Mathematisches
Forschungsinstitut
Oberwolfach

Mitglied der Leibniz-Gemeinschaft

Oberwolfach Preprints

OWP 2016 - 26

BAHAREH AFSHARI AND GRAHAM E. LEIGH

Finitary Proof Systems for Kozen's μ

Mathematisches Forschungsinstitut Oberwolfach gGmbH
Oberwolfach Preprints (OWP) ISSN 1864-7596

Oberwolfach Preprints (OWP)

Starting in 2007, the MFO publishes a preprint series which mainly contains research results related to a longer stay in Oberwolfach. In particular, this concerns the Research in Pairs-Programme (RiP) and the Oberwolfach-Leibniz-Fellows (OWLF), but this can also include an Oberwolfach Lecture, for example.

A preprint can have a size from 1 - 200 pages, and the MFO will publish it on its website as well as by hard copy. Every RiP group or Oberwolfach-Leibniz-Fellow may receive on request 30 free hard copies (DIN A4, black and white copy) by surface mail.

Of course, the full copy right is left to the authors. The MFO only needs the right to publish it on its website *www.mfo.de* as a documentation of the research work done at the MFO, which you are accepting by sending us your file.

In case of interest, please send a **pdf file** of your preprint by email to *rip@mfo.de* or *owlf@mfo.de*, respectively. The file should be sent to the MFO within 12 months after your stay as RiP or OWLF at the MFO.

There are no requirements for the format of the preprint, except that the introduction should contain a short appreciation and that the paper size (respectively format) should be DIN A4, "letter" or "article".

On the front page of the hard copies, which contains the logo of the MFO, title and authors, we shall add a running number (20XX - XX).

We cordially invite the researchers within the RiP or OWLF programme to make use of this offer and would like to thank you in advance for your cooperation.

Imprint:

Mathematisches Forschungsinstitut Oberwolfach gGmbH (MFO)
Schwarzwaldstrasse 9-11
77709 Oberwolfach-Walke
Germany

Tel +49 7834 979 50
Fax +49 7834 979 55
Email admin@mfo.de
URL www.mfo.de

The Oberwolfach Preprints (OWP, ISSN 1864-7596) are published by the MFO.

Finitary proof systems for Kozen's μ

Bahareh Afshari
Vienna University of Technology

Graham E. Leigh
University of Gothenburg

We present three finitary cut-free sequent calculi for the modal μ-calculus. Two of these derive annotated sequents in the style of Stirling's 'tableau proof system with names' (2014) and feature special inferences that discharge open assumptions. The third system is a variant of Kozen's axiomatisation in which cut is replaced by a strengthening of the ν-induction inference rule. Soundness and completeness for the three systems is proved by establishing a sequence of embeddings between the calculi, starting at Stirling's tableau-proofs and ending at the original axiomatisation of the μ-calculus due to Kozen. As a corollary we obtain a completeness proof for Kozen's axiomatisation which avoids the usual detour through automata or games.

This research was supported by the Mathematisches Forschungsinstitut Oberwolfach through the programmes "Oberwolfach Leibniz Fellows 2015" and "Research in Pairs 2016".

1 Introduction

Modal μ-calculus is the extension of propositional modal logic by quantifiers μ and ν that range over fixed points of propositional functions. More formally, the formulæ $\mu \mathsf{x} A$ and $\nu \mathsf{x} A$ are interpreted over directed labelled graphs as, respectively, the least and greatest fixed points of the (monotone) function formalising, at the semantic level, the mapping $\mathsf{x} \mapsto A(\mathsf{x})$. The quantifiers, combined with modal language, permit the expression of a variety of finite and infinite path quantification, giving the μ-calculus a second order flavour.

Since its inception in the early 1980s, the modal μ-calculus has become established as a central logic in computer science. On the one hand, the calculus is sufficiently rich to encompass many of the temporal logics used in systems verification, most prominently *computational tree logic* and *propositional dynamic logic*. On the other hand, despite its expressive power, the standard computational problems, such as validity and model checking, remain decidable.

The earliest deductive system for the modal μ-calculus, due to Kozen [10], is a Hilbert-style axiomatisation. A natural formulation of Kozen's axiomatisation as a (Tait-style) sequent calculus expands the usual sequent rules for (multi-)modal logic K by fixed point and induction inferences, and the logical rule of cut:

$$\frac{\Gamma, A(\sigma \mathsf{x} A(\mathsf{x}))}{\Gamma, \sigma \mathsf{x} A}\sigma \qquad\qquad \frac{\Gamma, A(\overline{\Gamma})}{\Gamma, \nu \mathsf{x} A(\mathsf{x})}\mathsf{ind} \qquad\qquad \frac{\Gamma, A \quad \Gamma, \overline{A}}{\Gamma}\mathsf{cut}$$

where σ ranges over the two quantifiers, $\overline{A}$ denotes the negation of A (in negation normal form), and $\overline{\Gamma}$ denotes the conjunction over negations of elements of Γ.[1]

The above proof system, henceforth denoted Koz, is known to be both sound and complete for the modal μ-calculus. Soundness was proved by Kozen in [10]; completeness was established later by Walukiewicz [16] utilising tableaux and infinite tree automata. It is, however, desirable to find a complete cut-free axiomatisation. A natural candidate is, of course, the subsystem Koz without cut, which we denote Koz^-. As Walukiewicz' proof makes essential use of the cut rule in Koz, and the result has, to date, proved surprisingly impervious to alternative approaches, completeness of Koz^- remains a significant open problem. Attention has thus shifted to providing alternative cut-free proof systems for the μ-calculus, such as the infinitary system $K_\omega(\mu)$ of [9] and, more recently, Stirling's 'tableau proof system with names' [13].

In this paper we prove completeness for the strengthening of Koz^-, denoted $\mathsf{Koz}^-_\mathsf{s}$, in which the induction rule is replaced by the inference

$$\frac{\Gamma, \nu \mathsf{x} A(\overline{\Gamma} \vee \mathsf{x})}{\Gamma, \nu \mathsf{x} A(\mathsf{x})}\mathsf{ind}_s$$

The new inference rule can be seen as combining the induction rule in Koz with two general fixed-point principles:

$$\nu \mathsf{x} \nu \mathsf{y} A(\mathsf{x}, \mathsf{y}) \leftrightarrow \nu \mathsf{x} A(\mathsf{x}, \mathsf{x}) \qquad\qquad \nu \mathsf{x} A(\mathsf{x} \vee \mathsf{x}) \leftrightarrow \nu \mathsf{x} A(\mathsf{x})$$

the first of which is referred to as the "golden lemma of μ-calculus" by Arnold and Niwinski [2].

[1] See Section 3 for full details.

The proof of completeness for $\mathsf{Koz}_\mathsf{s}^-$ is achieved by introducing two additional sequent calculi which are annotated in the style of Stirling [13]. These are utilised to establish a sequence of direct embeddings between calculi starting from a variant of Stirling's system, denoted Stir, and ending at $\mathsf{Koz}_\mathsf{s}^-$. Completeness of Stir and soundness of $\mathsf{Koz}_\mathsf{s}^-$ (both of which are relatively straightforward) will establish both properties for all systems.

As a consequence of our result we obtain a new, and constructive, completeness proof for Kozen's axiomatisation. The construction can be sketched as follows. Starting with a valid formula, one builds a proof in Stirling's calculus via a goal-orientated, deterministic tableau construction. In this system every formula and sequent is annotated by a word formed from 'names' for ν-variables and non-axiom leaves may be 'discharged' by a recurrence of the same sequent within the proof subject to certain conditions. Next we impose progressively stricter regularity conditions on these proofs by a process of systematic unfolding (identifying discharged leaves with their associated inner node). This leads to a proof whose structure induces a natural reading of annotations as describing applications of ind_s.

Outline of paper In the next section we fix the notation and definitions necessary for later work, including the sequent calculus that forms the basis for the systems utilised in this paper. Following this we briefly overview, in Section 3, the relevant literature on proof systems for the μ-calculus. The new proof systems we employ and the reductions between them form the content of Sections 4–7. We begin by establishing regularity conditions that can be imposed on Stirling proofs. These are subsequently exploited in Section 5 to interpret proofs in Stir in a closely related proof system Circ, referred to as *circular proofs*. Circ features an alternative rule for discharging open assumptions (see Definition 5.1) which permits a simpler and more transparent notion of proof avoiding the non-local 'repeat and reset' requirement present in Stir. The proof system and the reduction were previously outlined by the authors in [1] without detailed proof. In Section 6 a fresh inference rule, called *ν-closure*, is introduced (in Definition 6.1) which also discharges certain non-axiom leaves but captures the closure property for greatest fixed points more succinctly than either of the calculi Stir and Circ. Sequents in this system, which we denote Clo, are still annotated though in much simpler form than the calculi of the preceding sections. We prove that Circ-proofs can be translated into proofs in Clo. In Section 7 we provide an embedding of Clo into $\mathsf{Koz}_\mathsf{s}^-$, by translating annotated formulæ to 'plain' formulæ that make each discharged assumption a theorem of Koz^- and each instance of the ν-closure rule an application of ind_s. Section 8 treats the infinitary calculus $K_\omega(\mu)$ of [9]. We provide an embedding of Clo in $K_\omega(\mu)$ which, in combination with Studer's interpretation of $K_\omega(\mu)$-proofs as tableaux [14], implies the mutual embeddability of $K_\omega(\mu)$ and each of the annotated proof systems. With the embedding of Clo into Koz, proofs in $K_\omega(\mu)$ can be systematically transformed into Koz-proofs. We conclude with Section 9 where we discuss consequences and potential applications of our results.

2 Preliminaries

2.1 Syntax

The set of *μ-formulæ* is given by the grammar

$$A := p \mid \overline{p} \mid \mathsf{x} \mid A \wedge A \mid A \vee A \mid [\mathfrak{a}]A \mid \langle \mathfrak{a} \rangle A \mid \mu \mathsf{x} A \mid \nu \mathsf{x} A$$

where p ranges over a set *Prop* of *propositional constants*, $\mathfrak{a}$ over a set $\mathfrak{Act}$ of *actions* and x over a countably infinite set Var of *(formal) variables*. The sentential operators $\langle \mathfrak{a} \rangle$ and $[\mathfrak{a}]$ are referred to as *modalities*. An occurrence of a variable x in a μ-formula A is *bound* if it is within the scope of a quantifier $\mu\mathsf{x}$ or $\nu\mathsf{x}$, and is *free* otherwise. A μ-formula is *closed* if all variables are bound.

We define the following operations on μ-formulæ. Set $\bot = p \wedge \overline{p}$ and $\top = \overline{p} \vee p$ for some fixed $p \in \mathit{Prop}$ and define $A \to B = \overline{A} \vee B$ where $\overline{A}$ denotes the *dual* of A, given by

$$\begin{array}{llll} \overline{A \wedge B} = \overline{A} \vee \overline{B} & \overline{[\mathfrak{a}]A} = \langle \mathfrak{a} \rangle \overline{A} & \overline{\mu \mathsf{x} A} = \nu \mathsf{x} \overline{A} & \overline{\mathsf{x}} = \mathsf{x} \\ \overline{A \vee B} = \overline{A} \wedge \overline{B} & \overline{\langle \mathfrak{a} \rangle A} = [\mathfrak{a}]\overline{A} & \overline{\nu \mathsf{x} A} = \mu \mathsf{x} \overline{A} & \overline{\overline{p}} = p \end{array}$$

It is important to note that variables are never negated, so for instance $\overline{\mu\mathsf{x}(\mathsf{y} \vee [\mathfrak{a}]\mathsf{x})} = \nu\mathsf{x}(\mathsf{y} \wedge \langle \mathfrak{a} \rangle \mathsf{x})$. Given a μ-formula $A(\mathsf{x})$, possibly containing the variable x free, and a μ-formula B, we let $A(B)$ denote the result of replacing all free occurrences of x in A by B, renaming bound variables in A as necessary to avoid variable capture. As formulæ will always be taken *up to α*-equivalence, we assume that renaming of variables is unnecessary. Thus when writing $A(B)$ we implicitly mean that $A(\mathsf{x})$ is a μ-formula and no free occurrence of x in A is within the scope of a bound variable which is free in B. If $B = \sigma\mathsf{x}A(\mathsf{x})$ for some $\sigma \in \{\mu, \nu\}$, we call $A(B)$ the *unravelling* of $\sigma\mathsf{x}A$.

To reduce the use of parentheses we assume the binary connectives $\wedge$ and $\vee$ associate to the right, i.e. $A_1 \vee A_2 \vee \cdots \vee A_k$ abbreviates $A_1 \vee (A_2 \vee \cdots \vee A_k)$, and bind stronger than quantifiers and modalities. Appending a full stop to a quantifier expresses that it takes the largest possible scope, so $\mu\mathsf{x}.\, \langle \mathfrak{a} \rangle \mathsf{x} \vee \mathsf{x}$ represents the formula $\mu\mathsf{x}((\langle \mathfrak{a} \rangle \mathsf{x}) \vee \mathsf{x})$ whereas $\mu\mathsf{x}\langle \mathfrak{a} \rangle \mathsf{x} \vee \mathsf{x}$ should be read as $(\mu\mathsf{x}(\langle \mathfrak{a} \rangle \mathsf{x})) \vee \mathsf{x}$.

2.2 Semantics

Semantics for the modal μ-calculus is a direct extension of Kripke semantics for (multi-)modal logic incorporating variables and quantifiers. A *frame*, or *labelled transition system*, is a tuple $\mathscr{K} = \langle K, R, \lambda \rangle$ where $R \colon \mathfrak{Act} \to K \times K$ and $\lambda \colon \mathit{Prop} \to 2^K$. The set K is called the *domain* of $\mathscr{K}$. A *valuation* (over $\mathscr{K}$) is a function $v \colon \mathsf{Var} \to 2^K$.

Given a frame $\mathscr{K} = \langle K, R, \lambda \rangle$, μ-formula A and valuation v over $\mathscr{K}$, we define $\|A\|_v^{\mathscr{K}}$ by induction on A:

$$\begin{array}{rlrl} \|\mathsf{x}\|_v^{\mathscr{K}} &= v(\mathsf{x}) & & \\ \|p\|_v^{\mathscr{K}} &= \lambda(p) & \|[\mathfrak{a}]A\|_v^{\mathscr{K}} &= \{s \in K \mid \forall t \in K((s,t) \in R(\mathfrak{a}) \to t \in \|A\|_v^{\mathscr{K}})\} \\ \|\overline{p}\|_v^{\mathscr{K}} &= K \setminus \lambda(p) & \|\langle \mathfrak{a} \rangle A\|_v^{\mathscr{K}} &= \{s \in K \mid \exists t \in K((s,t) \in R(\mathfrak{a}) \wedge t \in \|A\|_v^{\mathscr{K}})\} \\ \|A \wedge B\|_v^{\mathscr{K}} &= \|A\|_v^{\mathscr{K}} \cap \|B\|_v^{\mathscr{K}} & \|\mu\mathsf{x}A\|_v^{\mathscr{K}} &= \bigcap\{S \subseteq K \mid \|A\|_{v[\mathsf{x} \mapsto S]}^{\mathscr{K}} \subseteq S\} \\ \|A \vee B\|_v^{\mathscr{K}} &= \|A\|_v^{\mathscr{K}} \cup \|B\|_v^{\mathscr{K}} & \|\nu\mathsf{x}A\|_v^{\mathscr{K}} &= \bigcup\{S \subseteq K \mid S \subseteq \|A\|_{v[\mathsf{x} \mapsto S]}^{\mathscr{K}}\} \end{array}$$

where $v[\mathsf{x} \mapsto S]$ denotes the valuation v' given by $v'(\mathsf{y}) = v(\mathsf{y})$ for every $\mathsf{y} \in \mathsf{Var} \setminus \{\mathsf{x}\}$, and $v'(\mathsf{x}) = S$. Since variable symbols appear only positively in μ-formulæ, when treated as functions of free variables, the semantics of μ-formulæ are *monotone*: if $v(\mathsf{x}) \subseteq T \subseteq K$ then for every formula A,

$$\|A\|_v^{\mathscr{K}} \subseteq \|A\|_{v[\mathsf{x}\mapsto T]}^{\mathscr{K}}.$$

An easy exercise shows that the semantics is closed under unravelling μ-formulæ, so for every $\sigma\mathsf{x}A(\mathsf{x})$ with $\sigma \in \{\mu, \nu\}$, frame $\mathscr{K}$ and valuation v,

$$\|\sigma\mathsf{x}A(\mathsf{x})\|_v^{\mathscr{K}} = \|A(\mathsf{x})\|_{v[\mathsf{x}\mapsto\|\sigma\mathsf{x}A\|_v^{\mathscr{K}}]}^{\mathscr{K}} = \|A(\sigma\mathsf{x}A)\|_v^{\mathscr{K}}. \tag{1}$$

In particular, for $\sigma = \mu$ (resp. $\sigma = \nu$), the Knaster-Tarski theorem implies $\|\sigma\mathsf{x}A\|_v^{\mathscr{K}}$ is the unique least (resp. greatest) set S ordered under inclusion satisfying the equation $S = \|A\|_{v[\mathsf{x}\mapsto S]}^{\mathscr{K}}$.

If v is a valuation over a frame $\mathscr{K} = \langle K, R, \lambda\rangle$, we let $\overline{v}$ denote the dual valuation which maps a variable x to the set $K \setminus v(\mathsf{x})$. It is not difficult to see that for every μ-formula A,

$$\|\overline{A}\|_v^{\mathscr{K}} = K \setminus \|A\|_{\overline{v}}^{\mathscr{K}}.$$

A formula A is *satisfiable* if there exists a frame $\mathscr{K}$ such that $\|A\|_v^{\mathscr{K}}$ is non-empty for some valuation v, and is *valid* if $\overline{A}$ is not satisfiable, i.e. $\|A\|_v^{\mathscr{K}}$ is the domain of $\mathscr{K}$ for every frame $\mathscr{K}$ and valuation v. Two μ-formulæ A and B are *equivalent* if both $A \to B$ and $B \to A$ are valid.

Example 2.1. Fix a formula $B = B(\mathsf{x}, \mathsf{y})$. We show that the μ-formula $\mu\mathsf{x}\nu\mathsf{y}B \to \nu\mathsf{y}\mu\mathsf{x}B$ is valid. The proof we present is due to Arnold and Niwinski [2]. Fix an arbitrary frame $\mathscr{K} = \langle K, R, \lambda\rangle$ and valuation v over $\mathscr{K}$. To simplify the presentation, for sets $X, Y \subseteq K$ and formula $C(\mathsf{x}, \mathsf{y})$, write $\|C(X, Y)\|_v^{\mathscr{K}}$ as shorthand for $\|C(\mathsf{x}, \mathsf{y})\|_{v[\mathsf{x}\mapsto X][\mathsf{y}\mapsto Y]}^{\mathscr{K}}$. In the following we drop explicit mention of $\mathscr{K}$ and v.

Let $S = \|\mu\mathsf{x}\nu\mathsf{y}B\|$. By definition $\|\nu\mathsf{y}\mu\mathsf{x}B\| = \bigcup\{Y \mid Y \subseteq \|\mu\mathsf{x}B(\mathsf{x}, Y)\|\}$ so $\mu\mathsf{x}\nu\mathsf{y}B \to \nu\mathsf{y}\mu\mathsf{x}B$ is valid if

$$S \subseteq \|\mu\mathsf{x}B(\mathsf{x}, S)\| \tag{2}$$

which we now prove. Let $T = \|\mu\mathsf{x}B(\mathsf{x}, S)\|$. By the fixed point property (1)

$$S = \|\nu\mathsf{y}B(S, \mathsf{y})\| = \|B(S, S)\|.$$

However, $T = \bigcap\{Y \mid \|B(S, Y)\| \subseteq Y\}$, so $T \subseteq S$ and by monotonicity

$$\|\nu\mathsf{y}B(T, \mathsf{y})\| \subseteq \|\nu\mathsf{y}B(S, \mathsf{y})\| = S$$
$$\|\nu\mathsf{y}B(T, \mathsf{y})\| = \|B(T, \|\nu\mathsf{y}B(T, \mathsf{y})\|)\| \subseteq \|B(T, S)\| = T.$$

Since $S = \bigcap\{Y \mid \|\nu\mathsf{y}B(Y, \mathsf{y})\| \subseteq Y\}$ we deduce $S \subseteq T$, i.e. (2).

2.3 Subsumption

Fix a formula A and let Var_A denote the variable symbols (bound or free) occurring in A. A induces a preorder $<_A$ on Var_A, called the *subsumption ordering for* A, given by $\mathsf{x} <_A \mathsf{y}$ if $\sigma \mathsf{y} B$ occurs as a sub-formula of A for some $\sigma \in \{\mu, \nu\}$ and B, and x is free in $\sigma \mathsf{y} B$. We call a formula A *locally well-named* if $<_A$ is irreflexive. Observe that every formula is α-equivalent to a locally well-named μ-formula, and that being locally well-named does not preclude a variable symbol being quantified in multiple places or contexts. Nevertheless, if $A = \mu \mathsf{x} B(\mathsf{x})$ is locally well-named then $A' = B(\mu \mathsf{x} B(\mathsf{x}))$ is locally well-named and $<_A = <_{A'}$.

It is convenient to assume that all formulæ considered are not only locally well-named but also pair-wise so, a condition that can be formalised thus. There exists a preorder $<$ on Var such that every finite preorder is embeddable in $<$.[2] Let us call a formula A *well-named* if $<_A$ is a sub-structure of $<$, i.e. $<_A = < \cap (\mathsf{Var}_A \times \mathsf{Var}_A)$. Every μ-formula is α-equivalent to a well-named formula. Henceforth we assume all μ-formulæ are well-named and write $\mathsf{x} \leq \mathsf{y}$ if either $\mathsf{x} < \mathsf{y}$ or $\mathsf{x} = \mathsf{y}$. Observe that $\overline{A}$ is well-named iff A is.

The next lemma captures the main use of well-named formulæ, namely that they are closed under unravelling fixed points expressions.

