiable, and one cannot find a model in the tableau, then one can construct for the formula another tableau-like structure called a *refutation*. In [10] a stronger axiomatisation was proposed and it was shown that:

> If there is a refutation for φ then $\neg\varphi$ is provable in the stronger system. (a)

This proof does not work for Kozen's axiomatisation and it does not look like any simple modification of the argument can help here.

It is also possible to look at Kozen's proof for aconjunctive formulas from the point of view of refutations. One can introduce a notion of *thin refutation*, which is a refutation where reductions of conjunctions are restricted. A slight extension of Kozen's arguments gives us:

> If there is a thin refutation for φ then $\neg\varphi$ is provable. (b)

Thin refutations suggest the notion of *weakly aconjunctive formulas*. These formulas have the property that every refutation for such a formula is thin. As the name indicates, all aconjunctive formulas are weakly aconjunctive. Below we will use both fact (b) and the notion of weakly aconjunctive formulas.

Let us now give an outline of the proof presented here. As we noted above it seems very hard to directly improve the statement (b) by trying to enlarge the class of refutations for which it holds. On the other hand, by fact (b) in order to show completeness it is enough to prove:

> For every formula φ there is a semantically equivalent aconjunctive formula $\hat{\varphi}$ such that $\varphi \Rightarrow \hat{\varphi}$ is provable. (c)

This cannot work because it is not true that every formula is equivalent to an aconjunctive formula. This obstacle can be avoided if we allow weakly aconjunctive formulas but still these formulas are not particularly easy to work with. It would certainly save us some work if we first tried to find some class of formulas with better properties.

We will prove a statement like (c) but instead of aconjunctive formulas we will use *disjunctive formulas* [2]. These formulas have several useful properties. First, tableaux for disjunctive formulas have very simple structure. Next, the proof of the fact that negation of every unsatisfiable disjunctive formula is provable is much easier than for weakly aconjunctive formulas (see Theorem 24). The third important property is that for every formula there is a semantically equivalent disjunctive formula. This last statement can be even strengthened as we will describe in the next paragraph.

The properties stated above suggest that instead of proving (c) we should try to prove:

> For every formula φ there is a semantically equivalent disjunctive formula $\hat{\varphi}$ such that $\varphi \Rightarrow \hat{\varphi}$ is provable. (d)

The tool we will use to construct a proof of $\varphi \Rightarrow \hat{\varphi}$ is *tableau equivalence*. As we have mentioned above, models for a formula can be constructed from a tableau for the formula. We will say that two tableaux are equivalent if they are essentially the same from the perspective of the model construction procedure. This induces equivalence on formulas which is stronger than semantical equivalence because there exist semantically equivalent formulas which do not have equivalent tableaux. Now it was shown in [2] that for every formula there is a disjunctive formula with an equivalent tableau. The use of tableau equivalence is important because it allows us to replace semantical equivalence with an equivalence which is much finer and syntactically defined.

Another important observation is that we can prove (d) in case φ is a weakly aconjunctive formula. This follows from:

> If α is a weakly aconjunctive formula, δ is a disjunctive formula, and the two formulas have equivalent tableaux then $\alpha \Rightarrow \delta$ is provable. (e)

Observe that already with this statement we increase the class of formulas which are known to be provable. We now know that some formulas of the form $\neg(\alpha \wedge \neg\delta)$ are provable, where $\neg\delta$ may not be a weakly aconjunctive formula.

Let us try to use (e) to prove (d) by induction on the structure of φ. This way we will see what we can do and where the problems are.

Suppose $\varphi = \nu X.\alpha(X)$. By induction assumption we have a disjunctive formula $\hat{\alpha}(X)$ and a proof of $\alpha(X) \Rightarrow \hat{\alpha}(X)$. Hence, $\nu X.\alpha(X) \Rightarrow \nu X.\hat{\alpha}(X)$ is provable. Because $\hat{\alpha}(X)$ is a disjunctive formula, $\nu X.\hat{\alpha}(X)$ is a weakly aconjunctive formula although it may not be a disjunctive formula. Let $\hat{\varphi}$ be a disjunctive formula with a tableau equivalent to a tableau for $\nu X.\hat{\alpha}(X)$. By (e) we have a proof of $\nu X.\hat{\alpha}(X) \Rightarrow \varphi$. So $\varphi \Rightarrow \hat{\varphi}$ is provable.

The problems come only in one case when $\varphi = \mu X.\alpha(X)$. This is because $\mu X.\hat{\alpha}(X)$ may not be a weakly aconjunctive formula. Fortunately, by the fixpoint rule, to prove $\mu X.\hat{\alpha}(X) \Rightarrow \hat{\varphi}$ it is enough to prove $\hat{\alpha}(\hat{\varphi}) \Rightarrow \hat{\varphi}$. Formula $\hat{\alpha}(\hat{\varphi})$ is weakly aconjunctive but this time we meet another problem. There may be no tableau for $\hat{\alpha}(\hat{\varphi})$ which is equivalent to a tableau for $\hat{\varphi}$. This should not come as a big surprise as the notion of tableau equivalence is very restrictive; it would be rather surprising if it worked all the way. We remedy this by introducing a weaker relation between tableaux which we call *tableau consequence*. It turns out to be a relation which refines semantical consequence.

We prove that there is a tableau for $\hat{\alpha}(\hat{\varphi})$ of which a tableau for φ is the consequence. On the other hand the notion of tableau consequence is still strong enough to show a statement similar to (e):

> If α is a weakly aconjunctive formula, δ is a disjunctive formula, and a tableau for δ is a consequence of a tableau for α then $\alpha \Rightarrow \delta$ is provable. (f)

This way we obtain a proof of $\hat{\alpha}(\hat{\varphi}) \Rightarrow \hat{\varphi}$ and hence also a proof of $\varphi \Rightarrow \hat{\varphi}$.

The plan of the paper is as follows. We start by defining the μ-calculus and some auxiliary notions like positive guarded formulas, binding function, or $(a \to \Psi)$ construct. In the next section we recall the results from [2] which we will need here. The notions of tableau equivalence and disjunctive formula are introduced there. Next, we present Kozen's axiomatisation of the logic and show some simple properties of it. The following section deals with weakly aconjunctive formulas. The next section is devoted to the properties of the tableau consequence relation. The last section gives the inductive proof of (d).

2. PRELIMINARY DEFINITIONS

Let $Prop = \{p, q, ...\}$ be a set of propositional letters, $Var = \{X, Y, ...\}$ be a set of variables, and $Act = \{a, b, ...\}$ be a set of actions. Formulas of the μ-calculus over these three sets are defined by the following grammar:

$$F := \top \mid \bot \mid Var \mid Prop \mid \neg F \mid F \vee F \mid F \wedge F \mid$$

$$\langle Act \rangle\, F \mid [Act]\, F \mid \mu\, Var.F \mid \nu\, Var.F.$$

Additionally we require that in formulas of the form $\mu X.\alpha(X)$ and $\nu X.\alpha(X)$, variable X occurs in $\alpha(X)$ only *positively*, i.e., under an even number of negations.

We will use σ to denote μ or ν. Formulas will be denoted by lowercase Greek letters. Uppercase Greek letters will denote finite sets of formulas. We write $\alpha \Rightarrow \beta$ for $\neg\alpha \vee \beta$. For a finite set of formulas Γ we denote by $\bigwedge \Gamma$ the conjunction of formulas in Γ. Similarly $\bigvee \Gamma$ denotes the disjunction of formulas in Γ. As usual the conjunction of the empty set is true and the disjunction of the empty set is false. Propositional constants, variables, and their negations will be called *literals*.

Formulas are interpreted in Kripke structures $\mathcal{M} = \langle S, R, \rho \rangle$, where S is a nonempty set of states, R: $Act \to \mathcal{P}(S \times S)$ is a function assigning a binary relation on S to each action in Act, and ρ: $Prop \to \mathcal{P}(S)$ is a function assigning a set of states to each propositional letter in $Prop$.

The meaning of a formula in a model is a set of states where it is true. For a given model $\mathcal{M}$ and a valuation V: $Var \to \mathcal{P}(S)$, the meaning of a formula α, denoted $\|\alpha\|_V^{\mathcal{M}}$, is defined by induction on the structure of α by the following clauses (we will omit superscript $\mathcal{M}$ when it causes no ambiguity):

$$\|\top\|_V = S \quad \|\bot\|_V = \emptyset$$

$$\|X\|_V = V(X)$$

$$\|p\|_V = \rho(p)$$

$$\|\neg\alpha\|_V = S - \|\alpha\|_V$$

$$\|\alpha \wedge \beta\|_V = \|\alpha\|_V \cap \|\beta\|_V$$

$$\|\alpha \vee \beta\|_V = \|\alpha\|_V \cup \|\beta\|_V$$

$$\|\langle a \rangle\, \alpha\|_V = \{s\colon \exists t.\, (s, t) \in R(a) \wedge t \in \|\alpha\|_V\}$$

$$\|[a]\, \alpha\|_V = \{s\colon \forall t. (s, t) \in R(a) \Rightarrow t \in \|\alpha\|_V\}$$

$$\|\mu X.\alpha(x)\|_V = \bigcap \{T \subseteq S : \|\alpha\|_{V[T/X]} \subseteq T\}$$

$$\|\nu X.\alpha(X)\|_V = \bigcup \{T \subseteq S : T \subseteq \|\alpha\|_{V[T/X]}\}.$$

Sometimes we will write $\mathcal{M}, s, V \models \alpha$ instead of $s \in \|\alpha\|_V^{\mathcal{M}}$.

Definition 1 (Positive, guarded formulas). We call a formula *positive* iff all negations in the formula appear only before propositional constants and free variables.

The variable X in $\mu X.\alpha(X)$ is called *guarded* iff every occurrence of X in α is in the scope of some modality operator: $\langle a \rangle$ or $[a]$. We say that a *formula is guarded* iff every bound variable in the formula is guarded.

Proposition 2 (Kozen). *Every formula is equivalent to a positive guarded formula.*

Proof. Let φ be a formula. We first show how to obtain an equivalent guarded formula. The proof proceeds by induction on the structure of the formula with the only nontrivial cases being fixpoint constructors. We present here the case for the least fixpoint. The case for the greatest fixpoint is similar.

Assume that $\varphi = \mu X.\alpha(X)$ and $\alpha(X)$ is a guarded formula. Suppose X is unguarded in some subformula of $\alpha(X)$ of the form $\sigma Y.\beta(Y, X)$. By the assumption, the variable Y is guarded in $\sigma Y.\beta(Y, X)$. We can use the equivalence $\sigma Y.\beta(Y, X) = \beta(\sigma Y.\beta(Y, X), X)$ to obtain a formula with all unguarded occurrences of X outside the fixpoint operator. This way we obtain a formula equivalent to $\alpha(X)$ with all unguarded occurrences of X not in the scope of a fixpoint operator.

Now using the laws of classical propositional logic we can transform this formula to a conjunctive normal form (considering fixpoint formulas and formulas of the form $\langle a\rangle\,\gamma$ and $[a]\,\gamma$ as propositional constants). This way we obtain a formula

$$(X \vee \alpha_1(X)) \wedge \cdots \wedge (X \vee \alpha_i(X)) \wedge \beta(X), \tag{1}$$

where all occurrences of X in $\alpha_1(X)$, ..., $a_i(X)$, $\beta(X)$ are guarded. Observe that some of $\alpha_j(X)$ may be just $\perp$ and $\beta(X)$ may be $\top$. The variable X occurs only positively in (1) because it did so in our original formula. Formula (1) is equivalent to

$$(X \vee (\alpha_1(X) \wedge \cdots \wedge \alpha_i(X))) \wedge \beta(X).$$

We will show that $\mu X.(X \vee \bar{\alpha}(X)) \wedge \beta(X)$ is equivalent to $\mu X.\bar{\alpha}(X) \wedge \beta(X)$. It is obvious that

$$(\mu X.\bar{\alpha}(X) \wedge \beta(X)) \Rightarrow (\mu X.(X \wedge \bar{\alpha}(X)) \wedge \beta(X)).$$

Let γ stand for $\mu X.\bar{\alpha}(X) \wedge \beta(X)$. To prove another implication it is enough to observe that γ is a pre-fixpoint of $\mu X.(X \vee \bar{\alpha}(X)) \vee \beta(X)$ as the following calculation shows:

$$\begin{aligned}
(\gamma \vee \bar{\alpha}(\gamma)) \wedge \beta(\gamma) &\Rightarrow \\
((\bar{\alpha}(\gamma) \wedge \beta(\gamma)) \vee \bar{\alpha}(\gamma)) \wedge \beta(\gamma) &\Rightarrow \\
\bar{\alpha}(\gamma) \wedge \beta(\gamma) &\Rightarrow \gamma.
\end{aligned}$$

If φ is a guarded formula then we use dualities of the μ-calculus,

$$\begin{aligned}
\neg(\alpha \vee \beta) &= \neg\alpha \wedge \neg\beta & \qquad \neg(\alpha \wedge \beta) &= \neg\alpha \vee \neg\beta \\
\neg\langle a\rangle\,\alpha &= [a]\neg\alpha & \qquad \neg[a]\,\alpha &= \langle a\rangle\neg\alpha \\
\neg\mu X.\alpha(X) &= \nu X\neg\alpha(\neg X) & \qquad \neg\nu X.\alpha(X) &= \mu X.\neg\alpha(\neg X),
\end{aligned}$$

to produce an equivalent positive formula. It is easy to see that it will be still a guarded formula. ∎

Next we introduce some tools which allow us to deal with occurrences of subformulas of a given formula. These tools are very similar to those used in [4] or [8]. We would like to have a different name (which will be a variable) for every fixpoint subformula of a given formula. We will also introduce a notion of a binding function which will associate subformulas to names.

Definition 3 (Binding). We call a formula *well-named* if every variable is bound at most once in the formula and free variables are distinct from bound variables. For a variable X bound in a well-named formula α there exists the unique subterm of α of the form $\sigma X.\beta(X)$, from now on called the *binding definition of X* *in* α and denoted $\mathscr{D}_\alpha(X)$. We will omit subscript α when it causes no ambiguity. We call X a *ν-variable* if $\sigma = \nu$, otherwise we call X a *μ-variable*.

The function $\mathscr{D}_\alpha$ assigning to every bound variable its *binding definition* in α will be called the *binding function* associated with α.

Remark. Every formula is equivalent to a well-named one which can be obtained by some consistent renaming of bound variables. The substitution of a formula β for all free occurrences of a variable X in α, denoted $\alpha[\beta/X]$, can be made modulo some consistent renaming of bound variables of β, so that the formula $\alpha[\beta/X]$ so obtained is still well-named.

DEFINITION 4 (Dependency order). Given a formula α, we define the *dependency order* $\leqslant_\alpha$ on the bound variables of α as the least partial order relation such that if X occurs free in $\mathscr{D}_\alpha(Y)$ then $X \leqslant_\alpha Y$. We will say that a bound variable Y *depends* on a bound variable X in α when $X \leqslant_\alpha Y$.

EXAMPLE. In case $\alpha = \mu X.\langle a \rangle X \vee \nu Y.\langle b \rangle Y$, variables X and Y are incomparable in the $\leqslant_\alpha$ ordering. On the other hand, if α is $\mu X.\nu Y.\langle a \rangle X \vee \mu Z.\langle a \rangle (Z \vee Y)$ then $X \leqslant_\alpha Z$.

DEFINITION 5 (Expansion). Let α be a formula with an associated binding function $\mathscr{D}_\alpha$. For every subformula β of α we define the *expansion* of β with respect to $\mathscr{D}_\alpha$ as

$$\langle\!\langle \beta \rangle\!\rangle_{\mathscr{D}_\alpha} = \beta[\mathscr{D}_\alpha(X_n)/X_n] \cdots [\mathscr{D}_\alpha(X_1)/X_1],$$

where the sequence $(X_1, X_2, ..., X_n)$ is a linear ordering of all bound variables of α compatible with the dependency partial order, i.e., if $X_i \leqslant_\alpha X_j$ then $i \leqslant j$.

DEFINITION 6. We extend the syntax of the μ-calculus by allowing a new construct of the form $(a \to \Psi)$, where a is an action and Ψ is a finite set of formulas. As for its semantics, we will consider such a formula to be an abbreviation of the formula $\bigwedge\{\langle a \rangle \psi : \psi \in \Psi\} \wedge [a] \bigvee \Psi$.

Remark. By itself the $(a \to \Psi)$ construction is nothing but a way to hide some conjunctions. This construct arises when one tries to find a notion of automata corresponding to the μ-calculus that is able to cope with potentially unbounded branching. In our case, we use this construction to provide a more symmetric rule for reducing modalities. It also makes the definition of a special conjunction (Definition 25) more natural. With this one construct it is possible to express both $[a]$ and $\langle a \rangle$ modalities. A formula $[a]\psi$ is equivalent to $(a \to \emptyset) \vee (a \to \{\psi\})$ and a formula $\langle a \rangle \psi$ is equivalent to $(a \to \{\psi, \top\})$. All the notions from this section such as guarded formula and binding function extend to formulas with this new construct.

Definition 7 (Terminal formulas). A formula of the form $(a \to \emptyset)$ will be called a *terminal formula* because its meaning is that there are not states reachable by action a from a given state.

Proviso. If not otherwise stated all formulas are assumed to be well named, positive, and guarded and to contain $(a \to \Psi)$ construct instead of $\langle a \rangle \psi$ and $[a] \psi$ modalities. By observations stated above, this is not a restriction as far as semantics is concerned. As we will mention later every formula is provably equivalent to a formula of this kind.

3. TABLEAU EQUIVALENCE AND DISJUNCTIVE FORMULAS

In this section we recall results from [2] which we are going to use later on. We define the notions of tableau and tableau equivalence. It turns out that if two tableaux are equivalent then the formulas in the roots of the tableaux are semantically equivalent. In spite of the fact that the implication in the other direction does not hold, tableau equivalence turns out to be a very handy tool. Next we define a notion of disjunctive formula. Some of the properties of these formulas are discussed in [2]. Here we will recall only one result: for a given tableau which can be presented as a finite graph one can construct a disjunctive formula with an equivalent tableau.

Definition 8 (Tableau rules). For a formula φ and its binding function $\mathscr{D}_\varphi$ we define the system of tableau rules $\mathscr{S}^\varphi$ parametrised by φ or rather its binding function. The system is presented in Fig. 1 (we use $\{\alpha, \Gamma\}$ as a shorthand for $\{\alpha\} \cup \Gamma$).

Remark. (1) We see applications of the rules as a process of reduction. Given a finite set of formulas Γ that we want to derive, we look for the rule the conclusion of which matches our set. Then we apply the rule and obtain the assumptions of the instance of the rule in which Γ is the conclusion.

