



HAL
open science

Linear logic with the least and greatest fixed points : truth semantics, complexity and a parallel syntax

Abhishek De

► **To cite this version:**

Abhishek De. Linear logic with the least and greatest fixed points : truth semantics, complexity and a parallel syntax. Computation and Language [cs.CL]. Université Paris Cité, 2022. English. NNT : 2022UNIP7128 . tel-04274800

HAL Id: tel-04274800

<https://theses.hal.science/tel-04274800v1>

Submitted on 8 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS CITÉ

École doctorale de Sciences Mathématiques de Paris-Centre (ED 386)
Institut de Recherche en Informatique Fondamentale (IRIF)

THÈSE DE DOCTORAT EN INFORMATIQUE
Dirigée par Alexis SAURIN

Linear logic with the least and greatest fixed points

Truth semantics, complexity and a parallel syntax

Par Abhishek DE

Présentée et soutenue publiquement le 1er décembre 2022

Devant un jury composé de:

Mme. Bahareh AFSHARI	Asst. Professor, University of Amsterdam	Rapportrice
M. Laurent REGNIER	Professeur, Université d'Aix-Marseille	Rapporteur
M. David BAELDE	Professeur, Université de Rennes	Examinateur
Mme. Amina DOUMANE	CR CNRS, LIP, ENS Lyon	Examinatrice
Mme. Delia KESNER*	Professeure, Université Paris Cité	Examinatrice
M. Olivier LAURENT	DR CNRS, LIP, ENS Lyon	Examinateur
M. Alexis SAURIN	CR CNRS, Université Paris Cité	Directeur de thèse
M. R. Ramanujam*	Visiting Professor, Azim Premji University	Membre Invité

*en visioconférence

Linear logic with the least and greatest fixed points

Truth semantics, complexity, and a parallel syntax

This page has been intentionally left blank.

To everyone fighting against imperialism across the world

Abstract. The subject of this thesis is the proof theory of linear logic with the least and greatest fixed points. In the literature, several systems have been studied for this language *viz.* the wellfounded system that relies on Park's induction rule, and systems that implicitly characterise induction such as the circular system and the non-wellfounded system. This thesis contributes to the theory of these systems with the ultimate goal of exactly capturing the provability relation of these systems and application of these objects in programming languages supporting (co)inductive reasoning.

This thesis contains three parts. In the first part, we recall the literature on linear logic and the main approaches to the proof theory of logics with fixed points. In the second part, we obtain truth semantics for the wellfounded system, devise new wellfounded infinitely branching systems, and compute the complexity of provability in circular and non-wellfounded systems. In the third part, we devise non-wellfounded proof-nets and study their dynamics.

Keywords. proof theory | linear logic | cyclic proofs | fixed points | proof nets | phase semantics | cut elimination

Résumé. Le sujet de cette thèse est la théorie de la preuve de la logique linéaire avec les plus petits et les plus grands points fixes. Plusieurs systèmes ont été étudiés dans la littérature pour ce langage : le système bien fondé qui repose sur la règle d'induction de Park, et des systèmes qui caractérisent implicitement l'induction comme le système circulaire et son extension non bien fondée. Cette thèse contribue à la théorie de ces systèmes avec pour but ultime de capturer exactement la relation de prouvabilité de ces systèmes et de permettre l'application de ces objets dans les langages de programmation supportant le raisonnement (co)inductif.

Cette thèse contient trois parties. Dans la première partie, nous rappelons la littérature sur la logique linéaire et les principales approches de la théorie de la preuve des logiques à points fixes. Dans la deuxième partie, nous obtenons une sémantique de vérité pour le système bien fondé, nous concevons de nouveaux systèmes infiniment ramifiés bien fondés, et nous calculons la complexité de la prouvabilité dans les systèmes circulaires et non bien fondés. Dans la troisième partie, nous concevons des réseaux de preuves non bien fondés et étudions leur dynamique.

