Cut elimination for Cyclic Proofs: A case study in temporal logic*

Bahareh Afshari
University of Gothenburg
Gothenburg, Sweden
bahareh.afshari@gu.se

Johannes Kloibhofer
University of Amsterdam
Amsterdam, The Netherlands
j.kloibhofer@uva.nl

We consider modal logic extended with the well-known temporal operator ‘eventually’ and provide a cut-elimination procedure for a cyclic sequent calculus that captures this fragment. The work showcases an adaptation of the reductive cut-elimination method to cyclic calculi. Notably, the proposed algorithm applies to a cyclic proof and directly outputs a cyclic cut-free proof without appealing to intermediate machinery for regularising the end proof.

1 Introduction

Non-wellfounded derivation systems underwent their first rigorous study in [9, 11, 2, 14]. In contrast to the traditional finite derivation trees, one allows for infinite branches that mimic the induction axioms required for capturing semantics of temporal operators (more generally, implicit/explicit fixpoints). To exclude vicious reasoning and capture the true semantics of temporal constructs, a so called *global validity condition* is placed on infinite branches. This condition necessitates a re-investigation of reductive cut elimination in the non-wellfounded setting. *Cyclic proofs* correspond to the regular fragment of non-wellfounded proofs. They are, in some respects, more amenable to this treatment as the global validity condition for a cyclic proof can be reformulated in terms of a condition posed on simple cycles in the graph representation of a cyclic proof.

With the advance of non-wellfounded and cyclic proof theory, cut elimination has also received considerable attention (see e.g. [5, 3, 13, 12]). The common approach proceeds in two steps: first cuts are pushed ‘up’ away from the root by the usual cut reductions to obtain, in the limit, a non-wellfounded derivation; in the second step the acquired limit structure is shown to satisfy the global validity condition. When it comes to cyclic systems, cut-elimination procedures that can *directly* produce cut-free cyclic proofs are rare. Although the aforementioned two-step approach can be applied, the resulting structure may not necessarily be a regular tree. For those calculi whose cyclic fragment does exhaust all validities one may invoke other machinery, such as automata, to find a cut-free cyclic proof.

We consider a simple temporal logic MLe with a complete cyclic sequent calculus GKe and describe a syntactic cut-elimination procedure for GKe that works directly on cyclic proofs. The logic MLe extends modal logic with the so called ‘eventually’ operator, the transitive closure of the diamond modality. Although MLe is expressively weak, this work can be seen as a starting point to apply cut elimination to cyclic proof systems for richer modal fixpoint logics.

Cuts occurring in a cyclic proof can be split into two categories: cuts that reside in a cycle and those that do not. We call them, respectively, *unimportant* and *important* cuts. We show that unimportant cuts are for most part harmless, in the sense that the cut formula does not interfere with the global validity

*This work was supported by the Dutch Research Council [OCENW.M20.048, 617.001.857] and Knut and Alice Wallenberg Foundation [2020.0199].

Submitted to:
FICS 2024

condition and as such, these cuts can be eliminated by the standard way of pushing them away from the root of the proof-tree. The main challenge posed is to eliminate important cuts, as in this case a 'successful' trace fulfilling the global validity condition can follow a cut formula.

The global validity condition can take different forms. A *trace-based condition* is formulated in terms of traces, i.e. paths of ancestry of formulas, on a branch. Such conditions are relatively easy to formulate but harder to work with. By annotating sequents with additional information one may succeed in formulating a *branch-based condition* that is not directly concerned with traces. Such a condition is easier to work with when performing operations such as projecting a proof or zipping two proofs that are critical in reductive cut elimination. Annotated proof systems were first introduced by Jungteerapanich [6] and Stirling [16] building on work of Walukiewicz [17]. There are general ways of obtaining annotated proof systems from cyclic proof systems [4, 7]. It is with this rich framework in mind that the presented case analysis has been carried out.

2 Syntax and Semantics of MLe

Let Prop be an infinite set of propositions. *Formulas* of MLe are defined inductively by

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \mathsf{F}\varphi$$

where $p \in \mathsf{Prop}$. The connectives $\vee, \rightarrow$ and $\Box$ are taken as abbreviations and can be defined as usual. To see MLe as a fragment of the modal μ-calculus [1] define $\mathsf{F}\varphi \equiv \mu x.(\varphi \vee \Diamond x)$.

Formulas are evaluated on Kripke models in the standard way:

Definition 1. A *Kripke model* is a tuple $\mathbb{S} = (S, R, V)$, where S is a non-empty set, R is a binary relation on S and V is a function $S \to \mathscr{P}(\mathsf{Prop})$.

Let $\mathbb{S} = (S, R, V)$ be a Kripke model, $s \in S$ and let R^* be the transitive closure of R. The relation $\Vdash$ is defined inductively by:

$\mathbb{S}, s \Vdash p$	$\Leftrightarrow$	$p \in V(s)$
$\mathbb{S}, s \Vdash \neg\varphi$	$\Leftrightarrow$	$\mathbb{S}, s \nVdash \varphi$
$\mathbb{S}, s \Vdash \varphi \wedge \psi$	$\Leftrightarrow$	$\mathbb{S}, s \Vdash \varphi$ and $\mathbb{S}, s \Vdash \psi$
$\mathbb{S}, s \Vdash \Diamond\varphi$	$\Leftrightarrow$	there exists $t \in S$ with sRt and $\mathbb{S}, t \Vdash \varphi$
$\mathbb{S}, s \Vdash \mathsf{F}\varphi$	$\Leftrightarrow$	there exists $t \in S$ with sR^*t and $\mathbb{S}, t \Vdash \varphi$

A formula φ is called *valid* if $\mathbb{S}, s \Vdash \varphi$ for every Kripke model $\mathbb{S} = (S, R, V)$ and every $s \in S$.

3 Proof System GKe

We present the cyclic annotated proof system GKe, which is sound and complete with respect to the valid formulas of MLe. The annotations allow us to formulate the global validity condition as a simple condition on paths. As MLe is a fragment of the alternation-free mu-calculus one bit of information per formula suffices [8]. Moreover, as in MLe traces are very well-behaved, in particular cyclic traces do not pass through disjunctions, at most one formula in focus in each sequent is sufficient [10].

An *annotated formula* is a pair (φ, a), usually denoted as φ^a, where $a \in \{f, u\}$. We call annotated formulas of the form φ^f *in focus* and φ^u *out of focus*. Let Γ and Δ be *sets* of annotated formulas, where every formula in Δ is out of focus and at most one formula in Γ of the form $\mathsf{F}\varphi$ or $\Diamond\mathsf{F}\varphi$ is in focus. Then $\Gamma \Rightarrow \Delta$ is called an *annotated sequent*. As every formula in Δ is out of focus, we omit the annotations in the right side of the annotated sequent. For a set of annotated formulas Γ we define $\Gamma^u = \{\varphi^u \mid \varphi^a \in \Gamma\}$. If it is clear from the context we will be sloppy and call annotated formulas just formulas and annotated sequents just sequents. We define $\Diamond\Delta = \{\Diamond\varphi \mid \varphi \in \Delta\}$.

Definition 2. Let φ and ψ be formulas. We write $\varphi \to_C \psi$ if ψ is either a direct subformula of φ or $\varphi \equiv \mathsf{F}\chi$ and $\psi \equiv \Diamond\mathsf{F}\chi$. The *closure* $\mathsf{Clos}(\Gamma \Rightarrow \Delta)$ of a sequent $\Gamma \Rightarrow \Delta$ is the least superset of $\Gamma \cup \Delta$ that is closed under $\to_C$.