(2) There is no rule for reducing formulas of the form $\langle a \rangle \varphi$ or $[a] \varphi$ because we assume that these formulas are replaced by equivalent formulas using the $(a \to \Psi)$ notation.

(3) The rule (*mod*) has as many assumptions as there are formulas in the sets Ψ, for which, $(a \to \Psi) \in \Gamma$. For example,

$$\frac{\{\varphi_1, \varphi_3\} \quad \{\varphi_2, \varphi_3\} \quad \{\varphi_1 \vee \varphi_2, \varphi_3\} \quad \{\psi_1\} \quad \{\psi_2\}}{\{(a \to \{\varphi_1, \varphi_2\}), (a \to \{\varphi_3\}), (b \to \{\psi_1, \psi_2\})\}}$$

is an instance of the rule. We will call a son labelled by an assumption obtained by *reducing* and action a a $\langle a \rangle$-*son*. In our example, if a node n of a tableau is labelled by the conclusion of the rule then its son labelled by $\{\varphi_1, \varphi_3\}$ is a $\langle a \rangle$-son of n and a son labelled by $\{\psi_1\}$ is a $\langle b \rangle$-son of n.

Definition 9 (Tableaux). *Tableau* for a formula φ is a pair $\langle T, L \rangle$, where T is a tree and L is a labelling function such that:

1. The root of T is labelled by $\{\varphi\}$.

2. The sons of any internal node n are created and labelled according to the rules of the system $\mathscr{S}^\varphi$. Additionally, we require that the rule (*mod*) is applied only when no other rule is applicable.

$$(and)\ \frac{\{\alpha,\beta,\Gamma\}}{\{\alpha\wedge\beta,\Gamma\}} \qquad (or)\ \frac{\{\alpha,\Gamma\}\quad\{\beta,\Gamma\}}{\{\alpha\vee\beta,\Gamma\}}$$

$$(\mu)\ \frac{\{\alpha(X),\Gamma\}}{\{\mu X.\alpha(X),\Gamma\}} \qquad (\nu)\ \frac{\{\alpha(X),\Gamma\}}{\{\nu X.\alpha(X),\Gamma\}}$$

$$(reg)\ \frac{\{\alpha(X),\Gamma\}}{\{X,\Gamma\}} \quad \text{whenever } X \text{ is a bound variable of } \varphi \text{ and } \mathcal{D}_\varphi(X)=\sigma X.\alpha(X)$$

$$(mod)\ \frac{\{\psi\}\cup\{\bigvee\theta : (a\to\theta)\in\Gamma,\theta\neq\Psi\}\ \text{ for every } (a\to\Psi)\in\Gamma,\psi\in\Psi}{\Gamma}$$

FIG. 1. The system $\mathcal{S}^\varphi$.

As our tableaux may be infinite we will be interested not only in the form of the leaves but also in the internal structure of tableaux. We are now going to distinguish some nodes of tableaux and define a notion of trace which captures the idea of a history of a regeneration of a formula.

DEFINITION 10 (Modal and choice nodes). Leaves and nodes where reduction of modalities is performed, i.e., the rule (*mod*) is used, will be called *modal nodes*. The root of the tableau and sons of modal nodes will be called *choice nodes*.

If φ is a guarded formula then the sequence of all the choice nodes on the path of a tableau for φ induces a partition of the path into finite intervals beginning in choice nodes and ending in modal nodes. We will say that a modal node m is *near* a choice node n iff they are both in the same interval, i.e., in the tableau there is a path from n to m without an application of the rule (*mod*). Observe that in some cases a choice node may be also a modal node.

DEFINITION 11 (Trace). Given a path $\mathcal{P}$ of a tableau $\mathcal{T}=\langle T,L\rangle$, a *trace* on $\mathcal{P}$ will be a function $\mathcal{T}\!r$ assigning a formula to every node in some initial segment of $\mathcal{P}$ (possibly to all of $\mathcal{P}$), satisfying the following conditions:

- If $\mathcal{T}\!r(m)$ is defined then $\mathcal{T}\!r(m)\in L(m)$.
- Let m be a node with $\mathcal{T}\!r(m)$ defined and let $n\in\mathcal{P}$ be a son of m. If a rule applied in m does not reduce the formula $\mathcal{T}\!r(m)$ then $\mathcal{T}\!r(n)=\mathcal{T}\!r(m)$. If $\mathcal{T}\!r(m)$ is reduced in m then $\mathcal{T}\!r(n)$ is one of the results of the reduction. This should be clear for all the rules except possibly for (*mod*). If m is a modal node and n is labelled by $\{\psi\}\cup\{\bigvee\theta\colon (a\to\theta)\in\Gamma,\ \theta\neq\Psi\}$ for some $(a\to\Psi)\in L(m)$ and $\psi\in\Psi$, then $\mathcal{T}\!r(n)=\psi$ if $\mathcal{T}\!r(m)=(a\to\Psi)$ and $\mathcal{T}\!r(n)=\bigvee\theta$ if $\mathcal{T}\!r(m)=(a\to\theta)$ for some $(a\to\theta)\in\Gamma$, $\theta\neq\Psi$. Traces from all other formulas end in the node m.

DEFINITION 12 (μ-trace). We say that there is a *regeneration* of a variable X on a trace $\mathcal{T}\!r$ on some path of a tableau for φ if for some node m and its son n on the path $\mathcal{T}\!r(m)=X$ and $\mathcal{T}\!r(n)=\alpha(X)$, where $\mathcal{D}_\varphi(X)=\sigma X.\alpha(X)$.

We call a trace a *μ-trace* iff it is an infinite trace (defined for the whole path) on which the smallest variable (with respect to $\leqslant_\varphi$ ordering) regenerated infinitely often is a μ-variable. Similarly, a trace will be called a *ν-trace* iff it is an infinite trace where the smallest variable which regenerates infinitely often is a ν-variable.

Remark. Every infinite trace is either a μ-trace or a ν-trace because all the rules except regenerations decrease the size of formulas and formulas are guarded; hence, every formula is eventually reduced.

3.1. Tableau Equivalence

The main result of this subsection is Theorem 19 which says that tableau equivalence refines semantical equivalence. We start by characterising satisfiability of a formula in a state of a structure by means of markings of a tableau of the formula.

Definition 13 (Marking). For a tableau $\mathscr{T} = \langle T, L \rangle$ we define its *marking* with respect to a structure $\mathscr{M} = \langle S, R, \rho \rangle$ and state s_0 to be a relation $M \subseteq S \times T$ satisfying the following conditions:

1. $(s_0, r) \in M$, where r is the root of T.
2. If some pair (s, m) belongs to M and a rule other than (mod) was applied in m, then for some son n of m, $(s, n) \in M$.
3. If $(s, m) \in M$ and the rule (mod) was applied in m then for every action a for which exists a formula of the form $(a \to \Psi)$ in $L(m)$:

   (a) for every $\langle a \rangle$-son n of m there exists a state t such that $(s, t) \in R(a)$ and $(t, n) \in M$.

   (b) for every state t, such that $(s, t) \in R(a)$, there exists a $\langle a \rangle$-son n of m such that $(t, n) \in M$.

Definition 14 (Consistent marking). Keeping the notation from Definition 13, we say that a marking M of $\mathscr{T}$ is *consistent* with respect to $\mathscr{M}$, s, Val iff it satisfies the following conditions:

local consistency for every modal node m and state t, if $(t, m) \in M$ then $\mathscr{M}, t, Val \models \Delta$, where Δ is the set of all the literals occurring in $L(m)$,

global consistency for every path $\mathscr{P} = n_0, n_1, \ldots$ of $\mathscr{T}$ such that for every $i = 0, 1, \ldots$, there exist s_i with $(s_i, n_i) \in M$ there is no μ-trace on $\mathscr{P}$.

Theorem 15. *A formula φ (satisfying our proviso at the end of Section 2) is satisfied in a structure $\mathscr{M}$, state s, and valuation Val iff there is a tableau $\mathscr{T}$ for φ and a marking M of $\mathscr{T}$ consistent with $\mathscr{M}$, s, Val.*

Proof. First we introduce notions of a signature and ν-signature similar to that considered by Streett and Emerson [9]. These notions come from the characterisation of fixpoint formulas by means of transfinite chains of approximations.

In order to describe these approximations we introduce two new constructs, $\mu^\tau X.\alpha(X)$ and $\nu^\tau X.\alpha(X)$, where τ is an ordinal, with the following semantics:

— $[\![\mu^0 X.\alpha(X)]\!] = \emptyset$, $[\![\nu^0 X.\alpha(X)]\!] = S$,

— $[\![\sigma^{\tau+1} X.\alpha(X)]\!] = \|\alpha(X)\|_{Val[[\![\sigma^\tau X.\alpha(X)]\!]/X]}$ (σ stands for μ or ν),

— $[\![\mu^\tau X.\alpha(X)]\!] = \bigcup_{\tau' < \tau} [\![\mu^{\tau'} X.\alpha(X)]\!]$, for τ limit ordinal,

— $[\![\nu^\tau X.\alpha(X)]\!] = \bigcap_{\tau' < \tau} [\![\nu^{\tau'} X.\alpha(X)]\!]$, for τ limit ordinal.

With these definitions we have:

$$[\![\mu X.\alpha(X)]\!] = \bigcup_{\tau} [\![\mu^{\tau} X.\alpha(X)]\!]$$

$$[\![\nu X.\alpha(X)]\!] = \bigcap_{\tau} [\![\nu^{\tau} X.\alpha(X)]\!].$$

We extend the notion of binding function from Section 2, by allowing values of the form $\sigma^{\tau} X.\alpha(X)$ (as before σ stands for μ or ν). The concept of expansion $\langle\!\langle \alpha \rangle\!\rangle_{\mathscr{D}}$ extends immediately.

DEFINITION 16. Let us take a formula β and a binding function $\mathscr{D}$ defined for all the variables occurring in β. Let $\preccurlyeq$ be a linearisation of the dependency order between the variables in the domain of $\mathscr{D}$. Let $U_1, U_2, ..., U_{d^{\mu}}$ ($V_1, ..., V_{d^{\nu}}$) be all the μ-variables (ν-variables, respectively) from the domain of $\mathscr{D}$ listed according to $\preccurlyeq$ ordering.

If the formula $\langle\!\langle \beta \rangle\!\rangle_{\mathscr{D}}$ is satisfied in a state s of a model $\mathscr{M}$ with a valuation Val then we can define a *signature* of β in s, $Sig(\beta, s)$, as the least, in the lexicographical ordering, sequence of ordinals $(\tau_1, ..., \tau_{d^{\mu}})$ such that $\mathscr{M}, s, Val \models \langle\!\langle \beta \rangle\!\rangle_{\mathscr{D}'}$, where $\mathscr{D}'$ is a binding function constructed from $\mathscr{D}$ by replacing, for each $i = 1, ..., d^{\mu}$, ith μ-variable definition $\mathscr{D}(U_i) = \mu X.\alpha_i(X)$ by $\mathscr{D}'(U_i) = \mu^{\tau_i} X.\alpha_i(X)$.

If the formula $\langle\!\langle \beta \rangle\!\rangle_{\mathscr{D}}$ is *not* satisfied in a state s of a model $\mathscr{M}$ with a valuation Val then we can define a *ν-signature* of β in s, ${}^{\nu}Sig(\beta, s)$, as the least, in the lexicographical ordering, sequence or ordinals $(\tau_1, ..., \tau_{d^{\nu}})$ such that $\mathscr{M}, s, Val \not\models \langle\!\langle \beta \rangle\!\rangle_{\mathscr{D}'}$, where $\mathscr{D}'$ is a definition list constructed from $\mathscr{D}$ by replacing, for each $i = 1, ..., d^{\nu}$, ith ν-variable definition $\mathscr{D}(V_i) = \nu X.\alpha_i(X)$ by $\mathscr{D}'(V_i) = \nu^{\tau_i} X.\alpha_i(X)$.

Remark. Of course, signature of a formula depends not only on a state but also on a valuation and a binding function. This is not taken into the account in our notation. These parameters will be always clear from the context.

It can be shown that signatures behave nicely with respect to formula reduction, namely:

LEMMA 17. *Let s be a state of a model $\mathscr{M}$, let Val be a valuation, let $\mathscr{D}$ be a definition list with some linear ordering $\preccurlyeq$ on its domain as in the definition of signature. For all formulas α, β, $\mu X.\alpha(X)$, $\nu X.\alpha(X)$ such that every variable occurring in them belongs to the domain of $\mathscr{D}$ the following holds*:

— *If $\mathscr{M}, s, Val \models \langle\!\langle \alpha \wedge \beta \rangle\!\rangle_{\mathscr{D}}$ then $Sig(\alpha \wedge \beta, s) = \max(Sig(\alpha, s), Sig(\beta, s))$.*

— *If $\mathscr{M}, s \models \langle\!\langle \alpha \vee \beta \rangle\!\rangle_{\mathscr{D}}$ then $Sig(\alpha \vee \beta, s) = Sig(\alpha, s)$ or $Sig(\alpha \vee \beta, s) = Sig(\beta, s)$.*

— *If $\mathscr{M}, s \models \langle\!\langle (a \to \Phi) \rangle\!\rangle_{\mathscr{D}}$ then* (i) *for every formula $\varphi \in \Phi$ there is a state t such that $(s, t) \in R(a)$ and $Sig(\varphi, t) \leqslant Sig((a \to \Phi), s)$,* (ii) *for every state t such that $(s, t) \in R(a)$, $Sig(\bigvee \Phi, t) \leqslant Sig((a \to \Phi), s)$.*

— *If $\mathscr{M}, s \models \langle\!\langle \nu X.\alpha(X) \rangle\!\rangle_{\mathscr{D}}$ and $\mathscr{D}(V) = \nu X.\alpha(X)$ then $Sig(\nu X.\alpha(X), s) = Sig(V, s)$.*

— *If* $\mathcal{M}, s \models \langle\!|\mu X.\alpha(X)|\!\rangle_{\mathcal{D}}$ *and* $\mathcal{D}(U_i) = \mu X.\alpha(X)$ *is the* i*th* (*in the* $\preccurlyeq$ *ordering*) μ*-variable in* $\mathcal{D}$ *then the prefixes of length* $i-1$ *of* $Sig(\mu X.\alpha(X), s)$ *and* $Sig(U_i, s)$ *are equal.*

— *If* $\mathcal{M}, s \models \langle\!|W|\!\rangle_{\mathcal{D}}$ *and* $\mathcal{D}(W) = \sigma X.\alpha(X)$ *then* $Sig(W, s) = Sig(\alpha(W), s)$ *if* W *is a* ν*-variable. If* W *is the* i*th* μ*-variable then the second signature is smaller and the difference is at the position* i.

Similarly for ν*-signatures but with interchanged role of* μ *with* ν*, conjunction with disjunction and the dual statement in* $(a \to \Phi)$ *case.*

Proof. We will consider only the last case. Suppose $\mathcal{M}, s \models \langle\!|U_i|\!\rangle_{\mathcal{D}}$, where U_i is the ith μ-variable in our linear ordering $\preccurlyeq$ of variables.

Let $\mathcal{D}(U_i) = \mu X.\alpha_i(X)$. Recall that $\preccurlyeq$ extends the dependency order; hence, only variables $\preccurlyeq$-smaller than U_i can appear in $\alpha_i(X)$. Let $Sig(U_i, s) = (\tau_1, ..., \tau_n)$ and let $\mathcal{D}'$ be a binding function obtained from $\mathcal{D}$ by replacing the jth μ-variable definition $\mathcal{D}(U_j) = \mu X.\alpha_j(X)$ by $\mathcal{D}(U_j) = \mu^{\tau_j} X.\alpha_j(X)$ for every $j = 1, ..., n$. Let us denote $\langle\!|\alpha_i(X)|\!\rangle_{\mathcal{D}'}$ by $\beta(X)$.

It should be clear that the signature of $\alpha_i(\mu X.\alpha_i(X))$ is the same as the signature of $\mu X.\alpha_i(X)$. This means that the signatures of U_i and $\alpha_i(U_i)$ are the same on positions smaller than i. From the definition of the signature we have $\mathcal{M}, s \models \mu^{\tau_i} X.\beta(X)$. Observe that τ_i must be a successor ordinal. Hence $\mathcal{M}, s \models \beta(\mu^{\tau_i - 1} X.\beta(X))$, which implies the thesis of the lemma. ∎

Proof of Theorem 15 $\Rightarrow$. Let us first focus on the left to right implication. Suppose that φ is satisfied in a state s of a structure $\mathcal{M}$ with a valuation Val. Let $\mathcal{T}$ be a tableau for φ. We will construct a consistent marking M of $\mathcal{T}$ with respect to $\mathcal{M}, s, Val$.

- We put the pair consisting of s and the root of $\mathcal{T}$ into M.
- If $(s, n) \in M$ and the unary rule was applied in n then we put (s, n') into M, for n' the son of n.
- Suppose $(s, n) \in M$ and the rule (or) was applied in n:

$$\frac{\{\alpha, \Gamma\} \quad \{\beta, \Gamma\}}{\{\alpha \vee \beta, \Gamma\}}.$$

Node n has two sons, n_α and n_β, labelled by the obtained assumptions. We put the pair (s, n_α) it into M if $Sig(\alpha, s) < Sig(\beta, s)$; otherwise we put (s, n_β) into M.

- Suppose $(s, n) \in M$ and the rule (mod) was applied in n. If for some $(a \to \Phi) \in L(n)$, $\varphi \in \Phi$ and t with $(s, t) \in R(a)$ we have $Sig(\varphi, t) \leqslant Sig((a \to \Phi), s)$ then we put the pair (t, n_φ) into M, where n_φ is the a-son of n containing φ.

Observe that by the construction for every $(s, n) \in M$ we have $\mathcal{M}, s, Val \models L(n)$. From this and Lemma 17 it follows that M is a marking. This observation also implies that M is a locally consistent marking.

Let us check the global consistency condition of the marking. Let $\mathcal{P} = n_1, n_2, ...$ be a path of $\mathcal{T}$ such that for every node n of $\mathcal{P}$ there is a state s of $\mathcal{M}$ with

$(s, n) \in M$. Construct a sequence of states $s_1, s_2, \ldots$ such that for all $i = 1, \ldots$, $(s_i, n_i) \in M$ and (s_i, n_i) is the reason for the pair (s_{i+1}, n_{i+1}) to be in M; in other words $(s_{i+1}, n_{i+1}) \in M$ because of one of the above rules for constructing the marking.