Mots-clés. théorie de la preuve | logique linéaire | preuve cyclique | points fixes | réseaux de preuves | sémantique des phases | élimination des coupures

Acknowledgements

I will begin, as is customary in the genre of thesis acknowledgement writing, with the jury. I am grateful to them for accepting to be on it and gracing the defence with their presence (offline and online). I am especially indebted to Laurent Regnier's personalised and detailed remarks and Bahareh Afshari's highly positive review. They were essential for the manuscript to reach its present form and for fortifying my confidence for the defence respectively. Amina, thanks for joining despite not being in the pink of your health. I have learnt an awful lot from your thesis and keep doing so. David, Delia you made the defence feel so comfy: what I assumed would feel like an exam felt like a rite of passage. Olivier, I am simply floored by your erudition and rigour. They are benchmarks I strive for. Jam, thanks for your warmth, wisdom, and above all, for the push that plunged me deep into logic. Finally, Alexis, no amount of gratitude can be enough for your guidance. Everything I know about linear logic is due to you. As a product of the first generation that could Google answers to homework assignments, you have taught me to suppress my urge to search for an answer in the literature before thinking about a solution. You have a unique way of looking at the most mundane taken-for-granted results of logic. Like an elder brother teaching his sibling to ride a bicycle, you have taken off the trainers at the right moment. I believe we can now ride as fellow cyclists on this cycle lane of cyclic proof theory (a term I know you don't like but I couldn't resist the alliteration).

I would like to thank my other collaborators: Luc, Anupam, and Farzad. I thoroughly enjoyed my short visit to Birmingham and I am grateful to Anupam for the invitation. Thanks, Gianluca, Sonia, and Marianna for the wonderful time. I will cherish the several friends I made and the discussions we had over beers in the several conferences, workshops, and summer schools I was fortunate to attend. From Prague to Swansea to Funchal, it has been an amazing experience. A big thanks to the people at IRIF secretariat (past and present), Ariella Brianni at FSMP, and Anne Mathurin at INRIA Paris for helping to dribble through the notorious French bureaucracy.

Back in CMI, I express my sincere gratitude to all my professors especially Madhavan for supervising what might be the world's first joint Master's thesis. A special thanks to Shiva who not only taught us linear algebra in his inimitable style but also awakened a political consciousness in many of us. A cocktail of Ambedkar, Shostakovich, V.I. Arnold, and Herzog is exactly what we need in those formative years. Going back even further, at Xavier's, I must acknowledge Saswata for introducing me to the RMO. Our collective problem-solving sessions are one of my most cherished memories of high school.

In Calcutta, Soviet books were still circulating in second-hand bookstores, flea markets, and personal collections as late as the early 2000s. I have learnt a lot of science from Mir Publisher books and I believe these books were instrumental in building my scientific temper in the Nehruvian sense of the phrase. As a tribute to their influence, this manuscript is typeset in *Literaturnaya*, the standard typeface for Soviet science books.

The real PhD is probably the friends we made along the way. There are many to name but I will try my best to enumerate. Firstly, I am thankful for the wonderful room-mates in 3026 at IRIF I had over the years, Remi, Farzad, Felix, Aymeric. Juliusz, thanks for being a breath of fresh air during the doldrums of an average day at the lab. Hopefully, one day you'll give a lecture in English or Bangla and we can pick up our debate on the finer points of the decline of the Abbasid Caliphate from where we left it. My mates at IRIF who did not forget me before going for lunch: Ranadeep, Rachid, Soumyajit, Alen, Easie. My fellow failures in French class, Anupa and Suha: in a parallel universe, we can perfectly recite Rimbaud. The larger Indian family in Paris: Samar and Afeef, KC, Chait, Amrita, Nirbhay, Hina, Madhuresh, Abhishek and Sindhura, Sharbat and Manvi, Ritam and Amrita, Dipanjana, ... I am grateful to Chait for introducing me to this family. A special mention of Mihir's

hospitality during the stressful week of my defence. My comrades: Com. (Manuj, Saurabh, Ritam, Adrija), *inquilaab zindabad!* I have learnt a lot from you and keep doing so. The CMI group in Europe: Sougata, Arnab, Ritam, Rajarshi, Debraj, Avinandan, Charles, Devesh, Thejaswini, Suman. The larger CMI support groups of G10 and S17 are too large to be mentioned individually but you know who you are and what you mean to me. Lots of love to Kazi, Debam, and Anirban for helping me keep my sanity during the Covid-19 lockdown. Here's to many more years of ISL matches in smoky Calcutta dive bars.

Finally, heartfelt gratitude to my parents for tolerating me all these years.

Birmingham, 2023

Preface to the thesis:

Linear logic with the least and greatest fixed points: truth semantics, complexity and a parallel syntax

When logic and proportion
Have fallen sloppy dead
And the White Knight is talking backwards
And the Red Queen's "Off with her head!"
Remember what the dormouse said:
Feed your head!
Feed your head!
Feed your head!