The rules of GKe for the propositional connectives and the modal operator are standard. The rules F_L and F_R for the eventually operator F stem from the identity $\mathsf{F}\varphi \equiv \varphi \vee \Diamond\mathsf{F}\varphi$. In F_L the formula $\mathsf{F}\varphi^a$ is called its *principal formula*. The rule F_L and the focus rules u and f are the only rules that change the annotations of formulas. The *discharge rule* D allows us to discharge leaves, that are labelled by sequents that already occur at one of its ancestors. To qualify as a GKe proof, this is only allowed if a certain success-condition is met. Each instance of D^{x} is labelled by a unique *discharge token* x taken from a fixed infinite set $\mathscr{D} = \{\mathsf{x}, \mathsf{y}, \mathsf{z}, \ldots\}$.

$$\mathsf{Ax}: \frac{}{p^u \Rightarrow p} \qquad \Diamond: \frac{\varphi^a \Rightarrow \Delta}{\Diamond\varphi^a \Rightarrow \Diamond\Delta} \qquad \wedge_L: \frac{\varphi^u, \psi^u, \Gamma \Rightarrow \Delta}{\varphi \wedge \psi^u, \Gamma \Rightarrow \Delta} \qquad \wedge_R: \frac{\Gamma \Rightarrow \Delta, \varphi \qquad \Gamma \Rightarrow \Delta, \psi}{\Gamma \Rightarrow \Delta, \varphi \wedge \psi}$$

$$\mathsf{F}_L: \frac{\Diamond\mathsf{F}\varphi^a, \Gamma \Rightarrow \Delta \qquad \varphi^u, \Gamma \Rightarrow \Delta}{\mathsf{F}\varphi^a, \Gamma \Rightarrow \Delta} \qquad \mathsf{F}_R: \frac{\Gamma \Rightarrow \Delta, \varphi, \Diamond\mathsf{F}\varphi}{\Gamma \Rightarrow \Delta, \mathsf{F}\varphi} \qquad \neg_L: \frac{\Gamma \Rightarrow \Delta, \varphi}{\neg\varphi^u, \Gamma \Rightarrow \Delta} \qquad \neg_R: \frac{\varphi^u, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg\varphi}$$

$$\mathsf{cut}: \frac{\Gamma_l \Rightarrow \Delta_l, \varphi \qquad \varphi^u, \Gamma_r \Rightarrow \Delta_r}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r} \qquad \mathsf{w}_L: \frac{\Gamma \Rightarrow \Delta}{\varphi^u, \Gamma \Rightarrow \Delta} \qquad \mathsf{w}_R: \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \varphi}$$

$$\begin{array}{c} [\Gamma \Rightarrow \Delta]^{\mathsf{x}} \\ \vdots \\ \mathsf{D}^{\mathsf{x}}: \dfrac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta} \end{array} \qquad \mathsf{u}: \frac{\varphi^f, \Gamma \Rightarrow \Delta}{\varphi^u, \Gamma \Rightarrow \Delta} \qquad \mathsf{f}: \frac{\varphi^u, \Gamma \Rightarrow \Delta}{\varphi^f, \Gamma \Rightarrow \Delta}$$

Figure 1: Rules of GKe

Definition 3. A GKe *derivation* $\pi = (T, P, \mathsf{S}, \mathsf{R})$ is a quadruple such that (T, P) is a, possibly infinite, tree with nodes T and binary relation P; S is a function that maps every node $u \in T$ to an annotated sequent; R is a function that maps every node $u \in T$ to either (i) the name of a rule in Figure 1 or (ii) a discharge token, such that (i) the specifications of the rules in Figure 1 are satisfied and (ii) every node labeled with a discharge token is a leaf.

For every leaf l that is labeled with a discharge token $\mathsf{x} \in \mathscr{D}$ there has to be a proper ancestor $c(l)$ of l that is labeled with D^{x} and such that l and $c(l)$ are labelled with the same sequent. In this situation we call l a *repeat leaf*, and $c(l)$ its *companion*.

A GKe derivation of a sequent $\Gamma \Rightarrow \Delta$ is a GKe derivation, where the root is labelled by $\Gamma \Rightarrow \Delta$.

Definition 4. Let $\pi = (T, P, \mathsf{S}, \mathsf{R})$ be a derivation. We define two trees we are interested in: (i) The usual proof tree $\mathscr{T}_\pi = (T, P)$ and (ii) the *proof tree with back edges* $\mathscr{T}_\pi^C = (T, P^C)$, where $P^C = P \cup \{(l, c(l)) \mid l \text{ is a repeat leaf}\}$ is the parent relation plus back-edges for each repeat leaf.

Definition 5. A *path* τ in a GKe derivation $\pi = (T, P, \mathsf{S}, \mathsf{R})$ is a path in $\mathscr{T}_\pi^C$. A path τ is called *successful* if it holds that

1. Every sequent on τ has a formula in focus and
2. The path τ passes through an application of F_L, where the principal formula is in focus.

We call a repeat leaf v in a GKe derivation π a *discharged leaf* if the path $\tau(v)$ in (T, P) from $c(v)$ to v is successful. A leaf is called *closed* if it is either a discharged leaf or labelled by Ax and *open* otherwise.

Definition 6. A GKe *proof* is a finite GKe derivation, where every leaf is closed.

Next lemma is usually a consequence of guardedness. For the GKe proof system this is immediate.

Lemma 7. *Let v be a discharged leaf in a* GKe *proof π. On $\tau(v)$ there is a node labelled by $\Diamond$.*

The soundness and completeness of GKe follows from [10] wherein a cyclic hypersequent calculi is given for a class of modal logics with the master modality characterised by frame conditions. In the case of no frame conditions, the calculus coincides with GKe, with the difference that the modalities $\Box$ and $\boxast$ are taken as principal instead of $\Diamond$ and F. Yet, as $\Box\varphi \equiv \neg\Diamond\neg\varphi$ and $\boxast\varphi \equiv \neg\mathsf{F}\neg\varphi$, the systems can be easily translated into each other.

Theorem 8 (Soundness and Completeness)**.** *There is a* GKe *proof of a sequent $\Gamma \Rightarrow \Delta$ iff $\bigwedge\Gamma \to \bigvee\Delta$ is valid in every Kripke model.*

4 Cut Elimination

We introduce a cut-elimination method based on reductive cut elimination tailored for cyclic proofs. The cut reductions we employ are the expected ones for the calculi; see Appendix A for full list.

The main difficulty in this approach are cut reductions, that alter successful paths. As an example consider the following principal cut reduction:

$$\dfrac{\dfrac{\Gamma \Rightarrow \Delta, \varphi, \Diamond \mathsf{F}\varphi}{\Gamma \Rightarrow \Delta, \mathsf{F}\varphi}\,\mathsf{F}_R \qquad \dfrac{\dfrac{\Diamond \mathsf{F}\varphi^f,\Gamma\Rightarrow\Delta \qquad \varphi^u,\Gamma\Rightarrow\Delta}{\mathsf{F}\varphi^f,\Gamma\Rightarrow\Delta}\,\mathsf{F}_L}{\mathsf{F}\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{u}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\longrightarrow \quad \dfrac{\dfrac{\Gamma\Rightarrow\Delta,\varphi,\Diamond\mathsf{F}\varphi \qquad \dfrac{\Diamond\mathsf{F}\varphi^f,\Gamma\Rightarrow\Delta}{\Diamond\mathsf{F}\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{u}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{cut} \qquad \varphi^u,\Gamma\Rightarrow\Delta}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

Note that in this reduction an F_L rule, where the principal formula is in focus, gets removed. Hence, paths that were successful before the reduction, could be unsuccessful afterwards. We define *unimportant cuts* in such a way, that in all cut reductions successful paths remain successful, i.e. the behaviour from above may not occur. Thus, for unimportant cuts it suffices to push all cuts upwards until a repeat is reached.

The treatment of *important cuts* is more complicated, as descendants of the cut-formulas could be in focus, and therefore being crucial in the discharge condition of repeats. In order to make these observation formal we first need some technical definitions.

4.1 Preliminary definitions

Let (G,E) be a graph. A *strongly connected subgraph* of G is a set of nodes $A \subseteq G$, such that from every node of A there is a path to every other node in A. A *cluster* of G is a maximally strongly connected subgraph of G. A cluster is called *trivial* if it consists of only one node.

Let π be a GKe proof. A *cluster* of π is a cluster of $\mathcal{T}_\pi^C$. Let $\mathcal{S}_\pi$ be the set of non-trivial clusters of π. We define a relation $\twoheadrightarrow$ on $\mathcal{S}_\pi$ as follows: $S_1 \twoheadrightarrow S_2$ if $S_1 \neq S_2$ and there are nodes $v_1 \in S_1, v_2 \in S_2$ such that there is a path from v_1 to v_2 in $\mathcal{T}_\pi^C$. The relation $\twoheadrightarrow$ is irreflexive, transitive and antisymmetric. We write $\mathsf{depth}(S)$ for the length of the longest path in $(\mathcal{S}_\pi, \twoheadrightarrow)$ starting from S.

Let v be a node in a proof π. We define the *depth* of v, written $\mathsf{depth}(v)$, to be $\max\{\mathsf{depth}(S) \mid S \in \mathcal{S}_\pi$ and there is a path from v to some $u \in S\}$ and 0 if this set is empty. The *depth* of a proof π, written $\mathsf{depth}(\pi)$, is the depth of its root.

The *component* of v, written $\mathsf{comp}(v)$ is the set of nodes $u \in \pi$, that are reachable from v in $\mathcal{T}_\pi^C$ with $\mathsf{depth}(u) = \mathsf{depth}(v)$. Note that $\mathsf{comp}(v)$ does not have to coincide with a cluster in π, but may contain multiple clusters. The component of the root is called the *root-component* and the cluster of the root is called the *root-cluster*. A descendant ψ at a node u of a formula φ at a node v is called a *component descendant* of φ if $u \in \mathsf{comp}(v)$.

Definition 9. Let C be an occurrence of a cut rule with cut-formula φ in a GKe proof π. We call C *important* if all formulas in the conclusion of C are out of focus and *unimportant* otherwise.