Suppose on the contrary that there is a μ-trace $\alpha_1, \alpha_2, \ldots$ on $\mathscr{P}$. Let U_k be the smallest variable regenerated infinitely often on this trace; of course it must be μ-variable. Consider the sequence of signatures: $\{Sig(s_i, \alpha_i)\}_{i \in \{1, \ldots\}}$. These are defined because $\mathscr{M}, s_i, Val \models L(n_i)$. From the assumption that U_k is the smallest variable regenerated infinitely often it follows that after some initial part of the sequence the signatures do not increase on positions $1, \ldots, k$. Moreover, each time k is regenerated the signature of a node where it happens decreases at the position k. This is a contradiction with the fact that lexicographic ordering on k-tuples of ordinals is well ordering.

Proof of Theorem 15 $\Leftarrow$. To prove the theorem in the direction from right to left let us assume that there is a tableau $\mathscr{T}$ for φ and its marking M consistent with respect to $\mathscr{M}, s, Val$. Assume conversely that $\mathscr{M}, s, Val \not\models \varphi$. We will show that this assumption leads to a contradiction. We will show that there must be a μ-trace on a path of $\mathscr{T}$ such that for every node n of it there is a state s with $(s, n) \in M$.

Suppose that we have constructed this hypothetical trace up to a node n, formula $\alpha \in L(n)$ is the last formula of it, and s is a state such that $(s, n) \in M$ and $\mathscr{M}, s, Val \not\models \alpha$. We proceed according to the rule which was applied in n.

- Suppose the rule is unary. If it was applied to α then the next element of the trace is the result of a reduction of α; otherwise, the next element is the formula α itself. In the case the (*and*) rule was applied to $\alpha = \gamma_1 \wedge \gamma_2$, choose γ_1 if ${}^{\nu}Sig(\gamma_1, s)$ is smaller than ${}^{\nu}Sig(\gamma_2, s)$ or γ_2 otherwise. It is clear that the new last element of the trace is not satisfied in s.

- If the rule (*or*) was applied in n then choose a son n' of n, s.t. $(s, n') \in M(n')$. The next element of the trace will be the result of a reduction of α which appears in n' or α itself if α was not reduced by this application of the rule.

- If the rule (*mod*) was applied in n then by the definition of a consistent marking α cannot be a literal or a terminal formula. Hence, it is a formula of the form $(a \to \Phi)$ with $\Phi \neq \emptyset$. In this case either:

  1. There is a formula $\varphi \in \Phi$ such that every t with $(s, t) \in R(a)$ we have $t \not\models \varphi$ and ${}^{\nu}Sig(\varphi, t) \leqslant {}^{\nu}Sig((a \to \Phi), s)$. In this case we choose a son n' of n labelled by $\{\varphi\} \cup \{\bigvee\theta\colon (a \to \theta) \in L(n), \theta \neq \Phi\}$. For the next state we take a state t such that $(t, n') \in M$.

  2. There is a state t, s.t. $(s, t) \in R(a)$ and $t \not\models \bigvee \Phi$ with ${}^{\nu}Sig(\bigvee \Phi, t) \leqslant {}^{\nu}Sig((a \to \Phi), s)$. In this case take a son n' of n such that $(t, n') \in M$. Our next formula is $\bigvee \Phi$ or some $\psi \in \Phi$ depending on which one appears in $L(n)$.

Using arguments similar to those in the proof of the left to right implication one can easily prove that the constructed trace must be a μ-trace. This contradicts our assumption about consistency of the marking. ∎

We are now going to define what it means for two tableaux to be equivalent. It occurs that we can abstract from the order of application of nonmodal rules, but the structure of a tree designated by modal nodes will be very important.

DEFINITION 18 (Tableau equivalence). We say that two tableaux $\mathscr{T}_1$ and $\mathscr{T}_2$ are *equivalent* iff there is a bijection $\mathscr{E}$ between the choice and modal nodes of $\mathscr{T}_1$ and $\mathscr{T}_2$ such that:

1. $\mathscr{E}$ maps the root of $\mathscr{T}_1$ to the root of $\mathscr{T}_2$; it maps choice nodes to choice nodes and modal nodes to modal nodes.

2. If n is a descendant of m then $\mathscr{E}(n)$ is a descendant of $\mathscr{E}(m)$. Moreover if for some action a, node n is a $\langle a \rangle$-son of a modal node m then $\mathscr{E}(n)$ is a $\langle a \rangle$-son of $\mathscr{E}(m)$.

3. For every modal node m, the sets of literals and terminal formulas (recall that these are formulas of the form $(a \to \emptyset)$) occurring in $L(m)$ and in $L(\mathscr{E}(m))$ are equal.

4. There is a μ-trace on a path $\mathscr{P}$ of $\mathscr{T}_1$ iff there is a μ-trace on a path of $\mathscr{T}_2$ determined by the image of $\mathscr{P}$ under $\mathscr{E}$.

The next theorem shows that tableau equivalence is a refinement of the semantical equivalence. It is quite easy to see that this is a strict refinement; there are semantically equivalent formulas which do not have equivalent tableaux.

THEOREM 19. *If two formulas (satisfying our proviso) have equivalent tableaux then they are semantically equivalent.*

Proof. Let α, β be two formulas and let $\mathscr{T}_1$, $\mathscr{T}_2$ be equivalent tableaux for α and β, respectively. Let $\mathscr{E}: \mathscr{T}_1 \to \mathscr{T}_2$ denote the bijection showing the equivalence of $\mathscr{T}_1$ and $\mathscr{T}_2$. We will show that for every structure $\mathscr{M}$, state s, and valuation Val, we have $\mathscr{M}, s, Val \models \alpha$ iff $\mathscr{M}, s, Val \models \beta$.

Suppose $\mathscr{M}, s, Val \models \alpha$. By Theorem 15 there is a consistent marking M_1 of $\mathscr{T}_1$ with respect to $\mathscr{M}, s, Val$. This marking determines a consistent marking M_2 of $\mathscr{T}_2$ defined as follows. First for every modal or choice node n of $\mathscr{T}_2$ we let $(s, n) \in M_2$ iff $(s, \mathscr{E}^{-1}(n)) \in M_1$. The marking of the internal nodes is uniquely determined by this choice. Directly from the definition of the equivalence it follows that M_2 is consistent with respect to $\mathscr{M}, s, Val$.

Observe that $\mathscr{E}^{-1}$ is also a function showing equivalence of $\mathscr{T}_2$ and $\mathscr{T}_1$; hence, there is a way of obtaining a consistent marking of $\mathscr{T}_1$ from a consistent marking of $\mathscr{T}_2$. ∎

3.2. Disjunctive Formulas

Here we define the notion of disjunctive formulas. The main theorem of this subsection shows that every formula is equivalent to some disjunctive formula.

DEFINITION 20 (Special conjunctions and disjunctive formulas). A conjunction $\alpha_1 \wedge \cdots \wedge \alpha_n$ is called *special* iff every α_i is either a literal or a formula of the form

$(a \to \Psi)$ and for every action a there is at most one conjunct of the form $(a \to \Psi)$ among $\alpha_1, ..., \alpha_n$.

The set of *disjunctive formulas*, $\mathscr{F}_d$, is the smallest set defined by the following clauses:

1. Every literal is a disjunctive formula.

2. If $\alpha, \beta \in \mathscr{F}_d$ then $\alpha \vee \beta \in \mathscr{F}_d$; if, moreover, X occurs only positively in α and not in the context $X \wedge \gamma$ for some γ, then $\mu X.\alpha$, $\nu X.\alpha \in \mathscr{F}_d$.

3. $(a \to \Psi) \in \mathscr{F}_d$ if $\Psi \subseteq \mathscr{F}_d$.

4. A special conjunction of disjunctive formulas is a disjunctive formula.

Remark. Modulo the order of application of (*and*) rules, disjunctive formulas have unique tableaux. Moreover, on every infinite path there is one and only one infinite trace.

A permutation of the order of application of (*and*) rules does not change the shape of a tableau. It just changes the order in which conjunctions are replaced with commas. On the other hand it matters when (*or*) rules are applied, as the following example shows:

$$\dfrac{\{(\alpha \vee \beta) \vee \gamma, \alpha \vee \beta\}}{\dfrac{\{\alpha, \alpha \vee \beta\}}{\{\alpha\} \quad \{\alpha, \beta\}} \quad \dfrac{\{\alpha, \alpha \vee \beta\}}{\{\gamma, \alpha\} \quad \{\gamma, \beta\}}}$$

$$\dfrac{\{(\alpha \vee \beta) \vee \gamma, \alpha \vee \beta\}}{\dfrac{\{(\alpha \vee \beta) \vee \gamma, \alpha\}}{} \quad \dfrac{\{(\alpha \vee \beta) \vee \gamma, \beta\}}{}}$$

$$\vdots$$

$$\{\alpha\} \quad \{\alpha, \beta\} \quad \{\alpha, \gamma\} \quad \{\alpha, \beta\} \quad \{\beta\} \quad \{\beta, \gamma\}.$$

DEFINITION 21 (Unwinding, regular tableau). Given a labelled graph with a source, $\mathscr{G} = \langle G_{\mathscr{G}}, L_{\mathscr{G}}, s_{\mathscr{G}} \rangle$ an *unwinding* of $\mathscr{G}$ is a tree whose nodes are finite paths of $G_{\mathscr{G}}$ starting from $s_{\mathscr{G}}$ and the label of such a finite path is the label of the last node in the path. A tableau is *regular* if it is an unwinding of a finite graph.

The following theorem shows that every formula is equivalent to a disjunctive formula. The theorem is even stronger and this stronger statement will be needed in the completeness proof.

THEOREM 22. *For every formula φ and every regular tableau $\mathscr{T}$ for φ there is a disjunctive formula φ with a tableau equivalent to $\mathscr{T}$.*

Proof. Let $\mathscr{T} = \langle T, L \rangle$ be a regular tableau for φ. Suppose $\mathscr{T}$ is the unwinding of a finite labelled graph with a source $\mathscr{G} = \langle G_{\mathscr{G}}, L_{\mathscr{G}}, s_{\mathscr{G}} \rangle$.

We first show that it is possible to find another finite graph $\mathscr{K} = \langle G_{\mathscr{K}}, L_{\mathscr{K}}, s_{\mathscr{K}} \rangle$ that also unwinds to $\mathscr{T}$ and that is a finite *tree with back edges*. A back edge is an edge leading from a node to one of its ancestors. We will still use tree-like terminology to trees with back edges. For example, we will say that one node is a son of the other if it is so in a tree obtained by forgetting about back edges. The most useful property of the graph $\mathscr{K}$ will be a special colouring of nodes described in the lemma.

LEMMA 22.1. *It is possible to construct a finite tree with back edges* $\mathscr{K} = \langle G_{\mathscr{K}}, L_{\mathscr{K}}, s_{\mathscr{K}} \rangle$ *satisfying the following conditions*:

(1) $\mathscr{K}$ *unwinds to* $\mathscr{T}$,

(2) *Every node to which a back edge points can be assigned the colour magenta or navy so that for every infinite path of the unwinding of* $\mathscr{K}$ *we have that there is a* μ*-trace on the path iff the closest to the root node of* $\mathscr{K}$ *appearing infinitely often on the path is coloured magenta.*

Proof. There are only finitely many possible labels of nodes of $\mathscr{T}$. Hence we can consider these labels as letters of our alphabet. It is easy to see that there is a Muller automaton on infinite words recognising those paths of $\mathscr{T}$ which have a μ-trace on them. We can assume that this automaton is deterministic.

From the results of Mostowski [6] it follows that there is an equivalent deterministic automaton $\mathscr{A}$ with so-called *parity conditions*. Parity conditions are given by a function Ω assigning to each state of the automaton a natural number. A run of the parity automaton is accepting if the smallest priority among priorities of the states appearing infinitely often on the run is even.

We can run our parity automaton $\mathscr{A}$ also on finite paths of $\mathscr{G}$ (the input is a sequence of the labels of the nodes). Let $\mathscr{A}(p)$ denote a state that $\mathscr{A}$ reaches after reading p. We define our tree with back edges $\mathscr{K} = \langle G_{\mathscr{K}}, L_{\mathscr{K}}, s_{\mathscr{K}} \rangle$ as follows:

(a) The vertices of the graph $G_{\mathscr{K}}$ are finite paths $s_0, ..., s_k$ of $G_{\mathscr{G}}$ with the following property: if $s_i = s_j$ for some $i \neq j$ $(i, j \in \{1, ..., k\})$ and s_i is a choice but not a modal node then either $\mathscr{A}(s_i) \neq \mathscr{A}(s_j)$ or there is $i < i' < j$ with $\Omega(\mathscr{A}(s_{i'})) < \Omega(\mathscr{A}(s_i))$.

(b) If p, p' are two vertices of $G_{\mathscr{K}}$ then there is an edge from p to p' if either (i) p' is one element longer than p or (ii) p' is a prefix of p ending in a node s_j such that ps_j is a path of $\mathscr{G}$ but not a vertex of $G_{\mathscr{K}}$.

It should be clear that the unwinding of $\mathscr{K}$ is $\mathscr{T}$. For arbitrary automaton $\mathscr{A}$, the graph $\mathscr{K}$ may not be finite. It may happen when in all choice nodes $\mathscr{A}$ is forced to take states with some big priority. But this is the only technical complication. To avoid it we can assume, without a loss of generality, that $\mathscr{A}$ is such that a state assigned to a choice node in an accepting run of $\mathscr{A}$ always has a smaller priority than any state assigned to nonchoice nodes. If $\mathscr{A}$ has this property then it is easy to show that $\mathscr{K}$ is finite as on every infinite path some choice and not modal node must appear infinitely often.

We colour a node p of $\mathscr{K}$ magenta if the priority of the state $\mathscr{A}(p)$ is even; otherwise we colour p navy. Let us check that this colouring satisfies condition (2) of the lemma. Let us take an infinite path $p_0, p_1, ...$ of the unwinding of $\mathscr{K}$. The sequence $\mathscr{A}(p_0), \mathscr{A}(p_1), ...$ is a run of $\mathscr{A}$ on the path. Let $\bar{p}$ be the closest to the root node of $\mathscr{K}$ appearing infinitely often on the path. It is not difficult to see that such a node exists. From the construction of $\mathscr{K}$ it follows that $\Omega(\mathscr{A}(\bar{p}))$ is the smallest among priorities of those states which appear infinitely often in $\mathscr{A}(p_0), \mathscr{A}(p_1), ...$. Hence, $\mathscr{A}(p_0), \mathscr{A}(p_1), ...$ is an accepting run iff $\Omega(\mathscr{A}(\bar{p}))$ is coloured magenta. By definition of $\mathscr{A}$, it happens iff there is a μ-trace on $p_0, p_1, ...$. ∎

From the tree with back edges $\mathscr{K}$ we are going to construct a disjunctive formula which has a tableau equivalent to $\mathscr{T}$. We start from the leaves of $\mathscr{K}$ and going to the top assign a formula $F(n)$ to each node n of $\mathscr{K}$ in the following way:

- If there are no edges going from n then in the label of n only literals and terminal formulas can occur. We let $F(n)$ be a conjunction of all the formulas appearing in the label of n plus the formula $\mu X_n.\top$.

- If there are edges going from n then we assume that every son of n has some formula assigned to it. It will be convenient to consider that a formula assigned to a son is also assigned to an edge leading from n to this son. There can be also back edges leading from n to some ancestors of n. Of course those ancestors have no formula assigned yet. To such a back edge from n to, say, m we assign the variable U_m (the index of the variable is the target of the edge). We first define an auxiliary formula γ depending on the rule which was applied in n.

  — If one of the rules (μ), (ν), (*cons*), or (*and*) was applied in n then γ is exactly the same as the formula assigned to the only edge leading from n.

  — If the rule (*or*) was applied in n then there are two edges leading from n which have been assigned formulas ψ_1 and ψ_2. We let $\gamma = \psi_1 \vee \psi_2$.

  — If the rule (*mod*) was applied then let Ψ_a be the set of all the formulas assigned to the edges leading from n to some node labelled by a result of reduction of the action a. We let γ be a conjunction of all the literals and terminal formulas appearing in $L(n)$ together with all the formulas of the form $(a \to \Psi_a)$. If n is not a choice node then we add a conjunct $\mu X_n.\top$ to γ.

If there are no back edges leading to n then $F(n)$ is just γ. Otherwise, let $F(n) = \sigma U_n.\gamma$, where σ is μ or ν depending on whether n was coloured magenta or navy, respectively.

We let $\hat{\varphi}$ be the formula assigned to the root of $\mathscr{K}$, i.e., $\varphi = F(n_0)$. Observe that $\hat{\varphi}$ has only one tableau, call it $\hat{\mathscr{T}}$. We show that it is equivalent to $\mathscr{T}$.

Remark. Strictly speaking the constructed formula φ is not a disjunctive formula. This is because of conjuncts $\mu X_n.\top$ added during the construction. These conjuncts are necessary to obtain the property from Observation 22.2 which in turn simplifies main arguments. To be completely formal we could allow such conjuncts in the definition of a disjunctive formula and say in the definition of the tableau equivalence that a presence or absence of $\top$ in the label does not matter. These changes would just obscure the notions and would not give any complications in any of the proofs. We have chosen to pretend that $\hat{\varphi}$ is a disjunctive formula and to forget about these inessential technicalities.

Let $\hat{\mathscr{T}}$ be a tableau for $\hat{\varphi}$. We will define an equivalence function $\mathscr{E}\colon \hat{\mathscr{T}} \to \mathscr{T}$. Recall that $\mathscr{T}$ is the unwinding of $\mathscr{K}$; hence, its nodes are paths of $\mathscr{K}$ and the label of such a node is the label of the last element of the path. Function $\mathscr{E}$ will have the property that for every modal or choice node $\hat{m}$ of $\hat{\mathscr{T}}$,

$$F(\text{last}(\mathscr{E}(\hat{m}))) = L(\hat{m}); \tag{2}$$

here $\text{last}(\mathscr{E}(\hat{m}))$ is the last element of the path $\mathscr{E}(\hat{m})$.

We construct $\mathscr{E}(\hat{m})$ by induction on the distance of $\hat{m}$ from the root. To the root of $\hat{\mathscr{T}}$ we assign the root of $\mathscr{T}$. If the root of $\hat{\mathscr{T}}$ is a modal node then by construction F, the root of $\mathscr{T}$ is also a modal node. Conversely, if the root of $\mathscr{T}$ is a modal node then by the condition (a) of the definition of $\mathscr{K}$ we know that there cannot be a back edge to the root in $\mathscr{K}$ (because it is both a modal and a choice node). Hence, by the construction F, the root of $\hat{\mathscr{T}}$ is also a modal node.

Suppose $\mathscr{E}$ is defined for some modal node $\hat{m}$. By the property (2) and the construction F there is a bijection between the sons of $\hat{m}$ and the sons of m. We extend F according to this bijection.