Grace Slick, *White Rabbit*

Taming the infinite and induction

Throughout the history of mathematics and logic, a lot of time has been devoted to comprehending the concept of infinity. It is intrinsically counter-intuitive because there is little material need for a concept of infinity in one's daily life (besides possibly theological). One might need the number 56 to convey how many goats they have or the number $\sqrt{2}$ to convey the distance between two corners of the town square or even the number π to convey the length of rope required to wrap around a tree, but not infinity. In fact, whenever we talk about something "infinite" in mathematics: there is a trade-off between rigour and clarity. Take *infinitesimals*, for example. The concept dominated mathematics since its conception by Newton/Leibnitz. Even now, high-school students are taught that in the following, one can safely cancel $(x - 1)$ in the numerator and denominator since we are not cancelling zeroes; rather something which is *infinitesimally* close to zero.

$$\lim_{x \rightarrow 1} \frac{x^2 - 1}{x - 1}$$

However, sceptics as early as Bishop Berkeley have pointed out that this is not rigorous [Ber34]. On the other hand, a rigorous ε/δ definition of the limits of sequences loses out on clarity and is the reason why they are not taught to high-school students. Consequently, a school of thought developed in the early 20th century that advocated *finitism*: representing the infinite by the finite.

But how far can we go armed with only finite tools? We should at least be able to reason over natural numbers, which is an infinite set ¹. We claim that, instinctively, we employ finitist tools to reason about them. For example, let us try to prove the assertion P defined as follows.

$P :=$ Every natural number is either even or odd.

¹The fact that infinite sets are well-defined is not derivable and needs to be stated explicitly as an axiom. The axiom of infinity in Zermelo-Frankel set theory asserts that there is at least one infinite set *viz.* the set of all natural numbers.

Observe that P is an infinite conjunction of smaller propositions P_n asserting that n is either even or odd. So, technically P has an infinite proof which is an infinite conjunction of the finite proofs of each P_n .

However, the response of any high-school student on looking at P will be to prove by mathematical induction. Immediately, we realise that we have to prove only finitely many things *viz.*

- 1 is odd;
- n is odd implies $n + 1$ is even; and
- n is even implies $n + 1$ is odd.

Suppose in the induction case, we have $n + 1$. By the induction hypothesis, n is either even or odd. If it is odd then by item (3), $n + 1$ is even; otherwise, n is even and by (4), $n + 1$ is odd. Combining these two, we have $n + 1$ is even or odd.

Naturally, finitists like Wittgenstein would champion the latter method rather than the former [WRK78]. Although induction is obvious to a working mathematician, it does not come for free in logic. While formalising number theory, it has to be included as a separate axiom (which are dictums like postulates of Euclidean geometry). In fact, Poincaré [Poi05] labelled induction as a genuine *synthetic a priori i.e.* it is not an obvious tautology but still a concept that is universally true.

Self-reference and circular proofs

One can imagine mathematical proofs as a game between two players: *prover* and *denier*. The denier questions something and the prover provides an argument and the game goes on like this. If at some point, the prover is unable to provide an argument, the statement we started with is false and if we end up with statements which have been a priori accepted to be true, the denier has nothing to question; he loses and the statement is deemed true.

Now imagine this game goes on forever. Who wins? Analogously, imagine a program running forever. Is it intrinsically wrong? Suppose it is printing the output of `ROOTTWO(0)` defined as follows.

```
function ROOTTWO( $n$ )
  compute  $v_n$  = the  $n$ th decimal place value of  $\sqrt{2}$  (By the long-division method)
  print  $v_n$ 
  return ROOTTWO( $n + 1$ )
end function
```

Figure 1: A pseudocode for the printing $\sqrt{2}$ in decimal representation

This is indeed a meaningful program despite it running for an infinite time. But obviously, not all programs running forever is meaningful (say it is because the programmer wrote the wrong exit condition). We need to be able to distinguish between these two and what one needs to check is if there has been some progress (*i.e.* one is not going around in circles). Therefore, the prover wins even if she and the denier debate for an infinite time as long as they are making progress in reaching a consensus. For example, the infinite proof of P mentioned above can be formulated as an infinite prover-denier game: at every round, the denier keeps providing a number n and the prover proves P_n . One can ensure that there is progress by asserting that the denier cannot ask about the same number twice.