The *rank* $\mathsf{rank}(\varphi)$ of a formula φ is the maximal nesting depth of F's in φ. The *rank* of a cut is the rank of the cut formula. The *depth* of a cut is the depth of its conclusion. The *sum-depth* of a cut is the sum of the depths of both premises. The *cut-rank*, *cut-depth* and *cut-sum-depth* of a proof π are the maximal rank, depth and sum-depth of a cut in π, respectively.

The next lemma justifies the definition of unimportant cuts. It implies, that pushing unimportant cuts upwards does not alter successful paths.

Lemma 10. *Let C be an unimportant cut in a GKe proof π. Every component descendant of the cut-formula of C is out of focus.*

Any node v in a non-trivial cluster of a GKe proof π has a formula in focus, as it is on the path $\tau(l)$ of a discharged leaf l. For nodes in a trivial cluster this is not necessarily the case. We can apply f and u rules in a certain way to minimize nodes with a formula in focus. By doing so, nodes with a formula in focus resemble the non-trivial clusters of the proof tree with back edges: Any node with a formula in focus is either in a non-trivial cluster or is labelled by f.

Formally we call a GKe proof *minimally focused* if (i) if v is labelled by u, then its child is labelled by D and (ii) if $\mathsf{depth}(v) < \mathsf{depth}(v')$ for a child v of v', then v is labelled by f. As every proof can be transformed to a minimally focused proof of the same sequent with the same cut-rank, cut-depth and cut-sum-depth, we will always assume that GKe proofs are minimally focused.

Definition 11. Let π be a GKe proof and v be a node in π. We define the GKe proof π_v as the subproof of π rooted at v, where recursively every open leaf l in π_v is replaced by $\pi_{c(l)}$.[1]

Note that π_v is well-defined, as $c(l)$ is a proper descendant of v for every open leaf l. Hence, at some point $\pi_{c(l)}$ has no open leaves. Importantly, $\mathsf{depth}(\pi_v) = \mathsf{depth}(v)$.

[1] In order to guarantee that D rules are labelled by unique discharge tokens, discharge tokens x are replaced by fresh discharge tokens, whenever a D^{x} rule is duplicated.

4.2 Important cuts

The strategy to eliminate an important cut of rank n is as follows. We first push the cut upwards, until it can not be reduced any more, those cuts we call *critical*. Then we take π_l and π_r, the proofs of the left and right premise of the cut, and '*zip*' them together, while deleting all descendants of the cut formula. This is done recursively, similarly to how a cut can be pushed 'up'. In order to find suitable repeats we also have to keep track of the already constructed cut-free subproof. The intermediate objects of this process are called *traversed proofs*, that correspond to proofs, where on every branch of the proof there is at most one cut of rank n.

Definition 12. Let π be a GKe proof and C an important cut in π. We call C *critical* if

1. the left premise of C is labelled by a u rule, i.e., the cut is of the form

$$\dfrac{\dfrac{\dfrac{\Sigma_l' \Rightarrow \Pi_l, \psi}{\Sigma_l' \Rightarrow \Pi_l, \psi}\ \mathsf{D}^{\mathsf{x}}}{\Sigma_l \Rightarrow \Pi_l, \psi}\ \mathsf{u} \qquad \psi^u, \Sigma_r \Rightarrow \Pi_r}{\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r}\ \mathsf{cut}$$

2. the cut formula ψ is of the form $\mathsf{F}\varphi$ or $\Diamond\mathsf{F}\varphi$.

Lemma 13. *Let π be a* GKe *proof with one important cut at the root. Then there is a* GKe *proof π' of the same sequent with the same cut-rank, cut-depth and sum-cut-depth, where all cuts are critical.*

Proof. We can apply cut reductions from Appendix A until all cuts satisfy point 1 of Definition 12. If ψ is not of the form $\mathsf{F}\varphi$ or $\Diamond\mathsf{F}\varphi$, then unfold D^{x}, i.e. apply a cut reduction to u and D^{x} on the left branch. Continue this process until all cuts are critical. In all cut reductions, except where $\mathsf{F}\varphi$ is principal, the syntactic size (i.e., number of symbols) of the cut formula is not increased, and in the reduction for $\Diamond$ the syntactic size of the cut formula decreases. Due to Lemma 7 on the path from a companion node to one of its discharged leaves there is a $\Diamond$ rule. Hence, the syntactic size of the cut formula ψ decreases until ψ is of the form $\mathsf{F}\varphi$ and after applying further cut reductions both conditions of Definition 12 are satisfied. □

Definition 14. An $\mathsf{F}\varphi$-*traversed proof* ρ of $\Sigma \Rightarrow \Pi$ is a derivation of $\Sigma \Rightarrow \Pi$, where all leaves are either closed or labelled by a sequent $\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r$ together with a triple (π_l, ψ, π_r), written as $[\pi_l]\psi[\pi_r]$, such that $\psi \equiv \mathsf{F}\varphi$ or $\psi \equiv \Diamond\mathsf{F}\varphi$, π_l is a proof of $\Gamma_l \Rightarrow \Delta_l, \psi$ and π_r is a proof of $\psi^a, \Gamma_r' \Rightarrow \Delta_r$, where $(\Gamma_r')^u = \Gamma_r$. If $\mathsf{F}\varphi$ is clear from the context we will just write *traversed proof*.

By applying a cut with cut-formula ψ at every open leaf, every traversed proof ρ corresponds to a GKe proof π, where on every branch of the proof there is at most one cut of rank $\mathsf{rank}(\psi)$. Hence, transforming a $\mathsf{F}\varphi$-traversed proof to a traversed proof without open leaves corresponds to eliminating cuts of rank $\mathsf{rank}(\mathsf{F}\varphi)$.

Lemma 15. *Let π be a* GKe *proof with only one critical cut at the root, i.e. π is of the form*

$$\dfrac{\dfrac{\overset{\pi_L}{\Sigma_l' \Rightarrow \Pi_l, \psi}}{\Sigma_l \Rightarrow \Pi_l, \psi}\ \mathsf{u} \qquad \overset{\pi_R}{\psi^u, \Sigma_r \Rightarrow \Pi_r}}{\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r}\ \mathsf{cut}$$

where π_L and π_R are cut-free and of depths k_L and k_R, respectively, and the cut formula ψ equals $\mathsf{F}\varphi$ or $\Diamond\mathsf{F}\varphi$. Then we can construct a proof π' of $\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$, where all cuts either have rank $\mathsf{rank}(\varphi)$ *or rank* $\mathsf{rank}(\mathsf{F}\varphi)$*, depth $\leq \max\{k_L, k_R\}$ and sum-depth $< k_L + k_R$.*

Proof. Let ρ_I be the $\mathsf{F}\varphi$-traversed proof of $\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$ consisting of one application of u and an open leaf labelled by $\Sigma'_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$ together with $[\pi_L]\psi[\pi_R]$, this we will denote by

$$\dfrac{\begin{array}{c}[\pi_L]\psi[\pi_R]\\ \Sigma'_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r\end{array}}{\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r}\ \mathsf{u}$$

Starting from ρ_I we will inductively transform the traversed proof, until we end up with a traversed proof of $\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$, where all leaves are closed. This will be done as follows: Let ρ be a traversed proof. If all leaves are closed we are done. Otherwise consider the leftmost open leaf v labelled by

$$\begin{array}{c}[\pi_l]\psi[\pi_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

Let the roots of π_l and π_r be labelled by the rules R_l and R_r, respectively. We transform ρ by a case distinction on R_l and R_r. We only show the crucial cases, the full proof can be found in Appendix B.

- If R_l is D^{x}, we make a case distinction:
  - If there is a node c in ρ, that is an ancestor of v and labelled by $\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r$, such that the path from c to v is successful, then insert a D^{z}-rule at c and let v be the discharged leaf $[\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r]^{\mathsf{z}}$ with fresh discharge token z.
  - Else let $\tilde{\pi}_l$ be the proof obtained from π'_l by replacing every discharged leaf labelled by x with π_l. We replace v by

$$\begin{array}{c}[\tilde{\pi}_l]\psi[\pi_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

- If R_r is D^{y}, then let $\tilde{\pi}_r$ be the proof obtained from π'_r by replacing every discharged leaf labelled by y with π_r. We replace v by

$$\begin{array}{c}[\pi_l]\psi[\tilde{\pi}_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

- If R_l is F_R and R_r is F_L, i.e. π_l and π_r are of the forms

$$\dfrac{\begin{array}{c}\pi'_l\\ \Gamma_l \Rightarrow \Delta_l, \varphi, \Diamond\mathsf{F}\varphi\end{array}}{\Gamma_l \Rightarrow \Delta_l, \mathsf{F}\varphi}\ \mathsf{F}_R \qquad\qquad \dfrac{\begin{array}{c}\pi'_r\\ \Diamond\mathsf{F}\varphi^a, \Gamma_r \Rightarrow \Delta_r\end{array} \quad \begin{array}{c}\pi^u_r\\ \varphi^u, \Gamma_r \Rightarrow \Delta_r\end{array}}{\mathsf{F}\varphi^a, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{F}_L$$

  then v is replaced by

$$\dfrac{\begin{array}{c}[\pi'_l]\Diamond\mathsf{F}\varphi[\pi'_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r, \varphi\end{array} \quad \begin{array}{c}\pi^u_r\\ \varphi^u, \Gamma_r \Rightarrow \Delta_r\end{array}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{cut}$$

  where the introduced cut has rank $\mathsf{rank}(\varphi)$.