Suppose $\mathscr{E}$ is defined for some choice node $\hat{m}$. If it is also a modal node then by the property (2) so is $\mathscr{E}(\hat{m})$. If m is not a modal node then by the construction of $\mathscr{K}$ node $\mathscr{E}(m)$ also cannot be a modal node. By the construction F there is a bijection between modal nodes near m and modal nodes near m. We use this bijection to extend $\mathscr{E}$. It can be checked that $\mathscr{E}$ satisfies clauses 1–3 of the definition of tableaux equivalence.

The truth of the clause 4 of the definition of the equivalence follows from properties of our colouring. Let $\hat{\mathscr{P}}$ be a path of $\mathscr{T}$. The path $\mathscr{E}(\hat{\mathscr{P}})$ is some path in $\mathscr{K}$. Let n be the closest to the root node of $\mathscr{E}(\hat{\mathscr{P}})$ which appears infinitely often on $\mathscr{E}(\hat{\mathscr{P}})$. By the construction F and property (2) we have that the smallest variable regenerated infinitely often on the unique trace on $\hat{\mathscr{P}}$ is U_n. This variable is a μ-variable if n is coloured magenta and it is a ν-variable otherwise. Hence, the smallest variable regenerated on the unique trace on $\hat{\mathscr{P}}$ is a μ-variable iff there is a μ-trace on $\mathscr{E}(\hat{\mathscr{P}})$.

The addition of $\mu X_n.\top$ components in the construction guarantees the following additional property:

Observation 22.2. For a modal or choice node $\hat{m}$ of $\hat{\mathscr{T}}$ and a set of modal or choice nodes $\hat{N}$ of $\hat{\mathscr{T}}$, if $\hat{L}(\hat{m}) \subseteq \bigcup_{n \in \hat{N}} \hat{L}(\hat{n})$ then $L(\mathscr{E}(\hat{m})) \subseteq \bigcup_{\hat{n} \in \hat{N}} L(\mathscr{E}(\hat{n}))$. In particular if $\hat{L}(\hat{m}) = \hat{L}(\hat{n})$ then $L(\mathscr{E}(\hat{m})) = L(\mathscr{E}(\hat{n}))$. ∎

4. THE SYSTEM

Here we present an axiomatisation of the μ-calculus proposed by Kozen [4] and show some simple properties of the system.

We adopt the original formulation of Kozen. The basic judgement of the system has the form $\alpha = \beta$ with the intended meaning that the two formulas are semantically equivalent. Judgement $\alpha \leqslant \beta$ is an abbreviation for $\alpha \wedge \beta = \alpha$. A formula α is *provable* if $\alpha = \top$ is provable.

The axiomatisation consists of the axioms and rules of equational logic (including substitution of equals by equals, i.e., cut rule) and the following axioms and rules:

$$(K1) \quad \text{axioms for Boolean algebra}$$

$$(K2) \quad \langle a \rangle \varphi \vee \langle a \rangle \psi = \langle a \rangle (\varphi \vee \psi)$$

$$(K3) \quad \langle a \rangle \varphi \wedge [a] \psi \leqslant \langle a \rangle (\varphi \wedge \psi)$$

$$(K4)\qquad \langle a\rangle \perp = \perp$$

$$(K5)\qquad \alpha(\mu X.\alpha(X)) \leqslant \mu X.\alpha(X)$$

$$(K6)\qquad \frac{\alpha(\varphi)\leqslant\varphi}{\mu X.\alpha(X)\leqslant\varphi}.$$

Because we have put ν and box directly into the language we have to define them by equivalences:

$$[a]\,\alpha = \neg\langle a\rangle\neg\alpha$$

$$\nu X.\alpha(X) = \neg\mu X.\neg\alpha(\neg X).$$

It was proved in [4] (or it easily follows from [4]) that the following rules are admissible:

$$(\langle\ \rangle)\qquad \frac{\psi\wedge\bigwedge\{\bigvee\theta\colon (a\to\theta)\in\Gamma,\ \theta\neq\Psi\}\leqslant\perp}{(a\to\Psi)\wedge\bigwedge\Gamma\leqslant\perp}\qquad \text{for some}\quad \psi\in\Psi$$

$$(\text{fix})\qquad \frac{\gamma\wedge\alpha(\mu X.\neg\gamma\wedge\alpha(X))\leqslant\perp}{\gamma\wedge\mu X.\alpha(x)\leqslant\perp}$$

$$(\text{mon})\qquad \frac{\alpha\leqslant\beta}{\varphi(\alpha)\leqslant\varphi(\beta)}\qquad X \text{ occurs only positively in } \varphi(X).$$

According to our proviso we restrict ourselves to well-named, positive, and guarded formulas. We must show that it is a harmless restriction as far as provability is concerned.

FACT 23. *Every formula is provably equivalent to a formula satisfying the proviso at the end of Section* 2.

Proof. Just observe that all the steps used in transforming a formula to a positive guarded form as described in Proposition 2 use provable equivalences. ∎

One of the nice properties of disjunctive formulas is that they are easy as far as provability is concerned.

THEOREM 24. *For every unsatisfiable disjunctive formula* α *the formula* $\neg\alpha$ *is provable.*

Proof. In [2] it was shown:

> A disjunctive formula α is satisfiable iff β obtained from α by replacing all μ-variables by $\perp$ and all ν-variables by $\top$ is satisfiable. (3)

We prove the theorem by induction on the size of α.

Suppose α is a special conjunction $\alpha_1\wedge\cdots\wedge\alpha_n$; we have two cases. If $\alpha_i=\perp$ or $\alpha_i=\neg\alpha_j$ for some $i, j\in\{1, ..., n\}$ then $\neg\alpha$ is easily provable. Otherwise one of the conjuncts must be of the form $(a\to\Psi)$ and one of the formulas from Ψ must be unsatisfiable. From the induction assumption and the rule ($\langle\ \rangle$) we obtain the proof of $\neg\alpha$.

If $\alpha = \gamma \vee \delta$ then by the induction assumption we have proofs of $\neg\gamma$ and $\neg\delta$. We can use propositional calculus laws to obtain the proof of $\neg\alpha$.

If $\alpha = \mu X.\gamma(X)$ then because this formula is unsatisfiable so is $\gamma(\bot)$. By induction assumption there is a proof of $\neg\gamma(\bot)$ and we can use the derivable rule:

$$\frac{\neg\gamma(\bot) = \top}{\neg\mu X.\gamma(X) = \top}.$$

If $\alpha = \nu X.\gamma(X)$ then we consider $\gamma(\top)$. It is, of course, a disjunctive formula. By (3), the formula $\gamma(\top)$ is satisfiable iff $\nu X.\gamma(X)$ is satisfiable. As the latter formula is not satisfiable we have by the induction assumption the proof $\neg\gamma(\top)$ and we can use the derivable rule:

$$\frac{\neg\gamma(\top) = \top}{\neg\nu.\gamma(X) = \top}. \quad \blacksquare$$

5. PROVABILITY FOR WEAKLY ACONJUNCTIVE FORMULAS

In this section we will consider a class of formulas for which the provability is easier than in the general case (although not as easy as for disjunctive formulas). We recall the notion of *aconjunctive formulas* [4] and propose its slight generalisation called *weakly aconjunctive formulas*. Our goal in the section is to obtain a generalisation of the main result from [4] which states that the negation of every unsatisfiable aconjunctive formula is provable. To do this we introduce the notion of *thin refutation*, which isolates the cases for which the original proof still goes through. It turns out that every refutation of a weakly aconjunctive formula is thin.

DEFINITION 25 (Weakly aconjunctive formulas). Let φ be a formula, $\mathscr{D}_\varphi$ be its binding function, and $\leqslant_\varphi$ be the dependency ordering (see Definitions 3 and 4).

— We say that a variable X is *active* in ψ, a subformula of φ, iff there is a variable Y appearing in ψ and $X \leqslant_\varphi Y$.

— Let X be a variable with its binding definition $\mathscr{D}_\varphi(X) = \mu X.\gamma(X)$. The variable X is called *aconjunctive* iff for all subformulas of γ of the form $\alpha \wedge \beta$ it is not the case that X is active in α as well as in β.

— A variable X as above is called *weakly aconjunctive* iff for all subformulas of γ of the form $\alpha \wedge \beta$ if X is active in both α and β then $\alpha \wedge \beta$ is a special conjunction as defined in Definition 20.

— A formula φ is called (*weakly*) *aconjunctive* iff all μ-variables in φ are (weakly) aconjunctive.

In the following we will be interested only in weakly aconjunctive formulas. The definition of aconjunctive formulas was restated just to compare the two notions. It also makes sense to compare weakly aconjunctive formulas with disjunctive formulas.

FACT 26. *Every disjunctive formula is a weakly aconjunctive formula.*

Proof. Directly from the fact that the only conjunctions in disjunctive formulas are special conjunctions. ∎

From the next observation it follows that all formulas appearing in a tableau for a weakly aconjunctive formula are weakly aconjunctive.

FACT 27. *Every formula appearing in a tableau for φ is a subformula of φ.*

The next proposition states some closure properties of the class of weakly aconjunctive formulas. Observe that weakly aconjunctive formulas are not closed under negation or under the least fixpoint operation.

PROPOSITION 28 (Composition). *If $\gamma(X)$ and δ are weakly aconjunctive formulas then $\gamma[\delta/X]$, $\nu X.\gamma(X)$, and $\delta \wedge \gamma(X)$ are also weakly aconjunctive formulas.*

Proof. As we consider only well-named formulas, when conjunction is formed we make sure that the bound variables in δ and $\gamma(X)$ are different. With this observation it should be easy to see that $\nu X.\gamma(X)$ and $\delta \wedge (X)$ are weakly aconjunctive.

Also while performing substitution $\gamma[\delta/X]$ we keep bound variables of δ distinct from the bound variables of γ. Let $\alpha = \gamma[\delta/X]$ and let Y be a μ-variable of α. This variable is bound either in γ or in δ. If it is a bound variable from γ then because no bound variable of γ is free in δ we have that for every $Y \leqslant_\alpha Z$, variable Z is a bound variable of γ. Hence, Y is weakly aconjunctive in α iff it was weakly aconjunctive in γ. For a similar reason every μ-variable of δ is weakly aconjunctive in α. ∎

Next we turn our attention to *refutations*. These will be tableau-like objects corresponding to unsatisfiability, in the sense that a formula has a refutation iff it is unsatisfiable. From our point of view the most interesting property is that from some refutations, called *thin*, we can construct a proof of the negation of the initial formula.

DEFINITION 29 (Refutations). A *refutation* for a formula φ is defined as tableau, but this time we modify system $\mathcal{S}^\varphi$ (presented in Fig. 1) by adding the explicit weakening rule

$$\frac{\Gamma}{\{\alpha, \Gamma\}}$$

and instead of (*mod*) rule we take ($\langle\,\rangle$) rule:

$$(\langle\,\rangle) \quad \frac{\{\psi\} \cup \{\bigvee \theta : (a \to \theta) \in \Gamma, \theta \neq \Psi\}}{\Gamma} \qquad (a \to \Psi) \in \Gamma, \psi \in \Psi,$$

This rule is similar to (*mod*) but has only one assumption. Additionally, we require that on every infinite path of the refutation there should be a μ-trace and every leaf of the refutation must be labelled by a set containing $\perp$ or some literal and its negation.

We call a refutation *thin* iff whenever a formula of the form $\alpha \wedge \beta$ is reduced in some node of the refutation and some variable is active in α as well as in β then either $\alpha \wedge \beta$ is a special conjunction or one of the conjuncts is immediately discarded by the use of the weakening rule.

The following is an easy consequence of Fact 27.

FACT 30. *Every refutation for a weakly aconjunctive formula is a thin refutation.*

It was shown in [7] that every unsatisfiable formula has a refutation. From this perspective the next theorem essentially says that one can prove the negation of an unsatisfiable weakly aconjunctive formula. The theorem is stated in greater generality because in Lemma 36 we deal with thin proof tableaux for formulas which may not be weakly aconjunctive.

THEOREM 31. *If a formula has a thin refutation then its negation is provable.*

Proof. The proof is a reformulation of Kozen's argument from [4]. Let $\mathscr{R}$ be a thin refutation for φ. We can assume that in $\mathscr{R}$ we reduce special conjunctions only when no other formula can be reduced by rules other than (*mod*). This restriction does not change the shape of the tableau. Let $\mathscr{D}$ be the binding function associated with φ and let $\leqslant_\varphi$ be the dependency ordering on the bound variables of φ. It will be convenient here to use some arbitrary linearisation of $\leqslant_\varphi$. We will write $<_\varphi$ for strictly less relation determined by this linearisation.

We will assign to every node m of $\mathscr{R}$ a formula which will contain some information about the path up to m. The information we are interested in is what variables were regenerated and in what nodes. To see what we mean consider a node m labelled $\{X, \Gamma\}$ and its son n labelled $\{\alpha(X), \Gamma\}$. The formula assigned to m will have the form $\neg(\gamma \wedge \mu X.\beta(X))$. Now, to remember the context in which X was regenerated we can use the rule (*fix*) and assign to n the formula $\neg(\gamma \wedge \beta(\mu X.\neg\gamma \wedge \beta(X)))$. If it ever happens that in some descendant o of n we regenerate X in the same context then we can use this recorded information in a sense that the formula assigned to o will be of the form $\neg(\gamma \wedge \neg\gamma \wedge \beta(\mu X.\neg\gamma \wedge \beta(X)))$; hence, it will be a provable formula. Summarising, we want two properties from our assignment of formulas:

1. If the formula assigned to a node is unprovable then the formula assigned to one of the sons is unprovable.

2. In some nodes use remembering so that on every path there is a variable regenerated in the context which is already recorded.

If $\neg\varphi$ is not provable then, by the first property, we can find an infinite path of $\mathscr{R}$, every node of which has associated an unprovable formula. By the second property, we obtain a contradiction because a formula associated with the node where some variable was regenerated for the second time in the same context is provable.

Unfortunately the second property is quite difficult to obtain. If we just used the remembering trick in every possible node, it could happen that we could get infinitely many different contexts. We have to be very careful about what information we remember and what we should forget. This is why the assignment of formulas to nodes is rather involved. We split it into two steps. First, starting from the root of $\mathscr{R}$ we assign a token list to every node; then we use this list assignment to define formulas.

We assume that we have a countable set of *tokens*. We can remove tokens from the list and we can add tokens to the right end of the list. Removed tokens are never used again. Each token has its own counter. We also assign a pair (formula, bound variable of φ) to every token on the list.

Let us first introduce some operations on labelled lists of tokens. We say that α *is replaceable* by β in some list of tokens if either of the conditions holds:

1. β does not appear in the labels of tokens in the list,
2. the smallest variable X_α such that (α, X_α) is the label of some token is smaller than the smallest variable X_β such that (β, X_β) is the label of some token,
3. variables X_α and X_β are the same but the token labelled (α, X_α) is to the left of the token labelled (β, X_β).

If α is replaceable by β then to *replace* α by β means, first, to delete all the tokens labelled (β, Y), for some variable Y, and next replace each label of the form (α, Z) by (β, Z). If α is not replaceable by β then we can *delete* β from the list by removing all the tokens labelled (β, Y) for some variable Y.

To the root of $\mathscr{R}$ we assign an empty list of tokens. A list of tokens for an internal node n is constructed from the list for its father m according to the following rules:

1. If the weakening rule was applied and some formula, say α, was deleted from the label of the node then delete α also from the token list.
2. Suppose the (*or*) rule is applied in m to $\alpha \vee \beta$ and, say, α is the result of the reduction which appears in the label of n. The token list for n is obtained by replacing $\alpha \vee \beta$ by α if $\alpha \vee \beta$ is replaceable by α.
3. Suppose in m we apply the rule

$$(\mathrm{reg})\quad \frac{\{\alpha(X), \Gamma\}}{\{X, \Gamma\}} \qquad \mathscr{D}(X) = \sigma X.\alpha(X).$$

If X is replaceable by $\alpha(X)$ then replace X by $\alpha(X)$. In case X is a μ-variable additionally increase the counter of the token now labelled $(\alpha(X), X)$ and set to 0 the counters of all the tokens to the right of it.

4. Suppose in m we apply the rule (μ) or (ν):

$$(\sigma)\quad \frac{\{\alpha(X), \Gamma\}}{\{\sigma X.\alpha(X), \Gamma\}} \qquad \mathscr{D}(X) = \sigma X.\alpha(X).$$

If $\sigma X.\alpha(X)$ is replaceable by $\alpha(X)$ then we replace $\sigma X.\alpha(X)$ by $\alpha(X)$. In case X is a μ-variable we additionally put a new token labelled $(\alpha(X), X)$ at the end of the list.

5. Suppose in m we apply the (*and*) rule to a formula $\alpha \vee \beta$ which is not a special conjunction. Because our refutation is thin we know that either (i) one of

the conjuncts is deleted or (ii) every variable is active in at most one of the formulas α or β. In the first case we proceed in exactly the same way as in the case of the (*or*) rule.

In the second case if $\alpha \wedge \beta$ is replaceable by α then for every token labelled $(\alpha \wedge \beta, Y)$, with Y active in α, replace its label by (α, Y). Similarly if $\alpha \wedge \beta$ is replaceable by β then for every token labelled $(\alpha \wedge \beta, Y)$, with Y active in β, replace its label by (β, Y).

6. Suppose we have zero or more applications of the (*and*) rule to special conjunctions followed by an application of the $(\langle\,\rangle)$ rule:

$$\dfrac{\dfrac{\{\gamma^i \wedge (a_1 \to \Psi^i_1) \wedge \cdots \wedge (a_k \to \Psi^i_k) : i=1, ..., l\}}{\vdots}}{\dfrac{\{\gamma^i, (a_1 \to \Psi^i_1), ..., (a_k \to \Psi^i_k) : i=1, ..., l\}}{\{\bigvee \Psi^1_i, ..., \bigvee \Psi^{j-1}_i, \psi_j, \bigvee \Psi^{j+1}_i, ..., \bigvee \Psi^l_i\}}}\,.$$

Consider indices $p=1, ..., l$ one by one, starting from the smallest. If $\bigvee \Psi^p_i$ is different from $\bigvee \Psi^q_i$ for all $q<p$ then for every token labelled

$$(\gamma^p \wedge (a_1 \to \Psi^p_1) \wedge \cdots \wedge (a_k \to \Psi^p_k), X) \tag{4}$$

for some X, replace this label by $(\bigvee \Psi^p_i, X)$. If $\bigvee \Psi^p_i = \bigvee \Psi^q_i$, for some $q<p$, then check if the formula (4) is replaceable by $(\bigvee \Psi^p_i, X)$; if so, perform the replacement. Similarly, for ψ_j check whether the formula (4), for $p=j$, is replaceable by ψ_j. If so, perform the replacement.