Lemma 2.1. *If $\sigma \mathsf{x} A(\mathsf{x})$ is a well-named μ-formulæ then $A(\sigma \mathsf{x} A(\mathsf{x}))$ is well-named.*

2.4 Threads

Definition 2.2 (Threads). A *thread* is a sequence of formulæ $\alpha = (A_i)_{i<N}$, where $N \leq \omega$, such that for every $n+1 < N$, one of the following conditions hold

1. $A_{n+1} = A_n$,
2. A_{n+1} is an immediate sub-formula of A_n,
3. $A_n = \sigma \mathsf{x} A(\mathsf{x})$ for some σ, x and A, and $A_{n+1} = A(A_n)$.

An *infinite thread* is as above with $N = \omega$. Let $\alpha = (A_i)_{i<\omega}$ be an infinite thread and fix $\sigma \in \{\mu, \nu\}$. A variable x *occurs infinitely often in* α *as* σ if for every $i < \omega$ there exists $i < j < \omega$ such that $A_j = \sigma \mathsf{x} A$ for some A. We call α a σ*-thread* if there is a variable x that i) occurs infinitely often in α as σ and ii) for all y that occurs infinitely often in α as μ, $\mathsf{x} < \mathsf{y}$. Given a sequence $\beta = (\Gamma_i)_{i<N}$ of sets of μ-formulæ, a *thread through* β is any thread $\alpha = (A_i)_{i<N}$ such that $A_i \in \Gamma_i$ for every $i < N$.

Lemma 2.3. *Let α be an infinite thread starting from a well-named formula such that at least one of condition 2 or 3 in Definition 2.2 is satisfied infinitely often. Then α is either a ν-thread or a μ-thread and not both.*

2.5 Sequent calculi

A *sequent* is a finite set of closed μ-formulæ. Sequents are denoted by Γ, Δ, etc., and formulæ are identified with singleton sequents. The union of two sequents is abbreviated by comma, writing Γ, Δ for $\Gamma \cup \Delta$ and Γ, A for $\Gamma \cup \{A\}$. Later we will also consider

[2]This is a consequence of Fraïssé's Theorem from Model Theory, cf [8, Theorem 7.1.2].

sequents with further structure and so to avoid confusion refer to sequents as *plain sequents*. We write $\bigvee\Gamma$ ($\bigwedge\Gamma$) to denote the disjunction (conjunction) over elements of Γ; the *dual* of Γ, denoted $\overline{\Gamma}$, is the μ-formula $\overline{\bigvee\Gamma}$. A sequent Γ is *valid* if the μ-formula $\bigvee\Gamma$ is valid.

In this paper, a *plain sequent calculus*, or simply a *calculus*, is a set of *inference rules* of the form

$$\dfrac{\begin{matrix}[\Delta]^{\dagger}\\ \vdots \\ \Gamma, A_1, \ldots, A_k\end{matrix}}{\Gamma, B_1, \ldots, B_k}\,\mathsf{dis}_{\dagger} \qquad \text{or} \qquad \dfrac{\Gamma, A_i \quad (i \in I)}{\Gamma, B}\,\mathsf{inf} \tag{3}$$

where Γ and Δ are (plain) sequents, A_i, B_i and B denote μ-formulæ, and I is an index set. In each of the two forms, the sequent(s) immediately above the rule is referred to as the *premise(s)*, and the lower sequent the *conclusion*. The set Δ in the left inference is the sequent *discharged* by the inference, and emphasised by the square bracket notation and annotation (in this case '†') clarifying the associated discharging inference. We refer to Γ as the set of *side formulæ*, the formulæ A_i occurring in the premises are the *active formulæ* and $B_1, \ldots, B_k, B$ in the conclusions are *principal*. An instance of the right-hand rule in which I is infinite is called an *infinitary* rule. All other inference rules of the form in (3) are *finitary*. An *axiom* is an inference rule with no premises.

Definition 2.4 (Derivations)**.** Let T be a plain sequent calculus comprising inference rules and axioms of the form in (3). A *pre-derivation* in T is a possibly infinite tree labelled by plain sequents locally consistent with the inference rules and axioms of T. A *derivation* is a pre-derivation in which every path is finite, i.e. the underlying tree is well-founded. The *conclusion* of a (pre-)derivation is the sequent labelling the root and an *assumption* is a leaf labelled by a sequent that is not an axiom of T. An assumption is *discharged* if it has been discharged by an inference in sense of (3) above, and is *open* otherwise. A *closed* (pre-)derivation is a (pre-)derivation in which every assumption is discharged.

Given a derivation π we identify paths in π with the sequence of sequents they induce and it is in this sense that we may talk about threads (through a path) in π.

For many of the calculi we present 'closed derivation' is synonymous with 'proof,' to be understood in the usual way. This is not true of all the proof systems we consider, however, and in general a proof in a calculus T will be a closed derivation in which every path to a discharged assumption satisfies some additional structural requirement. Specifically, each calculus T will be associated a particular class of (open) T-derivations called *(open) proofs*. Relative to a calculi T and the notion of open proof for T we may introduce the following abbreviations. If $\{\Gamma\} \cup \mathscr{A}$ is a set of sequents we write $\mathscr{A} \vdash_{\mathsf{T}} \Gamma$ if there exists an open proof in T with conclusion Γ and for which the open assumptions form a subset of $\mathscr{A}$, and write $\mathsf{T} \vdash \Gamma$ if $\emptyset \vdash_{\mathsf{T}} \Gamma$, i.e. there is a T-proof of Γ. T will be called *sound* if for every plain sequent Γ, $\mathsf{T} \vdash \Gamma$ implies $\bigvee\Gamma$ is valid, and *complete* if the converse holds. Finally, we often drop explicit reference to T if it is clear from the context which calculus is meant.

We begin by outlining a basic sequent calculus which we call *fixed point logic* that serves as a basis for all the calculi we present in this paper.

$$\mathsf{Ax1}\colon p, \bar{p} \qquad \frac{\Gamma, B, C}{\Gamma, B \vee C}\vee \qquad \frac{\Gamma, B \quad \Gamma, C}{\Gamma, B \wedge C}\wedge$$

$$\frac{\Gamma, A}{\langle \mathfrak{a} \rangle \Gamma, [\mathfrak{a}]A}\mathsf{mod} \qquad \frac{\Gamma, A(\sigma \mathsf{x} A(\mathsf{x}))}{\Gamma, \sigma \mathsf{x} A}\sigma \qquad \frac{\Gamma}{\Gamma, A}\mathsf{weak}$$

Figure 1: Rules and axioms of *fixed point logic*, Fix.

Definition 2.5 (Fixed point logic)**.** *Fixed point logic*, denoted Fix, is the sequent calculus comprising the seven inference rules and axioms listed in Figure 1. A Fix-*proof* is a finite closed Fix-derivation.

The class of Fix-proofs is not particularly interesting. Although sound, Fix is not complete for the μ-calculus, an easy consequence of the fact that Fix is sound with respect to interpreting the fixed point quantifiers as picking out any set satisfying the fixed point equation in (1). Complete proof systems can be obtained by extending Fix by further rules (and axioms) or relaxing the well-foundedness condition on proofs.

Remark on terminology Part of the motivation behind this work was a proof of completeness for Kozen's axiomatisation that avoids recourse to the theory of automata or parity games. This is achieved via a sequence of translations between sequent calculi. In order to succeed in our objective it is necessary that each embedding of one calculus into another does not appeal to 'indirect' arguments involving automata, games, or similar. A natural restriction in this direction would be that each embedding is *effective* or even *primitive recursive*. While this is certainly true of our completeness proof, Walukiewicz's proof is also readily seen as computable, so an effective embedding of a sound calculus T into Koz can, in theory, consist of replicating soundness and completeness for the two calculi, an approach that we want to avoid. However, Walukiewicz's completeness proof is non-constructive in nature, in contrast to our approach. Thus in the present paper an *embedding* from T to S is to be understood as a constructive, primitive recursive translation of T-proofs to S-proofs preserving the conclusion.

3 Background on proof systems for modal μ

3.1 Kozen's axiomatisation

The first proof system for the μ-calculus is due to Kozen. In [10], Kozen presents a Hilbert-style proof system for the calculus extending the basic (multi-)modal logic K by two axioms and one rule of inference:

$$A(\nu \mathsf{x} A(\mathsf{x})) \to \nu \mathsf{x} A(\mathsf{x}) \qquad \nu \mathsf{x} A \vee \mu \mathsf{x} \overline{A} \qquad \frac{B \to A(B)}{B \to \nu \mathsf{x} A(\mathsf{x})}$$

A natural presentation of Kozen's system as a sequent calculus is the extension of Fix by the rule ind and axiom Ax2 in Figure 2 and the rule cut in Figure 3. In the present paper we are concerned with variants of Kozen's system without the cut rule, and in order to more easily accommodate these cut-free calculi we also include a generalisation of the disjunction rule, $\vee_d$ in Figure 2, which although admissible in the presence of cut, is not

$$\frac{\Gamma, A(\overline{\Gamma})}{\Gamma, \nu \mathsf{x} A(\mathsf{x})}\,\mathsf{ind} \qquad \mathsf{Ax2}\colon \nu\mathsf{x}A, \mu\mathsf{x}\bar{A} \qquad \frac{\Gamma, A(B), A(C)}{\Gamma, A(B \vee C)}\,\vee_d$$

Figure 2: Additional rules present in Koz^-.

$$\frac{\Gamma, A \quad \Gamma, \overline{A}}{\Gamma}\,\mathsf{cut} \qquad \frac{\Gamma, \nu\mathsf{x}A(\overline{\Gamma} \vee \mathsf{x})}{\Gamma, \nu\mathsf{x}A(\mathsf{x})}\,\mathsf{ind}_s$$

Figure 3: Inference rules *cut* and *strengthened induction*.

obviously so in the systems lacking cut. Thus we define Koz^- to be the extension of Fix by the three rules and axioms in Figure 2 and represent Kozen's axiomatisation as the sequent calculus $\mathsf{Koz}^- + \mathsf{cut}$ which we denote Koz.

Proofs in Koz (and its subsystems) are closed finite derivations. Unlike Fix, a sound *and* complete proof system obtains.

Theorem 3.1. Koz *is sound and complete for the μ-calculus.*

Soundness for the calculus was established by Kozen in [10], as was completeness for a fragment known as the *aconjunctive* fragment. Completeness for the full calculus, however, was only proved much later, by Walukiewicz [16].

Completeness was (and still is) a significant result. To establish completeness, Walukiewicz isolated a class of μ-formulæ, called *disjunctive formulæ*,[3] that is, provably in Koz, equi-expressive with the full language, and for which weak completeness can be derived.[4] Weak completeness turns out to be trivial for duals of disjunctive formulæ as Koz-proofs can be directly read off from tableaux for duals of unsatisfiable disjunctive formulæ. In contrast, the proof of equi-expressivity of the disjunctive fragment and the full μ-calculus is extremely involved, and depends heavily on automata and game theoretic techniques, all operating on tableaux, to establish that every guarded μ-formula can be associated an equivalent disjunctive formula and that the equivalence is provable in Koz. Moreover, the equivalent disjunctive formula is typically exponentially larger than the starting μ-formula.

The upshot is that for a valid μ-formula A, the proof constructed by the above procedure comprises two parts combined by a cut: a simple Koz-proof for the dual of an (exponentially larger) equivalent disjunctive formulæ, A', and a proof of the sequent $A' \to A$. In particular, the resulting proof bears little relation to the Koz-derivations typically constructed by hand, wherein syntactic constraints such as disjunctiveness rarely materialise. To demonstrate the latter point we give a proof of the valid formula $\mu\mathsf{x}\nu\mathsf{y}B \to \nu\mathsf{y}\mu\mathsf{x}B$ of Example 2.1. First we make the following observation.

Lemma 3.2. *Let $A(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1})$ be a formula with at most the designated variables free.*

[3]See, for example, [16] or [3] for a definition of disjunctive formulæ (called *automaton normal form* in [3]).

[4]Weak completeness is the statement that $\mathsf{Koz} \vdash \overline{A}$ for every unsatisfiable μ-formula A.

If B_i and C_i are closed formulæ for each $i < k$, then

$$\{B_i, C_i\}_{i<k} \vdash_{\mathsf{Koz}^-} \overline{A}(B_0, \ldots, B_{k-1}), A(C_0, \ldots, C_{k-1}).$$

Proof. The proof proceeds by induction on A. We present the case $A = \nu \mathsf{x}_k A_0(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1}, \mathsf{x}_k)$. The remaining cases are straightforward. Let $B_k = \overline{A}(B_0, \ldots, B_{k-1})$. As the sequent $B_k, \overline{B_k}$ is an instance of $\mathsf{Ax2}$, the induction hypothesis implies

$$\{B_i, C_i\}_{i<k} \vdash_{\mathsf{Koz}^-} \overline{A_0}(B_0, \ldots, B_k), A_0(C_0, \ldots, C_{k-1}, \overline{B_k}),$$

whereby an application of μ yields

$$\{B_i, C_i\}_{i<k} \vdash_{\mathsf{Koz}^-} B_k, A_0(C_0, \ldots, C_{k-1}, \overline{B_k})$$

and an application of ind completes the proof. ☐

A particular instance of Lemma 3.2 worth noting is

Lemma 3.3. $\mathsf{Koz}^- \vdash \Gamma, \overline{\Gamma}$ *for every plain sequent* Γ.

Example 3.1. Recall the valid sequent $\{\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B\}$ from Example 2.1. Let $C = \nu\mathsf{x}\mu\mathsf{y}\overline{B}$ and $D = \nu\mathsf{y}\mu\mathsf{x}B$. The following derivation, which we denote π_{koz}, is the Koz-proof of this sequent motivated by the semantic validity argument:

$$C, \overline{C}$$
$$\vdots \ \text{Lemma 3.2}$$
$$\dfrac{\mu\mathsf{y}\overline{B}(C, \mathsf{y}), \nu\mathsf{y}B(\overline{C}, \mathsf{y})}{C, \nu\mathsf{y}B(\overline{C}, \mathsf{y})}\ \nu \qquad\qquad C, \overline{C}$$
$$\vdots \ \text{Lemma 3.2}$$
$$\dfrac{\dfrac{\dfrac{\overline{B}(C, C), B(\overline{C}, \nu\mathsf{y}B(\overline{C}, \mathsf{y}))}{\overline{B}(C, C), \nu\mathsf{y}B(\overline{C}, \mathsf{y})}\ \nu}{\overline{B}(C, C), \overline{C}}\ \mu}{\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \overline{C}}\ \mathsf{ind}$$
$$\vdots \ \text{Lemma 3.2}$$
$$\dfrac{\mu\mathsf{y}\overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mathsf{y}), \nu\mathsf{y}B(\overline{C}, \mathsf{y})}{\mu\mathsf{y}\overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mathsf{y}), \overline{C}}\ \mu \qquad \nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mu\mathsf{x}B(\mathsf{x}, \overline{C})$$
$$\vdots \ \text{Lemma 3.2}$$
$$\dfrac{\dfrac{\dfrac{\overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mu\mathsf{y}\overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mathsf{y})), B(\mu\mathsf{x}B(\mathsf{x}, \overline{C}), \overline{C})}{\mu\mathsf{y}\overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, C), \mathsf{y}), \mu\mathsf{x}B(\mathsf{x}, \overline{C})}\ \mu, \mu}{C, \mu\mathsf{x}B(\mathsf{x}, \overline{C})}\ \mathsf{ind}}{C, D}\ \mathsf{ind}$$

This example well demonstrates the non-triviality of generating Koz-proofs which can be attributed to the impredicative nature of the ν-induction rule. Note that the cut rule is not utilised in Lemma 3.2, so π_{koz} is in fact a proof in Koz^-.

3.2 Tableaux proofs

The first deductive system for the μ-calculus for which completeness was established is a system of ill-founded Fix derivations, referred to as *tableaux*, and is due to Niwinski and Walukiewicz [12]. Every formula induces a class of pre-derivations obtained by applying the proof rules of Fix in a 'bottom-up' fashion, systematically decomposing the formula into sequents of sub-formulæ. A cardinality argument shows that every infinite path of a pre-derivation must contain an infinitely repeating sequent and it so happens that validity can be characterised by a syntactic condition on the threads occurring along these paths.

Definition 3.4 (Tableaux)**.** Recall the definition of ν-thread from Section 2. A *tableau for* Γ *from assumptions* $\mathscr{A}$, written $\mathscr{A} \vdash_\infty \Gamma$, is a (possibly infinite) Fix pre-derivation π with conclusion Γ such that all open assumptions are in $\mathscr{A}$ and every infinite path through π contains an infinite ν-thread. When $\mathscr{A} = \emptyset$ we call π a *tableau.*

Theorem 3.5 (Niwinski and Walukiewicz [12])**.** *Let A be a guarded μ-formula. Then A is valid iff there exists a tableaux for A.*

Guardedness is the syntactic restriction on μ-formulæ requiring that for every sub-formula $\sigma \mathsf{x} B$ every occurrence of x in B is under the scope of a modality (in B). Every μ-formula is equivalent to a guarded formula and this equivalence is provable in Koz [10]. However, the restriction to guardedness in Theorem 3.5 turns out to be unnecessary:

Theorem 3.6. *A formula A of μ-calculus is valid iff there exists a tableaux for A.*

Theorem 3.6 is a corollary of the main result of Studer [14, Theorem 7.2] and also follows from Friedmann and Lange [6] on satisfiability tableaux for unguarded formulæ. Below we outline a direct proof of the 'left-to-right' direction of Theorem 3.6 as it will be important in pursuing completeness for *cut-free* proof systems where one cannot appeal to equivalences between guarded and unguarded formulæ.

Guardedness is an important restriction in the treatment of [12]. Given a sequent Γ containing guarded formulæ only, every decomposition of Γ by the inferences $\vee$, $\wedge$, μ and ν produces a finite tree with leaves labelled by sets of atoms and formulæ of the form $[\mathfrak{a}]A$ and $\langle\mathfrak{a}\rangle A$. If Γ is valid then applications of weak can reduce each leaf to either an axiom or a valid conclusion of mod, the premise of which will also be valid. Repeating this process ad infinitum completes the construction of the pre-derivation whence a second argument shows it to be a tableau.

If the starting valid sequent contains an unguarded formula then the decomposition by $\vee$, $\wedge$, μ and ν need not be finite and an infinite branch need not contain a ν-thread, as is the case, for example, with the valid sequent $\Gamma = \{[\mathfrak{a}]\bot, \langle\mathfrak{a}\rangle\top, \mu\mathsf{x}(B \wedge \mathsf{x})\}$. Nevertheless, a tableaux for Γ can be easily obtained from a tableaux for $\{[\mathfrak{a}]\bot, \langle\mathfrak{a}\rangle\top\}$ and an application of weak. A more 'deterministic' approach to constructing a tableaux for Γ is to apply reductions μ, ν, $\vee$ and $\wedge$ until a repetition occurs and an infinite μ-thread becomes apparent, after which the offending formula is removed by weak. As an example consider the following tableau where $\star$ marks the points at which the unguarded formula C may be safely eliminated to stop the creation of an infinite μ-thread, and where the right-hand path is continued by identifying the leaf labelled by A, B, C with the root.

$$
\begin{prooftree}
\AxiomC{$p, \overline{p}$}
\RightLabel{$\vee$}
\UnaryInfC{$\top$}
\RightLabel{weak}
\UnaryInfC{$\bot, \top$}
\RightLabel{mod}
\UnaryInfC{$[\mathfrak{b}]\bot, \langle\mathfrak{b}\rangle\top$}
\RightLabel{weak*}
\UnaryInfC{$[\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{b}\rangle\top, [\mathfrak{a}]C$}
\RightLabel{weak}
\UnaryInfC{$(\star)\ [\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{b}\rangle\top, [\mathfrak{a}]C, C$}
\RightLabel{$\mu, \vee$}
\UnaryInfC{$[\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{b}\rangle\top, C$}
\RightLabel{$\mu, \vee$}
\UnaryInfC{$A, \langle\mathfrak{b}\rangle\top, C$}
\AxiomC{$\vdots$}
\noLine
\UnaryInfC{A, B, C}
\RightLabel{mod}
\UnaryInfC{$\langle\mathfrak{a}\rangle A, \langle\mathfrak{a}\rangle B, [\mathfrak{a}]C$}
\RightLabel{weak}
\UnaryInfC{$[\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{a}\rangle B, [\mathfrak{a}]C$}
\RightLabel{weak}
\UnaryInfC{$(\star)\ [\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{a}\rangle B, [\mathfrak{a}]C, C$}
\RightLabel{$\mu, \vee$}
\UnaryInfC{$[\mathfrak{b}]\bot, \langle\mathfrak{a}\rangle A, \langle\mathfrak{a}\rangle B, C$}
\RightLabel{$\mu, \vee$}
\UnaryInfC{$A, \langle\mathfrak{a}\rangle B, C$}
\RightLabel{$\wedge$}
\BinaryInfC{$A, \langle\mathfrak{b}\rangle\top \wedge \langle\mathfrak{a}\rangle B, C$}
\RightLabel{ν}
\UnaryInfC{$\underbrace{\mu \mathsf{y}([\mathfrak{b}]\bot \vee \langle\mathfrak{a}\rangle\mathsf{y})}_{A}, \underbrace{\nu\mathsf{y}(\langle\mathfrak{b}\rangle\top \wedge \langle\mathfrak{a}\rangle\mathsf{y})}_{B}, \underbrace{\mu\mathsf{x}([\mathfrak{a}]\mathsf{x} \vee \mathsf{x})}_{C}$}
\end{prooftree}
$$

These observations suffice to prove the generalisation of Theorem 3.5 to unguarded formulæ, which we now present.

Proof of Theorem 3.6. Begin by assigning to every sequent Γ a pre-derivation π_Γ obtained as follows. Starting with conclusion Γ, apply the rules $\vee$, $\wedge$, μ and ν in a bottom-up fashion in any order except if there arises a path $\alpha = (m_i)_{i \leq N}$ in π_Γ with m_0 the root of π_Γ, and a thread $(A_i)_{i \leq N}$ though α such that i) $A_k = A_N = \mu \mathsf{x} A$ for some A and $k < N$, ii) A_k is principal in the sequent at m_k, and iii) if $A_i = \nu \mathsf{y} B$ for any $k < i < N$ then $\mathsf{x} < \mathsf{y}$. If this case arises then instead of one of $\vee$, $\wedge$, μ or ν, an application of weak is added with principal formula A_N. By construction the pre-derivation π_Γ is such that it contains no applications of mod, every infinite path contains an infinite ν-thread, and every leaf is a sequent containing only propositional constants or formulæ starting with a modality. Moreover, if Γ is a valid sequent then every leaf of π_Γ is labelled by a valid sequent.

