On the Proof Theory of the Constructive μ-Calculus

Gianluca Curzi[1], Graham E. Leigh[1], Luca Sauter[1]

[1]*Department of Philosophy, Linguistics, Theory of Science; University of Gothenburg*

In this work, we contribute to a general procedure for cut elimination in the context of ill-founded derivations. To achieve this, we create a proof system for constructive modal logic extended with fixed points and adapt a recent procedure by Curzi and Leigh [5] for this sequent calculus.

Modal logic presents an interesting case in the landscape of logical theory, partially because while there are interesting consequences resulting from enforcing an intuitionistic version of modal logic, there is no singular approach to achieve this. Depending on the nature of the underlying motivation, the constructed system satisfies different axioms and properties [6]. Instead of considering axioms that should be fulfilled by the intuitionistic modal logic [3], a different, proof-theoretically driven approach is to place the regular intuitionistic restriction of only one formula on the right hand side on the sequent calculus for classical modal logic [14]. This results in what is commonly called constructive modal logic, also referred to as CK.

The study of logics with least and greatest fixed points represents a significant trend in recent research. Many of these logics lie in the modal realm, with Kozen's modal μ-calculi [9] being a central part of this development. Notably, an underlying intuitionistic motivated logic has several implications for the added fixed points, for example see [4, 7].

When constructing a proof system for logics including fixed points, ill-founded proof systems have emerged as well suited to the requirements [11]. In such systems, proof trees with infinite height are allowed, but placed under an external constriction. This greatly simplifies the rules for the fixed points operators, which can now simply consist of the unfolding of said fixed point. To ensure soundness, appealing to local soundness is no longer sufficient. Instead, the external constriction most often takes the form of the global trace condition, also referred to as the progressivity condition. Hereby, the condition imposes that some feature occurs infinitely often on all infinite branches in a proof tree.

In light of this, it seems natural to examine the logic constructed by adding least and greatest fixed points to an intuitionistic modal logic [1]. The close alignment of CK with its proof theory and the prior, albeit limited, usage of its extension with fixed points lead us to the choice of CK as the basis. We will refer to this constructive μ-calculus as $\mu\mathsf{CK}$.

Little is known about $\mu\mathsf{CK}$. Prior to this work, this logic only finds mention in two papers. In a 2024 paper, Pacheco introduces game semantics for $\mu\mathsf{CK}$ and relates them to their Kripke-style semantic counterpart [12]. As far as we know, this is the only work in which this logic is the main focus. A (weak) fragment of $\mu\mathsf{CK}$ has a prominent role in the proof of completeness for the (classical) modal μ-calculus by Afshari et al [2] where the constraints on formal proofs implicit to intuitionistic logic bootstrap completeness for the classical logic.

Another motivation comes from developments towards a general cut elimination procedure for illfounded proof system in a recent paper by Curzi and Leigh [5]. The method presented there is developed and demonstrated on $\mu\mathsf{MALL}$, the multiplicative and additive fragment of linear logic extended with fixed points. The argument heavily uses the dualities of the connectives in $\mu\mathsf{MALL}$ and is based on an one-sided sequent calculus. To adapt this approach for other logics, many elements of the proof need to be adjusted.

Two major differences between $\mu\mathsf{MALL}$ and $\mu\mathsf{CK}$ are the inclusion of modalities and the missing interdefinability between $\Box$ and $\Diamond$, and μ and ν, which results in a two-sided sequent calculus. Therefore, $\mu\mathsf{CK}$ serves as a good case study for extending the strategy in [5] for different logics. Establishing the result of cut elimination for $\mu\mathsf{CK}$ would provide an example of how to achieve this for other intuitionistic logics, as well as for logics with modalities.

All together, our work defines $\mu\mathsf{CK}$, the underlying constructive modal logic and the fixed points, with its Kripke-style and game semantics. After providing an ill-founded proof system for it, we adapt the argument in [5] for this new setting to show general cut elimination for $\mu\mathsf{CK}$.