7. After the above steps we remove tokens which are either (i) labelled with pairs (α, Y) with Y not active in α or (ii) labelled with formulas not appearing in the label of the node.

OBSERVATION 31.1. *For every path $\mathscr{P}$ of $\mathscr{R}$ there is a counter which gets arbitrarily big on $\mathscr{P}$.*

Proof. As $\mathscr{R}$ is a refutation, there is a μ-trace $\mathscr{T}\!r$ on $\mathscr{P}$. Let X be the smallest variable regenerated on this trace. Let n_0 be a node of $\mathscr{P}$ where X is regenerated on $\mathscr{T}\!r$ and after which no variable smaller than X is regenerated on $\mathscr{T}\!r$.

Let t_0 be a token from the list for n_0 labelled (X, Y_0) for the smallest possible Y_0. We call t_0 the *support* of the trace in n_0. Such a support exists because (X, X) is in the label of some token. Because of the step 7 of the construction, we know that X depends on Y_0 (i.e., $Y_0 \leqslant_\varphi X$).

Suppose that n_1 is a node where t_0 is deleted. As X is active in $\mathscr{T}\!r(n_1)$, it can happen only because there was a token t_1 labelled $(\mathscr{T}\!r(n_1), Y_1)$ for Y_1 smaller than Y_0 or maybe $Y_0 = Y_1$ but t_1 was to the left of t_0. The new support for the trace becomes t_1. It should be clear that support can change only finitely many times.

Let m be a node after which the support does not change and where X is regenerated. Let t be a token labelled (X, X) on the list of m. From this point t is never deleted and its counter is increased every time X is regenerated.

If the counter of t is not unbounded then there is a token to the left of it in the list whose counter is being increased; i.e., the counter of the leftmost such token is unbounded. ∎

Next we assign a formula to every node of $\mathscr{R}$. To do this for every node n of $\mathscr{R}$ and every formula $\beta \in L(n)$ we define a binding function $\mathscr{D}_{n,\beta}$ depending on the token list for n. These binding functions will be obtained from $\mathscr{D}$ by modifications of one kind. For some μ-variables instead of $\mathscr{D}(X) = \mu X.\alpha(x)$ we will have $\mathscr{D}_{n,\beta}(X) = \mu X.\neg\gamma_1 \wedge \cdots \wedge \neg\gamma_k \wedge \alpha(X)$, where formulas $\gamma_1, ..., \gamma_k$ are determined in the following way:

> Consider ancestors of n up to the nearest node where a token now labelled (β, X) is added or its counter is reset to zero. Among these ancestors choose all $n_1, ..., n_k$ where the counter of the token was increased, then for $i = 1, ..., k$, $(*)$

$$\gamma_i = \bigwedge \{ \langle\!\lbrack \delta \rbrack\!\rangle_{\mathscr{D}_{n_i,\delta}} : \delta \in L(n_i), \delta \neq \alpha(X) \}.$$

The formula assigned to the node n is

$$\neg \bigwedge \{ \langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}} : \beta \in L(n) \}.$$

OBSERVATION 31.2. *If for some node m formula $\neg\bigwedge\{\langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{m,\beta}} : \beta \in L(m)\}$ is unprovable then there is a son n of m such that $\neg\bigwedge\{\langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}} : \beta \in L(n)\}$ is unprovable.* (*In case special conjunctions are reduced in m, the node n is not a son of m but a son of the modal node near m.*)

Proof. The proof is by cases depending on the rule which was applied in m. We will consider only one case when the rule applied in m is a regeneration of a μ-variable:

$$(\text{reg}) \quad \frac{\{\alpha(X), \Gamma\}}{\{X, \Gamma\}} \qquad \mathscr{D}(X) = \mu X.\alpha(X).$$

If X is not replaceable by $\alpha(X)$ then

$$\{ \langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}} : \beta \in \{\alpha(X), \gamma\} \} = \{ \langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{m,\beta}} : \beta \in \Gamma \}.$$

Otherwise, looking at the changes to the token list for the son n of m we can see that for every $\beta \in \Gamma$ and every variable Y, $\mathscr{D}_{n,\beta}(Y)$ is either $\mathscr{D}_{m,\beta}(Y)$ or $\mathscr{D}(Y)$. This implies that $\langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{m,\beta}} \leqslant \langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}}$ is provable for all $\beta \in \Gamma$.

By definition $\langle\!\lbrack X \rbrack\!\rangle_{\mathscr{D}_{n,\alpha(X)}}$ is of the form $\mu X.\gamma \wedge \alpha(X)$ and let us denote $\{\langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}}: \beta \in L(n), \beta \neq \alpha(X)\}$ by θ. We know that

$$\neg\Big(\bigwedge \theta \wedge (\mu X.\gamma \wedge \alpha(X))\Big)$$

is unprovable. By rule (*fix*)

$$\neg\Big(\bigwedge \theta \wedge \gamma \wedge \alpha\Big(\mu X.\neg \bigwedge \theta \wedge \gamma \wedge \alpha(X)\Big)\Big)$$

is unprovable, hence

$$\neg\Big(\bigwedge \theta \wedge \alpha\Big(\mu X.\neg \bigwedge \theta \wedge \gamma \wedge \alpha(X)\Big)\Big)$$

is unprovable. But $\alpha(\mu X.\neg\bigwedge \theta \wedge \gamma \wedge \alpha(X)) = \langle\!\lbrack \alpha(X) \rbrack\!\rangle_{\mathscr{D}_{n,\alpha(X)}}$. ∎

For the root n_0 of $\mathscr{R}$ we have $\mathscr{D}_{n_0,\varphi} = \mathscr{D}$. Using the assumption that $\neg\varphi$ is unprovable and the above observation we obtain an infinite path $\mathscr{P}$ of $\mathscr{R}$ such that for every node n of $\mathscr{P}$ the formula $\neg\bigwedge\{\langle\!\lbrack \beta \rbrack\!\rangle_{\mathscr{D}_{n,\beta}}: \beta \in L(n)\}$ is unprovable.

Let t be a token whose counter can be arbitrarily big on $\mathscr{P}$. Let X be a variable from the label of t and let $\mathscr{D}(X) = \mu X.\alpha(X)$ be its original definition. Because the counter of t is unbounded there must be two nodes n_1, n_2 such that (i) $L(n_1) = L(n_2)$, (ii) in both nodes the parts of the lists to the left of t are identical, (iii) t is labelled by $(\alpha(X), X)$, and (iv) the counter of t was increased and it was not reset between n_1 and n_2. Let us assume that n_2 is a descendant of n_1. We will show that $\neg\bigvee\{\langle\!\lbrack \gamma \rbrack\!\rangle_{\mathscr{D}_{n_2,\gamma}}: \gamma \in L(n_2)\}$ is provable.

As binding functions are established by (*) we have that

$$\mathscr{D}_{n_2,\delta} = \mathscr{D}_{n_1,\delta} \qquad \text{for every formula } \delta \in L(n_1), \quad \delta \neq \alpha(X). \tag{5}$$

This is because by (iii) and (iv) the counters of all the tokens to the right of t are 0 and all the counters to the left of t are the same in n_1 and n_2. Of course, the counter of t in n_1 is strictly smaller than in n_2.

We have:

$$\mathscr{D}_{n_1,\alpha(X)}(X) = \mu X.\neg\gamma_1 \wedge \cdots \wedge \neg\gamma_i \wedge \alpha(X)$$
$$\mathscr{D}_{n_2,\alpha(X)}(X) = \mu X.\neg\gamma_1 \wedge \cdots \wedge \neg\gamma_j \wedge \alpha(X),$$

where $j > i$ and formulas $\gamma_1, ..., \gamma_j$ are determined by the rule (*). We know that γ_i is $\bigwedge\{\langle\!\lbrack \delta \rbrack\!\rangle_{\mathscr{D}_{n_1,\delta}}: \delta \in L(n_1), \delta \neq \alpha(X)\}$ and by (5) it is the same as $\bigvee\{\langle\!\lbrack \delta \rbrack\!\rangle_{\mathscr{D}_{n_2,\delta}}: \delta \in L(n_2), \delta \neq \alpha(X)\}$. Finally we have that $\neg\bigvee\{\langle\!\lbrack \gamma \rbrack\!\rangle_{\mathscr{D}_{n_2,\gamma}}: \gamma \in L(n_2)\}$ is of the form

$$\neg\left(\neg\gamma_i \wedge \beta(\mu X.\neg\gamma_i \wedge \beta(X)) \wedge \bigwedge\{\langle\!\lbrack \delta \rbrack\!\rangle_{\mathscr{D}(n_2,\delta)}: \delta \neq \alpha(X), \delta \in L(n_2)\}\right)$$

which is just an instance of the propositional tautology $\neg(\neg\alpha \wedge \beta \wedge \alpha)$ — a contradiction with the choice of $\mathscr{P}$. ▮

6. TABLEAU CONSEQUENCE

In the completeness proof it will turn out that tableau equivalence is a too restrictive notion. Here we introduce a weaker notion called *tableau consequence*. It will turn out that this is a refinement of a semantic consequence relation. The main result of this section is Lemma 36, which shows that if a disjunctive formula δ has a tableau which is a consequence of a tableau for an aconjunctive formula α then $\alpha \leqslant \delta$ is provable.

DEFINITION 32 (Tableau consequence). Given a pair of tableaux $(\tilde{\mathscr{T}}, \mathscr{T})$, where $\tilde{\mathscr{T}} = \langle \tilde{\mathscr{T}}, \tilde{L} \rangle$ and $\mathscr{T} = \langle T, L \rangle$, we define a two player game $\mathscr{G}(\tilde{\mathscr{T}}, \mathscr{T})$ with the following rules.

1. The starting position is the pair of the roots of each tableaux.

2. Suppose a position of a play is $(\tilde{n}, n)$, both nodes being choice nodes of $\tilde{\mathscr{T}}$ and $\mathscr{T}$, respectively. Player I must choose a modal node $\tilde{m}$ near $\tilde{n}$ and player II must respond by choosing a modal node m near n. Node m must have the property that every literal and terminal formula from $L(m)$ appears in $\tilde{L}(\tilde{m})$.

3. Suppose a position of a play is $(\tilde{N}, N)$ with $\tilde{N}$, N being sets of choice nodes of $\tilde{\mathscr{T}}$ and $\mathscr{T}$, respectively. Player I must choose a modal node $\tilde{m}$ near some $\tilde{n} \in \tilde{N}$ and player II must respond with a modal node m near some $n \in N$, such that every literal and terminal formula from $L(m)$ appears in $\tilde{L}(\tilde{m})$.

4. Suppose a position consists of a pair of modal nodes $(\tilde{m}, m)$ from $\tilde{\mathscr{T}}$ and $\mathscr{T}$, respectively. Player I chooses some action a and has two possibilities afterwards. He can choose a $\langle a \rangle$-son n of m and player II then has to respond with a $\langle a \rangle$-son $\tilde{n}$ of $\tilde{m}$. Otherwise, player I can choose all $\langle a \rangle$-sons of m and player II must respond with the set of all $\langle a \rangle$-sons of $\tilde{m}$.

The game may end in a finite number of steps because one of the players cannot make a move. In this case, the opposite player wins. When the game has infinitely many steps we get as the result two infinite paths: $\tilde{\mathscr{P}}$ from $\tilde{\mathscr{T}}$ and $\mathscr{P}$ from $\mathscr{T}$. Player I wins if there is no μ-trace on $\tilde{\mathscr{P}}$ but there is a μ-trace on $\mathscr{P}$; otherwise, player II is the winner.

DEFINITION 33 (Strategy). A *strategy* $\mathscr{S}$ for the second player in the game $\mathscr{G}(\tilde{\mathscr{T}}, \mathscr{T})$ is a function assigning to a position consisting of two modal nodes $(\tilde{m}, m)$ and a son n of m a son $\mathscr{S}(\tilde{m}, n)$ of $\tilde{m}$ of the same type as n. If $(\tilde{n}, n)$ is a pair of choice nodes and $\tilde{m}$ is a modal node near $\tilde{n}$ then the strategy gives us a modal node $\mathscr{S}(\tilde{m}, n)$ near n. If a position consists of two sets $(\tilde{N}, N)$ then for every modal node $\tilde{m}$ near some $\tilde{n} \in \tilde{N}$ strategy $\mathscr{S}$ gives a modal node $\mathscr{S}(\tilde{m}, N)$ near some $n \in N$. A strategy is *winning* if it guarantees that player II wins no matter what the moves of player I are.

We will say that a tableau $\mathscr{T}$ is a *consequence* of a tableau $\tilde{\mathscr{T}}$ iff player II has a winning strategy in $\mathscr{G}(\tilde{\mathscr{T}}, \mathscr{T})$.

The definition of the game is based on the following intuition about tableaux. A tableau for a formula describes the semantics of a formula in an operational way. In order to satisfy formulas in a choice node n, we must provide a state which satisfies the label of one of the modal nodes near n. The sons of a modal node describe the transitions from a hypothetical state satisfying its label. Every $\langle a \rangle$-son describes an a-successor which is required. The set of all $\langle a \rangle$-sons puts a restriction on all possible a-successors of the node. In this way, a tableau of a formula describes all possible models of that formula.

The game is defined so that whenever player II has a winning strategy from a position $(\tilde{n}, n)$ then every model of the label of $\tilde{n}$, $\tilde{L}(\tilde{n})$, is also a model of the label of n, $L(n)$. If $\tilde{n}$ and n are both choice nodes then a model of $\tilde{L}(\tilde{n})$ must satisfy the label of one of the modal nodes near $\tilde{n}$. Hence, for every modal node near $\tilde{n}$ we must find a modal node near n whose label is implied by it. If $\tilde{n}$, n are modal nodes then every $\langle a \rangle$-son of n describes a state the existence of which is required in order to satisfy $L(n)$. We must show that existence of such a state is also required by $\tilde{L}(\tilde{n})$. The set of all the $\langle a \rangle$-sons represents general requirements, imposed by $L(n)$, on states reachable by action a. We must show that they are implied by the general requirements in $\tilde{L}(\tilde{n})$.

FACT 34. *Tableau consequence is transitive.*

Proof. If $\mathcal{T}_2$ is a consequence of $\mathcal{T}_1$ and $\mathcal{T}_3$ is a consequence of $\mathcal{T}_2$ then we can in some sense compose the winning strategies, $\mathcal{S}_{1,2}$ in $\mathcal{G}(\mathcal{T}_1, \mathcal{T}_2)$ and $\mathcal{S}_{2,3}$ in $\mathcal{G}(\mathcal{T}_2, \mathcal{T}_3)$, to obtain a winning strategy in $\mathcal{G}(\mathcal{T}_1, \mathcal{T}_3)$. For example, if the current position is a pair of choice nodes (n_1, n_3) then player I chooses a modal node m_1 near n_1. Strategy $\mathcal{S}_{1,2}$ gives us a modal node m_2 near n_2. We can consider m_2 as a move of player I in the game $\mathcal{G}(\mathcal{T}_2, \mathcal{T}_3)$. Hence the strategy $\mathcal{S}_{2,3}$ can give us a node n_3 near m_3. It is easy to check that the strategy defined in such a way is winning. ∎

The following lemma shows that tableau consequence is indeed weaker than equivalence.

LEMMA 35. *If two tableaux $\mathcal{T}_1$ and $\mathcal{T}_2$ are equivalent then $\mathcal{T}_1$ is a consequence of $\mathcal{T}_2$.*

Proof. Let $\mathcal{E}: \mathcal{T}_1 \to \mathcal{T}_2$ be an equivalence function. Consider the game $\mathcal{G}(\mathcal{T}_1, \mathcal{T}_2)$. The strategy for player II is to keep to positions of the form $(n, \mathcal{E}(n))$. The initial position is of this form. The strategy is defined by the following rules:

- If a position of a play is a pair of choice or modal nodes $(m, \mathcal{E}(m))$, then player I chooses some node n and player II replies by choosing $\mathcal{E}(n)$.
- If a position of a play is $(N, \mathcal{E}(N))$ with N being a set of choice nodes of $\mathcal{T}_1$ and $\mathcal{E}(N) = \{\mathcal{E}(n): n \in N\}$ then player I chooses a modal node m near some $n \in N$ and player II responds with the modal node $\mathcal{E}(m)$.

By the definition of the tableau equivalence this strategy is winning. ∎

In the next lemma we show how to use the fact that a tableau for a formula δ is a consequence of a tableau for a formula α to prove $\alpha \leqslant \delta$.

LEMMA 36. *Suppose that we have a weakly aconjunctive formula* α *and a disjunctive formula* δ. *If there is a tableau for* δ *which is a consequence of a tableau for* α *then* $\neg(\alpha \wedge \neg\delta)$ *is provable.*

Proof. Let $\mathscr{T}_\alpha = \langle T_\alpha, L_\alpha \rangle$ and $\mathscr{T}_\delta = \langle T_\delta, L_\delta \rangle$ be tableaux for α and δ, respectively, such that the second is a consequence of the first. Let $\mathscr{S}$ be a winning strategy for player II in the game $\mathscr{G}(\mathscr{T}_\alpha, \mathscr{T}_\delta)$. We will construct a thin refutation $\mathscr{R} = \langle T, L \rangle$ for $\alpha \wedge \neg\delta$.

To facilitate the construction we will define two correspondence functions $\mathscr{C}_\alpha$ and $\mathscr{C}_\delta$ which assign to every considered node of $\mathscr{R}$ (that is, not to all the nodes) a node of $\mathscr{T}_\alpha$ and $\mathscr{T}_\delta$, respectively. It will be always the case that:

$C1$ $\quad L(n) = L_\alpha(\mathscr{C}_\alpha(n)) \cup \{\neg \bigvee L_\delta(\mathscr{C}_\delta(n))\}$,

$C2$ $\quad$ strategy $\mathscr{S}$ is defined for the position $(\mathscr{C}_\alpha(n), \mathscr{C}_\delta(n))$.

Of course, the root of $\mathscr{R}$ will be labelled by $\{\alpha \wedge \neg\delta\}$. The next node, say m_0, will be labelled by $\{\alpha, \neg\delta\}$. We let $\mathscr{C}_\alpha(m_0)$ and $\mathscr{C}_\delta(m_0)$ be the roots of $\mathscr{T}_\alpha$ and $\mathscr{T}_\beta$, respectively. The next two observations show how to prolong $\mathscr{R}$.