Fix a valid μ-formula A. A pre-derivation π for A is constructed by recursion. Let $\pi_0 = \pi_{\{A\}}$. Given π_n, π_{n+1} is determined by inserting at every leaf of π_n applications of weak reducing the sequent to either an axiom or to an application of mod with a valid premise Δ, and in the latter case inserting a copy of π_Δ to this leaf. The limit of the construction is a closed Fix pre-derivation π for A. It remains to prove that π is a tableau, i.e. that all infinite paths contain a ν-thread. Fix an infinite path α in π. It follows that α is either already contained in π_n for some n, or α consists of infinitely many mod rules. In the former case α contains a ν-thread by design, whereas in the latter case every infinite thread through α passes through modalities infinitely often and since every sequent in α is valid one concludes that α contains a ν-thread by the usual argument. □

Lemma 3.7. *For every formula $A(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1})$ with at most the distinguished variables free, and for all sequences of closed formulæ $(B_i)_{i<k}$ and $(C_i)_{i<k}$ we have $\{B_i, C_i\}_{i<k} \vdash_\infty \overline{A}(B_0, \ldots, B_{k-1}), A(C_0, \ldots, C_{k-1})$.*

Proof. By induction on A. In the case $A = \nu \mathsf{x}_k A_0(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1}, \mathsf{x}_k)$ apply the induction hypothesis to $A_0(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1}, \mathsf{x}_k)$, picking $B_k = \overline{A}$ and $C_k = A$, and yielding a tableau with conclusion $\overline{A}, A$ from open assumptions $\{B_i, C_i\}$ for $i \leq k$. Identifying the final assumption with the root of the derivation and unravelling the result to an infinite tree yields the desired tableau. □

Example 3.2. Recall the valid sequent $\{\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B\}$ from Example 2.1. Let $C = \nu\mathsf{x}\mu\mathsf{y}\overline{B}$ and $D = \nu\mathsf{y}\mu\mathsf{x}B$. The following Fix pre-derivation π_∞ is a tableaux for $\{C, D\}$ where the missing inferences are provided by Lemma 3.7 and by identifying the two † and ‡ nodes.

$$
\begin{array}{c}
\dfrac{\begin{array}{c}\vdots\\ \mu\mathsf{y}\overline{B}(C,\mathsf{y}), \mu\mathsf{x}B(\mathsf{x},D)\quad(\dagger)\end{array}}{C, \mu\mathsf{x}B(\mathsf{x},D)}\,\nu
\qquad
\begin{array}{c}\vdots\\ \mu\mathsf{y}\overline{B}(C,\mathsf{y}), D\quad(\ddagger)\end{array}\\
\text{- -}\\
\vdots\ \text{Lemma 3.7}\\
\dfrac{\dfrac{\dfrac{\dfrac{\overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y})), B(\mu\mathsf{x}B(\mathsf{x},D), D)}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}), B(\mu\mathsf{x}B(\mathsf{x},D), D)}\,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}), \mu\mathsf{x}B(\mathsf{x},D)\ (\dagger)}\,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}), D\ (\ddagger)}\,\nu}{C, D}\,\nu
\end{array}
$$

3.3 Semi-formal systems

An alternative to the infinitely 'long' tableaux proofs is to consider infinitely 'wide' proofs, where one or more inference rules take infinitely many premises. Such systems, known as *semi-formal* systems or *infinitary* calculi, have been widely studied in the context of first and second order theories of arithmetic.

Jäger, Kretz and Studer [9], drawing on this background, define a sound and complete cut-free proof system for μ-calculus by adding an infinitary rule characterising the greatest fixed point. For each $n < \omega$, define a new 'quantifier' ν^n by $\nu^0\mathsf{x}A = \top$, and $\nu^{n+1}\mathsf{x}A(\mathsf{x}) = A(\nu^n\mathsf{x}A)$. The ν_ω inference rule is the following infinitary proof rule, the premises to which is a derivation of $\Gamma, \nu^n\mathsf{x}A$ for each n.

$$
\dfrac{\Gamma, \nu^0\mathsf{x}A \quad \Gamma, \nu^1\mathsf{x}A \quad \cdots}{\Gamma, \nu\mathsf{x}A}\,\nu_\omega
$$

Following [9] we call the extension of Fix by the above rule $K_\omega(\mu)$. A $K_\omega(\mu)$*-proof* is a closed well-founded derivation in $K_\omega(\mu)$. Note that $K_\omega(\mu)$-proofs are infinite objects due to the ω-branching of the ν_ω rule but well-foundedness implies a $K_\omega(\mu)$-proof has no infinite paths.

Soundness for $K_\omega(\mu)$ is an immediate consequence of the *finite model property* of the μ-calculus: if the premises are all valid sequents but the conclusion is not there exists a *finite* frame $\mathscr{K}$ such that $\mathscr{K} \not\models \nu\mathsf{x}A$ but $\mathscr{K} \models \nu^n\mathsf{x}A$ for every n, which yields a contradiction when n is at least the size of the domain of $\mathscr{K}$. Completeness is established via an infinitary version of the canonical model constructions used in proving completeness for modal logics, and techniques from infinitary proof theory.

The finite model property implies that for each instantiation of ν_ω only finitely many of the premises are needed to deduce the validity of the conclusion. It does not by itself, however, provide any bound, but since the μ-calculus satisfies also the small model property, a 'finitisation' of the ν_ω rule is possible. This latter property states that there is an exponential function $f\colon \omega \to \omega$ such that every satisfiable μ-formula has a model with domain of size at most $f(d(A))$ where $d(A)$ denotes the logical depth of the formula

A. Jäger, Kretz and Studer therefore also consider the following rule in place of ν_ω.

$$\frac{\Gamma, \nu^0 \mathsf{x} A \quad \cdots \quad \Gamma, \nu^k \mathsf{x} A}{\Gamma, \nu \mathsf{x} A} \nu_{<\omega} \text{ where } k = f(d(\bigvee \Gamma \wedge \nu \mathsf{x} A))$$

The resulting calculus, denoted $K_{<\omega}(\mu)$, is cut-free, sound and complete, and comprises only finitary inference rules. Nevertheless, the system is not entirely satisfactory as a finitary calculus in the traditional sense due to inferences of arbitrarily high arity: in general, the $\nu_{<\omega}$ rule deriving a sequent Γ has a number of premises exponential in the logical complexity of Γ.

3.4 Annotated proof systems

Stirling [13] introduces a 'tableau proof system with names' for the μ-calculus that captures the infinite ν-thread condition in tableaux within a finite tree by means of annotating formulæ and sequents with names for fixed point variables.

For each variable symbol x, fix an infinite set N_x of *names* for x. We assume $N_\mathsf{x} \cap N_\mathsf{y} = \emptyset$ if $\mathsf{x} \neq \mathsf{y}$ and use symbols x, y, z (also with indices) as names for the formal variables x, y and z respectively. Let $N := \bigcup_{\mathsf{x} \in \mathsf{Var}} N_\mathsf{x}$. Given a name $x \in N$, let N_x be the set N_x such that $x \in N_\mathsf{x}$. The subsumption ordering $<$ extends to apply to names in the obvious way: for $x, y \in N$, $x \leq y$ ($x < y$) if $x \in N_\mathsf{x}$ and $y \in N_\mathsf{y}$, and $\mathsf{x} \leq \mathsf{y}$ (resp. $\mathsf{x} < \mathsf{y}$). For a set $M \subseteq N$ of names, M^* is the set of finite words in M. For $a \in N^*$ and $\mathsf{x} \in \mathsf{Var}$ we write $a < \mathsf{x}$ ($a \leq \mathsf{x}$) iff a is a word in $\bigcup_{\mathsf{y} < \mathsf{x}} N_\mathsf{y}$ (resp. $\bigcup_{\mathsf{y} \leq \mathsf{x}} N_\mathsf{y}$). Let $\sqsubset$ denote the (reflexive) sub-word relation on N^*. Given words $a, b \in N^*$, we write $b{\restriction}_a$ for the largest sub-word of b comprising only names occurring in a.

In Stirling's proof system (and two further systems we introduce in later sections) sequents will be annotated by words from variable names. An *annotation* is a non-repeating word $a \in N^*$ that is increasing in $<$: if $a = x_0 \cdots x_{k-1}$ then for all $0 \leq i < j < k$, $x_i \leq x_j$ and $x_i \neq x_j$. An *annotated formula* is a pair (a, A), henceforth written A^a, where A is a closed μ-formula and $a \in N^*$ is an annotation consisting of names for variables occurring in A. An *annotated sequent* is a finite set of closed annotated formulæ $\{A_1^{a_1}, \ldots, A_n^{a_n}\}$ together with a finite word $a \in N^*$ without repetitions, called the *control*, written as $a \vdash A_1^{a_1}, \ldots, A_n^{a_n}$. We let Γ, Δ, Π range over finite sets of annotated formulæ and identify a plain sequent $\{A_1, \ldots, A_k\}$ with the annotated sequent $\epsilon \vdash A_1^\epsilon, \ldots, A_k^\epsilon$ in which the control and all annotations are the empty annotation.

The notion of an *inference rule*, *axiom*, *derivation*, *thread*, etc., for plain sequents generalises to annotated sequents in the obvious way. The notion of a *proof*, however, is not so robust as in each of the calculi we present the annotations and controls play slightly different roles, leading to different notions of proof.

Let *fixed point logic with names*, FixN, be the generalisation of Fix to annotated sequents with rules and axioms given in Figure 4. The notable restrictions on annotations are in the quantifier rules μ and ν, wherein the annotation b is required to comprise only names for x and variables subsuming x, and in exp, which permits expanding annotations by fresh variable names. FixN will form the underlying calculus for the annotated proof systems in this paper. The first complete proof system extending FixN that we present is due to Colin Stirling [13].

Definition 3.8 (Stirling proofs)**.** Let Stir denote the calculus extending FixN by the rule dis, and the rules ν_x and reset_x for each variable name $x \in N$, listed in Figure 5.

$$\mathsf{Ax1}\colon \epsilon \vdash p^\epsilon, \overline{p}^\epsilon \qquad \frac{a \vdash \Gamma, B^b, C^b}{a \vdash \Gamma, (B \vee C)^b}\vee \qquad \frac{a \vdash \Gamma, B^b \quad a \vdash \Gamma, C^b}{a \vdash \Gamma, (B \wedge C)^b}\wedge$$

$$(b \le \mathsf{x})\,\frac{a \vdash \Gamma, A(\sigma\mathsf{x}A(\mathsf{x}))^b}{a \vdash \Gamma, \sigma\mathsf{x}A^b}\sigma \qquad \frac{a \vdash \Gamma, A^b}{a \vdash \langle\mathfrak{a}\rangle\Gamma, [\mathfrak{a}]A^b}\mathsf{mod} \qquad \frac{a \vdash \Gamma}{a \vdash \Gamma, A^b}\mathsf{weak}$$

$$(\forall i \le k.\, a_i \sqsubset b_i \;\&\; b_i{\restriction_{a_0}} \sqsubset a_i)\,\frac{a_0 \vdash A_1^{a_1}, \dots, A_k^{a_k}}{b_0 \vdash A_1^{b_1}, \dots, A_k^{b_k}}\mathsf{exp}$$

Figure 4: Rules and axioms of *fixed-point logic with names*, FixN.

$$(x \in N_{\mathsf{x}} \;\&\; b \le \mathsf{x})\,\frac{ax \vdash \Gamma, A(\nu\mathsf{x}A)^{bx}}{a \vdash \Gamma, \nu\mathsf{x}A^b}\nu_x \qquad \begin{array}{c}[a \vdash \Gamma]^\dagger \\ \vdots \\ \dfrac{a \vdash \Gamma}{a \vdash \Gamma}\mathsf{dis}_\dagger\end{array}$$

$$(x_0, \dots, x_k \in N_x)\,\frac{a \vdash \Gamma, A_0^{bx}, \dots, A_k^{bx}}{a \vdash \Gamma, A_0^{bxx_0}, \dots, A_k^{bxx_k}}\mathsf{reset}_x$$

Figure 5: Additional inference rules for Stirling proofs. In the rules ν_x and reset_x, x may not occur in Γ.

An *open Stirling proof* is a finite derivation π in Stir equipped with a function $l \mapsto l^c$ mapping each discharged assumption to the conclusion of an associated instance of dis, fulfilling the conditions:

1. The control of every node is a non-repeating word in variable names;
2. For every annotated sequent $a \vdash \Gamma$ occurring in π and every annotated formula $B^b \in \Gamma$, $b \sqsubset a$;
3. For every discharged assumption l, there is a variable name x appearing in the control of every node on the path from l^c to l inclusive, and an application of reset_x on this path.

For each discharged assumption l the node l^c is called the *companion of* l and l a *companion leaf* of l^c. A *companion node* is a companion to some discharged assumption in π. If $\mathscr{A}$ is a set of annotated sequents, we write $\mathscr{A}, a \vdash_{\mathsf{Stir}} \Gamma$ if there is an open Stirling proof with conclusion $a \vdash \Gamma$ and open assumptions in $\mathscr{A}$. A *Stirling proof*, or Stir-*proof*, is a closed open Stirling proof, i.e. one with no open assumptions.

When presenting Stirling proofs, we illustrate the companion function c by labelling discharged assumptions and the corresponding application of dis (by symbols † and ‡) as in Figure 5. Note that in the presence of condition 2 above, the restriction accompanying applications of the expansion rule exp in Figure 4 simplifies to

$$(\forall i \le k.\, b_i{\restriction_{a_0}} = a_i)\,\frac{a_0 \vdash A_1^{a_1}, \dots, A_k^{a_k}}{b_0 \vdash A_1^{b_1}, \dots, A_k^{b_k}}\mathsf{exp}$$

The proof system Stir presented here differs slightly from [13] which was goal orientated with essentially deterministic inference rules when read 'bottom-up.' For instance, Stirling's system contains neither of the rules weak and exp: the former appears in a restricted form, called *thinning*, with premise $a \vdash \Gamma, B^b$ and conclusion $a \vdash \Gamma, B^b, B^c$ provided $b \prec_a c$ for a total ordering $\prec_a$ on annotations; the latter is instead incorporated directly into the other rules. It follows that each inference rule of Stirling's system can be simulated in Stir by a combination of an inference in Stir and (possibly) an application of exp and weak without affecting the requirements on proofs in Definition 3.8. Thus completeness for Stir is a corollary of the completeness proof in [13], using the argument in Theorem 3.6 to accommodate unguarded formulæ.

Theorem 3.9. *Let A be a closed well-named formula. If A is valid then $\mathsf{Stir} \vdash A$.*

Stirling's proof of Theorem 3.9 proceeds by building a partial FixN derivation, starting from the root, via a deterministic strategy for applying the inference rules. A cardinality argument ensures the construction terminates, yielding a Stirling proof iff the starting formula is valid.

Stirling also proves soundness of his proof system though what follows does not depend on it. We provide an alternative proof of soundness for Stir (and, as a consequence, for the proof system of [13]) based on embedding Stir-proofs in Koz.

We close the section with two examples of Stirling proofs: a generalisation of Lemma 3.7 to annotated sequents and the application of the above theorem to Example 3.2.

Lemma 3.10. *Let $a \in N^*$ be a non-repeating word and $b, c \sqsubset a$ annotations. Suppose $\nu\mathsf{x}_0 \cdots \nu\mathsf{x}_k A(\mathsf{x}_0, \ldots, \mathsf{x}_k)$ is a closed well-named formula. For each $i \leq k$, let b_i, c_i be the restriction of b and c respectively to names in $\bigcup_{\mathsf{y} \leq \mathsf{x}_i} N_{\mathsf{y}}$. For all sequences of closed formulæ $(B_i)_{i \leq k}$ and $(C_i)_{i \leq k}$,*

$$\{a \vdash B_i^{b_i}, C_i^{c_i}\}_{i \leq k}, a \vdash_{\mathsf{Stir}} \overline{A}(B_0, \ldots, B_k)^b, A(C_0, \ldots, C_k)^c.$$

Proof. As in Lemma 3.7, the proof proceeds by induction on the formula A. As before, the only non-trivial case is when A is quantified. Suppose $A(\mathsf{x}_0, \ldots, \mathsf{x}_k) = \nu\mathsf{y} A_0(\mathsf{x}_0, \ldots, \mathsf{x}_k, \mathsf{y})$. Without loss of generality we may assume $\mathsf{x}_i < \mathsf{y}$ for each $i \leq k$ and $a, b, c \leq \mathsf{y}$. Let

$$E = \overline{A}(B_0, \ldots, B_k) \qquad E_0(\mathsf{y}) = \overline{A_0}(B_0, \ldots, B_k, \mathsf{y})$$
$$F = A(C_0, \ldots, C_k) \qquad F_0(\mathsf{y}) = A_0(C_0, \ldots, C_k, \mathsf{y}).$$

Also, let $y, y' \in N_{\mathsf{y}}$ be names for y not occurring in a. Using the induction hypothesis, the derivation

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{a \vdash B_0^{b_0}, C_0^{c_0}}{ay \vdash B_0^{b_0}, C_0^{c_0}}\,\mathsf{exp} \quad \cdots \quad \dfrac{a \vdash B_k^{b_k}, C_k^{c_k}}{ay \vdash B_k^{b_k}, C_k^{c_k}}\,\mathsf{exp} \quad \dfrac{\dfrac{[ay \vdash E^b, F_0(F)^{cy}]^\dagger}{ayy' \vdash E^b, F_0(F)^{cyy'}}\,\mathsf{reset}_y}{ay \vdash E^b, F^{cy}}\,\nu_{y'}}{\vdots\ \text{I.H.}}}{ay \vdash E_0(E)^b, F_0(F)^{cy}}}{ay \vdash E^b, F_0(F)^{cy}}\,\mu}{\dfrac{ay \vdash E^b, F_0(F)^{cy}}{a \vdash E^b, F^c}\,\nu_y}\,\mathsf{dis}_\dagger \tag{4}$$

suffices for the lemma. □

Example 3.3. The tableaux π_∞ in Example 3.2 can be transformed into a Stirling proof π_{stir} of the annotated sequent $\epsilon \vdash \nu\mathsf{x}\mu\mathsf{y}\overline{B}^\epsilon, \nu\mathsf{y}\mu\mathsf{x}B^\epsilon$:

$$\dfrac{\dfrac{[xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y]^\dagger}{xyx' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{xx'}, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{reset}_x}{xy \vdash C^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\nu_{x'} \qquad \dfrac{\dfrac{[xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y]^\dagger}{xyy' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'}}\,\mathsf{reset}_y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, D^y}\,\nu_{y'}$$

$$\vdots \text{ (Lemma 3.10)}$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{xy \vdash \overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{dis}_\dagger}{x \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \nu\mathsf{y}\mu\mathsf{x}B^\epsilon}\,\nu_y}{\epsilon \vdash C^\epsilon, D^\epsilon}\,\nu_x$$

where $C = \nu\mathsf{x}\mu\mathsf{y}\overline{B}$, $D = \nu\mathsf{y}\mu\mathsf{x}B$, the missing inferences are provided by the previous lemma and the two leaves are discharged by the rule $\mathsf{dis}_\dagger$.

4 Unfolding Stirling proofs

In this section we establish two important closure properties of Stir-proofs, which we call *monotonicity* and *invariance*, that will prove critical for embedding Stir into other calculi. We begin with some preliminary observations.

Definition 4.1 (Leaf invariants)**.** Let π be a derivation in Stir and l a discharged assumptions in π. The *invariant* for l, denoted $inv_\pi(l)$, is the shortest initial segment of the control of l with the form ax such that x appears in the control of every node between l^c and l inclusive and there exists an application of reset_x at some node between l^c and l. If no such word exists, we set $inv_\pi(l) = \epsilon$. If $inv_\pi(l) = ax$, the conclusion to the first occurrence of reset_x on the path from l^c to l will be referred to as the *reset node for l*. If π is clear from the context we write $inv(l)$ in place of $inv_\pi(l)$.

We may rephrase the definition of Stir-proofs in terms of leaf invariants:

Lemma 4.2. *A closed derivation in* Stir *is a* Stir*-proof iff it satisfies conditions 1 and 2 in Definition 3.8 and every discharged leaf has non-trivial invariant.*

Proof. Consequence of the definition. □

Lemma 4.3. *Let π be an* Stir*-proof with plain conclusion.*

1. *If A^{axc} and B^{bxd} are annotated formulæ occurring in the same sequent in π then $a = b$.*

2. *If l is a discharged assumption in π then the control of every node between l^c and l is prefixed by the invariant of l.*

3. *If l_0 and l_1 are two leaves of π with the same invariant and associated the same companion node then their respective reset nodes are either the same or incomparable.*

Proof. 1 and 2 are proved by induction on π. The only non-trivial case to consider for 1 is an application of exp in π, which we may suppose has the form given in Figure 4. By assumption, the annotations $b_1, \ldots, b_k$ labelling formulæ in the conclusion have the property stated in 1. Suppose $a_i = c_i x d_i$ and $a_j = c_j x d_j$ for some $1 \leq i < j \leq k$. Since then x appears in the annotations b_i and b_j there is a single annotation e such that ex is a prefix of both b_i and b_j. By the condition on the application of exp, $c_i = e{\restriction_{a_0}} = c_j$.

3 is a consequence of the definition of leaf invariants and Lemma 4.2. □

As a discharged leaf and its associated companion are always labelled by the same annotated sequent, a Stirling proof can be unfolded by recursively replacing discharged assumptions by a fresh copy of their companion's sub-proof. Depending on how one associates companions to newly created leaves, different unfoldings of a single proof can be constructed. The following definition makes precise the operation of unfolding a leaf and selecting companions for new leaves, and the subsequent lemma establishes that Stirling proofs are closed under all forms of unfolding.