- If R_l is f of the form

$$\dfrac{\begin{array}{c}\pi'_l\\ \Gamma^u_l \Rightarrow \Delta_l, \psi\end{array}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{f}$$

then v is replaced by

$$\dfrac{\begin{array}{c}\pi_l' \\ \Gamma_l^u \Rightarrow \Delta_l, \psi\end{array} \qquad \dfrac{\begin{array}{c}\pi_r \\ \psi^a, \Gamma_r' \Rightarrow \Delta_r\end{array}}{\psi^u, \Gamma_r \Rightarrow \Delta_r}\,\mathsf{u}}{\dfrac{\Gamma_l^u, \Gamma_r \Rightarrow \Delta_l, \Delta_r}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\,\mathsf{f}}\,\mathsf{cut}$$

where the introduced cut has rank $\mathsf{rank}(\mathsf{F}\varphi)$, depth $\leq \max\{k_L, k_R\}$ and sum-depth lower than $k_L + k_R$, as the left premise of the cut is in the next component.

- If R_r is f, we do the analogous transformation as in the case above.
- If R_l or R_r is any non-principal rule, we push the open leaf 'up', similarly as a cut can be pushed away from the root.

Let ρ_i and ρ_j be traversed proofs. We write $\rho_i < \rho_j$ if ρ_j can be obtained from ρ_i by the above construction and $\rho_i \neq \rho_j$. If $\rho_i < \rho_j$, then ρ_i is a subproof of ρ_j, in the sense that ρ_j can be obtained from ρ_i by replacing some open leaves in ρ_i by traversed proofs and inserting nodes labelled by D. Thus, ρ_j consists of at least the nodes in ρ_i and we can identify nodes in ρ_i with nodes in ρ_j.

Let $\rho > \rho_I$ and let v be an open leaf in ρ labelled by

$$\begin{array}{c}[\pi_l]\,\psi[\pi_r] \\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

We can define a node $u_L(v) \in \pi_L$ with $\pi_l = \pi_{u_L(v)}$, $\mathsf{S}(u_L(v)) = \Gamma_l \Rightarrow \Delta_l, \psi$ and such that $u_L(v)$ is in the root-cluster of π_L or is labelled by f. Analogously we can define $u_R(v) \in \pi_R$. These definitions extend to nodes w below an open leaf v. Hence, nodes below an open leaf can be labelled by only finitely many distinct sequents.

Let τ be the path from the root of ρ to the open leaf v. Using the above definitions of $u_L(w)$ and $u_R(w)$ we obtain corresponding paths τ_L in π_L and τ_R in π_R.

Combining these observations, we can find a bound on the height of open leaves: As there are only finitely many distinct sequents on τ, at some point an open leaf v is reached, such that an ancestor c in ρ is labelled by the same sequent. The path $\tau(v)$ from c to v in ρ corresponds to a path $\tau_L(v)$ in π_L. If $\tau(v)$ is of some certain length the path $\tau_L(v)$ is successful. This implies that the path $\tau(v)$ is successful as well and the leaf v will be closed. A more detailed proof with explicit bounds can be found in Appendix B.

In conclusion, as every constructed tree is finitely branching, after finitely many steps a traversed proof ρ_T without open leaves is constructed. As every traversed proof without open leaves is a proof we have shown the lemma. □

Example 16. Consider the following proof π containing a crucial cut

$$\dfrac{\dfrac{\begin{array}{c}\pi_L \\ \mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}\Diamond p\end{array}}{\mathsf{F}\Diamond\Diamond p^u \Rightarrow \mathsf{F}\Diamond p}\,\mathsf{u} \qquad \dfrac{\begin{array}{c}\pi_R \\ \mathsf{F}\Diamond p^f \Rightarrow \mathsf{F}p\end{array}}{\mathsf{F}\Diamond p^u \Rightarrow \mathsf{F}p}\,\mathsf{u}}{\mathsf{F}\Diamond\Diamond p^u \Rightarrow \mathsf{F}p}\,\mathsf{cut}$$

where the two proofs π_L and π_R are

$$
\dfrac{\dfrac{\dfrac{\dfrac{[\mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}\Diamond p]^{\mathsf{x}}}{\Diamond\mathsf{F}\Diamond\Diamond p^f \Rightarrow \Diamond\mathsf{F}\Diamond p}\,\Diamond \qquad \overline{\overline{\Diamond\Diamond p^u \Rightarrow \Diamond\mathsf{F}\Diamond p}}}{a:\ \mathsf{F}\Diamond\Diamond p^f \Rightarrow \Diamond p, \Diamond\mathsf{F}\Diamond p}\,\mathsf{F}_L}{\mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}\Diamond p}\,\mathsf{F}_R}{\mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}\Diamond p}\,\mathsf{D}^{\mathsf{x}}
\qquad
\dfrac{\dfrac{\dfrac{\dfrac{[\mathsf{F}\Diamond p^f \Rightarrow \mathsf{F}p]^{\mathsf{y}}}{\Diamond\mathsf{F}\Diamond p^f \Rightarrow \Diamond\mathsf{F}p}\,\Diamond \qquad \overline{\overline{\Diamond p^u \Rightarrow \Diamond\mathsf{F}p}}}{b:\ \mathsf{F}\Diamond p^f \Rightarrow p, \Diamond\mathsf{F}p}\,\mathsf{F}_L}{\mathsf{F}\Diamond p^f \Rightarrow \mathsf{F}p}\,\mathsf{F}_R}{\mathsf{F}\Diamond p^f \Rightarrow \mathsf{F}p}\,\mathsf{D}^{\mathsf{y}}
$$

For the sake of readability we apply weakening implicitly and omit w rules in this example. The leaves marked by double lines can be proved easily without cuts or repeats.

Following the construction of Lemma 15 we define ρ_I as above. In the first steps, the proofs π_L and π_R get unfolded. Then the F_R rule in π_R is added to ρ, this results in the following traversed proof:

$$
\dfrac{\dfrac{\begin{array}{c}[\pi_a]\mathsf{F}\Diamond p[\pi_b]\\ \mathsf{F}\Diamond\Diamond p^f \Rightarrow p, \Diamond\mathsf{F}p\end{array}}{\mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}p}\,\mathsf{F}_R}{\mathsf{F}\Diamond\Diamond p^u \Rightarrow \mathsf{F}p}\,\mathsf{u}
$$

Now the rules R_l and R_R are F_R and F_L, respectively. Hence a cut with cut-formula $\Diamond p$ is introduced. Then the traversed proof gets further transformed into

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{v:\ \begin{array}{c}[\pi_L]\mathsf{F}\Diamond p[\pi_R]\\ \mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}p\end{array}}{\Diamond\mathsf{F}\Diamond\Diamond p^f \Rightarrow \Diamond\mathsf{F}p}\,\Diamond \qquad \overline{\overline{\Diamond\Diamond p^u, \Rightarrow \Diamond\mathsf{F}p}}}{\mathsf{F}\Diamond\Diamond p^f \Rightarrow \Diamond p, \Diamond\mathsf{F}p}\,\mathsf{F}_L \qquad \overline{\overline{\Diamond p^u \Rightarrow \Diamond\mathsf{F}p}}}{c:\ \mathsf{F}\Diamond\Diamond p^f \Rightarrow p, \Diamond\mathsf{F}p}\,\mathsf{cut}}{\mathsf{F}\Diamond\Diamond p^f \Rightarrow \mathsf{F}p}\,\mathsf{F}_R}{\mathsf{F}\Diamond\Diamond p^u \Rightarrow \mathsf{F}p}\,\mathsf{u}
$$

The root of π_L is labelled by D^{x} and the repeat condition for v is satisfied: The node c is labelled by the same sequent as v and the path from c to v is successful. Thus a D^{z} rule is inserted at c and v is discharged by z, which results in a GKe proof.

4.3 Unimportant cuts

Definition 17. Let π be a GKe derivation and v be a node in π. The *infinite unfolding of* $\mathsf{comp}(v)$ in π, written π^{*_v}, is obtained from π by recursively replacing every discharged leaf l, that is a component descendant of v, by π_l and removing nodes labelled by D^{x}, where no discharged leaf is labelled by x.