The syntax of μCK adheres to the expected structure. We follow the convention of using a capital letter X for variables and a lower case letter p for propositional symbols.

Definition. A *formula* is generated by the following grammar:

$$\varphi \coloneqq p \mid \bot \mid \top \mid X \mid (\varphi \vee \varphi) \mid (\varphi \wedge \varphi) \mid (\varphi \to \varphi) \mid \Box\varphi \mid \Diamond\varphi \mid \mu X.\varphi \mid \nu X.\varphi$$

Notable, $\Box, \Diamond, \mu$ and ν are all part of the syntax, since they cannot be defined in terms of each other. To ensure a sound definition for the fixed points, we restrict the fixed point variable to only occur positively in its associated fixed point. Since our syntax does not include negation, positivity refers to the formula occurring in the antecedent of an implication an even number of times.

The Kripke-style semantics for the CK-fragment follow the historic precedent, set in the introductory work for constructive modal logic by Wijesekera [14], with the slight adjustment of allowing for fallible worlds that satisfy $\bot$ [10]. The semantics for the fixed points are added in the usual way, as fixed points of monotone functions, which are well-defined by the Knaster-Tarski theorem. The latter, and both the introduction and the proof of equivalence with game semantics can be found in [12].

To form our ill-founded system for μCK, we adjust the finitary system from [14] by only allowing exactly one formula on the right hand side[1] and add the standard rules for fixed point unfoldings. Since they are symmetric, we present the general fixed point rules for $\sigma \in \{\mu, \nu\}$. Figure 1 shows the rules of our proof system.

$$\frac{\Gamma, \varphi, \psi, \Gamma' \Rightarrow \delta}{\Gamma, \psi, \varphi, \Gamma' \Rightarrow \delta}\ \mathsf{ex} \qquad \frac{\Gamma, \varphi, \varphi \Rightarrow \delta}{\Gamma, \varphi \Rightarrow \delta}\ \mathsf{co} \qquad \frac{\Gamma \Rightarrow \delta}{\Gamma, \varphi \Rightarrow \delta}\ \mathsf{wk} \qquad \frac{\Gamma \Rightarrow \varphi \quad \Lambda, \varphi \Rightarrow \delta}{\Gamma, \Lambda \Rightarrow \delta}\ \mathsf{cut}$$

$$\frac{}{\bot \Rightarrow \delta}\ \bot \qquad \frac{}{\Rightarrow \top}\ \top \qquad \frac{}{p \Rightarrow p}\ \mathsf{id} \qquad \frac{\Gamma \Rightarrow \varphi}{\Box\Gamma \Rightarrow \Box\varphi}\ \Box \qquad \frac{\Gamma, \varphi \Rightarrow \delta}{\Box\Gamma, \Diamond\varphi \Rightarrow \Diamond\delta}\ \Diamond$$

$$\frac{\Gamma, \varphi \Rightarrow \delta \quad \Gamma, \psi \Rightarrow \delta}{\Gamma, \varphi \vee \psi \Rightarrow \delta}\ \vee_{\mathsf{L}} \qquad \frac{\Gamma \Rightarrow \varphi_i}{\Gamma \Rightarrow \varphi_0 \vee \varphi_1}\ \vee_{\mathsf{R}} \qquad \frac{\Gamma, \varphi_i \Rightarrow \delta}{\Gamma, \varphi_0 \wedge \varphi_1 \Rightarrow \delta}\ \wedge_{\mathsf{L}}$$

$$\frac{\Gamma \Rightarrow \varphi \quad \Gamma \Rightarrow \psi}{\Gamma \Rightarrow \varphi \wedge \psi}\ \wedge_{\mathsf{R}} \qquad \frac{\Gamma \Rightarrow \varphi \quad \Gamma, \psi \Rightarrow \delta}{\Gamma, \varphi \to \psi \Rightarrow \delta}\ \to_{\mathsf{L}} \qquad \frac{\Gamma, \varphi \Rightarrow \psi}{\Gamma \Rightarrow \varphi \to \psi}\ \to_{\mathsf{R}}$$