Definition 4.4 (Unfolding Stir-proofs)**.** Let π be a Stirling proof and M a non-empty set of discharged assumptions in π whose companion nodes form a set of pairwise incomparable nodes. Let O be a set of discharged leaves in π with the property that for every $o \in O$ there exists $m \in M$ such that $m^c \leq o^c < m$. We define the *O-unfolding of π at M* to be the closed derivation π' given by replacing each leaf $l \in M$ by a copy of the sub-proof of π at l^c in which assumptions in O are left open in the sub-proof whenever possible. This condition is made precise by the definition of the companion function c' of π'. Let $\hat{\cdot} \colon \pi' \to \pi$ be the function projecting π' back to π given by: if $n > m$ for some $m \in M$ then $\hat{n}$ is the node of π from which n was copied in the formation of π'; if $n \in M$ then $\hat{n} = n^c$; otherwise, $\hat{n} = n$. Fix an arbitrary discharged assumption l in π'. If l is a leaf in π then l has the same companion as $\hat{l}$ in π, i.e. $l^{c'} = \hat{l}^c$. Otherwise, $l > m$ for some $m \in M$ and $l^{c'}$ is defined according to the choice of O and position of $\hat{l}^c$:

- if $\hat{l}^c < m^c$ then $l^{c'} = \hat{l}^c$,
- if $m^c \leq \hat{l}^c < m$ and $\hat{l} \in O$ then also $l^{c'} = \hat{l}^c$,
- otherwise, $m^c \leq \hat{l}^c$ and $l^{c'}$ is chosen to be the unique node $m \leq o < l$ such that $\hat{o} = \hat{l}^c$.

An *unfolding* of π is the O-unfolding of π at M for some choice O and M.

If π' is an unfolding of π then every node that is a leaf in both π and π' is assigned the same companion in both derivations. We may therefore always assume the function assigning companions to leaves in π' extends the companion function for π and thus uniformly denote by l^c the companion of a node l in either proof. Moreover, since π is a sub-structure of π' we may assume a single accessibility relation, $<$, on nodes in both trees.

Example 4.1. Consider the Stir-proof π_{stir} of Example 3.3. Let l_0, l_1 name respectively the left and right leaf of π_{stir} and m mark the conclusion to the single dis rule in π_{stir}. There are twelve unfoldings of π_{stir}, given by picking $O, M \subseteq \{l_0, l_1\}$ (with M non-empty). The result with $O = \{l_0\}$ and $M = \{l_1\}$, for instance, is given in Figure 6 and obtained by replacing the leaf l_1 by a copy of the sub-proof of π_{stir} rooted at the conclusion to

$$
\dfrac{\dfrac{\left[xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y\right]^\dagger}{xyx' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{xx'}, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{reset}_x}{xy \vdash C^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\nu_{x'}
\qquad
\dfrac{\dfrac{\left[xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y\right]^\ddagger}{xyy' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'}}\,\mathsf{reset}_y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, D^y}\,\nu_{y'}
$$

$$
\begin{array}{c}
\text{- - - - - - - - - - - - - - - - - - - -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{xy \vdash \overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{dis}_\ddagger}{xyy' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^{yy'}}\,\mathsf{reset}_y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, D^y}\,\nu_{y'}
\end{array}
$$

$$
\dfrac{\dfrac{\left[xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y\right]^\dagger}{xyx' \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{xx'}, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{reset}_x}{xy \vdash C^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\nu_{x'}
$$

$$
\begin{array}{c}
\text{- - - - - - - - - - - - - - - - - - - -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{xy \vdash \overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, B(\mu\mathsf{x}B(\mathsf{x},D),D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mu}{xy \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x},D)^y}\,\mathsf{dis}_\dagger}{x \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^x, \nu\mathsf{y}\mu\mathsf{x}B^\epsilon}\,\nu_y}{\epsilon \vdash C^\epsilon, D^\epsilon}\,\nu_x
\end{array}
$$

Figure 6: The $\{l_0\}$-unfolding of π_{stir} at $\{l_1\}$, π^1_{stir}.

$\mathsf{dis}_\dagger$. The fresh copy of l_0 introduced above l_1 is assigned the original instance of dis as its companion. With the same M but $O = \{l_0, l_1\}$, the result is the same proof but in which all leaves are discharged by the instance of $\mathsf{dis}_\dagger$. With $M = \{l_0\}$ and $O = \{l_1\}$ the unfolding is defined in a symmetric manner, with l_0 unravelled and the fresh copy of l_1 inserted linked to the discharge rule near the root.

There are two extremes of unfolding π at a set of leaves M, the $\emptyset$-unfolding and the L-unfolding where L is the set of all discharged assumptions o for which there exists $m \in M$ such that $m^c \le o^c < m$. The former duplicates also the companion function, assigning fresh companions to new leaves if possible. The latter uses the most conservative choice of companion function possible, setting $l^c = \hat{l}^c$ for every discharged assumption l satisfying $\hat{l}^c < l$.

Lemma 4.5. *Every unfolding of a* Stir*-proof is a* Stir*-proof.*

Proof. Let π be a Stirling proof and π' be the O-unfolding of π at M for some choice of O and M. It suffices to show that every discharged assumption in π' has non-trivial invariant. Without loss of generality assume $M \subseteq c(m)$ for some companion node m in π. Fix a node $m' \in M$ and an arbitrary non-axiom leaf $l > m'$ in π'. Let l^c denote the companion of l in π' and $\hat{l}$ the projection of l into π. If $m' \le l^c < l$ then $inv_{\pi'}(l) = inv_\pi(\hat{l})$ by definition. Otherwise, $l^c < m' < l$ and it follows that l^c and m are comparable nodes, and l^c is the companion (in π) of the leaf $\hat{l}$. So $m \le l^c < m'$ or $l^c < m < \hat{l}$, whence Lemma 4.3(2) implies $inv_\pi(m')$ and $inv_\pi(\hat{l})$ are both prefixes of the control of one of l^c or m. Thus $inv_\pi(m')$ and $inv_\pi(\hat{l})$ are comparable and the shorter of the two is the invariant for l in π'. □

The next definition pins down two important regularity conditions on Stirling proofs.

The definition of an invariant Stir-proof is taken directly from [1]. There, also a version of monotonicity was formulated which, to simplify a technicality arising in Theorem 4.9 below, we strengthen slightly in the following definition.

Definition 4.6 (Invariant and monotone Stirling proofs)**.** A Stirling proof π is *invariant* if any two assumptions m, n with the same companion have the same invariant, and is *monotone* if for all pairs of assumptions m, n in π satisfying $m^c < n^c < m$ the invariant of m is a prefix of the invariant of n and $inv_\pi(m) = inv_\pi(n)$ implies the reset node for m is strictly below n^c. If π is an invariant Stir-proof, the function inv_π can be extended to companion nodes by setting $inv_\pi(l^c) = inv_\pi(l)$ for every assumption l.

Example 4.2. The Stirling proofs constructed by Lemma 3.10 are monotone and invariant. The Stir-proof π_{stir} in Example 3.3 is also monotone but not invariant as the two assumption leaves associated to $\mathsf{dis}_\dagger$ have distinct invariants, namely x for l_0 and xy for l_1. For each $i \in \{0, 1\}$, let π^i_{stir} be the $\{l_{1-i}\}$-unfolding of π_{stir} at $\{l_1\}$. π^0_{stir} is neither monotone nor invariant, but π^1_{stir} is both monotone and invariant.

In an invariant monotone proof the partial ordering of companion nodes by their invariants closely matches their ordering in the proof: for companion nodes $m < n$ in a monotone Stir-proof π, $inv_\pi(m)$ is a prefix of $inv_\pi(n)$ if $n < l$ for some $l \in c_\pi(m)$. There are many unfoldings that preserve invariance. An important class is given by the next lemma, which refines Lemma 4.5 above. In particular it shows that every $\emptyset$-unfolding of a monotone invariant proof is monotone invariant.

Lemma 4.7. *Let π be a Stirling proof, M_0 a non-empty set of pairwise incomparable companion nodes in π, $M \subseteq \bigcup\{c_\pi(m) \mid m \in M_0\}$ and O a set of assumptions in π. Let π' be the O-unfolding at M of π. For every $m \in M$ and assumption $l > m$ in π',*

1. *if $l^c \geq m$ then $inv_{\pi'}(l) = inv_\pi(\hat{l})$,*
2. *if $l^c < m$ then $inv_{\pi'}(l)$ is the longest common prefix of $inv_\pi(\hat{l})$ and $inv_\pi(m)$,*
3. *if π is invariant and for every $m \in M$ and $o \in O$, $inv_\pi(o)$ is a prefix of $inv_\pi(m)$, then $inv_{\pi'}(l) = inv_\pi(\hat{l})$ for every assumption l in π', and hence π' is invariant.*

Proof. 1 and 2 are established in the proof of Lemma 4.5; 3 is a special case of 1 and 2. ☐

The next two results constitute the main contribution of the present section, that the restriction to invariant monotone Stir-proofs is also complete for μ-calculus. It is established by repeatedly unfolding a given Stir-proof until first invariance, and then monotonicity, is obtained.

Lemma 4.8. *For every Stirling proof there exists an invariant Stirling proof with the same conclusion.*

Proof. Let π be a Stirling proof. Let m be a maximal node in π that fails the invariance property, i.e. there exist leaves $n_1, n_2 \in c_\pi(m)$ such that $inv_\pi(n_1) \neq inv_\pi(n_2)$. Define an equivalence relation on $c_\pi(m)$ by $n \sim_\pi n'$ if $inv_\pi(n) = inv_\pi(n')$, and let $L_0, \ldots, L_k$ enumerate the equivalence classes such that $inv(n)$ is a proper prefix of $inv(n')$ if $n \in L_i$, $n' \in L_j$ and $i < j$. Such an enumeration is possible as the invariant of every leaf with companion m is a prefix of the control at m. By assumption, $k > 0$. We prove there exists an unfolding π' of π such that no node in π' strictly above m fails the invariance property, and if m fails the invariance property then the equivalence relation $\sim_{\pi'}$ on $c_{\pi'}(m)$ has fewer equivalence classes.

Set $O = \bigcup_{i<k} L_i$ and let π' be the O-unfolding of π at L_k. Notice that the invariant of every assumption in O is a proper prefix of the invariant of nodes in L_k. By Lemma 4.7(1) it follows that every companion node in π' that is strictly above a node in L_k fulfils the invariance property. Let $L'_0, \dots, L'_{k'}$ enumerate the $\sim_{\pi'}$-equivalence classes of $c_{\pi'}(m)$. Applying Lemma 4.7(3) we conclude that $k' < k$ and for every $i < k$ there exists $j \le k'$ such that $L_i \subseteq L'_j$. □

Theorem 4.9. *Let* Γ *be a plain sequent.* $\mathsf{Stir} \vdash \Gamma$ *iff there exists an invariant monotone Stirling proof of* Γ*.*

Proof. By the previous lemma it suffice to prove that every invariant Stir-proof can be transformed into an invariant monotone Stir-proof. For each invariant proof π define a relation $\lhd_\pi$ between nodes in π that witnesses any failure of monotonicity: $m \lhd_\pi n$ iff m and n are both companion nodes and there exists $m' \in c_\pi(m)$ such that i) $m < n < m'$, ii) $inv_\pi(n)$ is a prefix of $inv_\pi(m)$, and iii) if $inv_\pi(n) = inv_\pi(m)$ then the reset associated to m' is above n. Observe that $\lhd_\pi$ is irreflexive and asymmetric, and that transitivity holds in some cases:

$$m \lhd_\pi o \lhd_\pi n \text{ and } n < m' \text{ for some } m' \in c(m) \text{ implies } m \lhd_\pi n.$$

Furthermore, for companion nodes m, n and o, the following closure property holds for $\lhd_\pi$:

$$\text{if } m < o < n \text{ and } inv_\pi(o) \text{ is a prefix of } inv_\pi(n) \text{ then } m \lhd_\pi n \text{ implies } m \lhd_\pi o. \tag{5}$$

Let the *rank* of a companion node m, written $rk_\pi(m)$, be the length of the longest $\lhd_\pi$-chain starting from m, i.e. the largest $r \ge 0$ for which there exist $m = m_0 < m_1 < \dots < m_r$ such that $m_i \lhd_\pi m_{i+1}$ for every $i < r$. If every node in a proof π has zero rank then π is monotone. For a set M of companion nodes in π define $c_\pi(M) = \bigcup_{m \in M} c_\pi(m)$, and for $\prec$ being $<$ or $\lhd_\pi$, write $n \prec M$ ($M \prec n$) if there exists $m \in M$ such that $n \prec m$ (resp. $m \prec n$).

The proof of the theorem now proceeds by induction on the maximal rank of nodes and a subsidiary induction on the number of nodes of maximal rank. Assume π is a closed invariant Stir-proof of Γ and let $r+1$ be the maximal rank among companion nodes in π. Fix a maximal companion node m_0 in π with rank $r+1$. We construct a sequence $\pi = \pi_0, \pi_1, \dots, \pi_k$ of invariant proofs such that for each $i < k$, π_{i+1} is an unfolding of π_i at a set of nodes whose companions are at or above m_0.

The particular sequence of proofs we consider is defined as follows. We write $\lhd_i$, inv_i, c_i and rk_i for $\lhd_{\pi_i}$, inv_{π_i}, c_{π_i} and rk_{π_i} respectively. Let $M_0 = \{m_0\}$. Suppose π_i has been defined and M_i is a set of companion nodes in π_i with non-zero rank. Define $O_{i+1} = \bigcup\{c_i(o) \mid M_i \lhd_i o\}$ and $M_{i+1} = \{m \in c_i(M_i) \mid \exists n(M_i \lhd_i n < m)\}$ which denote, respectively, the companion leaves to nodes $\lhd_i$-above M_i, and the set of companion leaves to a node in M_i witnessing non-zero rank of elements of M_i. If M_{i+1} is non-empty, set π_{i+1} to be the O_{i+1}-unfolding of π_i at M_{i+1}; otherwise we let $k = i$, which ends the sequence.

We claim that for each $i < k$ if $\hat{\cdot}\colon n \mapsto \hat{n}$ denotes the projection function from $\pi_{i+1} \to \pi_i$,

i) $c_{i+1}(M_i) = c_i(M_i) \setminus M_{i+1}$ and for every $j \leq i$, π_{i+1} differs from π_j only at nodes above M_{j+1};

ii) for every assumption l in π_{i+1}, $\hat{l}$ is an assumption in π_i and $inv_{i+1}(l) = inv_i(\hat{l})$ and hence π_{i+1} is invariant;

iii) $inv_{i+1}(m) = inv_{i+1}(m')$ for all $m, m' \in c_{i+1}(M_{i+1})$;

iv) $n \lhd_{i+1} o$ and $\hat{n} < \hat{o} < c_i(\hat{n})$ implies $\hat{n} \lhd_i \hat{o}$;

v) $n \lhd_{i+1} o$ implies $\hat{n} \lhd_i \hat{o}$;

vi) for every node n in π_{i+1}, $rk_{i+1}(n) \leq rk_i(\hat{n})$;

vii) for every node n in π_{i+1}, if $rk_{i+1}(n) = r + 1$ and $n \geq m_0$ then $n \in M_{i+1}$;

viii) for every $m \in M_i$, $|M_{i+1} \cap c_i(m)| \leq |c_0(m_0)| - i$.

We prove (i)–(vii) by simultaneous induction on $i < k$. Note that criterion (viii) implies the sequence of unfoldings terminates, with $k \leq |c_0(m_0)|$. Since M_{i+1} is empty iff $rk_i(m) = 0$ for every $m \in M_i$, (vi) and (vii) imply that π_k has fewer nodes of rank $r + 1$ than π.

Criterion (i) is a direct consequence of the definition of π_{i+1}, noting that $M_{i+1} \subseteq c_i(M_i)$ and since $\lhd_i$ is irreflexive, $c_i(M_i) \cap O_{i+1} = \emptyset$. To see (ii), observe that by the definition of $\lhd_i$, for any $o \in O_{i+1}$ there exists $m \in M_{i+1}$ such that the invariant of o in π_i is a prefix of the invariant of m in π_i, and apply Lemma 4.7(3) to the induction hypothesis for (iii) and the fact that π_i is invariant. (iii) then follows from (ii) and the induction hypothesis for (iii).

Concerning (iv), assume $n \lhd_{i+1} o$ and $\hat{n} < \hat{o} < c_i(\hat{n})$. There exists an assumption $l \in c_{i+1}(n)$ such that $o < l$ and $\hat{o} < \hat{l}$. Let l_r be the reset node for l. By (ii) $inv_{i+1}(l) = inv_i(\hat{l}) = inv_{i+1}(\hat{l})$, and $inv_{i+1}(n) = inv_i(\hat{n})$ is a prefix of $inv_{i+1}(o) = inv_i(\hat{o})$. If $inv_{i+1}(n)$ is a proper prefix of $inv_{i+1}(o)$ then $\hat{n} \lhd_i \hat{o}$. Otherwise the two invariants are identical and $n \lhd_{i+1} o$ implies $o < l_r$. If l is an assumption in π_i then $\hat{l} = l$ and $\hat{n} \lhd_i \hat{o}$ is immediate. Otherwise, let $m \in M_{i+1}$ be such that $m < l$. We perform a case distinction on the relative positions of m and o in π_{i+1}. By assumption, $\hat{n} < \hat{o} < c_i(\hat{n})$. If $m < o$ then $\hat{o} < \hat{l}_r < \hat{l} \in c_i(\hat{n})$ and $\hat{l}_r$ is the reset node for $\hat{l}$ in π_i as $inv_i(\hat{l}) = inv_{i+1}(l)$, so $\hat{n} \lhd_i \hat{o}$. If, instead, $o \leq m$ then $\hat{n} = n < o = \hat{o} < c_i(n)$, so Lemma 4.3(3) and the fact that $o < l_r$ implies $\hat{n} \lhd_i \hat{o}$.

We now turn to (v). First observe that for every $m \in M_i$, $rk_{i+1}(m) = 0$, for if $m \lhd_{i+1} o < m' \in c_{i+1}(m)$ then since $c_{i+1}(m) \subseteq c_i(m)$, we have $m \lhd_i o < m'$, meaning $m' \in M_{i+1}$ and contradicting (i). Thus suppose $n \lhd_{i+1} o$ and $n \notin M_i$. By (iv) it suffice to assume either $\hat{n} \not< \hat{o}$ or $\hat{n} < \hat{o} \not< c_i(\hat{n})$, from which it follows that there exists $m \in M_{i+1}$ such that $\hat{m} < n < m \leq o$. In particular, $\hat{n} = n$ and $c_i(n) \subseteq O_{i+1}$, i.e. $\hat{m} \lhd_i n$. Let $p > m$ be such that $\hat{p} = n$. We consider three cases: $m = o$, $m < o < p$ or $p < o$. In the first case $\hat{m} \lhd_i n \lhd_{i+1} m$, and a contradiction is reached by considering the position of the reset node for m in π_i. If $m < o < p$ then $\hat{m} < \hat{o} < n < o$ and, since $inv_i(\hat{o})$ is a prefix of $inv_i(n)$, equation (5) implies $\hat{m} \lhd_i \hat{o}$. But then $c_i(\hat{o}) \subseteq O_{i+1}$, which contradicts the fact that o is a companion node. This leaves the case $p < o$ from which we immediately deduce $\hat{n} \lhd_i \hat{o}$ by (iv).

(vi) is a simple proof by induction through π_{i+1} (starting from assumptions) appealing to (v); (vi) together with the fact that $rk_{i+1}(m) = 0$ for every $m \in M_i$ combine to prove (vii). Finally, to see (viii) fix $m \in M_i$. The result is trivial if $i = 0$, so assume $i = j + 1 > 0$. Let $m' \in M_j$ be such that $m \in c_j(m')$. The induction hypothesis implies $|M_i \cap c_j(m')| \leq |c_0(m_0)| - j$. By (v), the projection function $\hat{\cdot}$ injectively maps the set $M_{i+1} \cap c_i(m)$ into $M_i \cap c_j(m')$. Assume $l \in M_{i+1} \cap c_i(m)$ is such that $\hat{l} = m$. By the choice of M_{i+1} there exists a companion node n such that $m \lhd_{i+1} n < l$, and (v) implies $m' \lhd_i \hat{n} < m$, so $n \in O_i$ which contradicts that n is a companion node in π_{i+1}. Thus $|M_{i+1} \cap c_i(m)| < |M_i \cap c_j(m')|$ and we are done. □

Turning a Stir-proof into an invariant monotone proof via the above construction involves a substantial blow-up in size. Given a proof of height h the construction yields an invariant proof of height bounded by 2^h_h. Starting with an invariant proof of height h and maximal rank $r+1$, we obtain an invariant proof with rank r and height also bounded by 2^h_h. Since the transformation to invariant proof does not increase the rank of the starting derivation, the two bounds may be combined to deduce that the construction of invariant monotone proofs involves no worse than hyper-hyper-exponential increase in the height of Stirling proofs. However, a more efficient procedure for this transformation may be possible by interleaving the steps necessary for achieving monotonicity and invariance.

Before we conclude the present section, we observe the following fact about 'unused' reset rules in Stirling proofs.

Lemma 4.10. *For every monotone invariant* Stir*-proof there is a monotone invariant* Stir*-proof with the same conclusion such that for every node m, if m is the conclusion of an instance of* reset *then m is the rest node for some assumption $l > m$.*

Proof. The idea of the proof is to replace each 'unused' instance of reset by an application of exp. Fix a Stir-proof π and suppose

$$\frac{a \vdash \Gamma, A_0^{bx}, \ldots, A_k^{bx}}{a \vdash \Gamma, A_0^{bxx_0}, \ldots, A_k^{bxx_k}}\,\mathsf{reset}_x$$

is a reset occurring in π which is not the reset node associated to an assumption. By definition, $x_i \in N_x$ for each $i \leq k$ and x does not occur in Γ. We prove that this instance of reset can be removed at the cost of inserting instances of exp throughout π. Let m mark the conclusion of the rule in π. By condition 2 in the definition of Stir-proofs, for each $i \leq k$, x_i occurs in a, and it is for this reason that the rule cannot be directly replaced by an instance of exp. By Lemma 4.3(1), x_i occurs neither in Γ nor in the annotation b. Let $a' \sqsubset a$ be the result of removing from a the name x_i for each $i \leq k$. We replace the instance of reset_x by the corresponding instance of exp:

$$\frac{a' \vdash \Gamma, A_0^{bx}, \ldots, A_k^{bx}}{a \vdash \Gamma, A_0^{bxx_0}, \ldots, A_k^{bxx_k}}\,\mathsf{exp}$$

and recursively remove names $x_0, \ldots, x_k$ from the control of nodes working upwards through π. For the result to not be a Stirling proof at least one of two scenarios must now occur: i) there is an assumption $l > l^c > m$ such that the control at l contains some name from $x_0, \ldots, x_k$ which the control at l^c does not; ii) there is an assumption $l > m > l^c$ such that the control at l^c contains some name from $x_0, \ldots, x_k$ which the control at l does not. For l of type (i), simply insert a further instance of exp at l removing the offending names; for assumptions of type (ii), insert, immediately below l^c, an instance of exp. The latter case may injure other assumptions/companions above l^c in a similar manner, which are fixed by the addition of further applications of exp. Note the operation preserves monotonicity and invariance. ☐

5 Circular proofs

Stirling proofs, although cut-free, are not well-suited as a sequent calculus for the μ-calculus. Aside from requiring sequents to be annotated, the conditions on what constitutes a proof are non-local: an application of the discharge rule depends not on the form of the sequent being discharged but on properties of the whole path to an associated assumption. In the present section we provide a modification of Stir which permits a more natural definition of proof.