OBSERVATION 36.1. *Suppose we have already constructed* $\mathscr{R}$ *up to a node* m; $\mathscr{C}_\alpha(m)$, $\mathscr{C}_\delta(m)$ *are choice nodes of appropriate tableaux and satisfy* $C1$, $C2$. *We can construct a finite part of* $\mathscr{R}$ *and define for each leaf* n *of the constructed part* $\mathscr{C}_\alpha(n)$ *and* $\mathscr{C}_\delta(n)$ *so that* (i) $\mathscr{C}_\delta(n) = \mathscr{S}(\mathscr{C}_\alpha(n), \mathscr{C}_\delta(m))$, (ii) *conditions* $C1$ *and* $C2$ *are satisfied, and* (iii) *traces from* m *to* n *are reflected. This last property means that the traces from* m *to* n *are exactly the traces from* $\mathscr{C}_\alpha(m)$ *to* $\mathscr{C}_\alpha(n)$ *with the exception of the trace from* $\neg\bigwedge L_\delta(\mathscr{C}_\delta(m))$ *to* $\neg\bigwedge L_\delta(\mathscr{C}_\delta(n))$, *which corresponds to the* (*unique*) *negated trace from* $\mathscr{C}_\delta(m)$ *to* $\mathscr{C}_\delta(n)$.

Proof. By assumption $L(m) = L_\alpha(\mathscr{C}_\alpha(m)) \cup \{\neg\gamma\}$ as $\mathscr{C}_\delta(m)$ is labelled by one formula because δ is a disjunctive formula. The idea of the construction is represented in Fig. 2.

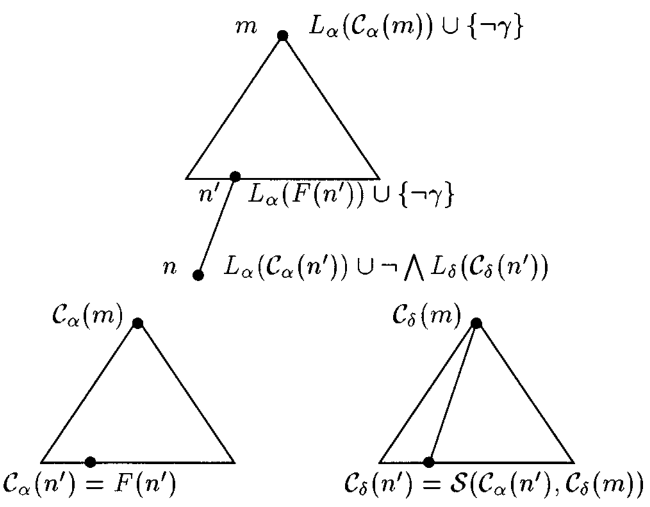

FIG. 2. Construction of $\mathscr{R}$.

From m we apply as long as possible rules other than $(\langle\ \rangle)$ and weakening to all the formulas in $L(m)$ except $\neg\gamma$. We apply them in the same order as they were applied from $\mathscr{C}_\alpha(m)$. This way we obtain a finite tree rooted in m. This tree is isomorphic to the part of $\mathscr{T}_\alpha$ between $\mathscr{C}_\alpha(m)$ and nearest modal nodes. Denoting this isomorphism F we have the property that for every leaf n' of this part $L(n') = L_\alpha(F(n')) \cup \{\neg\gamma\}$. Set $\mathscr{C}_\alpha(n') = F(n')$. Strategy $\mathscr{S}$ gives us a node $\mathscr{C}_\delta(n')$ which is a reply of player *II* to choosing $\mathscr{C}_\alpha(n')$ by player *I*. From the definition of the game it follows that $\mathscr{C}_\delta(n')$ is a modal node near $\mathscr{C}_\delta(m)$. Let us look at the path from $\mathscr{C}_\delta(m)$ to $\mathscr{C}_\delta(n')$ in $\mathscr{T}_\delta$. Because δ is a disjunctive formula, on this path first only (σ), (reg), and (or) rules may be applied and then we have zero or more applications of the (and) rule. Let us apply dual rules to $\neg\gamma$ (dual to (μ) is (ν), (reg) is self-dual). When it comes to an application of the (or) rule in $\mathscr{T}_\delta$, apply the (and) rule followed by weakening to leave only the conjunct which appears on the path to $\mathscr{C}_\delta(n')$. This way we make sure that the resulting tableau will be thin.

After these reductions we get a node n which is labelled by $L(\mathscr{C}_\alpha(n')) \cup \{\neg \bigwedge L_\delta(\mathscr{C}_\delta(n'))\}$. Setting $\mathscr{C}_\alpha(n) = \mathscr{C}_\alpha(n')$ and $\mathscr{C}_\delta(n) = \mathscr{C}_\delta(n')$ establishes the conditions C_1 and C_2. Finally trace reflection follows directly from the construction. ∎

OBSERVATION 36.2. *Suppose we have constructed $\mathscr{R}$ up to a node m. Assume that $\mathscr{C}_\alpha(m)$ and $\mathscr{C}_\delta(m)$ are modal nodes and $C1$, $C2$ are satisfied. We can construct a finite part of $\mathscr{R}$ and define $\mathscr{C}_\alpha(n)$, $\mathscr{C}_\delta(n)$ for every leaf n of this part in such a way that* (i) *position $(\mathscr{C}_\alpha(n), \mathscr{C}_\delta(n))$ is reachable from $(\mathscr{C}_\alpha(m), \mathscr{C}_\delta(m))$ when player II plays according to $\mathscr{S}$,* (ii) *conditions $C1$, $C2$ are satisfied, and* (iii) *traces are reflected.*

Proof. Let $\gamma = \bigwedge L_\delta(\mathscr{C}_\delta(m)) = \bigwedge\{\gamma_1, ..., \gamma_l\}$ and $\Gamma = L_\alpha(\mathscr{C}_\alpha(m))$. By $C1$ we have $L(m) = \{\neg\gamma\} \cup \Gamma$. Node $\mathscr{C}_\delta(m)$ is a modal node; hence, every γ_i is either a literal or a formula of the form $(a \to \Phi)$. When we negate γ we obtain a disjunction of negations of such formulas. Let us apply the (or) rule to eliminate these disjunctions. This way we obtain new leaves $m_1, ..., m_l$. For every $i \in 1, ..., l$ node m_i is labelled by $\{\neg\gamma_i\} \cup \Gamma$. We use γ_i to decide what rule to apply in m_i.

If γ_i is a literal or a terminal formula then we are done because γ_i appears in Γ. This follows directly from $C2$ and the definition of the game.

If γ_i is of the form $(a \to \Phi)$ with $\Phi \neq \emptyset$ then negated it becomes

$$\bigvee\{[a]\neg\varphi : \varphi \in \Phi\} \vee \langle a\rangle \bigwedge\{\neg\varphi : \varphi \in \Phi\}$$

or rather the translation of this formula to $(a \to \theta)$ notation. We apply disjunction rules as long as possible. This way we obtain a part of a tree. Each leaf u of this part is labelled by Γ and one of these disjuncts.

- Suppose this disjunct is $(a \to \emptyset)$. As $\Phi \neq \emptyset$, there is a $\langle a\rangle$-son of $\mathscr{C}_\delta(m)$ so there is, by the definition of a strategy, a $\langle a\rangle$-son of $\mathscr{C}_\alpha(m)$. Hence there is $(a \to \theta) \in \Gamma$ with $\theta \neq \emptyset$. Apply the $(\langle\,\rangle)$ rule to $(a \to \theta)$ and some $\psi \in \theta$. Then because $(a \to \emptyset) \in L(u)$ the son will contain $\bot$.

- If it is $(a \to \{\neg\varphi\})$ for some $\varphi \in \Phi$ then let us consult the strategy in the case when player *I* chooses the $\langle a\rangle$-son u_δ of $\mathscr{C}_\delta(m)$ labelled by $\{\varphi\}$. Strategy $\mathscr{S}$ gives us in this case $\langle a\rangle$-son u_α of $\mathscr{C}_\alpha(m)$. This son is labelled by

$\{\psi\} \cup \{\bigvee \theta\colon (a \to \theta) \in \Gamma, \theta \neq \Psi\}$ for some $(a \to \Psi) \in \Gamma$ and $\psi \in \Psi$. We apply the $(\langle\ \rangle)$ rule to $(a \to \Psi)$ in $L(u)$ and obtain a son u' of u labelled $\{\psi\} \cup \{\bigvee \theta\colon (a \to \theta) \in \Gamma, \theta \neq \Psi\} \cup \{\neg\varphi\}$. We let $\mathscr{C}_\alpha(u') = u_\alpha$ and $\mathscr{C}_\delta(u') = u_\delta$.

- If it is $(a \to \{\bigwedge\{\neg\varphi\colon \varphi \in \Phi\}, \top\})$ then we apply the $(\langle\ \rangle)$ rule to this formula and obtain a son u' of u labelled

$$\left\{\bigwedge \{\neg\varphi\colon \varphi \in \Phi\}\right\} \cup \left\{\bigvee \theta\colon (a \to \theta) \in \Gamma\right\}.$$

The construction from this point is presented in Fig. 3.

Let us choose one formula $\bigvee \Psi \in \{\bigvee\theta\colon (a \to \theta) \in \Gamma\}$ and apply (or) rules to it. This way we obtain a part of $\mathscr{R}$ each leaf of which is labelled by the set

$$\left\{\bigwedge\{\neg\varphi\colon \varphi \in \Phi\}\right\} \cup \{\psi\} \left\{\bigvee \theta\colon (a \to \theta) \in \Gamma, \theta \notin \Psi\right\}$$

for some $\psi \in \Psi$. Let o be one of such leaves and let o_α be a $\langle a\rangle$-son of $\mathscr{C}_\alpha(m)$ labelled by $\{\psi\} \cup \{\bigvee \theta\colon (a \to \theta) \in \Gamma, \theta \neq \Psi\}$. As in the previous observation apply rules other than (mod) and weakening to this set to obtain a finite part of a tree isomorphic to the part of $\mathscr{T}_\alpha$ between o_α and the nearest modal nodes. Let F denote this isomorphism. Let n' be a leaf of this part. We have $L(n') = L_\alpha(F(n')) \cup \bigwedge\{\neg\varphi\colon \varphi \in \Phi\}$. Now it is time to consult the strategy.

Let player I choose N_δ, the set of all $\langle a\rangle$-sons of $\mathscr{C}_\delta(m)$. Player II responds with N_α being the set of $\langle a\rangle$-sons of $\mathscr{C}_\alpha(m)$. For every modal node n_α near $o_\alpha \in N_\alpha$ the strategy $\mathscr{S}$ gives us a modal node $n_\delta = \mathscr{S}(n_\alpha, N_\delta)$. Let $n_\delta = \mathscr{S}(F(n'), N_\delta)$ and let o_δ be a choice node on the path to n_δ. We have $L_\delta(o_\delta) = \{\varphi\}$ for some $\varphi \in \Phi$. Apply

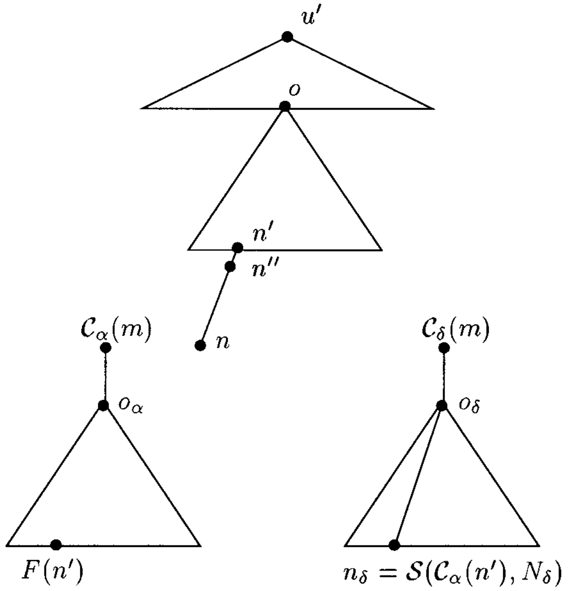

FIG. 3. Construction a part of $\mathscr{R}$.

the (*and*) rule followed by a contraction to obtain a node n'' labelled $\{\neg\varphi\} \cup L_\alpha(F(n'))$. Then reduce $\neg\varphi$ in n'' as in the proof of the previous observation. We arrive at a node n labelled $\{\neg\bigwedge L_\delta(n_\delta)\} \cup L_\alpha(F(n'))$. Let $\mathscr{C}_\alpha(n) = F(n')$ and $\mathscr{C}_\delta(n) = n_\delta$. ∎

The above two observations describe $\mathscr{R}$ completely. All the leaves are labelled by sets containing $\bot$ or some literal and its negation. For every infinite path $\mathscr{P}$ we have two possibilities. There may be a μ-trace on a path of $\mathscr{T}_\alpha$ designated by the image of $\mathscr{P}$ under $\mathscr{C}_\alpha$. In this case, by trace reflection, there is also a μ-trace on $\mathscr{P}$. If there is no μ-trace on $\mathscr{C}_\alpha(\mathscr{P})$ then there cannot be a μ-trace on $\mathscr{C}_\delta(\mathscr{P})$ because we were choosing our moves according to the strategy $\mathscr{S}$. Hence there is a ν-trace on $\mathscr{C}_\delta(\mathscr{P})$ which negated in $\mathscr{R}$ becomes a μ-trace.

This shows that $\mathscr{R}$ is a refutation. $\mathscr{R}$ is also a thin refutation because α is a weakly aconjunctive formula and whenever we reduce a conjunction coming from $\neg\delta$ we leave only one of the conjuncts. Hence by Theorem 31 the formula $\neg(\alpha \wedge \neg\delta)$ is provable. ∎

7. COMPLETENESS

Our main goal is:

THEOREM 37 (Completeness). *For every unsatisfiable formula φ formula $\neg\varphi$ is provable.*

Having Theorem 24 to prove completeness it is enough to show that for every unsatisfiable formula φ there is a disjunctive unsatisfiable formula $\hat{\varphi}$ such that $\varphi \leqslant \hat{\varphi}$ is provable. Of course we could just take $\hat{\varphi}$ to be $\bot$ but then the proof of this fact would be exactly as difficult as showing completeness. So in general we will look for more complicated formulas than $\bot$. Because we will prove this fact by induction on $\hat{\varphi}$ we clearly need to consider also satisfiable formulas. From these considerations it follows that we need:

THEOREM 38. *For every positive, guarded formula φ there is a semantically equivalent disjunctive formula $\hat{\varphi}$ such that $\varphi \leqslant \hat{\varphi}$ is provable. Moreover, if a variable occurs only positively in φ then it occurs only positively in $\hat{\varphi}$.*

Before proving this theorem let us show how to use it in the completeness proof.

Proof (Completeness). Let φ be an unsatisfiable formula. By Proposition 23 we may assume that φ satisfies our proviso from Section 2. From Theorem 38 it follows that there is a disjunctive formula $\hat{\varphi}$ equivalent to φ and $\varphi \leqslant \hat{\varphi}$ is provable. Hence it is enough to show that $\neg\varphi$ is provable. But this follows from Theorem 24. ∎

The rest of this section is devoted to the proof of Theorem 38.

Proof (Theorem 38). The proof is by induction on the structure of the formula φ.

Case: *φ is a literal.* In this case $\hat{\varphi}$ is just φ.

Case: $\varphi = \alpha \vee \beta$. By the induction assumption there are disjunctive formulas $\hat{\alpha}$, $\hat{\beta}$ equivalent to α and β, respectively. We let $\widehat{\alpha \vee \beta}$ be $\hat{\alpha} \vee \hat{\beta}$. Because $\alpha \leqslant \hat{\alpha}$ and $\beta \leqslant \hat{\beta}$ are provable, $\alpha \vee \beta \leqslant (\hat{\alpha} \vee \hat{\beta})$ is also provable.

Case: $\varphi = (a \rightarrow \Phi)$. This case is very similar to the previous one.

Case: $\varphi = \mu X.\alpha(X)$. The proof of this case will take a significant part of this section. Fortunately the tools developed here can be also used for the remaining cases.

By the induction assumption there is a disjunctive formula $\hat{\alpha}(X)$ equivalent to $\alpha(X)$. It is easy to see that $\mu X.\alpha(X)$ is semantically equivalent to $\mu X.\hat{\alpha}(X)$ and $\mu X.\alpha(X) \leqslant \mu X.\hat{\alpha}(X)$ is provable. Unfortunately $\mu X.\hat{\alpha}(X)$ may not be a disjunctive or even a weakly aconjunctive formula. This is because X may occur in a context $X \wedge \gamma$ for some γ. Therefore we have to construct $\hat{\varphi}$ from scratch.

By Theorem 22 there is a disjunctive formula $\hat{\varphi}$ which has the tableau equivalent to some tableau $\mathscr{T}$ for $\mu X.\hat{\alpha}(X)$. By Theorem 19 the two formulas are equivalent. We are left to show that $\mu X.\hat{\alpha}(X) \leqslant \hat{\varphi}$ is provable in Kozen's system. To do this it is enough to prove $\hat{\alpha}(\hat{\varphi}) \leqslant \hat{\varphi}$ and then use the rule (K6). Now, it is possible to show that if α and δ have equivalent tableaux, α is weakly aconjunctive, and δ is disjunctive then $\alpha \leqslant \delta$ is provable. Unfortunately the notion of tableau equivalence is too strong for us because there may be no tableau for $\hat{\alpha}(\hat{\varphi})$ equivalent to a tableau for $\hat{\varphi}$. It turns out that tableau consequence is what we need.

LEMMA 39. *The tableau $\mathscr{T}$ for $\mu X.\hat{\alpha}(X)$ is a consequence of the tableau $\hat{\mathscr{T}}$ for $\hat{\alpha}(\hat{\varphi})$.*

Proof. Let $\mathscr{T} = \langle T, L \rangle$, $\tilde{\mathscr{T}} = \langle \tilde{T}, \tilde{L} \rangle$ and let $\hat{\mathscr{T}} = \langle \hat{T}, \hat{L} \rangle$ be the tableau for φ. Recall that $\hat{\mathscr{T}}$ was constructed from $\mathscr{T}$ using Theorem 22. Hence we can assume that $\hat{\mathscr{T}}$ satisfies the properties from Observation 22.2. As $\mathscr{T}$ and $\hat{\mathscr{T}}$ are equivalent we have an equivalence function $\mathscr{E}\colon \mathscr{T} \rightarrow \hat{\mathscr{T}}$. Because the tableaux for $\mu X.\hat{\alpha}(X)$ and $\hat{\alpha}(\mu X.\hat{\alpha}(X))$ differ just by one application of the fixpoint rule in the root we will denote by $\mathscr{T}$ also the tableau for $\hat{\alpha}(\mu X.\hat{\alpha}(X))$.