Lemma 18. *Let π be a* GKe *proof of cut-rank n and cut-depth $k+1$, such that all cuts are unimportant and in the root-cluster. Then we can transform π into a* GKe *proof π' of the same sequent with cut-rank n and cut-depth $\leq k$.*

Proof. Let $\Gamma \Rightarrow \Delta$ be the sequent at the root r of π and $m = |\mathsf{Clos}(\Gamma \Rightarrow \Delta)|$. Let π^{*_r} be the infinite unfolding of $\mathsf{comp}(r)$ of π and let A be $\mathsf{comp}(r)$ in π^{*_r}. Note that A is infinite.

We want to push the cuts (also the ones with cut-depth $\leq k$, to avoid cut reductions with cuts) in A in π^{*r} upwards, until the set of nodes A_0 of height at most $8^m + 1$ in A do not contain cuts.

This can be done inductively: Choose a root-most node v labelled by cut, take a lowest descendant v' of v, that is labelled by cut, such that both premises are not labelled by cut (That could also be v itself). This can always be found as π is a proof and hence a repeat does not only consist of cut rules. Apply a suitable cut reduction defined in Appendix A to v'. At some point this will be a non-principal cut reduction, because π is a proof and the cut-formula is out of focus: There are infinitely many F_L rules, where the principal formula is in focus. Thus the height of the cut-free sub-proof increases. Continue this process until A_0 is cut-free. All cuts outside of A have cut-depth $\leq k$, otherwise they would be in A. Because all cut-reductions defined in Appendix A preserve the cut-rank, all cuts have cut-rank $\leq n$.

As A_0 is cut-free, each formula in A_0 is in $\mathsf{Clos}(\Gamma \Rightarrow \Delta)$, where it could occur or not occur in Γ and Δ, and be in focus or out of focus in Γ. Thus, there are at most 8^m many distinct sequents in A_0 and on each branch in A_0 there is a node v such that an ancestor $c(v)$ of v is labelled by the same sequent. For each such branch choose the root-most such node v, insert a D^{x} rule at $c(v)$ and let v be a discharged leaf labelled by x with fresh discharge token x. All sequents in A_0 have a formula in focus. Because of Lemma 10 all cut-reductions are not affecting formulas in focus and F_L-rules, where the principal formula is in focus, remain. The path from $c(v)$ to v is successful and we obtain a proof with cut-rank n and cut-depth $\leq k$. □

4.4 Cut elimination

As in the finitary case, we want to prove cut elimination by first showing cut admissibility. Yet in the context of cyclic proofs it does not suffice to eliminate one cut at the root of the proof, but rather we transform a proof with cuts only in the root-cluster to a cut-free proof. If the root-cluster is trivial this coincides with the usual notion of cut admissibility.

Lemma 19 (Cut admissibility)**.** *Let π be a* GKe *proof with cuts only in the root-cluster. Then we can transform π into a cut-free* GKe *proof π' of the same sequent.*

Proof. By triple induction on the cut-rank n, the cut-depth k and the sum-cut-depth m of π. If the root-cluster is trivial, there is one important cut at the root of π. Otherwise $k > 0$ and all cuts are unimportant.

In the first case Lemma 15 together with Lemma 13 yields a proof π_0, where all cuts have cut-rank $< n$ or rank n, depth $\leq k$ and sum-depth $< m$. In the second case Lemma 18 yields a proof π_1 with cut-rank n and cut-depth $< k$. In both cases we inductively apply the induction hypothesis to every subproof with cuts only in the root-cluster to obtain a cut-free proof π' of the same sequent. □

Theorem 20 (Cut elimination)**.** *We can transform every* GKe *proof π into a cut-free* GKe *proof π' of the same sequent.*

Proof. By induction on the number of clusters with cuts in π. Let π_0 be a subproof of π with cuts only in the root-cluster. Lemma 19 gives a cut-free proof π_0' of the same sequent as π_0. By substituting π_0' in π we obtain a proof of the same sequent, where the number of clusters with cuts is reduced. □

5 Conclusion

We have presented a syntactic cut-elimination procedure for the cyclic proof system GKe. To our knowledge, this is the first cut-elimination method that works directly on cyclic proofs without a detour on infinitary proofs. This work is a case study in the sense that we believe the introduced method can be generalised to other annotated, cyclic proof systems for fragments of the modal μ-calculus. As a next step we are working on applying the strategy to a proof system for Propositional Dynamic Logic [1] . Other interesting candidates are the focus system for the alternation-free fragment introduced in [8] and the annotated proof systems for the modal μ-calculus [6, 16, 4].

References

[1] Julian Bradfield & Colin Stirling (2007): *Modal mu-calculi*. *Handbook of Modal Logic* 3, Elsevier, pp. 721–756, doi:10.1016/S1570-2464(07)80015-2.

[2] James Brotherston (2006): *Sequent calculus proof systems for inductive definitions*. Ph.D. thesis. Available at `https://era.ed.ac.uk/handle/1842/1458`.

[3] Anupam Das & Damien Pous (2018): *Non-Wellfounded Proof Theory For (Kleene+Action)(Algebras+Lattices)*. doi:10.4230/LIPICS.CSL.2018.19.

[4] Maurice Dekker, Johannes Kloibhofer, Johannes Marti & Yde Venema (2023): *Proof Systems for the Modal μ-Calculus Obtained by Determinizing Automata*. In: *TABLEAUX 2023, Prague, Czech Republic*, *Lecture Notes in Computer Science* 14278, Springer, pp. 242–259, doi:10.1007/978-3-031-43513-3_14.

[5] Jérôme Fortier & Luigi Santocanale (2013): *Cuts for circular proofs: semantics and cut-elimination*. In: *Computer Science Logic 2013 (CSL 2013)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 23, pp. 248–262, doi:10.4230/LIPIcs.CSL.2013.248.

[6] Natthapong Jungteerapanich (2010): *Tableau systems for the modal μ-calculus*. Ph.D. thesis, School of Informatics; The University of Edinburgh. Available at `http://hdl.handle.net/1842/4208`.

[7] Graham E. Leigh & Dominik Wehr (2023): *From GTC to Reset: Generating Reset Proof Systems from Cyclic Proof Systems*. arXiv:2301.07544.

[8] Johannes Marti & Yde Venema (2021): *A Focus System for the Alternation-Free μ-Calculus*. In: *TABLEAUX 2021, Proceedings*, Lecture Notes in Computer Science, Springer, pp. 371–388, doi:10.1007/978-3-030-86059-2_22.

[9] Damian Niwinski & Igor Walukiewicz (1996): *Games for the mu-Calculus*. *Theor. Comput. Sci.* 163(1&2), pp. 99–116, doi:10.1016/0304-3975(95)00136-0.

[10] Jan Rooduijn (2021): *Cyclic Hypersequent Calculi for Some Modal Logics with the Master Modality*. In: *TABLEAUX 2021*, Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 354–370, doi:10.1007/978-3-030-86059-2_21.

[11] Luigi Santocanale (2002): *A Calculus of Circular Proofs and Its Categorical Semantics*. In: *Foundations of Software Science and Computation Structures*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 357–371, doi:10.1007/3-540-45931-6_25.

[12] Alexis Saurin (2023): *A Linear Perspective on Cut-Elimination for Non-wellfounded Sequent Calculi with Least and Greatest Fixed-Points*. In: *TABLEAUX 2023, Prague, Czech Republic*, *Lecture Notes in Computer Science* 14278, Springer, pp. 203–222, doi:10.1007/978-3-031-43513-3_12.

[13] Yury Savateev & Daniyar Shamkanov (2020): *Non-well-founded Proofs for the Grzegorcyk Modal Logic*. *The Review of Symbolic Logic* 14(1), pp. 22–50, doi:10.1017/s1755020319000510.

[14] Kurt Schütte (1977): *Proof Theory*. Springer-Verlag, Berlin, Heidelberg.

[15] Daniyar S. Shamkanov (2014): *Circular proofs for the Gödel-Löb provability logic*. *Mathematical Notes* 96(3-4), pp. 575–585, doi:10.1134/s0001434614090326.

[16] Colin Stirling (2014): *A Tableau Proof System with Names for Modal Mu-calculus*. In: *HOWARD-60. A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, *EPiC Series in Computing* 42, EasyChair, pp. 306–318, doi:10.29007/lwqm.

[17] Igor Walukiewicz (1993): *On Completeness of the mu-calculus*. In: *Proceedings of the Eighth Annual Symposium on Logic in Computer Science (LICS '93), Montreal, Canada, June 19-23, 1993*, IEEE Computer Society, pp. 136–146, doi:10.1109/LICS.1993.287593.

A Cut reductions

For readability we state the cut reductions for a simplified cut rule, where $\Gamma_l = \Gamma_r$ and $\Delta_l = \Delta_r$. This can be generalized in the obvious way.