$$\frac{\Gamma, \varphi(\sigma X.\varphi(X)) \Rightarrow \delta}{\Gamma, \sigma X.\varphi(X) \Rightarrow \delta}\ \sigma_{\mathsf{L}} \qquad \frac{\Gamma \Rightarrow \varphi(\sigma X.\varphi(X))}{\Gamma \Rightarrow \sigma X.\varphi(X)}\ \sigma_{\mathsf{R}}$$

Figure 1: The sequent calculus of μCK

To ensure soundness, we add the global trace condition. A trace is a sequence of formulas, starting at some formula in a branch and following the ancestors of this formula upwards through the branch. A progressing derivation or proof is a derivation where for every infinite branch, there exists a trace through that branch such that along this trace, a least fixed point unfolds infinitely often on the left hand side of the sequent, or a greatest fixed points unfolds infinitely often on the right hand side.

Towards cut elimination, specifying an infinitary cut reduction procedure for ill-founded proof systems, even along the infinite branches, is straightforward. Cut elimination is then obtained as a limit of such an infinitary procedure. The challenge lies in showing that such a limit cut-free derivation is both well defined and also still progressing. To illustrate this issue, consider a proof where the trace showing progressivity along some infinite branch begins in a cut formula. After applying the cut reductions, this cut formula and its ancestors will not be included in the cut-free derivation. Since the originial trace does no longer exist, what is now witnessing the progressivity of the derivation along this infinite branch?

[1] This adjusted sequent calculus can also be found in [6], where it is referred to as LCK

The cut elimination argument from [5] and adapted here is based on reducibility candidates, a concept introduced by Tait [13] and Girard [8] in the context of lambda calculus. Simply put, reducibility candidates model sets of objects that normalize in some way. In our case, they serve as a proof-theoretic interpretation of μCK formulas, consisting of certain derivations of the given formula that reduce under this cut reduction procedure to a sound and cut-free derivation. We write $[\![\varphi]\!]$ as the reducibility candidate interpretation of the formula φ. This interpretation is defined by an induction on the structure of φ and is of a syntactic nature. As an example, we give the definition of the interpretation of $\Diamond\varphi$:

$$[\![\Diamond\varphi]\!] := \left\{ \begin{array}{c} d \\ \Gamma, \psi \Rightarrow \varphi \\ \hline \Box\Gamma, \Diamond\psi \Rightarrow \Diamond\varphi \end{array} \Diamond \;\middle|\; d \in [\![\varphi]\!] \right\}^*$$

The interpretation of $\Diamond\varphi$ is constructed by taking derivation in the interpretation of φ and applying the $\Diamond$-rule to them. The $*$-operator closes the definition under cut reduction, meaning that every derivation that reduces to a derivation belonging in this set, also belong to the set. We can also establish the following lemma from the definition of this interpretation:

Lemma (Normalization)**.** Every derivation d in a reducibility interpretation $[\![\varphi]\!]$ can be rewritten by a cut reduction procedure into a progressing and cut-free proof.

We aim to show that every derivation is contained in some reducibility candidate. First, we establish the correct behaviour of our proof system in regards to the interpretation with the following lemma:

Lemma (Soundness)**.** The rules of the sequent calculus of μCK are sound with respect to the interpretation based on reducibility candidates.

We can prove this by a straightforward case distinction on the last rule of a derivation. Next, we want to show our main theorem, whose immediate corollary, together with the normalization lemma, is that every derivation reduces to one without cuts:

Theorem. Every proof d with conclusion $\Rightarrow \varphi$ is in the reducibility interpretation of φ, $[\![\varphi]\!]$.