By assumption $\hat{\varphi}$ and $\hat{\alpha}(X)$ are disjunctive formulas. We will use $\beta[\mu X.\hat{\alpha}(X)/\hat{\varphi}]$ and $\beta[\hat{\varphi}/\mu X.\hat{\alpha}(X)]$ to stand for the obvious replacements; it will be always the case that no free variable in $\mu X.\hat{\alpha}(X)$ or $\hat{\varphi}$ is bound by the context β. From Fact 27 we obtain:

OBSERVATION 39.1. *For every node $\tilde{n}$ of $\tilde{\mathscr{T}}$, every formula in $\tilde{L}(\tilde{n})$ is either a disjunctive formula or of the form $\delta(\hat{\varphi})$ with $\delta(X)$ being a disjunctive formula.*

As the first step, for every mode $\tilde{m}$ of $\tilde{\mathscr{T}}$ we will define two functions:

$$p_{\tilde{m}}\colon \tilde{L}(\tilde{m}) \rightarrow \mathbb{N} \cup \{\infty\} \qquad \widehat{nd}_{\tilde{m}}\colon \mathrm{Ran}(p_{\tilde{m}}) \cap \mathbb{N}\mathscr{P}(\tilde{\mathscr{T}}).$$

The first function assigns a priority which is a natural number or ∞ to every formula in $\tilde{L}(\tilde{m})$. The function $\widehat{nd}_{\tilde{m}}$ assigns sets of nodes of $\hat{\mathscr{T}}$ to finite priorities in the range of $p_{\tilde{m}}$. Sometimes we will identify a singleton set $\{m\}$ with the element m. For example, we will write $\hat{L}(\widehat{nd}_{\tilde{m}}(q))$ when $\widehat{nd}_{\tilde{m}}(q)$ is a singleton.

These two functions will satisfy the following condition which we call $I1$.

- if $\widehat{nd}_{\tilde{m}}(q)$ is a singleton then

$$p_{\tilde{m}}^{-1}(q) \subseteq \hat{L}(\widehat{nd}_{\tilde{m}}(q)) \subseteq \bigcup_{q' \leqslant q} p_{\tilde{m}}^{-1}(q');$$

- if $\widehat{nd}_{\tilde{m}}(q)$ is not a singleton then

$$p_{\tilde{m}}^{-1}(q) = \left\{ \bigvee \{\psi \colon \{\psi\} = \hat{L}(m), m \in \widehat{nd}_{\tilde{m}}(q)\} \right\}; \qquad \text{(I1)}$$

- if $\tilde{m}$ is a modal node then $\widehat{nd}_{\tilde{m}}(q)$ is a singleton for all $q \in \operatorname{Ran}(p_{\tilde{m}}) \cap \mathbb{N}$

The idea behind these two functions comes from considering $\tilde{\mathscr{T}}$ to be some kind of composition of $\mathscr{T}$ and many copies of $\hat{\mathscr{T}}$. To see what we mean consider a part of a path of $\tilde{\mathscr{T}}$ which is represented in the middle of Fig. 4. To the left of it we have put a corresponding path of $\mathscr{T}$ and to the right we have represented a part of $\hat{\mathscr{T}}$. Arrows represent traces.

The label of a node of $\tilde{\mathscr{T}}$ can be divided into a set of formulas to which there is no trace going through φ and the rest which have such a trace. Every formula γ of the first kind corresponds to a formula $\gamma[\mu X.\hat{\alpha}(X)/\hat{\varphi}]$ of $\mathscr{T}$. These formulas will have priority ∞. In our figure they are represented by Γ with indices. For a formula δ of the second type, there is the earliest occurrence of $\hat{\varphi}$ from which there is a trace to δ. This occurrence determines the priority of the formula and the whole trace determines the node of $\hat{\mathscr{T}}$. We use priorities when contraction occurs. Consider, for example, a situation represented in the last two nodes of the part of the path in Fig. 4. Formula δ_1^1 is reduced and becomes the same as δ_2^0. Nevertheless the node associated with δ_2^0 may be different than the son of the node associated with δ_1^1 because the histories of reductions of these formulas could have been different. The priority tells us that the path from the son of δ_1^1 should not be followed but the path from the left occurrence of δ_2^0 should continue. We can say that the trace

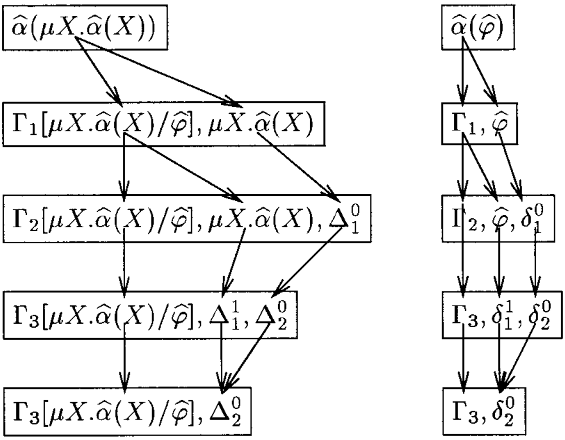

FIG. 4. Decomposition of a path of $\tilde{\mathscr{T}}$.

jumps from δ_1^1 to the left occurrence of δ_2^0. Arranging formulas this way we have that every trace on a path of $\tilde{\mathcal{T}}$ is either some trace on the corresponding path of $\mathcal{T}$ or is eventually (after finitely many jumps) a trace on a path of $\hat{\mathcal{T}}$ designated by the $\widehat{nd}$ function.

The functions $p_{\tilde{m}}$ and $\widehat{nd}_{\tilde{m}}$ will be defined by simultaneous induction on the distance of $\tilde{m}$ from the root. For the root $\tilde{r}$ of $\tilde{\mathcal{T}}$ we let $p_{\tilde{r}}(\hat{\alpha}(\hat{\varphi}))=\infty$. The induction step is handled by the following two observations.

Observation 39.2. *Suppose $\tilde{m}\in\tilde{T}$ is a modal node, both $p_{\tilde{m}}$ and $nd_{\tilde{m}}$ are defined and satisfy condition $I1$. For every son $\tilde{n}$ of $\tilde{m}$ we can define $p_{\tilde{n}}$ and $nd_{\tilde{n}}$ so that $I1$ will be satisfied.*

Proof. Let $\{(a\to\theta_1), ..., (a\to\theta_k)\}$ be all the formulas from $\tilde{L}(\tilde{m})$ having the form $(a\to\theta)$.

Let $\tilde{n}$ be a $\langle a\rangle$-son of $\tilde{m}$ and to keep indexing simple say it is labelled by $\{\varphi_1\}\cup\{\bigvee\theta_2, ..., \bigvee\theta_k\}$ for some $\varphi_1\in\theta_1$.

If $\varphi_1\neq\bigvee\theta_i$ for all $i=2, ..., k$ then set $p_{\tilde{n}}(\bigvee\theta_i)=p_{\tilde{m}}(a\to\theta_i)$ and let $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\bigvee\theta_i))$ be the set of all $\langle a\rangle$-sons of $nd_{\tilde{m}}(p_{\tilde{m}}(a\to\theta_i))$. For φ_1 let $p_{\tilde{n}}(\varphi_1)=p_{\tilde{m}}(a\to\theta_1)$ and let $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\varphi_1))$ be the son of $\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(a\to\theta_1))$ which is labelled by $\{\varphi_1\}$. Of course we set $\widehat{nd}_{\tilde{n}}$ only whenever the priority is finite.

Suppose now that $\varphi_1=\bigvee\theta_i$ for some $i=2, ..., k$. We must decide whether to treat this formula as φ_1 or as $\bigvee\theta_i$.

- If $p_{\tilde{m}}(a\to\theta_i)<p_{\tilde{m}}(a\to\theta_1)$ then let $p_{\tilde{n}}(\bigvee\theta_i)=p_{\tilde{m}}(a\to\theta_i)$ and let $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\bigvee\theta_i))$ be the set of all the $\langle a\rangle$-sons of $\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(a\to\{\theta_i\}))$.

- If $p_{\tilde{m}}(a\to\theta_i)>p_{\tilde{m}}(a\to\theta_1)$ then let $p_{\tilde{n}}(\varphi_1)=p_{\tilde{m}}(a\to\theta_1)$ and let $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\varphi_1))$ be the son of $\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(a\to\theta_1))$ which is labelled by φ_1.

With all $\bigvee\theta_j$ for $j=2, ..., k$, $j\neq i$ we proceed as before. ∎

Observation 39.3. *Suppose $\tilde{m}\in\tilde{T}$ is not a modal node, both $p_{\tilde{m}}$ and $nd_{\tilde{m}}$ are defined and satisfy condition $I1$. For every son $\tilde{n}$ of $\tilde{m}$ we can define $p_{\tilde{n}}$ and $nd_{\tilde{n}}$ so that $I1$ will be satisfied.*

Proof. If $\tilde{m}$ is not a modal node then only one formula, say β, is reduced by the rule applied in $\tilde{m}$. Let $q=p_{\tilde{m}}(\beta)$. Let $\tilde{n}$ be a son of $\tilde{m}$ and let $\gamma\in L(\tilde{n})$ be one of the formulas obtained by reducing β. We have several cases depending on the type of formula β.

- If $\widehat{nd}_{\tilde{m}}(q)$ is not a singleton then $\beta=\bigvee\{\psi\colon\{\psi\}=\hat{L}(\hat{m}), \hat{m}\in\widehat{nd}_{\tilde{m}}(q)\}$. In this case γ is a disjunct of β. Let $p_{\tilde{n}}(\gamma)=q$ and let $\widehat{nd}_{\tilde{n}}(q)$ be an appropriate subset of $\widehat{nd}_{\tilde{m}}(q)$. For every $\delta\in\tilde{L}(\tilde{n})$, $\delta\neq\gamma$ let $p_{\tilde{n}}(\delta)=p_{\tilde{m}}(\delta)$ and $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\delta))=\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(\delta))$.

- If $\beta=\mu X.\hat{\alpha}(X)$ then we let $p_{\tilde{n}}(\gamma)$ be the smallest priority not in the range of $p_{\tilde{m}}$ and set $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\gamma))$ to the root of $\hat{\mathcal{T}}$. For every $\delta\in\tilde{L}(\tilde{n})$, $\delta\neq\gamma$ we proceed as before.

- If not the previous cases, $\gamma\in\tilde{L}(\tilde{m})$ and $p_{\tilde{m}}(\gamma)>p_{\tilde{m}}(\beta)$ then for every $\delta\in\tilde{L}(\tilde{n})$ let $p_{\tilde{n}}(\delta)=p_{\tilde{m}}(\delta)$ and $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\delta))=\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(\delta))$.

- If not the previous cases, $\gamma \notin \tilde{L}(\tilde{m})$ or $p_{\tilde{m}}(\gamma) \leqslant p_{\tilde{m}}(\beta)$ then $p_{\tilde{n}}(\gamma) = p_{\tilde{m}}(\beta)$ and $\widehat{nd}_{\tilde{n}}(p_{\tilde{n}}(\gamma))$ is the son of $\widehat{nd}_{\tilde{m}}(p_{\tilde{m}}(\beta))$ containing γ. For all $\delta \in \tilde{L}(\tilde{n})$, $\delta \neq \gamma$ we proceed as in the first case. ∎

The next step is to define a winning strategy in the game $\mathscr{G}(\tilde{\mathscr{T}}, \mathscr{T})$. We will write $nd_{\tilde{m}}(q)$ for $\mathscr{E}^{-1}(\widehat{nd}_{\tilde{m}}(q))$ (recall that $\mathscr{E}: \mathscr{T} \to \tilde{\mathscr{T}}$ is the equivalence function). All positions $(\tilde{m}, m)$ reachable in a game played according to the strategy will have the following property.

If $\tilde{m}$ is a choice or modal node then (whenever defined) $nd_{\tilde{m}}(q)$ is either a singleton or a set of all the $\langle a \rangle$-sons of some node. Let $L_{\tilde{m}}(q)$ stand for $L(n)$ if $\{n\} = nd_{\tilde{m}}(q)$ and let $L_{\tilde{m}}(q) = \{\bigvee \theta_1, ..., \bigvee \theta_k\}$ if $nd_{\tilde{m}}(q)$ is the set of all $\langle a \rangle$-sons of some node m and $\{(a \to \theta_1), ..., (a \to \theta_n)\} \subseteq L(m)$ is the set of all the formulas of the form $(a \to \theta)$ in $L(m)$. With this definition we have the property

$$L(m) \subseteq \bigcup_{q \in \operatorname{Ran}(p_{\tilde{m}}) \cap \mathbb{N}} L_{\tilde{m}}(q) \cup \{\psi[\mu X.\hat{\alpha}(X)/\hat{\varphi}]: \psi \in p_{\tilde{m}}^{-1}(\infty)\}. \tag{I2}$$

Condition $I2$ allows us to define a function $p_{\tilde{m}, m}: L(m) \to \mathbb{N} \cup \{\infty\}$. For every $\gamma \in L(m)$ let $p_{\tilde{m}, m}(\gamma)$ be the smallest priority q such that $\gamma \in L_{\tilde{m}}(q)$; in case there is no such q let $p_{\tilde{m}, m}(\gamma) = \infty$.

The strategy is described in the next three observations.

Observation 39.4. *Assume the game is in a position $(\tilde{m}, m)$ consisting of two modal nodes and the condition $I2$ is satisfied. Suppose player I chooses a $\langle a \rangle$-son n of m. We can find a $\langle a \rangle$-son $\tilde{n}$ of $\tilde{m}$ so that* (i) *condition $I2$ will be satisfied and* (ii) *if $\beta \in L(n)$ is obtained from $\alpha \in L(m)$ then $p_{\tilde{n}, m}(\beta) \leqslant p_{\tilde{m}, m}(\alpha)$.*

Proof. Let n be a $\langle a \rangle$-son of m. It is labelled by

$$\{\xi\} \cup \left\{\bigvee \theta: (a \to \theta) \in L(m), \theta \neq \Xi\right\}$$

for some $(a \to \Xi) \in L(m)$ and $\xi \in \Xi$. Let $q = p_{\tilde{m}, m}(a \to \Xi)$. If $q = \infty$ we take a son of $\tilde{m}$ labelled

$$\{\xi[\hat{\varphi}/\mu.\hat{\alpha}(X)]\} \cup \left\{\bigvee \theta: (a \to \theta) \in \tilde{L}(\tilde{m}), \theta \neq \Xi[\hat{\varphi}/\mu X.\hat{\alpha}(X)]\right\}.$$

The case when $q \in \mathbb{N}$ is represented in Fig. 5.

$$\begin{tikzcd} m \arrow[d] & & nd_{\tilde{m}}(q) \arrow[d] & & \mathcal{E}(nd_{\tilde{m}}(q)) \arrow[d] & & \tilde{m} \arrow[d] \\ n & \Longrightarrow & n' & \overset{\mathcal{E}}{\Longrightarrow} & \mathcal{E}(n') & \Longrightarrow & \tilde{n} \end{tikzcd}$$

FIG. 5. Finding a $\langle a \rangle$-son of $\tilde{m}$ for a $\langle a \rangle$-son of m.

By $I2$ there is a $\langle a\rangle$-son n' of $nd_{\tilde{m}}(q)$ labelled

$$\{\xi\}\cup\left\{\bigvee\theta\colon (a\to\theta)\in L(nd_{\tilde{m}}(q)),\,\theta\neq\Xi\right\}.$$

Using equivalence $\mathscr{E}$ we get $\langle a\rangle$-son $\mathscr{E}(n')$ of $\widehat{nd}_{\tilde{m}}(q)=\mathscr{E}(nd_{\tilde{m}}(q))$. It is labelled by some $\{\psi\}$ for $\psi\in\Psi$ and $(a\to\Psi)\in\hat{L}(\widehat{nd}_{\tilde{m}}(q))$. By $I1$ we can take as $\tilde{n}$ a $\langle a\rangle$-son of $\tilde{m}$ labelled

$$\{\psi\}\cup\left\{\bigvee\theta\colon (a\to\theta)\in\tilde{L}(\tilde{m}),\,\theta\neq\Psi\right\}.$$

It is quite straightforward to show that (i) and (ii) are satisfied if we use the fact that for every two choice nodes $\hat{n}_1$, $\hat{n}_2$ of $\hat{\mathscr{T}}$, whenever $\hat{L}(\hat{n}_1)=\hat{L}(\hat{n}_2)$ then $L(\mathscr{E}^{-1}(\hat{n}_1))=L(\mathscr{E}^{-1}(\hat{n}_2))$. We can assume this property by Observation 22.2. ∎

Observation 39.5. *Assume the game is in a position $(\tilde{m}, m)$ consisting of two choice nodes and the condition $I2$ is satisfied. Suppose player I chooses a modal node $\tilde{n}$ near $\tilde{m}$. We can find a modal node n near m so that* (i) *the condition $I2$ will be satisfied and* (ii) *the traces from m to n will be preserved. Preservation of traces means that whenever there is a trace from $\alpha\in L(m)$ to $\beta\in L(n)$ and Y is the smallest, in $\leqslant_{\mu X.\hat{\alpha}(X)}$ ordering, variable regenerated on the trace then either*:

- *$p_{\tilde{m},m}(\alpha)>p_{\tilde{n},n}(\beta)$ or*
- *$p_{\tilde{m},m}(\alpha)=p_{\tilde{n},n}(\beta)=q$ and when $q\in\mathbb{N}$ there is a trace from $\alpha\in L(nd_{\tilde{m}}(q))$ to $\beta\in L(nd_{\tilde{n}}(q))$ or when $q=\infty$ there is a trace from $\alpha[\hat{\varphi}/\mu X.\hat{\alpha}(X)]\in L(\tilde{m})$ to $\beta[\hat{\varphi}/\mu X.\hat{\alpha}(X)]\in L(\tilde{n})$. In both cases Y is the smallest regenerated variable on the trace.*

Proof. We will find n with the required properties by constructing a path to it from m. In some sense n is determined by $\tilde{n}$ and all $nd_{\tilde{n}}(q)$ for $q\in\mathbb{N}$. For every $q\in\operatorname{Ran}(p_{\tilde{m}})\cap\mathbb{N}$ let $\theta_q=\{\psi\colon\{\psi\}=\hat{L}(\hat{m}),\,\hat{m}\in\widehat{nd}_{\tilde{m}}(q)\}$. By $I1$ we have

$$\tilde{L}(\tilde{m})=\left\{\bigvee\theta_q\colon q\in\operatorname{Ran}(p_{\tilde{m}})\cap\mathbb{N}\right\}\cup p_{\tilde{m}}^{-1}(\infty).$$

On the path to $\tilde{n}$ there is a node $\tilde{s}$ where exactly one disjunct is chosen from each θ_q. Say its label is

$$\tilde{L}(\tilde{s})=\{\psi_q\colon q\in\operatorname{Ran}(p_{\tilde{m}})\cap\mathbb{N}\}\cup p_{\tilde{m}}^{-1}(\infty)$$

for some $\psi_q\in\theta_q$, $(q\in\operatorname{Ran}(p_{\tilde{m}})\cap\mathbb{N})$. For this node we can define $\widehat{nd}_{\tilde{s}}(q)$ to be a node from $\widehat{nd}_{\tilde{m}}(q)$ labelled $\{\psi_q\}$. As before we define $nd_{\tilde{s}}(q)=\mathscr{E}^{-1}(\widehat{nd}_{\tilde{s}}(q))$.