Partition each set N_{x} of variable names into two infinite sets: a set $N_{\mathsf{x}}^{\mathcal{A}}$ of *assumption names* (denoted $\hat{x}$, $\hat{y}$, etc.) and a set $N_{\mathsf{x}}^{\mathcal{V}}$, still called *variable names*. Let $N^{\mathcal{A}} = \bigcup_{\mathsf{x}\in\mathsf{Var}} N_{\mathsf{x}}^{\mathcal{A}}$; similarly for $N^{\mathcal{V}}$.

Definition 5.1 (Circular proofs)**.** Circ is the extension of FixN by the inference rules $\mathsf{dis}_{\hat{x}}$ for $\hat{x} \in N^{\mathcal{A}}$ and ν_x for $x \in N^{\mathcal{V}}$ given by

$$\begin{array}{c}[b\hat{x} \vdash \Gamma, A_0^{a_0\hat{x}x_0}, \ldots, A_k^{a_k\hat{x}x_k}]^{\hat{x}} \\ \vdots \\ (x_0,\ldots,x_k \in N_{\hat{x}}^{\mathcal{V}})\dfrac{b\hat{x} \vdash \Gamma, A_0^{a_0\hat{x}}, \ldots, A_k^{a_k\hat{x}}}{b \vdash \Gamma, A_0^{a_0}, \ldots, A_k^{a_k}}\mathsf{dis}_{\hat{x}}\end{array} \qquad (a \leq x \in N_{\mathsf{x}}^{\mathcal{V}})\dfrac{b \vdash \Gamma, A(\nu\mathsf{x}A)^{ax}}{b \vdash \Gamma, \nu\mathsf{x}A^a}\nu_x$$

with the restriction that in the ν_x inference, x may not appear in the side formulæ Γ. A Circ-*proof* is a finite closed derivation π satisfying:

1. the control of every node in π is a non-repeating word of assumption names,
2. there is at most one instance of $\mathsf{dis}_{\hat{x}}$ rule in π for each $\hat{x} \in N^{\mathcal{A}}$.

Given a finite set of annotated formulæ Γ we write $\mathsf{Circ} \vdash \Gamma$ if there exists a Circ-proof with conclusion $\epsilon \vdash \Gamma$.

The idea behind this new form of discharge rule in Circ is to separate what is essentially two distinct roles of variable names in Stirling proofs. On the one hand, names control the invariants associated to assumptions via the $x \in N$ for which a reset_x rule is applied. On the other, names record (via annotations) the unravelling of ν-quantifiers that are necessary for an application of reset. In Stir-proofs a given name may be utilised in both forms simultaneously: there may be assumptions l_1, l_2 with, say, $l_1^c < l_2^c < l_1$ and a name x which occurs in the invariant for l_1 but is eliminated on the path from l_2^c to l_2 by a reset rule. Although the restriction to monotone invariant Stir-proofs mitigates this particular problem, it will prove helpful to rephrase Stirling proofs within a framework that explicitly separates the two roles. In Circ, variable names are used solely to record ν-regenerations along threads and are excluded from the control; assumption names have the single purpose of providing invariants.

Lemma 5.2. *For every closed formula A and words $a \in N^* \cap N^{\mathcal{A}}$ and $b, c \leq \mathsf{Var}_A$ such that $a \vdash \overline{A}^b, A^c$ is an annotated sequent, there is a Circ-proof with conclusion $a \vdash \overline{A}^b, A^c$.*

Proof. The proofs constructed in Lemma 3.10 can be readily converted into Circ-proofs for this lemma. For example, the proof in (4) on page 16 with $A = \nu \mathsf{y} A_0(\mathsf{y})$, becomes

$$\dfrac{\dfrac{\dfrac{\dfrac{[a\hat{y} \vdash \overline{A}^b, A_0(A)^{c\hat{y}y}]^{\hat{y}}}{a\hat{y} \vdash \overline{A}^b, A^{c\hat{y}}}\nu_y \quad \vdots \text{ I.H.}}{\dfrac{a\hat{y} \vdash \overline{A_0}(\overline{A})^b, A_0(A)^{c\hat{y}}}{a\hat{y} \vdash \overline{A}^b, A_0(A)^{c\hat{y}}}\mu}}{a \vdash \overline{A}^b, A_0(A)^c}\mathsf{dis}_{\hat{y}}}{a \vdash \overline{A}^b, A^c}\nu$$

□

Fix a set $X \subseteq N^{\mathcal{V}}$ of variable names and let $\Gamma = \{A_1^{a_1}, \ldots, A_k^{a_k}\}$ and $\Delta = \{A_1^{b_1}, \ldots, A_k^{b_k}\}$ be two sets of annotated formulæ with the same underlying plain formulæ. We call Γ an *X-expansion* of Δ if for every $1 \le i \le k$ either $b_i = a_i$ or $b_i = b\hat{x}b'$ for some $\hat{x} \in N^{\mathcal{A}}$ and there exists $x \in X \cap N_{\hat{x}}$ such that $a_i = b\hat{x}xb'$.

The next lemma constitutes an important step in embedding Stir into Circ.

Lemma 5.3 (Expansion Lemma). *Let π be a circular proof of $a \vdash \Gamma$, possibly with open assumptions, and x a formal variable symbol. For every set $X \subseteq N_{\mathsf{x}}^{\mathcal{V}}$ of variable names not occurring in π and every X-expansion Γ' of Γ there exists a circular proof π' of $a \vdash \Gamma'$ such that every open assumption of π' is an X-expansion of an open assumption in π. Moreover, π' is such that if*

$$\dfrac{c_0 \vdash \Pi', A^c}{b_0 \vdash \Pi, A^b}\mathsf{exp}$$

is an instance of exp *in π' and b contains a name $x \in X$ not occurring in c, then b has the form $d\hat{x}xd'$ with $\hat{x} \in N_{\mathsf{x}}^{\mathcal{A}}$ and $c \sqsubset dd'$.*

Proof. The proof proceeds by induction on π via a case distinction on the last applied inference. For most inferences the desired circular proof π' follows directly from the induction hypothesis. The interesting case is if the last inference in π is an instance of $\mathsf{dis}_{\hat{y}}$ for an assumption name $\hat{y} \in N^{\mathcal{A}}$. Suppose therefore that π is an open proof witnessing $\mathscr{A}, a \vdash \Gamma$ for some finite set $\mathscr{A} = \{c_i \vdash \Lambda_i \mid i \in I\}$ of annotated sequents and $X \subseteq N_{\mathsf{x}}^{\mathcal{V}}$ is a finite set of variable names not occurring in π. We may assume Γ has the form $A_0^{a_0}, \ldots, A_k^{a_k}, \Delta$ and π ends with an application of $\mathsf{dis}_{\hat{y}}$ with $A_0, \ldots, A_k$ principal:

$$\pi \left\{ \begin{array}{c} c_i \vdash \Lambda_i \ (i \in I) \qquad [a\hat{y} \vdash A_0^{a_0\hat{y}y_0}, \ldots, A_k^{a_k\hat{y}y_k}, \Delta]^{\hat{y}} \\ \vdots \qquad\qquad\qquad\qquad \vdots \\ \text{- - - - - - - - - - - - - - - -} \\ \vdots\ \hat{\pi} \\ \dfrac{a\hat{y} \vdash A_0^{a_0\hat{y}}, \ldots, A_k^{a_k\hat{y}}, \Delta}{a \vdash A_0^{a_0}, \ldots, A_k^{a_k}, \Delta}\mathsf{dis}_{\hat{y}} \end{array} \right.$$

Let $\hat{\pi}$ be the immediate sub-proof of π whose open assumptions are each either sequents in $\mathscr{A}$ or of the form $a\hat{y} \vdash A_0^{a_0\hat{y}y_0}, \ldots, A_k^{a_k\hat{y}y_k}, \Delta$.

Let $\Delta = \{A^{a_{k+1}}_{k+1}, \dots, A^{a_{k+l}}_{k+l}\}$, $\alpha = (a_i)_{i \le k+l}$ and let $\vec{y} = (\hat{y}, \dots, \hat{y})$ be of length $k+1$. Given sequences of annotations $\gamma = (c_0, \dots, c_r)$, $\beta = (b_0, \dots, b_{r+s})$ and $\delta = (d_i)_{i \le k+l}$ define

$$\beta \circ \gamma = (b_0 c_0, \dots, b_r c_r, b_{r+1}, \dots, b_{r+s}), \qquad\qquad \Gamma^{\delta} = \{A_i^{d_i} \mid i \le l+k\}$$

so in particular $\Gamma = \Gamma^{\alpha}$ and the conclusion of $\hat{\pi}$ is $a\hat{y} \vdash \Gamma^{\alpha \circ \vec{y}}$. If β and γ each have length $k+l+1$ and Γ^{β} is an X-expansion of Γ^{γ} we call β an *X-expansion* of γ. The induction hypothesis yields, for every X-expansion β of α, a proof $\hat{\pi}_{\beta}$ with conclusion $a\hat{y} \vdash \Gamma^{\beta \circ \vec{y}}$ such that open assumptions in $\hat{\pi}_{\beta}$ are X-expansions of open assumptions in $\hat{\pi}$. Let $\mathscr{X}$ be the finite set of X-expansions of α.

Fix an X-expansion β of α. To obtain desired proof π' in the case $\Gamma' = \Gamma^{\beta}$ we first define, for every $\delta \in \mathscr{X}$ and $\mathscr{Y} \subseteq \mathscr{X}$, a circular proof $\hat{\pi}_{\delta,\mathscr{Y}}$ with conclusion $a\hat{y} \vdash \Gamma^{\delta \circ \vec{y}}$ such that every open assumption in $\hat{\pi}_{\delta,\mathscr{Y}}$ is either the X-expansion of an open assumption in π or has the form $a\hat{y} \vdash \Gamma^{\eta \circ (\hat{y}y_0, \dots, \hat{y}y_k)}$ for some $\eta \in \mathscr{Y}$ and some $y_0, \dots, y_k \in N^{\mathcal{V}}_{\hat{y}}$. Then π' will be defined as $\hat{\pi}_{\beta,\emptyset}$ in which the assumption name $\hat{y}$ is removed from all sequents at which it appears. Set $\hat{\pi}_{\delta,\mathscr{X}} = \hat{\pi}_{\delta}$. The definition of $\hat{\pi}_{\delta,\mathscr{Y}}$ for $\mathscr{Y} \subsetneq \mathscr{X}$ has two cases depending on whether $\delta \in \mathscr{Y}$. Suppose $\delta \notin \mathscr{Y}$ and let $\hat{x} \in N_{\hat{y}} \cap N^{\mathcal{A}}$ be a fresh assumption variable not occurring in $\hat{\pi}_{\delta,\mathscr{Y}\cup\{\delta\}}$ that names the same variable as $\hat{y}$ and let $\vec{x} = (\hat{x}, \dots, \hat{x})$ be of length $k+1$. Define $\hat{\pi}_{\delta,\mathscr{Y}}$ to be the proof obtained from $\hat{\pi}_{\delta,\mathscr{Y}\cup\{\delta\}}$ by: i) replacing all occurrences of $\hat{y}$ by $\hat{y}\hat{x}$; ii) inserting an application of $\mathsf{dis}_{\hat{x}}$ at the root discharging every open assumptions of the form $a\hat{y}\hat{x} \vdash \Gamma^{\delta \circ (\hat{y}\hat{x}x_0, \dots, \hat{y}\hat{x}x_k)}$ for $x_0, \dots, x_k \in N_{\hat{x}}$; iii) appending to each remaining open assumption an application of exp removing all occurrences of $\hat{x}$:

$$\hat{\pi}_{\delta,\mathscr{Y}} \left\{ \begin{array}{c} (\eta \ne \delta)\,\dfrac{a\hat{y} \vdash \Gamma^{\eta \circ \vec{y} \circ (y_0,\dots,y_k)}}{a\hat{y}\hat{x} \vdash \Gamma^{\eta \circ \vec{y} \circ \vec{x} \circ (y_0,\dots,y_k)}}\,\mathsf{exp} \qquad (\Delta \ne \Gamma)\,\dfrac{c' \vdash \Delta'}{c \vdash \Delta}\,\mathsf{exp} \qquad \left[a\hat{y}\hat{x} \vdash \Gamma^{\delta \circ \vec{y} \circ \vec{x} \circ (x'_0,\dots,x'_k)}\right]^{\hat{x}} \\ \vdots \qquad\qquad\qquad \vdots \qquad\qquad\qquad \vdots \\ \text{- -} \\ \vdots\ \hat{\pi}_{\delta,\mathscr{Y}\cup\{\delta\}} \\ \dfrac{a\hat{y}\hat{x} \vdash \Gamma^{\delta \circ \vec{y} \circ \vec{x}}}{a\hat{y} \vdash \Gamma^{\gamma \circ \vec{y}}}\,\mathsf{dis}_{\hat{x}} \end{array} \right.$$

where in the above $c' \vdash \Delta'$ denotes the result of removing $\hat{x}$ from the control and annotations of the sequent $c \vdash \Delta$. For $\delta \in \mathscr{Y}$, the proof $\hat{\pi}_{\delta,\mathscr{Y}}$ is defined in terms of $\{\hat{\pi}_{\gamma,\mathscr{Y}} \mid \gamma \in \mathscr{X} \setminus \mathscr{Y}\}$ as the proof derived from $\hat{\pi}_{\delta}$ by attaching, at any leaf $a\hat{y} \vdash \Gamma^{\gamma \circ \vec{y}}$ for $\gamma \notin \mathscr{Y}$, a copy of $\hat{\pi}_{\gamma,\mathscr{Y}}$:

$$\begin{array}{c} \begin{array}{cc} & \vdots\ \hat{\pi}_{\gamma,\mathscr{Y}} \\ a\hat{y} \vdash \Gamma^{\gamma \circ \vec{y} \circ (y_0,\dots,y_l)}\ (\gamma \in \mathscr{Y}) \qquad & (\gamma \notin \mathscr{Y})\,\dfrac{a\hat{y} \vdash \Gamma^{\gamma \circ \vec{y}}}{a\hat{y} \vdash \Gamma^{\gamma \circ \vec{y} \circ (y'_0,\dots,y'_l)}}\,\mathsf{exp} \\ \vdots & \vdots \end{array} \\ \text{- -} \\ \vdots\ \hat{\pi}_{\delta} \\ a\hat{y} \vdash \Gamma^{\delta \circ \vec{y}} \end{array}$$

As mentioned above, set π' to be the proof $\hat{\pi}_{\beta,\emptyset}$ in which the assumption name $\hat{y}$ is removed from all sequents at which it appears. □

Theorem 5.4. *If* $\mathsf{Stir} \vdash \Pi$ *and* Π *is a plain sequent then* $\mathsf{Circ} \vdash \Pi$.

Proof. Fix a closed Stir-proof π with plain conclusion and let ρ denote the root of π. By Theorem 4.9 we may assume π is invariant and monotone. Moreover, we may assume that every conclusion to an instance of dis is the companion to some assumption, and by Lemma 4.10, every instance of reset in π is the reset node for some assumption.

We construct, for all pairs of nodes $n \leq m$ in π where n is a companion node, a circular proof π^n_m by recursion on m. The interpretation of π will then be the circular proof π^ρ_ρ where ρ is the root of π. The role of the parameter n is to distinguish between the two ways of proceeding when m marks a reset node in π, which are needed to ensure that an open assumption has the appropriate form when the recursion reaches its companion node. Specifically, if m marks a companion node in π with premise m' then π^n_m is defined in terms of both $\pi^n_{m'}$ and $\pi^m_{m'}$.

The properties we require of π^n_m are:

1. the conclusion of π^n_m is the sequent labelling m in π, where every name in the control is considered an assumption name;
2. there exists a function f from open assumptions in π^n_m to non-axiom leaves in π such that every open assumption l in π^n_m can be written in the form

$$bx \vdash \Gamma, A_1^{axx_1}, \ldots, A_r^{axx_r}, A_{r+1}^{ax}, \ldots, A_s^{ax} \tag{6}$$

   where x is an assumption name not occurring in Γ, $x_1, \ldots, x_r \in N_x \cap N^{\mathcal{V}}$ are variable names naming the same ν-variable as x, and
   a) the sequent at $f(l)$ has the form $bxb' \vdash \Gamma, A_1^{axa_1}, \ldots, A_s^{axa_s}$ for words b', $a_1, \ldots, a_s$,
   b) the invariant of $f(l)$ is bx,
   c) if the reset node for $f(l)$ in π is $\geq m$ then
      - $f(l) \in c_\pi(n)$ implies $r = s$;
      - $f(l) \notin c_\pi(n)$ implies that for every $1 \leq i \leq r$ there exists either a node $o < l$ in π^n_m or a node $f(l)^c < o < m$ in π such that x_i does not occur in the sequent at o.

Criterion 2 stipulates the necessary information regarding the 'repeat and reset' condition in Stirling proofs so that the discharge rule from Stir can be interpreted as a discharge rule in Circ. The definition of π^n_m proceeds via a case distinction on the inference rule applied to derive m in π:

- axiom case. π^n_m is taken to be the same sequent.
- assumption. In this case $n < m$ and m has the form

$$bxb' \vdash \Gamma, A_1^{axx_1a_1}, \ldots, A_r^{axx_ra_r}, A_{r+1}^{axa_{r+1}}, \ldots, A_s^{axa_s}$$

  where $bx = inv(m)$, $x_1, \ldots, x_r$ all name the same variable as x, x does not appear in Γ and $x < a_i$ for each $r < i \leq s$. The proof π^n_m is the open derivation

$$\dfrac{bx \vdash \Gamma, A_1^{axx_1}, \ldots, A_r^{axx_r}, A_{r+1}^{ax}, \ldots, A_s^{ax}}{bxb' \vdash \Gamma, A_1^{axx_1a_1}, \ldots, A_r^{axx_ra_r}, A_{r+1}^{axa_{r+1}}, \ldots, A_s^{axa_s}}\,\mathsf{exp}$$

which trivially satisfies the necessary requirements.

- exp, $\vee$, $\wedge$, μ, ν, mod or weak. Immediate via induction hypothesis.

- ν_y for $y \in N_{\mathsf{y}}$. Suppose m and its unique successor m' are the sequents $b \vdash \Delta, (\nu \mathsf{y} A(\mathsf{y}))^a$ and $by \vdash \Delta, A(\nu \mathsf{y} A)^{ay}$ respectively for some formula $A(\mathsf{y})$ and annotation $a \sqsubset b$. We define π^n_m to be the result of removing the name y from the control (but not annotations) of all sequents in $\pi^n_{m'}$ and inserting an application of ν_y at the root.

- reset_y for $y \in N$. Let m' be the unique successor to m in π. The sequents at m and m' have the form $aya' \vdash \Delta, B_0^{byy_0}, \ldots, B_k^{byy_k}$ and $aya' \vdash \Delta, B_0^{by}, \ldots, B_k^{by}$ respectively where y does not appear in Δ and $y_i \in N_y$ for each $i \leq k$. By assumption m is the reset node for some open assumption l in the sub-proof of π rooted at m: let $n_0 = l^c$.

  Let $\pi^n_{m'}$ and its associate function f' be given satisfying 1 and 2. By the Expansion Lemma there exists, for each $Y \subseteq \{0, \ldots, k\}$, a circular proof $\pi^n_{m',Y}$ with conclusion $aya' \vdash \Delta, B_0^{byc_0}, \ldots, B_k^{byc_k}$ where $c_i = y_i$ if $i \in Y$ and $c_i = \epsilon$ otherwise. All open assumptions in $\pi^n_{m',Y}$ are $\{y_i \mid i \in Y\}$-expansions of open assumptions of $\pi^n_{m'}$, so f' naturally induces a function f^Y from assumptions in $\pi^n_{m',Y}$ to assumptions in π that fulfils all the requirements needed of π^n_m except, possibly, the case of 2c for an open assumption l such that $f^Y(l) \in c_\pi(n_0)$. Satisfying this final requirement depends on the appropriate choice of Y. If $n = n_0$ set $Y = \{0, \ldots, k\}$. Otherwise, let Y be the subset recording the variable names that do not appear in the control of some sequent between n_0 and m in π. Without loss of generality assume $Y = \{k', k'+1, \ldots, k\}$, and define π^n_m to be $\pi^n_{m',Y}$ with an application of exp attached at the root that inserts the name y_i for each $i < k'$:

$$\pi^n_m \left\{ \begin{array}{c} \pi^n_{m',Y} \\ \vdots \\ \dfrac{aya' \vdash \Delta, B_0^{by}, \ldots, B_{k'-1}^{by}, B_{k'}^{byy_{k'}}, \ldots, B_k^{byy_k}}{aya' \vdash \Delta, B_0^{byy_0}, \ldots, B_{k'-1}^{byy_{k'-1}}, B_{k'}^{byy_{k'}}, \ldots, B_k^{byy_k}}\,\mathsf{exp} \end{array} \right.$$

  In either case define $f = f^Y$. To see that condition 2c is now fulfilled, suppose $f(l) \in c_\pi(n_0)$. Then m is the reset node associated to $f(l)$ in π and the invariant of $f(l)$ is ay. Let $r \leq s$ be such that the sequent at $f(l)$ has the form

$$ay \vdash \Gamma, A_1^{ayy'_1}, \ldots, A_r^{ayy'_r}, A_{r+1}^{ay}, \ldots, A_s^{ay}$$

  where $y'_i \in N_y \cap N^{\mathcal{V}}$ for each $1 \leq i \leq r$ and y does not occur in Γ. If $n = n_0$ then the choice of $Y = \{0, \ldots, k\}$ in the formation of π^n_m implies $r = s$ by the second part of the Expansion Lemma. Assume therefore $n \neq n_0$ and fix $i \in \{1, \ldots, r\}$. If there exists a sequent on the path from n_0 to m in π at which y'_i does not occur, pick o to be this node. Otherwise, y'_i must occur in the sequent at m, so y'_i appears in the root of π^n_m and $y'_i = y_j$ for some $j < k'$, whence we let o be the conclusion of $\pi^n_{m',Y}$.