Principal cut reductions:

$$\dfrac{\dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi,\Diamond\mathsf{F}\varphi\end{matrix}}{\Gamma\Rightarrow\Delta,\mathsf{F}\varphi}\,\mathsf{F}_R \qquad \dfrac{\begin{matrix}\pi_1\\ \Diamond\mathsf{F}\varphi^u,\Gamma\Rightarrow\Delta\end{matrix} \quad \begin{matrix}\pi_2\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\mathsf{F}\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{F}_L}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\longrightarrow \quad \dfrac{\dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi,\Diamond\mathsf{F}\varphi\end{matrix} \quad \begin{matrix}\pi_1\\ \Diamond\mathsf{F}\varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{cut} \qquad \begin{matrix}\pi_2\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\dfrac{\dfrac{\begin{matrix}\pi_0\\ \gamma^a\Rightarrow\Delta,\varphi\end{matrix}}{\Diamond\gamma^a\Rightarrow\Diamond\Delta,\Diamond\varphi}\,\Diamond \qquad \dfrac{\begin{matrix}\pi_1\\ \varphi^u\Rightarrow\Delta\end{matrix}}{\Diamond\varphi^u\Rightarrow\Diamond\Delta}\,\Diamond}{\Diamond\gamma^a\Rightarrow\Diamond\Delta}\,\mathsf{cut} \quad \longrightarrow \quad \dfrac{\dfrac{\begin{matrix}\pi_0\\ \gamma^a\Rightarrow\Delta,\varphi\end{matrix} \quad \begin{matrix}\pi_1\\ \varphi^u\Rightarrow\Delta\end{matrix}}{\gamma^a\Rightarrow\Delta}\,\mathsf{cut}}{\Diamond\gamma^a\Rightarrow\Diamond\Delta}\,\Diamond$$

$$\dfrac{\dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix} \quad \begin{matrix}\pi_1\\ \Gamma\Rightarrow\Delta,\psi\end{matrix}}{\Gamma\Rightarrow\Delta,\varphi\wedge\psi}\,\wedge_R \qquad \dfrac{\begin{matrix}\pi_2\\ \varphi^u,\psi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\varphi\wedge\psi^u,\Gamma\Rightarrow\Delta}\,\wedge_L}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\longrightarrow \quad \dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix} \qquad \dfrac{\begin{matrix}\pi_1\\ \Gamma\Rightarrow\Delta,\psi\end{matrix} \quad \begin{matrix}\pi_2\\ \varphi^u,\psi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{cut}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\dfrac{\dfrac{\begin{matrix}\pi_0\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma\Rightarrow\Delta,\neg\varphi}\,\neg_R \qquad \dfrac{\begin{matrix}\pi_1\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix}}{\neg\varphi^u,\Gamma\Rightarrow\Delta}\,\neg_L}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \quad \longrightarrow \quad \dfrac{\begin{matrix}\pi_1\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix} \quad \begin{matrix}\pi_0\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

Trivial principal cut reductions:

$$\dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,p\end{matrix} \quad p^u\Rightarrow p}{\Gamma\Rightarrow\Delta,p}\,\mathsf{cut} \quad \longrightarrow \quad \begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,p\end{matrix}$$

$$\dfrac{p^u\Rightarrow p \quad \begin{matrix}\pi_1\\ p^u,\Gamma\Rightarrow\Delta\end{matrix}}{p^u,\Gamma\Rightarrow\Delta}\,\mathsf{cut} \quad \longrightarrow \quad \begin{matrix}\pi_1\\ p^u,\Gamma\Rightarrow\Delta\end{matrix}$$

$$\dfrac{\dfrac{\overset{\pi_0}{\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{w}_R \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \overset{\pi_0}{\Gamma\Rightarrow\Delta}$$

$$\dfrac{\overset{\pi_0}{\Gamma\Rightarrow\Delta,\varphi} \qquad \dfrac{\overset{\pi_1}{\Gamma\Rightarrow\Delta}}{\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{w}_L}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \overset{\pi_1}{\Gamma\Rightarrow\Delta}$$

Cut reductions for u, D^{x} and f:
We push u and f rules ‘upwards’ away from the root and unfold D^{x} rules. The presented reductions are analogous, if the right premise of the cut is labelled by u, D^{x} or f.

$$v:\ \dfrac{\dfrac{\dfrac{\overset{\pi_0}{\Gamma'\Rightarrow\Delta,\varphi}}{\Gamma'\Rightarrow\Delta,\varphi}\,\mathsf{D}^{\mathsf{x}}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{u} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \dfrac{\overset{\pi_0'}{\Gamma\Rightarrow\Delta,\varphi} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

where π_0' is obtained from π_0 by (i) unfocusing sequents up to D rules and leaves labelled by x and (ii) replacing every discharged leaf labelled by x with π_v, where v is the left premise of the cut rule.[2]

$$v:\ \dfrac{\dfrac{\overset{\pi_0}{\Gamma\Rightarrow\Delta,\varphi}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{D}^{\mathsf{x}} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \dfrac{\overset{\pi_0'}{\Gamma\Rightarrow\Delta,\varphi} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

where π_0' is obtained from π_0 by replacing every discharged leaf labelled by x with π_v, where v is the left premise of the cut rule.

$$\dfrac{\dfrac{\dfrac{\overset{\pi_0}{\Gamma_0\Rightarrow\Delta_0}}{\Gamma^u\Rightarrow\Delta,\varphi}\,\mathsf{u}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{f} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \dfrac{\dfrac{\overset{\pi_0'}{\Gamma^u\Rightarrow\Delta,\varphi}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{f} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

where π_0' is defined as above. For a rule R different from u we do the following:

$$\dfrac{\dfrac{\dfrac{\overset{\pi_0}{\Gamma_0\Rightarrow\Delta_0}}{\Gamma^u\Rightarrow\Delta,\varphi}\,\mathsf{R}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{f} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut} \qquad\longrightarrow\qquad \dfrac{\dfrac{\dfrac{\overset{\pi_0}{\Gamma_0\Rightarrow\Delta_0}}{\Gamma_0'\Rightarrow\Delta_0}\,\mathsf{f}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{R} \qquad \overset{\pi_1}{\varphi^u,\Gamma\Rightarrow\Delta}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

[2]Here and in the following cut reductions discharge tokens y are replaced by fresh discharge tokens, whenever a D^{y} rule is duplicated.

Non-principal cut-reductions for a rule R different from cut, $\Diamond$, u, f and D:

$$\dfrac{\begin{matrix}\pi_0\\ \Gamma \Rightarrow \Delta, \varphi\end{matrix} \qquad \dfrac{\begin{matrix}\pi_1\\ \varphi^u,\Gamma_1\Rightarrow\Delta_1\end{matrix} \quad \cdots \quad \begin{matrix}\pi_n\\ \varphi^u,\Gamma_n\Rightarrow\Delta_n\end{matrix}}{\varphi^u,\Gamma\Rightarrow\Delta}\,\mathsf{R}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\longrightarrow \quad \dfrac{\dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix} \quad \begin{matrix}\pi_1\\ \varphi^u,\Gamma_1\Rightarrow\Delta_1\end{matrix}}{\Gamma_1,\Gamma\Rightarrow\Delta_1,\Delta}\,\mathsf{cut} \quad \cdots \quad \dfrac{\begin{matrix}\pi_0\\ \Gamma\Rightarrow\Delta,\varphi\end{matrix} \quad \begin{matrix}\pi_n\\ \varphi^u,\Gamma_n\Rightarrow\Delta_n\end{matrix}}{\Gamma_n,\Gamma\Rightarrow\Delta_n,\Delta}\,\mathsf{cut}}{\Gamma\Rightarrow\Delta}\,\mathsf{R}$$

$$\dfrac{\dfrac{\begin{matrix}\pi_1\\ \Gamma_1\Rightarrow\Delta_1,\varphi\end{matrix} \quad \cdots \quad \begin{matrix}\pi_n\\ \Gamma_n\Rightarrow\Delta_n,\varphi\end{matrix}}{\Gamma\Rightarrow\Delta,\varphi}\,\mathsf{R} \qquad \begin{matrix}\pi_0\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma\Rightarrow\Delta}\,\mathsf{cut}$$

$$\longrightarrow \quad \dfrac{\dfrac{\begin{matrix}\pi_1\\ \Gamma_1\Rightarrow\Delta_1,\varphi\end{matrix} \quad \begin{matrix}\pi_0\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma_1,\Gamma\Rightarrow\Delta_1,\Delta}\,\mathsf{cut} \quad \cdots \quad \dfrac{\begin{matrix}\pi_n\\ \Gamma_n\Rightarrow\Delta_n,\varphi\end{matrix} \quad \begin{matrix}\pi_0\\ \varphi^u,\Gamma\Rightarrow\Delta\end{matrix}}{\Gamma_n,\Gamma\Rightarrow\Delta_n,\Delta}\,\mathsf{cut}}{\Gamma\Rightarrow\Delta}\,\mathsf{R}$$