Switching to the other tableau, by $I2$ we have

$$L(m)\subseteq\bigcup_{q\in\mathbb{N}}L_{\tilde{m}}(q)\cup\{\psi[\mu X.\hat{\alpha}(x)/\hat{\varphi}]\colon\psi\in p_{\tilde{m}}^{-1}(\infty)\}.$$

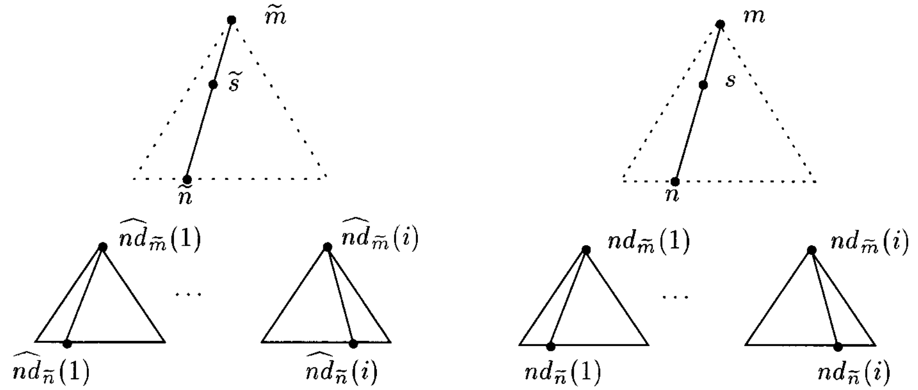

FIG. 6. Finding a modal node n for the modal node $\tilde{n}$.

If $nd_{\tilde{m}}(q)$ is a singleton then $\bigvee\theta_q = \psi_q$ and $L(nd_{\tilde{m}}(q)) = L(nd_{\tilde{s}}(q))$. Otherwise $L_{\tilde{m}}(q) = \{\bigvee\Delta_1, ..., \bigvee\Delta_k\}$ and $L(nd_{\tilde{s}}(q)) = \{\bigvee\Delta_1, ..., \delta_i, ...\bigvee\Delta_k\}$ for some $i = 1, ..., k$ and $\delta_i \in \Delta_i$. First apply (or) rules from m to obtain a node s such that

$$L(s) \subseteq \bigcup_{q \in \mathbb{N}} L(nd_{\tilde{s}}(q)) \cup \{\psi[\mu X.\hat{\alpha}(X)/\hat{\varphi}] \colon \psi \in p_{\tilde{m}}^{-1}(\infty)\}. \tag{6}$$

The obtained situation and the rest of the construction is represented in Fig. 6. From s we will construct a path choosing one node at the time. For every considered node o we will define a priority function $p_{\tilde{s}, o} \colon L(o) \to \mathbb{N} \cup \{\infty\}$. We will assume that for every considered node o and every $\psi \in L(o)$:

If $p_{\tilde{s}, o}(\psi) = \infty$ then $\psi[\varphi/\mu X.\alpha(X)]$ appears on the path from $\tilde{m}$ to $\tilde{n}$, otherwise if $p_{\tilde{s}, o}(\phi) = q \in \mathbb{N}$ then ψ appears on the path from $nd_{\tilde{m}}(q)$ to $nd_{\tilde{n}}(q)$ $\quad (I3)$

Function $p_{\tilde{s}, s}$ is defined using (6) by letting $p_{\tilde{s}, s}(\psi)$ be the smallest q such that $\varphi \in L(nd_{\tilde{s}}(q))$ or ∞ if there is no such q. Actually it may happen that s does not satisfy $I3$ or rather $I3$ does not make sense because $\mathrm{Ran}(p_{\tilde{n}}) \neq \mathrm{Ran}(p_{\tilde{m}})$. Let us extend $\widehat{nd}_{\tilde{m}}$ and $\widehat{nd}_{\tilde{n}}$. The only element which can appear in $\mathrm{Ran}(p_{\tilde{n}}) \backslash \mathrm{Ran}(p_{\tilde{m}})$ is the smallest priority q which does not appear in $\mathrm{Ran}(p_{\tilde{m}})$. We take care of this by extending the definition of $\widehat{nd}_{\tilde{m}}$ and letting $\widehat{nd}_{\tilde{m}}(q)$ be the root of $\hat{\mathscr{T}}$. Let $q \in \mathrm{Ran}(p_{\tilde{m}}) \backslash \mathrm{Ran}(p_{\tilde{n}})$ and let $\Gamma \subseteq \tilde{L}(\tilde{n})$ be the set of all the formulas to which there is a trace from the unique formula of priority q in $\tilde{L}(\tilde{m})$. Γ is the label of some modal node n near $\widehat{nd}_{\tilde{m}}(q)$. Let $\widehat{nd}_{\tilde{n}}(q) = n$. As $\Gamma \subseteq \bigcup\{\hat{L}(\widehat{nd}_{\tilde{n}}(q')) \colon q' \in \mathrm{Ran}(p_{\tilde{n}}), q' < q\}$ we know by Observation 22.2 that

$$nd_{\tilde{n}}(q) \subseteq \bigcup\{L(nd_{\tilde{n}}(q')) \colon q' \in \mathrm{Ran}(p_{\tilde{n}}), q' < q\}. \tag{7}$$

With these extensions, the condition $I3$ is satisfied for the node s and we may proceed with the construction of the path.

• If ψ is not a disjunction then there is only one son o' of o. Let ψ' be the result of reducing ψ. For every $\beta \in L(o')$, $\beta \neq \psi'$ we let $p_{\tilde{s},o'}(\beta) = p_{\tilde{s},o}(\beta)$. If $\psi' \notin L(o)$ then let $p_{\tilde{s},o'}(\psi') = p_{\tilde{s},o}(\psi)$; otherwise let $p_{\tilde{s},o'}(\psi') = \min\{p_{\tilde{s},o}(\psi), p_{\tilde{s},o}(\psi')\}$. If $\psi \neq \mu X.\hat{\alpha}(X)$ then function $p_{\tilde{s},o'}$ satisfies the condition $I3$. If $\psi = \mu X.\hat{\alpha}(X)$ then $\psi' = X$ and letting $p_{\tilde{s},o'}(X) = \infty$ would be unsound with respect to $I3$. We let $p_{\tilde{s},o'}(X)$ be the smallest priority q not in $\mathrm{Ran}(p_{\tilde{n}})$. This is sound as $nd_{\tilde{n}}(q)$ is a modal node near the root of $\mathscr{T}$.

• If $\psi = \alpha \vee \beta$ then o has two sons o_1, o_2 and we have to choose one of them. If $p_{\tilde{s},o}(\psi) = \infty$ then $\psi[\hat{\varphi}/\mu X.\hat{\alpha}(X)]$ is on the path from $\tilde{n}$ to $\tilde{m}$; otherwise ψ appears on the path from $nd_{\tilde{m}}(p_{\tilde{s},o}(\psi))$ to $nd_{\tilde{n}}(p_{\tilde{s},o}(\psi))$. We choose a son of o with the same disjunct as the one appearing on the appropriate path. For the chosen o' we define $p_{\tilde{s},o'}$ as in the case of an unary rule. It should be easy to check that for so-defined o' and $p_{\tilde{s},o'}$ the condition $I3$ holds.

Repeating this procedure we arrive at a modal node n near m. Let us check that condition $I2$ holds. Suppose $\psi \in L(n)$ and $q = p_{\tilde{s},n}(\psi)$. Because n is a modal node, ψ can be reducible only by application of (mod) rule. By $I3$ if $q = \infty$ then $\psi[\hat{\varphi}/\mu X.\hat{\alpha}(X)] \in \tilde{L}(\tilde{n})$; otherwise $q \in \mathbb{N}$ and $\psi \in L(nd_{\tilde{n}}(q))$. In the latter case either $q \in \mathrm{Ran}(p_{\tilde{n}})$ or by (7) we have $q' \in \mathrm{Ran}(p_{\tilde{n}})$, $q' \leqslant q$ with $\psi \in nd_{\tilde{n}}(q')$.

Finally it is easy to see that the traces are preserved. ∎

OBSERVATION 39.6. *Suppose a position in the game is* $(\tilde{N}, N)$ *for* $\tilde{N}$ *being a set of all the* $\langle a \rangle$*-sons of some node* $\tilde{m}$ *and* N *being a set of all the* $\langle a \rangle$*-sons of some node* m. *Suppose also that* $I2$ *holds for the pair* $(\tilde{m}, m)$. *For every modal node* $\tilde{o}$ *near some* $\tilde{n} \in \tilde{N}$ *we can find a modal node* o *near some* $n \in N$ *so that* (i) *condition* $I2$ *holds for position* $(\tilde{o}, o)$ *and* (ii) *the traces from* m *to* o *are preserved.*

Proof. For every $q \in \mathrm{Ran}(p_{\tilde{n}}) \cap \mathbb{N}$ let $\theta_q = \bigvee\{\psi : \{\psi\} = L(\tilde{n}), \tilde{n} \in \widehat{nd}_{\tilde{n}}(q)\}$. By $I1$ we have

$$\tilde{L}(\tilde{n}) = \left\{\bigvee \theta_q : q \in \mathrm{Ran}(p_{\tilde{n}}) \cap \mathbb{N}\right\} \cup p_{\tilde{n}}^{-1}(\infty).$$

On the way to $\tilde{o}$ we reach a node $\tilde{s}$ where exactly one formula ψ_q is chosen from each θ_q.

$$\tilde{L}(\tilde{s}) = \{\psi_q : q \in \mathrm{Ran}(p_{\tilde{m}}) \cap \mathbb{N}\} \cup p_{\tilde{n}}^{-1}(\infty)$$

For this node we can define $\widehat{nd}_{\tilde{s}}(q)$ to be a node from $\widehat{nd}_{\tilde{m}}(q)$ labelled $\{\psi_q\}$. As before we define $nd_{\tilde{s}}(q) = \mathscr{E}^{-1}(\widehat{nd}_{\tilde{s}}(q))$.

Now $L(nd_{\tilde{s}}(q)) = \{\delta_1\} \cup \{\bigvee \Delta_2, ..., \bigvee \Delta_k\}$ and $L(nd_{\tilde{n}}(q))$ is either the same set of formulas or it is $\{\bigvee \Delta_1, ..., \bigvee \Delta_k\}$ for some $\Delta_1 \ni \delta_1$.

We can find a choice node $n \in N$ and a descendant s of n such that

$$L(s) \subseteq \bigcup_{q \in \mathbb{N}} L(nd_{\tilde{s}}(q)) \cup \{\psi[\mu X.\hat{\alpha}(X)/\hat{\varphi}] : \psi \in p_{\tilde{m}}^{-1}(\infty)\}.$$

From this point we can repeat the arguments from the previous observation. ∎

Finally we show that the defined strategy is winning. Let us take some play of $\mathcal{G}(\tilde{\mathcal{T}}, \mathcal{T})$ where II plays according to the strategy. By the three observations above player II can always make a move so II cannot lose in a finite number of steps. Assume that the play was infinite. The result of the play is two paths $\tilde{\mathcal{P}} = \{\tilde{n}_0, \tilde{n}_1, ...\}$ of $\tilde{\mathcal{T}}$ and $\mathcal{P} = \{n_0, n_1, ...\}$ of $\mathcal{T}$. If there is no μ-trace on $\mathcal{P}$ then player II wins, so assume that there is a μ-trace on $\mathcal{P}$.

By condition $I2$ for every choice or modal node n_i of $\mathcal{P}$ we can define priority $p_{\tilde{n}_i, n_i}$. By trace preservation this priority cannot increase, hence after some index j it is constant, say equal q.

If $q = \infty$ then $j = 0$ and for every $k \geqslant 0$ by $I2$ we have $T(n_k)[\hat{\varphi}/\mu X.\hat{\alpha}(X)] \in \tilde{L}(\tilde{n}_k)$. By trace preservation we obtain a μ-trace on $\tilde{\mathcal{P}}$ going throughout these formulas.

If $q \in \mathbb{N}$ then for $k \geqslant j$, by $I2$, we have $T(n_k) \in L(nd_{\tilde{n}_k}(q))$ which by trace preservation gives us a μ-trace on the path $\mathcal{P}' = \{nd_{\tilde{n}_j}(q), nd_{\tilde{n}_{j+1}}(q), ...\}$ of $\mathcal{T}$. By equivalence $\mathcal{E}$ we have a μ-trace on the path $\hat{\mathcal{P}} = \{\widehat{nd}_{\tilde{n}_j}(q), \widehat{nd}_{\tilde{n}_{j+1}}(q), ...\}$ of $\hat{\mathcal{T}}$. By condition $I1$ we have a μ-trace on $\tilde{\mathcal{P}}$ which means that player II wins. ∎

We now summarise the case of the proof of Theorem 38 for $\varphi = \mu X.\alpha(X)$. By the induction assumption we have a disjunctive formula $\hat{\alpha}(X)$ equivalent to $\alpha(X)$ and know that $\alpha(X) \leqslant \hat{\alpha}(X)$ is provable. By Theorem 22, we a obtain a disjunctive formula $\hat{\varphi}$ which has a tableau $\hat{\mathcal{T}}$ equivalent to some tableau $\mathcal{T}$ for $\mu X.\hat{\alpha}(X)$. By Theorem 19, formula $\hat{\varphi}$ is equivalent to φ. By Lemma 39, $\mathcal{T}$ is a consequence of $\tilde{\mathcal{T}}$. By Lemma 35, $\hat{\mathcal{T}}$ is a consequence of $\mathcal{T}$. Hence, as the consequence relation is transitive, $\hat{\mathcal{T}}$ is a consequence of $\tilde{\mathcal{T}}$. Now by Proposition 28, $\hat{\alpha}(\hat{\varphi})$ is a weakly aconjunctive formula and $\hat{\varphi}$ is by definition a disjunctive formula. By Lemma 36, $\hat{\alpha}(\hat{\varphi}) \leqslant \hat{\varphi}$ is provable. Then $\mu X.\hat{\alpha}(X) \leqslant \hat{\varphi}$ is provable by rule (K6) and $\varphi \leqslant \mu X.\hat{\alpha}(X)$ is provable by the induction assumption. Hence $\varphi \leqslant \hat{\varphi}$ is provable.

Case: $\varphi = \nu X.\alpha(X)$. By the induction assumption we have an equivalent disjunctive formula $\hat{\alpha}(X)$ and $\alpha(X) \leqslant \hat{\alpha}(X)$ is provable. By Theorem 22, we obtain a disjunctive formula $\hat{\varphi}$ which has a tableau $\hat{\mathcal{T}}$ equivalent to some tableau $\mathcal{T}$ for $\nu X.\hat{\alpha}(X)$. Fortunately, by Proposition 28, $\nu X.\hat{\alpha}(X)$ is a weakly aconjunctive formula and by Lemma 35 $\hat{\mathcal{T}}$ is a consequence of $\mathcal{T}$. Hence we can use Lemma 36 to show that $\nu X.\alpha(X) \leqslant \hat{\varphi}$ is provable.

Case: $\varphi = \alpha \wedge \beta$. By the induction assumption there are disjunctive formulas $\hat{\alpha}, \hat{\beta}$ equivalent to α and β, respectively, and such that both $\alpha \leqslant \hat{\alpha}$ and $\beta \leqslant \hat{\beta}$ are provable. Hence $\alpha \wedge \beta \leqslant \hat{\alpha} \wedge \hat{\beta}$ is provable. By Theorem 22, there is a disjunctive formula $\hat{\varphi}$ which has a tableau equivalent to some tableau for $\hat{\alpha} \wedge \hat{\beta}$. Because, by Proposition 28, $\hat{\alpha} \wedge \hat{\beta}$ is a weakly aconjunctive formula, we can, as in the case before, use Lemma 36 to show that $\alpha \wedge \beta \leqslant \hat{\varphi}$ is provable. ∎

ACKNOWLEDGMENTS

I thank Dexter Kozen for the introducing me to the μ-calculus and sharing with me his understanding of the subject. I am grateful to Damian Niwiński for many important comments and remarks on this work. I thank David Janin, Nils Klarlund, Dexter Kozen, and Roger Villemaire for stimulating discussions. Finally, I thank the referees for their numerous suggestions.

Final manuscript received March 1, 1999

REFERENCES

1 Harel, D., Kozen, D., and Parikh, R. (1982), Process logic: Expressiveness, decidability and completeness, *J. Comput. System Sci.* **25**, 144–201.

2. Janin, D., and Walukiewicz, I. (1995), Automata for the μ-calculus and related results, *in* "MFCS '95," Lecture Notes in Computer Science, Vol. 969, pp. 552–562, Springer-Verlag, Berlin.

3. Kaivola, R. (1995), Axiomatising linear time mu-calculus, *in* "CONCUR '95," Lecture Notes in Computer Sciences, Vol. 962, pp. 423–437, Springer-Verlag, Berlin.

4. Kozen, D. (1983), Results on the propositional mu-calculus, *Theoret. Comput. Sci.* **27**, 333–354.

5. Kozen, D. (1988), A finite model theorem for the propositional μ-calculus, *Studia Logica* **47**, (3) 234–241.

6. Mostowski, A. W. (1984), Regular expressions for infinite trees and a standard form of automata, *in* "Fifth Symposium on Computation Theory" (A. Skowron, Ed.), Lecture Notes in Computer Science, Vol. 208, pp. 157–168, Springer-Verlag, Berlin.

7. Niwiński, D., and Walukiewicz, I. (February 1994), Games for μ-calculus, Technical Report, TR 94-03(192), Institute of Informatics, Warsav University. To appear in *Theoret. Comput. Sci.* Available from: http://www.daimi.aau.dk/igw/home.html.

8. Stirling, C. S., and Walker, D. J. (1991), Local model checking in the modal mu-calculus, *Theoret. Comput. Sci.* **89**, 161–171.

9. Streett, R. S., and Emerson, E. A. (1989), An automata theoretic procedure for the propositional mu-calculus, *Inform. and Comput.* **81**, 249–264.

10. Walukiewicz, I. (1993), On completeness of the μ-calculus, *in* "LICS '93," pp. 136–146.