- dis. Suppose the sequent at m is

$$bxb' \vdash \Gamma, A_1^{axx_1a_1}, \ldots, A_r^{axx_ra_r}, A_{r+1}^{axa_{r+1}}, \ldots, A_s^{axa_s} \tag{7}$$

and $inv(m) = bx$, x does not appear in Γ, $x_i \in N_x$ for each $1 \leq i \leq r$, and for each $r < i \leq s$, a_i does not start with a name in N_x. Let m' be the unique successor to m in π and let f' be the function fulfilling requirement 2 for $\pi^n_{m'}$ given by the induction hypothesis. It suffices to construct a circular proof of $bx \vdash \Gamma, A_1^{ax}, \ldots, A_s^{ax}$ satisfying requirement 2 since then a Circ-derivation with conclusion (7) can be obtain by inserting an application of exp at the root. Set $L = \{l \mid f'(l) \in c_\pi(m)\}$. Fix a fresh assumption variable $\hat{x}$ and let $\hat{\pi}^n_m$ denote the open circular proof of $bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}}, \ldots, A_s^{ax\hat{x}}$ obtained from $\pi^n_{m'}$ by i) recursively removing b', $x_1, \ldots, x_r$ and $a_1, \ldots, a_s$ from the root upwards, ii) replacing x throughout by $x\hat{x}$, and iii) inserting applications of exp eliminating $\hat{x}$ at every open assumption $l \notin L$. Monotonicity of π ensures $\hat{\pi}^n_m$ has the same underlying structure and inference rules as $\pi^n_{m'}$ and condition 2c of the induction hypothesis implies that the operation leaves all open assumptions unchanged except for assumptions in L, for which the control and annotations are affected by the substitution $x \mapsto x\hat{x}$. By the above observation $\hat{\pi}^n_m$ clearly fulfils condition 2 for any open assumption $l \notin L$. Let $l \in L$ be an open assumption in $\hat{\pi}^n_m$ and suppose l has the form given in (6). If $r = s$ then this assumption can be discharged by appending an instance of $\mathsf{dis}_{\hat{x}}$ to the root of $\hat{\pi}^n_m$. Indeed, if all assumptions from L have this form then we define π^n_m as the result of inserting an application of $\mathsf{dis}_{\hat{x}}$ at the root of $\hat{\pi}^n_m$:

$$
\cfrac{\cfrac{\begin{array}{c}\left[bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}x_1}, \ldots, A_s^{ax\hat{x}x_s}\right]^{\hat{x}} \\ \vdots\ \hat{\pi}^n_m \\ bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}}, \ldots, A_s^{ax\hat{x}}\end{array}}{bx \vdash \Gamma, A_1^{ax}, \ldots, A_s^{ax}}\ \mathsf{dis}_{\hat{x}}}{bxb' \vdash \Gamma, A_1^{axa_1}, \ldots, A_s^{axa_s}}\ \mathsf{exp}
$$

Otherwise there are assumptions in L that do not have the correct form to be discharged by $\mathsf{dis}_{\hat{x}}$, which are instead replaced by a copy of the proof $\hat{\pi}^m_m$:

$$
\cfrac{\cfrac{\begin{array}{c}\begin{array}{ccc}\begin{array}{c}\left[bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}x_1}, \ldots, A_s^{ax\hat{x}x_s}\right]^{\hat{x}} \\ \vdots\end{array} & & \begin{array}{c}\left[bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}x_1}, \ldots, A_s^{ax\hat{x}x_s}\right]^{\hat{x}} \\ \vdots\ \hat{\pi}^m_m \\ (r < s)\ \cfrac{bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}}, \ldots, A_s^{ax\hat{x}}}{bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}x_1}, \ldots, A_r^{ax\hat{x}x_r}, A_{r+1}^{ax\hat{x}}, \ldots, A_s^{ax\hat{x}}}\ \mathsf{exp} \\ \vdots\end{array}\end{array} \\ \text{- - - - - - - - - - - - - - - - - - - -} \\ \vdots\ \hat{\pi}^n_m \\ bx\hat{x} \vdash \Gamma, A_1^{ax\hat{x}}, \ldots, A_s^{ax\hat{x}}\end{array}}{bx \vdash \Gamma, A_1^{ax}, \ldots, A_s^{ax}}\ \mathsf{dis}_{\hat{x}}}{bxb' \vdash \Gamma, A_1^{axa_1}, \ldots, A_s^{axa_s}}\ \mathsf{exp}
$$

Condition 2 for π^m_m ensures that π^n_m satisfies all the necessary requirements. □

Example 5.1. Recall the monotone invariant Stirling proof π^1_{stir} in Figure 6. Let $C = \nu \mathsf{x} \mu \mathsf{y} \overline{B}$ and $D = \nu \mathsf{y} \mu \mathsf{x} B$. The circular proof obtained by following the construction of Theorem 5.4 is the proof π_{circ} in Figure 7. The omitted inferences in π_{circ} are the result of applying the transformation of the previous theorem to derivations provided by Lemma 3.10, which are in essence the proofs given by Lemma 5.2. The proof has the same underlying structure as π^1_{stir} with instances of the reset rule removed and the

$$
\dfrac{\dfrac{\left[\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}x'}, \mu\mathsf{x}B(\mathsf{x},D)^{y}\right]^{\hat{x}}}{\hat{x}\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}x'}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'\hat{y}}}\,\mathsf{exp}}{\hat{x}\hat{y} \vdash C^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'\hat{y}}}\,\nu_{x'}
\qquad
\dfrac{\left[\hat{x}\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'\hat{y}y''}\right]^{\hat{y}}}{\hat{x}\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, D^{yy'\hat{y}}}\,\nu_{y''}
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\hat{x}\hat{y} \vdash \overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y}))^{x\hat{x}}, B(\mu\mathsf{x}B(\mathsf{x},D),D)^{yy'\hat{y}}}{\hat{x}\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'\hat{y}}}\,\mu,\mu}{\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{yy'}}\,\mathsf{dis}_{\hat{y}}}{\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, D^{y}}\,\nu_{y'}
\end{array}
$$

$$
\dfrac{\left[\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}x'}, \mu\mathsf{x}B(\mathsf{x},D)^{y}\right]^{\hat{x}}}{\hat{x} \vdash C^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{y}}\,\nu_{x'}
\qquad
\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, D^{y}
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\dfrac{\hat{x} \vdash \overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y}))^{x\hat{x}}, B(\mu\mathsf{x}B(\mathsf{x},D),D)^{y}}{\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{y}}\,\mu,\mu}{\epsilon \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \mu\mathsf{x}B(\mathsf{x},D)^{y}}\,\mathsf{dis}_{\hat{x}}}{\epsilon \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \nu\mathsf{y}\mu\mathsf{x}B}\,\nu_{y}}{\epsilon \vdash \nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B}\,\nu_{x}
\end{array}
$$

Figure 7: The circular proof π_{circ} obtained by applying Theorem 5.4 to π^1_{stir}.

$$
\dfrac{\left[\hat{y}\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}x'}, \mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}y}\right]^{\hat{x}}}{\hat{y}\hat{x} \vdash C^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}y}}\,\nu_{x'}
\qquad
\dfrac{\left[\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, D^{\hat{y}y}\right]^{\hat{y}}}{\hat{y}\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, D^{\hat{y}y}}\,\mathsf{exp}
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\hat{y}\hat{x} \vdash \overline{B}(C, \mu\mathsf{y}\overline{B}(C,\mathsf{y}))^{x\hat{x}}, B(\mu\mathsf{x}B(\mathsf{x},D),D)^{\hat{y}y}}{\hat{y}\hat{x} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x\hat{x}}, \mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}y}}\,\mu,\mu}{\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}y}}\,\mathsf{dis}_{\hat{x}}}{\hat{y} \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \nu\mathsf{y}\mu\mathsf{x}B^{\hat{y}}}\,\nu_{y}}{\epsilon \vdash \mu\mathsf{y}\overline{B}(C,\mathsf{y})^{x}, \nu\mathsf{y}\mu\mathsf{x}B}\,\mathsf{dis}_{\hat{y}}}{\epsilon \vdash \nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B}\,\nu_{x}
\end{array}
$$

Figure 8: A simplification of π_{circ}.

relevant annotations propagating upwards. In particular, for this example no further unravellings of the form described in the proof of the Expansion Lemma are required.

Note that π_{circ} is not optimal. Shorter Circ-proofs of the sequent can be obtained by moving the application of the $\mathsf{dis}_{\hat{y}}$ inference either immediately below the $\nu_{y'}$ inference, or below the ν_y rule at the root. In the latter case, the proof in Figure 8 obtains. However, this proof could not result from a Stirling proof due to the form of the reset rules.

The proof of Lemma 5.3 and Theorem 5.4 make heavy use of the fact that, like Stirling proofs, circular proofs can be unfolded by identifying closed assumptions and their associated discharge rules. This operation becomes highly relevant in the next section where we prove that the discharge and annotated ν inferences can be combined into a single inference rule capturing regenerating ν-threads. We conclude this section with an analysis of the unravelling of circular proofs.

Let a *pre-proof* be a closed pre-derivation satisfying the two requirements of Circ-proofs

$$\left.\begin{array}{c} [b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}x_0}, \dots, A_k^{a_k\hat{x}x_k}]^{\hat{x}} \\ \vdots\, \pi_{\hat{x}} \\ \dfrac{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}}, \dots, A_k^{a_k\hat{x}}}{b \vdash \Delta, A_0^{a_0}, \dots, A_k^{a_k}}\,\mathsf{dis}_{\hat{x}} \\ \vdots \\ a \vdash \Gamma \end{array}\right\}\pi \mapsto \left.\begin{array}{c} \vdots\, \pi_{\hat{x}} \\ \dfrac{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}}, \dots, A_k^{a_k\hat{x}}}{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}x_0}, \dots, A_k^{a_k\hat{x}x_k}}\,\mathsf{exp} \\ \vdots\, \pi_{\hat{x}} \\ \dfrac{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}}, \dots, A_k^{a_k\hat{x}}}{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}x_0}, \dots, A_k^{a_k\hat{x}x_k}}\,\mathsf{exp} \\ \vdots\, \pi_{\hat{x}} \\ \dfrac{b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}}, \dots, A_k^{a_k\hat{x}}}{b \vdash \Delta, A_0^{a_0}, \dots, A_k^{a_k}}\,\mathsf{dis}_{\hat{x}} \\ \vdots \\ a \vdash \Gamma \end{array}\right\}\pi^*$$

Figure 9: Unravelling a circular proof π to a pre-proof π^*. $\pi_{\hat{x}}$ denotes the sub-proof of π rooted at the premise of $\mathsf{dis}_{\hat{x}}$.

in Definition 5.1, i.e. a closed Circ-proof with the finiteness condition relaxed.

Definition 5.5 (Unravelling circular proofs)**.** Fix a circular proof π with conclusion $a \vdash \Gamma$. For each assumption name $\hat{x}$ in π not in a, let $\pi_{\hat{x}}$ denote the sub-proof of π rooted at the premise of the instance of $\mathsf{dis}_{\hat{x}}$. The *unravelling of* π, denoted π^*, is the *pre-proof* given by recursively replacing every discharged assumption of the form $b\hat{x} \vdash \Delta, A_0^{a_0\hat{x}x_1}, \dots, A_k^{a_k\hat{x}x_k}$ by a fresh copy of $\pi_{\hat{x}}$ (i.e. renaming any discharged assumption names) with its conclusion expanded to the same sequent by an application of exp, as illustrated in Figure 9.

The unravelling of a Circ-proof π is essentially the $\emptyset$-unfolding from Stirling proofs iterated to the limit. The paths in π^* are in bijection with the finite sequences of nodes in π that trace the unravelling process. We refer to the latter sequences as paths *through* π.

Definition 5.6 (Paths through Circ-proofs)**.** Let π be a Circ-proof. A *path through* π is a sequence $\alpha = (\alpha_i)_{i<N}$ ($N \leq \omega$) of nodes in π such that α_0 is the root of π and for every $i < N$ exactly one of the following conditions hold.

1. α_i is an axiom or open assumption and $N = i + 1$;
2. α_i is a discharged assumption and α_{i+1} is the premise of the companion instance of dis;
3. α_i is not a leaf and α_{i+1} is an immediate successor of α_i.

A path $\alpha = (\alpha_i)_{i<N}$ is *infinite* if $N = \omega$ and *finite* otherwise, and the set of infinite paths through π is denoted $\mathfrak{P}_\pi$.

Given an infinite path $\alpha = (\alpha_i)_{i<\omega}$ through π, denote by $\alpha^* = (\alpha_i^*)_{i<\omega}$ the corresponding path through the unravelling of π. Notice that nodes in the unravelling of π

are uniquely named. In particular, if $\alpha = (\alpha_i)_{i<\omega}$ and $\beta = (\beta_i)_{i<\omega}$ are paths through the same circular proof π, and $\alpha_s^* = \beta_t^*$ for some $s, t < \omega$ then $s = t$ and $\alpha_i = \beta_i$ for all $i \leq s$. Without loss of generality we may therefore identify nodes in π^* with finite paths through π.

The next lemma confirms that pre-proofs obtained by unravelling a circular proof are indeed tableaux.

Lemma 5.7. *Let π be a circular proof. For every infinite path $\alpha = (\alpha_i)_{i<\omega}$ through π there exists $s < \omega$, formal variable x and name $\hat{x} \in N_{\mathsf{x}}^{\mathcal{A}}$ such that*

1. *α_s is the premise of an application of $\mathsf{dis}_{\hat{x}}$,*
2. *$\alpha_s = \alpha_j$ for infinitely many j,*
3. *for all $j \geq s$, α_j is a node in the sub-proof of π rooted at α_s and the control of α_s is a prefix of the control of α_j.*

Moreover, x and $\hat{x}$ are uniquely determined by α, and there exists a ν-thread $(B_i^{b_i})_{i<\omega}$ such that for every $i \geq s$, $B_i^{b_i}$ occurs in the sequent at α_i and $b_s\hat{x}$ is a prefix of b_i.

Proof. Fix a circular proof π and let $\alpha = (\alpha_i)_{i<\omega}$ be a path through π. The first part of the lemma is straightforward, as is the fact that x and $\hat{x}$ are uniquely determined. We prove that α contains a ν-thread. Let $s < \omega$, x and $\hat{x}$ be as above, and suppose the instance of $\mathsf{dis}_{\hat{x}}$ with premise α_s has the form

$$\dfrac{\begin{array}{c}\vdots \\ b\hat{x} \vdash \Gamma, A_0^{a\hat{x}}, \ldots, A_k^{a\hat{x}}\end{array}}{b \vdash \Gamma, A_0^{a}, \ldots, A_k^{a}}\,\mathsf{dis}_{\hat{x}}$$

Let $s < s_0 < s_1 < \ldots$ enumerate the occurrences of the node α_s in α and let $\mathfrak{T}_i$ denote the collection of finite threads $(B_j^{b_j})_{j \leq s_i}$ through α such that $B_{s_i}^{b_{s_i}} \in \{A_0^{a\hat{x}}, \ldots, A_k^{a\hat{x}}\}$. $\bigcup_{i<\omega} \mathfrak{T}_i$ forms an infinite, finitely branching tree under the prefix ordering so, by König's Lemma, contains an infinite branch which is a ν-thread through α. □

The above lemma concisely captures the motivation for Circ and its advantage over Stirling proofs. Dropping requirement 3, the lemma also holds for unravellings of Stirling proofs and is not difficult to prove. For monotone invariant Stir-proofs a weakened form of 3 is provable replacing 'control of α_s' by 'invariant of α_s' (the latter being well-defined by invariance). However, in the next section we make essential use of Lemma 5.7 and, particularly, requirement 3.

6 Circular proofs with ν-closure

We now turn to the task of finding an alternative form for discharge rules that more succinctly describes the closure properties for the ν-quantifier. This will be in the form of a fresh rule of inference, called *ν-closure*, that generalises the annotated ν inference by permitting the discharging of assumptions. The ν-closure rule will replace both the annotated ν-rules and discharge rules of the preceding sections, and although sequents in the new calculus, denoted Clo, are still annotated, they will be in much simpler form than either Stir or Circ.

Definition 6.1 (Proofs with ν-closure)**.** Let Clo be the proof system expanding FixN by the ν-clo inference

$$\dfrac{\begin{matrix}[\vdash \Gamma, \nu \mathsf{x} A^{ax}]^{x} \\ \vdots \\ \vdash \Gamma, A(\nu \mathsf{x} A)^{ax}\end{matrix}}{\vdash \Gamma, \nu \mathsf{x} A^{a}}\ \nu\text{-}\mathsf{clo}_x \qquad (a \leq x \in N_{\mathsf{x}})$$

with the restriction that x does not appear in Γ. A Clo-proof is a closed derivation using the rules of FixN and ν-clo in which all sequents are annotated sequents with empty control and there is at most one use of ν-clo_x for each $x \in N$.

Recall that FixN, and hence also Clo, does not contain the rules ν_x present in either Stir or Circ. The rule is admissible in Clo however as it corresponds to an application of ν-clo_x that does not discharge any assumptions.

Since controls in Clo-proofs are empty, the restriction on applications of exp in Clo reduces to checking the sub-word relation only, i.e. for the instance of exp given in Figure 4, the condition becomes $a_i \sqsubset b_i$ for every $i \leq k$. To simplify presentation, from now on we identify a finite set Γ of closed annotated μ-formulæ with the annotated sequent $\epsilon \vdash \Gamma$.

Lemma 6.2. *For every formula $A(\mathsf{x}_0, \ldots, \mathsf{x}_{k-1})$ with free variables among x_0, …, x_{k-1}, all closed formulæ B_i, C_i for $i < k$ and all annotations b, c such that $b, c \leq \mathsf{x}$ for each ν-variable x in A,*

$$\{B_i^b, C_i^c\}_{i<k} \vdash_{\mathsf{Clo}} \overline{A}(B_0, \ldots, B_{k-1})^b, A(C_0, \ldots, C_{k-1})^c.$$

Proof. We deal with the quantifier case and assume $k = 1$, the other cases and the generalisation to $k > 0$ are straightforward and follow the same principles. Let $A(\mathsf{y}) = \nu\mathsf{x} A_0(\mathsf{x}, \mathsf{y})$ and fix two closed formulæ B and C and annotations b, c satisfying the hypothesis of the lemma. Arguing by induction, assume

$$\{\overline{A}(B)^b, A(C)^{cx}\}, \{B^b, C^{cx}\} \vdash_{\mathsf{Clo}} \overline{A}_0(\overline{A}(B), B)^b, A_0(A(C), C)^{cx}.$$

where x is a fresh name for x. An application of μ yields

$$\{\overline{A}(B)^b, A(C)^{cx}\}, \{B^b, C^{cx}\} \vdash_{\mathsf{Clo}} \overline{A}(B)^b, A_0(A(C), C)^{cx}$$

and an application of ν-clo at the root and exp at the remaining open assumptions implies $\{B^b, C^c\} \vdash_{\mathsf{Clo}} \overline{A}(B)^b, A(C)^c$. □

Theorem 6.3. *Let Γ be a plain sequent. If $\mathsf{Circ} \vdash \Gamma$ then $\mathsf{Clo} \vdash \Gamma$.*

Proof. Fix a circular proof π with conclusion Γ and let π^* be the unravelling of π as described in Definition 5.5. We prove that a finite sub-tree of π^* encodes a Clo-proof. Let $\mathfrak{P}_\pi$ be the set of paths through the circular proof π. Recall that each path $\alpha = (\alpha_i)_{i<\omega} \in \mathfrak{P}_\pi$ induces a path α^* in π^* which we identify with the infinite sequence $(\alpha_i^*)_{i<\omega}$ where $\alpha_i^* = (\alpha_j)_{j \leq i}$.

We begin by defining functions $f, g \colon \mathfrak{P}_\pi \to \omega$ that select appropriate positions for instances of ν-clo and associated assumptions. Fix $\alpha = (m_i)_{i<\omega} \in \mathfrak{P}_\pi$ and let $s < \omega$, x, $\hat{x} \in N_{\mathsf{x}}$ and $\beta = (B_i^{b_i})_{i<\omega}$ be as given by Lemma 5.7. Since β is a thread through α there exists $s \leq t < t' < \omega$, $a \leq \mathsf{x}$, $x \in N_{\mathsf{x}}$ and a formula A such that

1. α_t is the conclusion of an application of ν_x in which $B_t^{b_t} = \nu\mathsf{x}A^{a\hat{x}}$ is the principal formula,

2. $(\alpha_t, B_t, b_t) = (\alpha_{t'}, B_{t'}, b_{t'})$.