B Important Cuts

Lemma 21. *Let S be a non-trivial cluster of size n in a proof π and $\tau = v_1v_2...$ be a path τ in S. Let $l(\tau)$ be the* length *of the path τ. Then,*

1. *if $l(\tau) \geq n$, then τ is successful,*
2. *if $l(\tau) \geq n$, then there is a companion node on τ,*
3. *if $l(\tau) \geq n$, then there is a node labelled by $\Diamond$ on τ,*
4. *if $l(\tau) \geq 2k\cdot n$ then there are k indices $1 \leq m_1 < \cdots < m_k \leq 2k\cdot n$ such that $v_{m_1},...,v_{m_k}$ are companion nodes and the paths $v_{m_i}v_{m_i+1}...v_{m_j}$ are successful for all $i < j$.*

Proof. Every node of τ is in S, hence it has a formula in focus. Thus for 1 we only have to show that τ passes through an application of F_L, where the principal formula is in focus. If every node of S is passed, then 1, 2 and 3 are satisfied. Otherwise there has to be a node that is passed twice. This is only possible if there is a discharged leaf u such that every node of $\tau(u)$ is in τ. Thus 1, 2 and 3 are satisfied. 4 follows by combining 2 and 1. □

Lemma 15. *Let π be a GKe proof with only one critical cut at the root, i.e. the cut formula ψ equals $\mathsf{F}\varphi$ or $\Diamond\mathsf{F}\varphi$ and π is of the form*

$$\dfrac{\dfrac{\begin{matrix}\pi_L\\ \Sigma_l'\Rightarrow\Pi_l,\psi\end{matrix}}{\Sigma_l\Rightarrow\Pi_l,\psi}\,\mathsf{u} \qquad \begin{matrix}\pi_R\\ \psi^u,\Sigma_r\Rightarrow\Pi_r\end{matrix}}{\Sigma_l,\Sigma_r\Rightarrow\Pi_l,\Pi_r}\,\mathsf{cut}$$

where π_L and π_R are cut-free and of depths k_L and k_R, respectively. Then we can construct a proof π' of $\Sigma_l,\Sigma_r \Rightarrow \Pi_l,\Pi_r$, where all cuts either have rank $\mathsf{rank}(\varphi)$ or rank $\mathsf{rank}(\mathsf{F}\varphi)$, depth $\leq \max\{k_L,k_R\}$ and sum-depth $< k_L + k_R$.

Proof. Let ρ_I be the $\mathsf{F}\varphi$-traversed proof of $\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$ consisting of one application of u and an open leaf labelled by $\Sigma_l', \Sigma_r \Rightarrow \Pi_l, \Pi_r$ together with $[\pi_L]\psi[\pi_R]$, this we will denote by

$$\dfrac{\begin{array}{c}[\pi_L]\psi[\pi_R]\\ \Sigma_l', \Sigma_r \Rightarrow \Pi_l, \Pi_r\end{array}}{\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r}\ \mathsf{u}$$

Starting from ρ_I we will inductively transform the traversed proof, until we end up with a traversed proof of $\Sigma_l, \Sigma_r \Rightarrow \Pi_l, \Pi_r$, where all leaves are closed. This will be done as follows: Let ρ be a traversed proof. If all leaves are closed we are done. Otherwise consider the leftmost open leaf v labelled by

$$\begin{array}{c}[\pi_l]\psi[\pi_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

Let the roots of π_l and π_r be labelled by the rules R_l and R_r, respectively. We transform ρ by a case distinction on R_l and R_r.

- If R_l is D^{x}, then π_l has the form

$$\dfrac{\begin{array}{c}\pi_l'\\ \Gamma_l \Rightarrow \Delta_l, \psi\end{array}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{D}^{\mathsf{x}}$$

  We make a case distinction:

  - If there is a node c in ρ, that is an ancestor of v and labelled by $\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r$, such that the path from c to v is successful, then insert a D^{z}-rule at c and let v be the discharged leaf $[\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r]^{\mathsf{z}}$ with fresh discharge token z.
  - Else let $\tilde{\pi}_l$ be the proof obtained from π_l' by replacing every discharged leaf labelled by x with π_l.[3]
    We replace v by

$$\begin{array}{c}[\tilde{\pi}_l]\psi[\pi_r]\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{array}$$

- If R_l is F_R with principal formula $\mathsf{F}\varphi$ we make a case distinction on R_r:
  - If R_r is F_L with principal formula $\mathsf{F}\varphi$, the proofs π_l and π_r have the form

$$\dfrac{\begin{array}{c}\pi_l'\\ \Gamma_l \Rightarrow \Delta_l, \varphi, \Diamond\mathsf{F}\varphi\end{array}}{\Gamma_l \Rightarrow \Delta_l, \mathsf{F}\varphi}\ \mathsf{F}_R \qquad\qquad \dfrac{\begin{array}{c}\pi_r'\\ \Diamond\mathsf{F}\varphi^a, \Gamma_r \Rightarrow \Delta_r\end{array} \quad \begin{array}{c}\pi_r^u\\ \varphi^u, \Gamma_r \Rightarrow \Delta_r\end{array}}{\mathsf{F}\varphi^a, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{F}_L$$

    and v is replaced by

$$\dfrac{\begin{array}{c}[\pi_l']\Diamond\mathsf{F}\varphi[\pi_r']\\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r, \varphi\end{array} \quad \begin{array}{c}\pi_r^u\\ \varphi^u, \Gamma_r \Rightarrow \Delta_r\end{array}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{cut}$$

    where the introduced cut has rank $\mathsf{rank}(\varphi)$.

[3] Discharge tokens y are replaced by fresh discharge tokens, whenever a D^{y} rule is duplicated.

– If R_r is u of the form

$$\dfrac{\begin{matrix}\pi_r' \\ \psi^a, \Gamma_r' \Rightarrow \Delta_r\end{matrix}}{\psi^u, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{u}$$

then v is replaced by

$$\begin{matrix}[\pi_l]\,\psi[\pi_r'] \\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{matrix}$$

– If R_r is D^{y} of the form

$$\dfrac{\begin{matrix}\pi_r' \\ \psi^a, \Gamma_r \Rightarrow \Delta_r\end{matrix}}{\psi^a, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{D}^{\mathsf{y}}$$

then let $\tilde{\pi}_r$ be the proof obtained from π_r' by replacing every discharged leaf labelled by y with π_r.[4] We replace v by

$$\begin{matrix}[\pi_l]\,\psi[\tilde{\pi}_r] \\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{matrix}$$

– If R_r is f of the form

$$\dfrac{\begin{matrix}\pi_r' \\ \psi^u, \Gamma_r \Rightarrow \Delta_r\end{matrix}}{\psi^a, \Gamma_r' \Rightarrow \Delta_r}\ \mathsf{f}$$

then v is replaced by

$$\dfrac{\begin{matrix}\pi_l \\ \Gamma_l \Rightarrow \Delta_l, \psi\end{matrix} \qquad \begin{matrix}\pi_r' \\ \psi^u, \Gamma_r \Rightarrow \Delta_r\end{matrix}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{cut}$$

where the introduced cut has rank $\mathsf{rank}(\mathsf{F}\varphi)$, depth $\leq \max\{k_L, k_R\}$ and sum-depth lower than $k_L + k_R$, as the right premise of the cut is in the next component.

– If R_r is w of the form

$$\dfrac{\begin{matrix}\pi_r' \\ \Gamma_r' \Rightarrow \Delta_r\end{matrix}}{\psi^u, \Gamma_r' \Rightarrow \Delta_r}\ \mathsf{w}$$

then v is replaced by

$$\dfrac{\dfrac{\begin{matrix}\pi_r' \\ \Gamma_r' \Rightarrow \Delta_r\end{matrix}}{\Gamma_r \Rightarrow \Delta_r}\ \mathsf{u}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{w}$$

[4] Discharge tokens z are replaced by fresh discharge tokens, whenever a D^{z} rule is duplicated.