As is common in arguments involving fixed points, the proof strategy of this theorem will hinge on ordinal approximates of fixed points. Additionally, we argue by contraposition. Assuming a derivation d does not belong to a suitable reducibility candidate and repeatedly applying the soundness lemma yields an infinite branch in d. Since d is a proof and therefore progressing, there is a trace along this branch with the correct fixed point unfoldings. From these, using the ordinal approximations, we can construct an infinitely descending chain of ordinals, which contracts the well-foundedness of ordinals.

This proof method is non-constructive. Our cut elimination argument only states the existence of a cut reduction procedure with a cut-free limit derivation without stating precisely which. This also avoids the problem of showing the progressivity of the limit derivation.

A natural continuation of this work would be the adaptation of the second cut elimination argument in [5]. With a stricter variation of progressivity placed on derivations and some additional results, the authors extend the ideas behind the approach presented here. Essentially, this variation encodes well-behaved cut reduction in the infinitary case into the progressivity condition. Showing that both this and the global trace condition agree on every derivation results in an explicit procedure for cut elimination.

Furthermore, in [2] several results are established for a fragment of μCK essentially corresponding to the implication free fragment. Most notably, cut elimination, a normal form theorem and completeness are shown for this fragment. While in our work we lift cut elimination to the entirety of μCK, a possible future endeavour consists of finding a precise formulation of the normal form theorem which can be applied to the entirety of μCK.

References

[1] Bahareh Afshari and Lide Grotenhuis. "Intuitionistic μ-Calculus with the Lewis Arrow". In: *Automated Reasoning with Analytic Tableaux and Related Methods*. Springer Nature Switzerland, 2026, pages 374–392. https://doi.org/10.1007/978-3-032-06085-3_20.

[2] Bahareh Afshari, Graham E. Leigh, and Guillermo Menèndez Turata. "Demystifying μ". In: *Fundamenta Informaticae* Volume 194, Issue 2: Fixed Points in Computer Science 2023 (2025).

[3] Milan Božić and Kosta Došen. "Models for Normal Intuitionistic Modal Logics". In: *Studia Logica: An International Journal for Symbolic Logic* 43.3 (1984), pages 217–245.

[4] Pierre Clairambault. "Least and Greatest Fixpoints in Game Semantics". In: *Foundations of Software Science and Computational Structures*. Edited by Luca de Alfaro. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pages 16–31.

[5] Gianluca Curzi and Graham E. Leigh. *Making progress: Reducibility Candidates and Cut Elimination in the Ill-founded Realm*. 2026. https://arxiv.org/abs/2602.01299.

[6] Anupam Das and Sonia Marin. "On Intuitionistic Diamonds (and Lack Thereof)". In: *Automated Reasoning with Analytic Tableaux and Related Methods*. Springer Nature Switzerland, 2023, pages 283–301.

[7] Sebastian Enqvist. *Computation by infinite descent made explicit*. 2025. https://arxiv.org/abs/2506.22206.

[8] J.Y. Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Éditeur inconnu, 1972.

[9] Dexter Kozen. "Results on the propositional μ-calculus". In: *Theoretical Computer Science* 27.3 (1983). Special Issue Ninth International Colloquium on Automata, Languages and Programming (ICALP) Aarhus, Summer 1982, pages 333–354.

[10] Michael Mendler and Valeria de Paiva. "Constructive CK for Contexts". In: *Context Representation and Reasoning (CRR-2005)*. 2005.

[11] Damian Niwiński and Igor Walukiewicz. "Games for the μ-calculus". In: *Theoretical Computer Science* 163.1 (1996), pages 99–116.

[12] Leonardo Pacheco. *Game semantics for the constructive μ-calculus*. 2024. https://arxiv.org/abs/2308.16697.

[13] W. W. Tait. "Intensional Interpretations of Functionals of Finite Type I". In: *The Journal of Symbolic Logic* 32.2 (1967), pages 198–212.

[14] Duminda Wijesekera. "Constructive modal logics I". In: *Annals of Pure and Applied Logic* 50.3 (1990), pages 271–301.