Fix (t', t) minimally with respect to lexicographical ordering (i.e. $(i, j) < (i', j')$ if $i < i'$, or $i = i'$ and $j < j'$) satisfying 1 and 2 above and define $f(\alpha) = t$ and $g(\alpha) = t'$. In particular, the two functions are continuous in their input: if $\alpha = (\alpha_i)_{i<\omega}$ and $\gamma = (\gamma_i)_{i<\omega}$ are both paths through π such that $\alpha_i = \gamma_i$ for every $i < g(\alpha)$ then $f(\alpha) = f(\gamma)$ and $g(\alpha) = g(\gamma)$. Let $G = \{\alpha^*_{g(\alpha)} \mid \alpha \in \mathfrak{P}_\pi\}$ be the set of π^* nodes distinguished by g, and similarly $F = \{\alpha^*_{f(\alpha)} \mid \alpha \in \mathfrak{P}_\pi\}$. By the choice of g, G is a set of pairwise incomparable nodes in π^* and every node in F is below some node in G. Let π_G denote the sub-tree of π^* truncated at the frontier G, and define a companion relation on π_G such that a leaf $l = \alpha^*_{g(\alpha)} \in G$ is associated companion node $l^c = \alpha^*_{f(\alpha)}$.

By the choice of f, for every leaf $l \in \pi_G$, l^c is the conclusion to an application of ν_{x_l} for some $x_l \in N^{\mathcal{V}}$ and there is a thread through the path from l^c to l mapping the principal formula at l^c to an identical formula in l. Along each such thread allow the variable name x_l to persist to l by adjusting each application of the exp rule affecting the thread and inserting fresh instances of exp on diverging paths to revert the annotation and eliminate x_l. By removing any remaining dis inferences, along with all assumption names in π_G, and replacing each ν_y inference by ν-clo_y (discharging the assigned companion leaves in G), a Clo-proof with conclusion Γ results. □

Example 6.1. Define the width of a discharge rule dis_x to be the number of formulæ in the premise that have x in their annotation. For example, the instance of dis_x given in Definition 5.1 has width $k + 1$ and the two discharge rules in the proof π_{circ} in Figure 7 each have width 1. In order to obtain Theorem 6.3 it suffices for each discharge rule in a circular proof to be unravelled a number of times equal to its width. In the case of π_{circ} this means that unravelling each of the discharge rules once yields a Circ-proof π'_{circ} containing a Clo-proof.

We now describe the construction of a Clo-proof from π_{circ}. Before proceeding observe that the inference ν_x in π_{circ} sits below the $\mathsf{dis}_{\hat{x}}$ rule which affects the same thread. As a result, during the embedding into Clo the ν_x rule will be replaced by an instance of ν-clo *without* associated assumptions, and as such the variable name x occurring throughout π_{circ} plays no further role and ν_x may be dropped in favour of an (unannotated) ν rule. Likewise, the instances of ν_y and $\nu_{y'}$ are below the corresponding $\mathsf{dis}_{\hat{y}}$ inference and their annotations may also be eliminated. Note that these changes are purely cosmetic as the variables x, y and y' will play no part in the transformation of π_{circ} into a Clo-proof. To further ease presentation, we also drop displaying the control from the sequents in the Circ-derivation. The version of π_{circ} we will work with is therefore

$$
\dfrac{\dfrac{\left[\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)\right]^{\hat{x}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mathsf{exp}}{C^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\nu_x
\qquad
\dfrac{\left[\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}y}\right]^{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \nu \mathsf{y}\mu \mathsf{x} B^{\hat{y}}}\,\nu_y
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu \mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}}, B(\mu \mathsf{x} B(\mathsf{x},D),D)^{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mu,\mu}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\mathsf{dis}_{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \nu \mathsf{y}\mu \mathsf{x} B}\,\nu
\end{array}
$$

$$
\dfrac{\left[\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)\right]^{\hat{x}}}{C^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\nu_x
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu \mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}}, B(\mu \mathsf{x} B(\mathsf{x},D),D)}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\mu,\mu}{\mu \mathsf{y}\overline{B}(C,\mathsf{y}), \mu \mathsf{x} B(\mathsf{x},D)}\,\mathsf{dis}_{\hat{x}}}{\nu \mathsf{x}\mu \mathsf{y}\overline{B}, \nu \mathsf{y}\mu \mathsf{x} B}\,\nu,\nu
\end{array}
$$

Unravelling the $\mathsf{dis}_{\hat{x}}$ and $\mathsf{dis}_{\hat{y}}$ rules yields the proof π'_{circ}

$$
\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi'\\ \mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)\end{array}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mathsf{exp}}{C^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\nu_x
\qquad
\dfrac{\left[\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}y'}\right]^{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \nu \mathsf{y}\mu \mathsf{x} B^{\hat{y}}}\,\nu_{y'}
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu \mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}}, B(\mu \mathsf{x} B(\mathsf{x},D),D)^{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mu,\mu}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}y}}\,\mathsf{exp}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \nu \mathsf{y}\mu \mathsf{x} B^{\hat{y}}}\,\nu_y
\end{array}
$$

$$
\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi'\\ \mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)\end{array}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mathsf{exp}}{C^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\nu_x
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu \mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}}, B(\mu \mathsf{x} B(\mathsf{x},D),D)^{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)^{\hat{y}}}\,\mu,\mu}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\mathsf{dis}_{\hat{y}}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \nu \mathsf{y}\mu \mathsf{x} B}\,\nu
\end{array}
$$

$$
\dfrac{\dfrac{\begin{array}{c}\vdots\,\pi'\\ \mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)\end{array}}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x}, \mu \mathsf{x} B(\mathsf{x},D)}\,\mathsf{exp}}{C^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\nu_x
$$

$$
\begin{array}{c}
\text{- -}\\
\vdots\\
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu \mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}}, B(\mu \mathsf{x} B(\mathsf{x},D),D)}{\mu \mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}}, \mu \mathsf{x} B(\mathsf{x},D)}\,\mu,\mu}{\mu \mathsf{y}\overline{B}(C,\mathsf{y}), \mu \mathsf{x} B(\mathsf{x},D)}\,\mathsf{dis}_{\hat{x}}}{\nu \mathsf{x}\mu \mathsf{y}\overline{B}, \nu \mathsf{y}\mu \mathsf{x} B}\,\nu,\nu
\end{array}
$$

where π' denotes the sub-proof of π_{circ} rooted at the premise to the $\mathsf{dis}_{\hat{x}}$ rule in which the instance of $\mathsf{dis}_{\hat{y}}$ has been unravelled (and $\hat{y}$ renamed $\hat{y}'$) and the conclusion suitably expanded, namely,

$$
\dfrac{\dfrac{\left[\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x'},\mu\mathsf{x}B(\mathsf{x},D)\right]^{\hat{x}}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x'},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}\,\mathsf{exp}}{C^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}\,\nu_{x'}
\qquad
\dfrac{\left[\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'y'''}\right]^{\hat{y}'}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},D^{\hat{y}'}}\,\nu_{y'''}
$$

$$\vdots$$

$$
\dfrac{\left[\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x'},\mu\mathsf{x}B(\mathsf{x},D)\right]^{\hat{x}}}{\dfrac{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x'},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}{C^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}\,\nu_{x'}}\,\mathsf{exp}
\qquad
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu\mathsf{y}B(C,\mathsf{y}))^{\hat{x}},B(\mu\mathsf{x}B(\mathsf{x},D),D)^{\hat{y}'}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'y''}}\,\mathsf{exp}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},D^{\hat{y}'}}\,\nu_{y''}
$$

$$\vdots$$

$$
\dfrac{\left[\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}x'},\mu\mathsf{x}B(\mathsf{x},D)\right]^{\hat{x}}}{C^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)}\,\nu_{x'}
\qquad
\dfrac{\dfrac{\dfrac{\overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}},B(\mu\mathsf{x}B(\mathsf{x},D),D)^{\hat{y}'}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)^{\hat{y}'}}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)}\,\mathsf{dis}_{\hat{y}'}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},D}\,\nu
$$

$$\vdots$$

$$
\dfrac{\overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y}))^{\hat{x}},B(\mu\mathsf{x}B(\mathsf{x},D),D)}{\mu\mathsf{y}\overline{B}(C,\mathsf{y})^{\hat{x}},\mu\mathsf{x}B(\mathsf{x},D)}\,\mu,\mu
$$

The next step in the translation is to isolate the ν-threads satisfying properties 1 and 2 in the proof of Theorem 6.3. The set labelled G marking the discharged assumptions in the final Clo-proof will comprise the conclusions to the three $\nu_{y'''}$ and nine $\nu_{x'}$ inferences in π'_{circ} as well as the conclusion to the $\nu_{y'}$ inference. The corresponding companion nodes are the three instances of $\nu_{x'}$ and $\nu_{y''}$ and the single instance of ν_y. Finally, we expand the annotations along the threads, replace all remaining annotated ν inferences by ν-clo inferences and remove the remaining discharge rules. The resulting Clo-proof, which we name π_{clo}, is:

$$
\dfrac{\overset{\vdots\;\pi'_{clo}}{C,\mu\mathsf{x}B(\mathsf{x},D)}}{C,\mu\mathsf{x}B(\mathsf{x},D)^{y}}\,\mathsf{exp}
\qquad
\left[\mu\mathsf{y}\overline{B}(C,\mathsf{y}),D^{y}\right]^{y}
$$

$$\vdots$$

$$
\overset{\vdots\;\pi'_{clo}}{C,\mu\mathsf{x}B(\mathsf{x},D)}
\qquad
\dfrac{\dfrac{\overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y})),B(\mu\mathsf{x}B(\mathsf{x},D),D)^{y}}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}),\mu\mathsf{x}B(\mathsf{x},D)^{y}}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}),D}\,\nu\text{-}\mathsf{clo}_y
$$

$$\vdots$$

$$
\overset{\vdots\;\pi'_{clo}}{C,\mu\mathsf{x}B(\mathsf{x},D)}
\qquad
\dfrac{\dfrac{\overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y})),B(\mu\mathsf{x}B(\mathsf{x},D),D)}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}),\mu\mathsf{x}B(\mathsf{x},D)}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}),D}\,\nu
$$

$$\vdots$$

$$
\dfrac{\dfrac{\overline{B}(C,\mu\mathsf{y}\overline{B}(C,\mathsf{y})),B(\mu\mathsf{x}B(\mathsf{x},D),D)}{\mu\mathsf{y}\overline{B}(C,\mathsf{y}),\mu\mathsf{x}B(\mathsf{x},D)}\,\mu,\mu}{\nu\mathsf{x}\mu\mathsf{y}\overline{B},\nu\mathsf{y}\mu\mathsf{x}B}\,\nu,\nu
$$

where π'_{clo} is

$$
\cfrac{\cfrac{\big[C^x, \mu\mathsf{x}B(\mathsf{x}, D)\big]^x}{C^x, \mu\mathsf{x}B(\mathsf{x}, D)^{y'}}\,\mathsf{exp} \qquad \big[\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, D^{y'}\big]^{y'}}{\cfrac{\vdots}{\cfrac{\cfrac{\overline{B}(C, \mu\mathsf{y}\overline{B}(C, \mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x}, D), D)^{y'}}{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x}, D)^{y'}}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, D}\,\nu\text{-}\mathsf{clo}_{y'}}}
$$

$$
\cfrac{\big[C^x, \mu\mathsf{x}B(\mathsf{x}, D)\big]^x \qquad \mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, D}{\cfrac{\vdots}{\cfrac{\cfrac{\overline{B}(C, \mu\mathsf{y}\overline{B}(C, \mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x}, D), D)}{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x}, D)}\,\mu,\mu}{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, D}\,\nu}}
$$

$$
\cfrac{\big[C^x, \mu\mathsf{x}B(\mathsf{x}, D)\big]^x \qquad \mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, D}{\cfrac{\vdots}{\cfrac{\cfrac{\overline{B}(C, \mu\mathsf{y}\overline{B}(C, \mathsf{y}))^x, B(\mu\mathsf{x}B(\mathsf{x}, D), D)}{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x, \mu\mathsf{x}B(\mathsf{x}, D)}\,\mu,\mu}{C, \mu\mathsf{x}B(\mathsf{x}, D)}\,\nu\text{-}\mathsf{clo}_x}}
$$

As in Example 5.1 the above proof is not optimal. We leave it as an exercise to the reader to find a more concise version of π_{clo}.

7 Strengthened induction and the embedding into Koz

The final step is to embed the annotated proof system Clo into the plain sequent calculus $\mathsf{Koz}^-_\mathsf{s}$, the variant of Koz^- with the strengthened induction rule ind_s of Figure 3 in place of ind, and to embed the latter calculus within Koz. Note that $\mathsf{Koz}^-_\mathsf{s}$ may be viewed as an extension of Koz^- since the induction rule ind is admissible in $\mathsf{Koz}^-_\mathsf{s}$ via the following combination of inferences.

$$
\cfrac{\cfrac{\cfrac{\cfrac{\Gamma, A(\overline{\Gamma})}{\Gamma, A(\overline{\Gamma}), A(\nu\mathsf{x}A(\overline{\Gamma} \vee \mathsf{x}))}\,\mathsf{weak}}{\Gamma, A(\overline{\Gamma} \vee \nu\mathsf{x}A(\overline{\Gamma} \vee \mathsf{x}))}\,\vee_d}{\Gamma, \nu\mathsf{x}A(\overline{\Gamma} \vee \mathsf{x})}\,\nu}{\Gamma, \nu\mathsf{x}A(\mathsf{x})}\,\mathsf{ind}_s
$$

Theorem 7.1. *Let A be a closed well-named formula. If* $\mathsf{Clo} \vdash A$ *then* $\mathsf{Koz}^-_\mathsf{s} \vdash A$.

Given a closed circular proof we will define a translation $*\colon A^a \mapsto A^{a*}$ of annotated formulæ into plain formulæ such that the $*$-translation of each assumption in π is a valid sequent derivable in Koz^-, each inference rule of FixN is admissible in $\mathsf{Fix} + \vee_d$, and each instance of ν-closure is admissible in $\mathsf{Koz}^-_\mathsf{s}$. As we will see, the definition of $*$ depends only on the contexts in which variable names are utilised in the circular proof.

Proof of Theorem 7.1. Fix an assignment $\Phi = \{A_x \mid x \in N\}$ of closed μ-formulæ to variable names. We define the associated translation $*$ induced by Φ by structural induction. Variables and other atoms will be unchanged by $*$. As a result, given an annotated formula $B(\mathsf{x})^b$ and a plain formula C, we may write $B^{b*}(C)$ for the result of replacing all occurrences of x by C in B^{b*}. For atoms, modal and logical connectives define

$$
(B \vee C)^{a*} = B^{a*} \vee C^{a*} \qquad ([\mathfrak{a}]B)^{a*} = [\mathfrak{a}]B^{a*}
$$
$$
(B \wedge C)^{a*} = B^{a*} \wedge C^{a*} \qquad (\langle\mathfrak{a}\rangle B)^{a*} = \langle\mathfrak{a}\rangle B^{a*}
$$

For quantified formulæ, suppose $\mathsf{x} \in \mathsf{Var}$ and $a = bx_1 \cdots x_k c$ is an annotation where $x_1, \ldots, x_k \in N_\mathsf{x}$ and $b < \mathsf{x} < c$. Then for $B = B(\mathsf{x})$ we set

$$(\mu\mathsf{x}B)^{a*} = \mu\mathsf{x}\, B^{b*} \qquad (\nu\mathsf{x}B)^{a*} = A_{x_1} \vee \cdots \vee A_{x_k} \vee \nu\mathsf{x}.\, B^{b*}(A_{x_1} \vee \cdots \vee A_{x_k} \vee \mathsf{x})$$

Given a set Γ of annotated formulæ we set $\Gamma^* = \{A^{a*} \mid A^a \in \Gamma\}$.

Observe that regardless of the choice of Φ (provided variable names are associated *closed* μ-formulæ), the interpretation described above translates the inferences of FixN to Koz^--derivations. The only non-trivial case is the translations of ν and exp inferences: the former simply becomes an application of the ν fixed point rule in Fix and the latter rule becomes a series of $\vee^d$ inferences. Since $*$ is also the identity on plain formulæ (i.e. $A^{\epsilon *} = A$ for every μ-formula A), it remains to show how an appropriate choice of Φ allows the interpretation of discharged assumptions as derivable sequents and instances of ν-clo as admissible rules in $\mathsf{Koz}_\mathsf{s}^-$.

Fix a closed Clo-proof π with plain conclusion Γ, say. Without loss of generality we may assume each variable name occurring in π is uniquely associated exactly one instance of the ν-clo rule that uses this name. Let N_π denote the set of variable names used in π. By recursion through the ν-clo rules in π we define the assignment $\Phi_\pi = \{A_x \mid x \in N_\pi\}$: for $x \in N_\pi$, set $A_x = \overline{\Delta^*}$ where Δ is the set of side formulæ of the instance of ν-clo_x in π. This is well-defined since a) the conclusion to an instance of ν-clo_x in π may only contain a name y if ν-clo_y is closer to the root than ν-clo_x; and b) the definition of B^{b*} depends only on the choice of A_x for x appearing in b.

We prove by induction through π that for each annotated sequent Δ in π, $\mathsf{Koz}_\mathsf{s}^- \vdash \Delta^*$. As mentioned above, we need only concern ourselves with discharged assumptions and instances of ν-clo. For each $x \in N_\pi$ let Γ_x denote the set of side formulæ to the instance of ν-clo_x in π, so $A_x = \overline{\Gamma_x^*}$. If Δ is a discharged assumption in π then $\Delta = \Gamma_x \cup \{\nu\mathsf{x}A^{ax}\}$ for some $x \in N_\mathsf{x}$ and $a \leq \mathsf{x}$, and $(\nu\mathsf{x}A)^{ax*} = \overline{\Gamma_x^*} \vee B$ for some closed μ-formula B. Lemma 3.3 implies $\mathsf{Koz}^- \vdash \Gamma_x^*, \overline{\Gamma_x^*}$, whence an application of weak and $\vee$ yields $\mathsf{Koz}^- \vdash \Delta^*$. For interpreting ν-clo, suppose $\Delta = \Gamma_x \cup \{\nu\mathsf{x}A(\mathsf{x})^a\}$ and, by the induction hypothesis, the sequent $\Gamma_x^*, A(\nu\mathsf{x}A(\mathsf{x}))^{ax*}$ is derivable in $\mathsf{Koz}_\mathsf{s}^-$, where $x \in N_\mathsf{x}$ and $a \leq \mathsf{x}$. Let $b < \mathsf{x}$ and $x_1, \ldots, x_k \in N_\mathsf{x}$ be such that $a = bx_1 \cdots x_k$ and set

$$B(\mathsf{x}) = \overline{\Gamma_{x_k}^*} \vee \cdots \vee \overline{\Gamma_{x_1}^*} \vee \mathsf{x} \qquad\qquad A'(\mathsf{x}) = A^{b*}(B(\mathsf{x})).$$

Then

$$\begin{aligned} (\nu\mathsf{x}A)^{a*} &= B(\nu\mathsf{x}A') & A(\nu\mathsf{x}A)^{ax*} &= A^{b*}((\nu\mathsf{x}A)^{ax*}) \\ (\nu\mathsf{x}A)^{ax*} &= B(\overline{\Gamma_x^*} \vee \nu\mathsf{x}.\, A'(\overline{\Gamma_x^*} \vee \mathsf{x})) & &= A'(\overline{\Gamma_x^*} \vee \nu\mathsf{x}.\, A'(\overline{\Gamma_x^*} \vee \mathsf{x})) \end{aligned}$$

and the following derivation interprets ν-clo_x:

$$\cfrac{\cfrac{\cfrac{\Gamma_x^*, A'(\overline{\Gamma_x^*} \vee \nu\mathsf{x}.\, A'(\overline{\Gamma_x^*} \vee \mathsf{x}))}{\Gamma_x^*, \nu\mathsf{x}A'(\overline{\Gamma_x^*} \vee \mathsf{x})}\,\nu}{\Gamma_x^*, \nu\mathsf{x}A'}\,\mathsf{ind}_s}{\Gamma_x^*, B(\nu\mathsf{x}A')}\,\mathsf{weak} + \vee^* \tag{8}$$

Thus $\mathsf{Koz}_\mathsf{s}^- \vdash \Delta^*$ is derivable for each sequent $\vdash \Delta$ occurring in π and since $\Gamma^* = \Gamma$, $\mathsf{Koz}_\mathsf{s}^- \vdash \Gamma$. $\square$

Example 7.1. Applying Theorem 7.1 to the proof π_{clo} of Example 6.1 yields a cut-free proof of the sequent $\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B$ in $\mathsf{Koz}_\mathsf{s}^-$. Recall the formulæ $C = \nu\mathsf{x}\mu\mathsf{y}\overline{B}$ and $D = \nu\mathsf{y}\mu\mathsf{x}B$. The interpretation of annotated formulæ is given by

$$\begin{array}{lll}
\Gamma_x = \{\mu\mathsf{x}B(\mathsf{x}, D)^\epsilon\} & \Gamma_y = \{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^\epsilon\} & \Gamma_{y'} = \{\mu\mathsf{y}\overline{B}(C, \mathsf{y})^x\} \\
\overline{\Gamma_x^*} = \nu\mathsf{x}\overline{B}(\mathsf{x}, \overline{D}) & \overline{\Gamma_y^*} = \nu\mathsf{y}B(\overline{C}, \mathsf{y}) & \overline{\Gamma_{y'}^*} = \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y})
\end{array}$$

$$\begin{aligned}
C^{x*} &= \nu\mathsf{x}\overline{B}(\mathsf{x}, \overline{D}) \vee \nu\mathsf{x}\mu\mathsf{y}.\, \overline{B}(\nu\mathsf{x}\overline{B}(\mathsf{x}, \overline{D}) \vee \mathsf{x}, \mathsf{y}) \\
D^{y*} &= \nu\mathsf{y}B(\overline{C}, \mathsf{y}) \vee \nu\mathsf{y}\mu\mathsf{x}B(\mathsf{x}, \nu\mathsf{y}B(\overline{C}, \mathsf{y}) \vee \mathsf{y}) \\
D^{y'*} &= \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y}) \vee \nu\mathsf{y}\mu\mathsf{x}B(\mathsf{x}, \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y}) \vee \mathsf{y})
\end{aligned}$$

Because x, y and y' are the only annotations occurring in π_{clo}, the above cases provide the interpretation of all other formulæ: if F^c is any other annotated formula in the proof then F^{c*} is the result of replacing in F, C^{x*} for C, D^{y*} for D or $D^{y'*}$ for $D^{y'}$, depending whether c is x, y or y'.