- If R_r is any other rule, then the proof has the form

$$\dfrac{\begin{matrix}\pi_r^1 \\ \psi^a, \Gamma_r^1 \Rightarrow \Delta_r^1\end{matrix} \quad \dots \quad \begin{matrix}\pi_r^n \\ \psi^a, \Gamma_r^n \Rightarrow \Delta_r^n\end{matrix}}{\psi^a, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{R}$$

and v is replaced by

$$\dfrac{\begin{matrix}[\pi_l]\psi[\pi_r^1] \\ \Gamma_l, \Gamma_r^1 \Rightarrow \Delta_l, \Delta_r^1\end{matrix} \quad \dots \quad \begin{matrix}[\pi_l]\psi[\pi_r^n] \\ \Gamma_l, \Gamma_r^n \Rightarrow \Delta_l, \Delta_r^n\end{matrix}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{R}$$

- If R_l is $\Diamond$, then we make a case distinction on R_r:
  - If R_r is $\Diamond$, the proofs have the form

$$\dfrac{\begin{matrix}\pi_l' \\ \chi^f \Rightarrow \Delta_l, \mathsf{F}\varphi\end{matrix}}{\Diamond\chi^f \Rightarrow \Diamond\Delta_l, \Diamond\mathsf{F}\varphi}\ \Diamond \qquad\qquad \dfrac{\begin{matrix}\pi_r' \\ \mathsf{F}\varphi^a \Rightarrow \Delta_r\end{matrix}}{\Diamond\mathsf{F}\varphi^a \Rightarrow \Diamond\Delta_r}\ \Diamond$$

  Then v is replaced by

$$\dfrac{\begin{matrix}[\pi_l']\mathsf{F}\varphi[\pi_r'] \\ \chi^f \Rightarrow \Delta_l, \Delta_r\end{matrix}}{\Diamond\chi^f \Rightarrow \Diamond\Delta_l, \Diamond\Delta_r}\ \Diamond$$

  - If R_r is any other rule, we do the same as in the previous case.
- If R_l is f of the form

$$\dfrac{\begin{matrix}\pi_l' \\ \Gamma_l^u \Rightarrow \Delta_l, \psi\end{matrix}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{f}$$

then v is replaced by

$$\dfrac{\dfrac{\begin{matrix}\pi_l' \\ \Gamma_l^u \Rightarrow \Delta_l, \psi\end{matrix} \qquad \dfrac{\begin{matrix}\pi_r \\ \psi^a, \Gamma_r' \Rightarrow \Delta_r\end{matrix}}{\psi^u, \Gamma_r \Rightarrow \Delta_r}\ \mathsf{u}}{\Gamma_l^u, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{cut}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{f}$$

where the introduced cut has rank $\mathsf{rank}(\mathsf{F}\varphi)$, depth $\leq \max\{k_L, k_R\}$ and sum-depth lower than $k_L + k_R$, as the left premise of the cut is in the next component.

- If R_l is w of the form

$$\dfrac{\begin{matrix}\pi_l' \\ \Gamma_l \Rightarrow \Delta_l\end{matrix}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{w}$$

then v is replaced by

$$\dfrac{\begin{matrix}\pi_l' \\ \Gamma_l \Rightarrow \Delta_l\end{matrix}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{w}$$

- If R_l is any other rule, π_l has the form

$$\dfrac{\begin{matrix}\pi_l^1 \\ \Gamma_l^1 \Rightarrow \Delta_l^1, \psi\end{matrix} \quad \dots \quad \begin{matrix}\pi_l^n \\ \Gamma_l^n \Rightarrow \Delta_l^n, \psi\end{matrix}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{R}$$

Then v is replaced by

$$\dfrac{\begin{matrix}[\pi_l^1]\,\psi[\pi_r] \\ \Gamma_l^1, \Gamma_r \Rightarrow \Delta_l^1, \Delta_r\end{matrix} \quad \dots \quad \begin{matrix}[\pi_l^n]\,\psi[\pi_r] \\ \Gamma_l^n, \Gamma_r \Rightarrow \Delta_l^n, \Delta_r\end{matrix}}{\Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r}\ \mathsf{R}$$

Let ρ_i and ρ_j be traversed proofs. We write $\rho_i < \rho_j$ if ρ_j can be obtained from ρ_i by the above construction and $\rho_i \neq \rho_j$. It holds that $<$ is irreflexive, antisymmetric and transitive. Moreover, if $\rho_i < \rho_j$, then ρ_i is a subproof of ρ_j, in the sense that ρ_j can be obtained from ρ_i by replacing some open leaves in ρ_i by traversed proofs and inserting nodes labelled by D. Thus, ρ_j consists of at least the nodes in ρ_i and we can identify nodes in ρ_i with nodes in ρ_j.

From now on, whenever we speak about a traversed proof ρ we mean a traversed proof $\rho > \rho_I$.

Let $\rho > \rho_I$ and let v be an open leaf in ρ labelled by

$$\begin{matrix}[\pi_l]\,\psi[\pi_r] \\ \Gamma_l, \Gamma_r \Rightarrow \Delta_l, \Delta_r\end{matrix}$$

We first define nodes $u_L(v) \in \pi_L$ and $u_R(v) \in \pi_R$ with

1. $\pi_l = \pi_{u_L(v)}$ and $\pi_r = \pi_{u_R(v)}$,
2. $\mathsf{S}(u_L(v)) = \Gamma_l \Rightarrow \Delta_l, \psi$ and $\mathsf{S}(u_R(v)) = \psi^a, \Gamma_r' \Rightarrow \Delta_r$,
3. $u_L(v)$ is in the root-cluster of π_L or is labelled by f and $u_R(v)$ is in the root-component of π_R or is labelled by f.

The nodes $u_L(v)$ and $u_R(v)$ are defined by recursion on the construction. For ρ_I define $u_L(v)$ to be the root of π_L and $u_R(v)$ to be the root of π_R. For the recursion step we follow the case distinction. For example, let ρ' be obtained from ρ with leftmost open leaves v' and v, respectively, where the last applied rules in π_l and π_r are $\diamond$. Then $u_L(v')$ is the child of $u_L(v)$ and $u_R(v')$ is the child of $u_R(v)$.

3 follows, as π is minimally focused, hence every path through π_L, which leaves the component, has to pass a f rule, but then the corresponding leaf in ρ gets closed.

This definition extends to nodes w in a traversed proof ρ below an open leaf. For such nodes w and w' it moreover holds, that if w is the parent of w', then either

1. $u_L(w) = u_L(w')$,
2. $u_L(w)$ is the parent of $u_L(w')$ or

3. $u_L(w)$ is an ancestor of $u_L(w')$, where all but one node on the ancestor path are labelled by D.

The same holds for the nodes $u_R(w)$ and $u_R(w')$.

Let τ be the path from the root of ρ to the open leaf v. By the above definition we can define corresponding paths τ_L in π_L and τ_R in π_R.

Let n_L be the size (i.e. the number of nodes) of the root-cluster of π_L and n_R be the size of the root-component of π_R. By the above argumentation nodes in τ can only be labelled by at most $n = 4 \cdot n_L \cdot n_R$ distinct sequents. Moreover, if $l(\tau) = k \cdot n_R$, then $l(\tau_L) \geq k$. This holds as in every step of the construction either π_l or π_r is transformed. Whenever π_l is transformed case 1 above can not be the case. But if $l(\tau) \geq n_R$, π_l has to be transformed, as otherwise $l(\tau_R) \geq n_R$. Due to Lemma 21.3 there has to be a $\Diamond$ rule on τ_R, in which case π_l gets transformed as well.

We claim that every open leaf in a traversed proof ρ has height at most $5 \cdot n_L^2 \cdot n_R^2$. For suppose that v is an open leaf of height more than $4 \cdot n_L^2 \cdot n_R^2$. Let τ be the path from the root of ρ to v and let τ_L the corresponding path in π_L. It holds that $l(\tau_L) \leq 5 \cdot n_L^2 \cdot n_R$. Using Lemma 21.4 we obtain $5 \cdot n_L \cdot n_R$ companion nodes $c_1, c_2, ...$ in τ_L. Hence there are as many traversed proofs with leftmost open leafs $v_1, v_2, ...$ with $u_L(v_j) = c_j$ for all j. This means that there are $5 \cdot n_L \cdot n_R$ traversed proofs $\rho_1 < \rho_2 < \cdots < \rho$, with left-most open leaves $v_1, v_2, ...$, which are *repeat leaves*, i.e. π_l is of the form

$$\dfrac{\begin{array}{c}\pi_l' \\ \Gamma_l \Rightarrow \Delta_l, \psi\end{array}}{\Gamma_l \Rightarrow \Delta_l, \psi}\ \mathsf{D}^{\times}$$

By the above argumentation the nodes $v_1, v_2, ...$ are labelled by at most $4 \cdot n_L \cdot n_R$ many distinct sequents. Hence there are $i < j$ such that $\mathsf{S}(v_i) = \mathsf{S}(v_j)$. Due to the choice of the companion nodes in Lemma 21.4 the path from c_i to c_j in τ_L is successful. Note that, if $u_L(w)$ is labelled by F_L, where the principal formula is in focus, then so is w. Therefore, and because all sequents in τ have a formula in focus, it follows that the path from v_i to v_j in τ is successful as well. Hence, the leaf gets discharged in the next step of transforming ρ_j. This contradicts the fact that v in ρ is an open leaf of height more than $5 \cdot n_l^2 \cdot n_r^2$.

In conclusion, as every constructed tree is finitely branching, after finitely many steps a traversed proof ρ_T without open leaves is constructed. As every traversed proof without open leaves is a proof we have shown the lemma. $\square$