Under the interpretation, discharged assumptions become derivable sequents and instances of ν-clo are replaced by the sequent of inferences in (8). For example, each of the three instances of ν-$\mathsf{clo}_{y'}$ becomes

$$\dfrac{\dfrac{\mu\mathsf{y}\overline{B}(C^{x*}, \mathsf{y}), \overbrace{\mu\mathsf{x}B(\mathsf{x}, \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y}) \vee \nu\mathsf{y}\mu\mathsf{x}B(\mathsf{x}, \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y}) \vee \mathsf{y}))}^{\mu\mathsf{x}B(\mathsf{x}, D)^{y'*}}}{\mu\mathsf{y}\overline{B}(C^{x*}, \mathsf{y}), \nu\mathsf{y}\mu\mathsf{x}B(\mathsf{x}, \nu\mathsf{y}B(\overline{C^{x*}}, \mathsf{y}) \vee \mathsf{y})}\,\nu}{\mu\mathsf{y}\overline{B}(C^{x*}, \mathsf{y}), \nu\mathsf{y}\mu\mathsf{x}B}\,\mathsf{ind}_s$$

The result is a $\mathsf{Koz}_\mathsf{s}^-$ proof of $\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B$.

Example 7.2. In Example 5.1 (Figure 8) a simplification of π_{circ} was described, obtained by applying a discharge rule closer to the root of the proof. Following this proof through the embeddings from Circ to Clo to $\mathsf{Koz}_\mathsf{s}^-$ yields another proof of the same sequent, namely

$$\dfrac{\overline{\overline{\mu\mathsf{x}B(\mathsf{x}, D_y)}, \mu\mathsf{x}B(\mathsf{x}, D_y)}}{C_{xx'}, \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\mathsf{weak}, \vee \qquad \dfrac{\dfrac{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), \nu\mathsf{y}B(\overline{C_x}, \mathsf{y})}{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), D_y}\,\mathsf{weak}, \vee}{\mu\mathsf{y}\overline{B}(C_{xx'}, \mathsf{y}), D_y}\,\mathsf{weak}, \vee_d$$

$$\vdots$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\overline{B}(C_{xx'}, \mu\mathsf{y}\overline{B}(C_{xx'}, \mathsf{y})), B(\mu\mathsf{x}B(\mathsf{x}, D_y), D_y)}{\mu\mathsf{y}\overline{B}(C_{xx'}, \mathsf{y}), \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\mu, \mu}{\nu\mathsf{x}\mu\mathsf{y}\overline{B}(\overline{\mu\mathsf{x}B(\mathsf{x}, D_y)} \vee \mathsf{x}, \mathsf{y}), \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\nu}{\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\mathsf{ind}_s}{C_x, \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\mathsf{weak}, \vee^* \qquad \dfrac{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), \nu\mathsf{y}B(\overline{C_x}, \mathsf{y})}{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), D_y}\,\mathsf{weak}, \vee$$

$$\vdots$$

$$\dfrac{\dfrac{\dfrac{\overline{B}(C, \mu\mathsf{y}\overline{B}(C_x, \mathsf{y})), B(\mu\mathsf{x}B(\mathsf{x}, D_y), D_y)}{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), \mu\mathsf{x}B(\mathsf{x}, D_y)}\,\mu, \mu}{\mu\mathsf{y}\overline{B}(C_x, \mathsf{y}), \nu\mathsf{y}\mu\mathsf{x}B}\,\nu, \mathsf{ind}_s}{\nu\mathsf{x}\mu\mathsf{y}\overline{B}, \nu\mathsf{y}\mu\mathsf{x}B}\,\nu, \mathsf{ind}_s$$

where

$$
\begin{aligned}
C_x &= \overline{\nu \mathsf{y} \mu \mathsf{x} B} \vee \nu \mathsf{x} \mu \mathsf{y}.\, \overline{B}(\overline{\nu \mathsf{y} \mu \mathsf{x} B} \vee \mathsf{x}, \mathsf{y}) \\
C_{xx'} &= \overline{\mu \mathsf{x} B(\mathsf{x}, D_y)} \vee \nu \mathsf{x} \mu \mathsf{y}.\, \overline{B}(\overline{\mu \mathsf{x} B(\mathsf{x}, D_y)} \vee \mathsf{x}, \mathsf{y}) \\
D_y &= \nu \mathsf{y} B(\overline{C_x}, \mathsf{y}) \vee \nu \mathsf{y} \mu \mathsf{x}.\, B(\mathsf{x}, \nu \mathsf{y} B(\overline{C_x}, \mathsf{y}) \vee \mathsf{y}).
\end{aligned}
$$

We complete the section with an embedding of $\mathsf{Koz}_\mathsf{s}^-$ into Koz. As the latter system is sound, this yields a soundness proof for $\mathsf{Koz}_\mathsf{s}^-$ and as such also for Clo and Circ.

Theorem 7.2. *If* $\mathsf{Koz}_\mathsf{s}^- \vdash \Gamma$ *then* $\mathsf{Koz} \vdash \Gamma$.

Proof. We first prove $\mathsf{Koz}^- \vdash \overline{\nu \mathsf{y} \nu \mathsf{x}\, B(\mathsf{y} \vee \mathsf{x})}, \nu \mathsf{x}\, B(\mathsf{x})$ for every formula $B(\mathsf{x})$ with at most x free. Fix a formula $B(\mathsf{x})$ and let $C = \nu \mathsf{y} \nu \mathsf{x}\, B(\mathsf{y} \vee \mathsf{x})$. The Koz^--proof

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\mathsf{Ax2}\colon \overline{C}, C \qquad \dfrac{\mathsf{Ax2}\colon \overline{\nu \mathsf{x} B(C \vee \mathsf{x})}, \nu \mathsf{x} B(C \vee \mathsf{x})}{\overline{\nu \mathsf{x} B(C \vee \mathsf{x})}, C}\,\nu}{\overline{C \vee \nu \mathsf{x} B(C \vee \mathsf{x})}, C}\,\wedge \\ \vdots\ \text{Lemma 3.2} \\ \overline{B(C \vee \nu \mathsf{x} B(C \vee \mathsf{x}))}, B(C)}{\overline{\nu \mathsf{x} B(C \vee \mathsf{x})}, B(C)}\,\mu}{\overline{\nu \mathsf{y} \nu \mathsf{x} B(\mathsf{y} \vee \mathsf{x})}, B(C)}\,\mu}{\overline{\nu \mathsf{y} \nu \mathsf{x} B(\mathsf{y} \vee \mathsf{x})}, \nu \mathsf{x} B}\,\mathsf{ind}}{}
$$

derives the desired sequent. The following derivation then establishes the admissibility of ind_s in Koz:

$$
\dfrac{\dfrac{\Gamma, \nu \mathsf{x} B(\overline{\Gamma} \vee \mathsf{x})}{\Gamma, \nu \mathsf{y} \nu \mathsf{x} B(\mathsf{y} \vee \mathsf{x})}\,\mathsf{ind} \qquad \overline{\nu \mathsf{y} \nu \mathsf{x} B(\mathsf{y} \vee \mathsf{x})}, \nu \mathsf{x} B(\mathsf{x})}{\Gamma, \nu \mathsf{x} B(\mathsf{x})}\,\mathsf{cut}
$$

□

Example 7.3. The previous two examples each describe a proof of $\mu \mathsf{x} \nu \mathsf{y} B \to \nu \mathsf{y} \mu \mathsf{x} B$ in $\mathsf{Koz}_\mathsf{s}^-$ which can be readily transformed into Koz-proofs by replacing each instance of ind_s by ind and cut. The resulting Koz-proofs have essentially the same structure as the $\mathsf{Koz}_\mathsf{s}^-$-proofs from which they are obtained.

8 Semi-formal systems

As discussed in Section 3.3, the system $K_\omega(\mu)$ of [9], which extends Fix by the infinitary ν-rule

$$
\dfrac{\Gamma, \nu^0 \mathsf{x} A \quad \Gamma, \nu^1 \mathsf{x} A \quad \cdots}{\Gamma, \nu^\omega \mathsf{x} A}\,\nu_\omega
$$

is sound and complete. In the present section we provide an alternative proof of completeness for $K_\omega(\mu)$ via an embedding of Clo into $K_\omega(\mu)$, a result that will establish the connection between $K_\omega(\mu)$ and the proof systems introduced in this paper, including Koz.

Recall the quantifiers ν^k, defined by $\nu^0 \mathsf{x} A = \top$ and $\nu^{k+1} \mathsf{x} A(\mathsf{x}) = A(\nu^k \mathsf{x} A)$; set $\nu^\omega \mathsf{x} A = \nu \mathsf{x} A$. Let $K'_\omega(\mu)$ be the extension of $K_\omega(\mu)$ by the rule

$$(n \leq m \leq \omega)\,\frac{\Gamma, A(\nu^m \mathsf{x} B)}{\Gamma, A(\nu^n \mathsf{x} B)}\,\mathsf{approx}$$

This rule is sound, and is admissible in $K_\omega(\mu)$ by a straightforward induction. Nevertheless, it is convenient to have the rule present during the embedding.

Definition 8.1. An *approximation* is a function $f\colon N^* \to \omega \cup \{\omega\}$ anti-monotone in the sub-word relation, i.e. for all $a, b \in N^*$, if $a \sqsubset b$ then $f(b) \leq f(a)$.

Approximations provide a natural interpretation of annotated formulæ as plain formulæ in which greatest fixed points are replaced by the approximant induced by the annotation. For each annotated formula A^a and approximation f, define $A^{f,a}$ recursively by

$$(A \wedge B)^{f,a} = A^{f,a} \wedge B^{f,a} \qquad ([\mathfrak{a}]A)^{f,a} = [\mathfrak{a}]A^{f,a} \qquad p^{f,a} = p$$
$$(A \vee B)^{f,a} = A^{f,a} \vee B^{f,a} \qquad (\langle\mathfrak{a}\rangle A)^{f,a} = \langle\mathfrak{a}\rangle A^{f,a} \qquad \mathsf{x}^{f,a} = \mathsf{x};$$

and for quantified formulæ, if $a = bcd$ where $b < \mathsf{x} < d$ and $c \in N^*_{\mathsf{x}}$, then

$$(\mu\mathsf{x}A)^{f,a} = \mu\mathsf{x}(A^{f,b}) \qquad (\nu\mathsf{x}A)^{f,a} = \nu^{f(bc)}\mathsf{x}(A^{f,b})$$

The restriction to anti-monotone functions ensures that for any approximation f and annotations $a \sqsubset b$ the formula $A^{f,a} \to A^{f,b}$ is valid.

Given a set Γ of annotated formulæ, define $\Gamma^f = \{A^{f,a} \mid A^a \in \Gamma\}$. For a set $\mathscr{A}$ of annotated sequents, define $\mathscr{A}^f = \{\Delta^f \mid \Delta \in \mathscr{A}\}$.

Lemma 8.2. *Suppose $\mathscr{A} \cup \{\Gamma\}$ is a set of annotated sequents with empty control such that for every $\Delta \in \mathscr{A}$ the annotations occurring in Δ are sub-words of annotations in Γ. If $\mathscr{A} \vdash_{\mathsf{Clo}} \Gamma$ then for every approximation f, $\mathscr{A}^f \vdash_{K'_\omega(\mu)} \Gamma^f$.*

Proof. Fix an approximation f. The proof proceeds by induction on the length of the (finite) derivation $\mathscr{A} \vdash_{\mathsf{Clo}} \Gamma$. In case Γ is an axiom or $\Gamma \in \mathscr{A}$, the result is trivial. The propositional and modal rules, and μ, translate to the corresponding plain rules. Moreover, instances of the ν rule become applications of approx and anti-monotonicity ensures that applications of exp become a series of applications of approx. The remaining rule to consider is ν-clo. Suppose the final rule applied is ν-clo of the form

$$\begin{array}{c} [\Gamma, \nu\mathsf{x}A^{ax}]^x \qquad \{\Delta \mid \Delta \in \mathscr{A}\} \\ \vdots \qquad\qquad\qquad \vdots \\ \hdashline \vdots \\ \dfrac{\Gamma, A(\nu\mathsf{x}A^{ax})}{\Gamma, \nu\mathsf{x}A^a}\,\nu\text{-}\mathsf{clo}_x \end{array}$$

where $x \in N_{\mathsf{x}}$. Given a word b, let b^- be the sub-word of b in which x has been removed if it occurs. For each $n \leq \omega$, define f_n to be the approximation

$$f_n(b) = \begin{cases} \min\{n, f(b^-)\}, & \text{if } x \text{ occurs in } b, \\ f(b), & \text{otherwise.} \end{cases}$$

Since x does not appear in Γ nor in any $\Delta \in \mathscr{A}$, it follows that f_n and f agree on all sub-words of annotations occurring in these sequents, i.e. $\Delta^{f_n} = \Delta^f$ for all $\Delta \in \mathscr{A} \cup \{\Gamma\}$.

Moreover, $f_{n+1}(b) \leq f_n(b) + 1$ for every b and n. Let $a = a_0 a_1$ where $a_0 < \mathsf{x} \leq a_1$. By the induction hypothesis, for every $n \leq \omega$,

$$\Gamma^f \cup \{\nu^{f_n(ax)}\mathsf{x}.\ A^{f,a_0}\}, \mathscr{A}^f \vdash_{K'_\omega(\mu)} \Gamma^f, \nu^{f_{n+1}(ax)}\mathsf{x}.\ A^{f,a_0}.$$

Inducting on $n < \omega$ we deduce

$$\mathscr{A}^f \vdash_{K'_\omega(\mu)} \Gamma^f, \nu^{f_n(ax)}\mathsf{x}.\ A^{f,a_0} \tag{9}$$

for every $n < \omega$. If $f(a) < \omega$ then pick $n = f(a)$ in (9). Otherwise, an application of ν_ω with sub-derivation (9) for each $n < \omega$ yields the desired derivation. ☐

Theorem 8.3. *For every plain sequent* Γ*,* $\mathsf{Clo} \vdash \Gamma$ *implies* $K_\omega(\mu) \vdash \Gamma$*.*

Proof. Suppose $\mathsf{Clo} \vdash \Gamma$ and let f be an approximation such that $f(\epsilon) = \omega$. Recall that Γ represents the annotated sequent $\{A^\epsilon \mid A \in \Gamma\}$. Then $\Gamma^f = \Gamma$ and Lemma 8.2 implies $K'_\omega(\mu) \vdash \Gamma$. As approx is admissible in $K_\omega(\mu)$ we obtain $K_\omega(\mu) \vdash \Gamma$. ☐

The converse is also possible. Studer, in [14], describes a transformation of $K_\omega(\mu)$-proofs into tableaux. Combining Studer's argument with the completeness proof for Stir yields an embedding of $K_\omega(\mu)$ in Stir.

Theorem 8.4. *For every plain sequent* Γ*, if* $K_\omega(\mu) \vdash \Gamma$ *then* $\mathsf{Stir} \vdash \Gamma$*.*

9 Conclusion

Combining the results of previous sections we obtain

Theorem 9.1. *Let* A *be a closed well-named formula. The following are equivalent.*

1. *A is valid.* *2.* $\mathsf{Stir} \vdash A$. *3.* $\mathsf{Circ} \vdash A$. *4.* $\mathsf{Clo} \vdash A$.

5. $\mathsf{Koz}^-_\mathsf{s} \vdash A$. *6.* $\mathsf{Koz} \vdash A$. *7.* $K_\omega(\mu) \vdash A$.

Moreover, there exists a constructive, primitive recursive transformation between the proof systems in the direction 2 to 6 and from Clo *to* $K_\omega(\mu)$*, and a recursive transformation of recursive* $K_\omega(\mu)$*-proofs to* Stir*.*

Proof. Theorems 5.4, 6.3, 7.1 and 7.2 provide the embeddings $2 \Rightarrow 3 \Rightarrow 4 \Rightarrow 5 \Rightarrow 6$ for arbitrary well-named A. For guarded formulæ, $1 \Rightarrow 2$ is proved in [13], and Theorem 3.9 generalises this to unguarded formulæ. The embedding $4 \Rightarrow 7$ is provided by Theorem 8.3, $7 \Rightarrow 2$ by Theorem 8.4, and $6 \Rightarrow 1$ by the soundness of Kozen's axiomatisation [10]. ☐

Theorem 9.1 settles a number of questions regarding proof systems for μ-calculus which we summarise below.

The system $\mathsf{Koz}^-_\mathsf{s}$ is a natural variant of Kozen's original axiomatisation wherein cut is dropped in favour of the strengthened induction rule

$$\frac{\Gamma, \nu\mathsf{x}A(\overline{\Gamma} \vee \mathsf{x})}{\Gamma, \nu\mathsf{x}A(\mathsf{x})}\,\mathsf{ind}_s$$

Specifically, the calculus derives only plain sequents and all inference rules have a fixed arity. As such, $\mathsf{Koz}_\mathsf{s}^-$ marks (as far as the authors are aware) the first finitary sound and complete cut-free proof system for the modal μ-calculus. Furthermore, the completeness of $\mathsf{Koz}_\mathsf{s}^-$ reduces the long standing open problem of whether Kozen's axiomatisation without cut is complete to whether the strengthened induction rule ind_s is admissible. As the proof of Theorem 7.2 demonstrates, this is in turn equivalent to the admissibility of the rule

$$\frac{\Gamma, \nu \mathsf{y} \nu \mathsf{x} B(\mathsf{y} \vee \mathsf{x})}{\Gamma, \nu \mathsf{x} B(\mathsf{x})}$$

which permits contracting like quantifiers in simple contexts.

The embeddings $2 \Rightarrow 3 \Rightarrow 4 \Rightarrow 5 \Rightarrow 6$ of Theorem 9.1 close the gap between Stirling tableau proofs and Kozen's axiomatisation, thereby answering the question posed by Stirling in [13] and providing a new proof of completeness for Kozen's original axiomatisation. The argument is purely proof-theoretic and constructive, and provides a computational approach to obtaining proofs in Koz. The approach involves a hyper-hyper-exponential blow-up in the size of the proof, which is due to the transformation of Stirling proofs to invariant monotone Stirling proofs (Theorem 4.9). Nevertheless, this is not always witnessed: in our running example of the valid formula $\mu \mathsf{x} \nu \mathsf{y} B \to \nu \mathsf{y} \mu \mathsf{x} B$ (Examples 3.2, 3.3, 5.1, 6.1, 7.1 and 7.3) the transformation between proof systems remains feasible and the resulting Koz-proof can be written down by hand.

The finitary sequent calculus Clo is also of independent interest as it is inter-translatable with the infinitary system $K_\omega(\mu)$ of [9]. In particular, via this finitary system, an embedding from $K_\omega(\mu)$ into Koz is obtained. It would be interesting to see whether Studer's interpretation of $K_\omega(\mu)$-proofs as tableaux can be adapted to a direct embedding of $K_\omega(\mu)$ in Clo. Such an embedding would serve to further clarify the relation of $K_\omega(\mu)$ to the systems $\mathsf{Koz}_\mathsf{s}^-$ and Koz.

The cut-free finitary proof systems introduced and the constructive completeness proof of Kozen's axiomatisation obtained in this article open the door to a number of avenues of research. One such application is to the problem of interpolation. It is known that uniform interpolation holds for the μ-calculus [5]. The proof of this result is indirect and uses μ-automata to show the μ-calculus interprets bisimulation quantifiers. Typically however, constructive proofs of interpolation can be obtained from cut-free calculi. Another application of interest is to the problem of cut elimination. While there are proof systems that are known to be cut-free complete, effective cut-elimination has only been established for small fragments of μ-calculus such as *propositional dynamic logic* [7], *common knowledge* [4] and the *one variable fragment* [11]. It may prove more viable to study effective cut elimination in the context of one of the annotated proofs systems, such as Clo, due to their more predicative nature. Finally, it would be interesting to see whether the techniques of this paper can be generalised to yield cut-free calculi for Venema's coalgebraic fixed point logic [15].

References

[1] B. Afshari and G. E. Leigh. 'Circular proofs for the modal mu-calculus'. In: *Proceedings of Applied Mathematics and Mechanics*. Vol. 16. 1. Wiley-VCH Verlag, 2016, pp. 893–894.

[2] A. Arnold and D. Niwinski. *Rudiments of μ-calculus*. Vol. 146. Studies in Logic. North Holland, 2001.

[3] J. Bradfield and C. Stirling. 'Modal mu-calculi'. In: *Handbook of Modal Logic*. Ed. by P. Blackburn, J. van Benthem and F. Wolter. Elsevier, 2007, pp. 721–756.

[4] K. Brünnler and T. Studer. 'Syntactic cut-elimination for common knowledge'. In: *Annals of Pure and Applied Logic* 160.1 (2009), pp. 82–95.

[5] G. D'Agostino and M. Hollenberg. 'Logical questions concerning the μ-calculus'. In: *The Journal of Symbolic Logic* 65.1 (2000), pp. 310–332.

[6] O. Friedmann and M. Lange. 'Deciding the unguarded modal μ-calculus'. In: *Journal of Applied Non-Classical Logics* 23.4 (2013), pp. 353–371.

[7] B. Hill and F. Poggiolesi. 'A contraction-free and cut-free sequent calculus for propositional dynamic logic.' In: *Studia Logica* 94.1 (2010), pp. 47–72.

[8] W. Hodges. *Model Theory*. Cambridge: Cambridge University Press, 1993.

[9] G. Jäger, M. Kretz and T. Studer. 'Canonical completeness of infinitary μ'. In: *Journal of Logic and Algebraic Programming* 76 (2008), pp. 270–292.

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

[11] G. Mints and T. Studer. 'Cut-elimination for the mu-calculus with one variable'. In: *Fixed Points in Computer Science 2012*. Vol. 77. EPTCS. Open Publishing Association, 2012, pp. 47–54.

[12] D. Niwinski and I. Walukiewicz. 'Games for the mu-calculus'. In: *Theoretical Computer Science* 163.1&2 (1996), pp. 99–116.

[13] C. Stirling. 'A tableau proof system with names for modal mu-calculus'. In: *HOWARD-60*. Ed. by Voronkov and Korovina. 2014, pp. 306–318.

[14] T. Studer. 'On the proof theory of the modal mu-calculus'. In: *Studia Logica* 89 (2008), pp. 343–363.

[15] Y. Venema. 'Automata and fixed point logic: A coalgebraic perspective'. In: *Information and Computation* 204.4 (2006), pp. 637–678.

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