Recursion is a way to finitely represent infinite behaviour. For example, the program² in Figure 1 is a finite representation of the infinite decimal expansion of $\sqrt{2}$. Circular proofs are a finite presentation of the infinite proofs presented above. They do so by way of self-reference (just like recursive function call themselves). A circular proof of P will go as follows.

We first prove the same three things as before *viz.* 1 is odd, n is odd implies $n + 1$ is even, and n is even implies $n + 1$ is odd. Then, we have the following³:

²Technically it is a *corecursive* program.

³This is a semi-formal presentation of a formal proof. See [BS10] for a formal proof.

$$\frac{\frac{\frac{m \text{ is even or odd}}{m \text{ is odd or } m + 1 \text{ is odd}}{n = 1 \text{ is even or odd}}}{n = m + 1 \text{ is even or odd}}}{n \text{ is even or odd}} \begin{array}{l} (\because m \text{ even} \Rightarrow m + 1 \text{ odd}) \\ (\because m \text{ odd} \Rightarrow m + 1 \text{ even}) \\ (\because 1 \text{ is odd}) \end{array}$$

Clearly, there are a lot of similarities between proofs by induction and circular proofs. A natural question at this point is the following.

Is induction as powerful as circular proofs?

Self-reference is at the heart of several logical fallacies and developments of modern logic. Take the *Liar's paradox*. Let S be the following sentence. Is it true or false?

$$S := S \text{ is false}$$

Now, consider *Russell's paradox*. Let X be a set defined as follows. Does $X \in X$?

$$X := \{Y \mid Y \notin Y\}$$

Note that in both cases, the negation was crucial (replacing false with true in the definition of S and replacing \notin by \in in the definition of X removes the apparent contradiction). In circular proofs, the progress condition guarantees this positivity and ensures consistency.

Circular proofs have deep roots in the history of logic and mathematical reasoning: starting with Euclid's [Euc56] heuristic of infinite descent through the more rigorous studies of Fermat (notably his August 1659 letter to Carcavi [dF94]). A systematic investigation of the connection between circular proofs and reasoning by infinite descent has been carried out by Brotherston and Simpson [Bro06, BS07, BS11].

The semantic notion of truth

Russell's paradox triggered a foundational crisis. Russell and Whitehead spent a lot of time [DPPDD09] trying to resolve it. Their efforts were the earliest examples of types (similar to what computer scientists call datatypes) and started modern set theory.

In the 1900 International Congress of Mathematics, Hilbert announced his list of twenty-three problems. The second problem in the list addressed not only the looming question of the foundational crisis but was also a manifestation of his finitist ideology: he hoped to establish, by purely finite combinatorial methods, that there exist no contradictions in mathematics.

In 1931, Gödel [Göd31] dashed Hilbert's dreams: he proved that a powerful enough mathematical theory could not establish its *own* consistency. Again, the problem was self-reference; we emphasise that *consistency* of a system could be proved just not within itself. Put very informally, to prove the consistency of a certain amount of mathematics, one must always use "more" mathematics. Gentzen's proof technique [Gen36] for consistency was ground-breaking since it only manipulated formal proofs and nothing else. The foundational crisis may not be as grave anymore hence Gentzen's proof could have been just a "dusty trinket displayed in the cabinet of mathematical curiosity"⁴ but the opposite happened: it opened the doors to a whole new world of logic called proof theory.

A key novelty of Gödel's theorem was to encode number theoretic functions as numbers themselves in a sound way. For Gödel, the sole notion of the truth of a formula was provability. In particular, incompleteness was expressed in terms of provability: a formula is neither provable nor unprovable. There is, however, a semantic or external notion of truth: for example, the liar's paradox is a semantic fallacy since syntactically (*i.e.* in the case of natural languages, grammatically) the sentence makes perfect sense.

A pervasive school of thought, championed by Tarski, within logic is to consider an external meaning of truth. It has been very successful over the years and has come up as one of the major branches of

⁴Melliès [Mel09]

logic, model theory⁵. Essentially one interprets a language \mathcal{L} in another metalanguage \mathcal{L}' and obtains a semantic notion of truth. **Tarski's undefinability theorem** [Tar33] shows that the Gödel encoding cannot be done for semantic notions of truth *i.e.* no sufficiently rich language (*viz.* a language capable of expressing negation and self-reference) can represent its own semantics.

To complete the story of self-reference in logic, Turing applied Gödel's technique (*viz.* the idea of encoding functions as natural numbers) to mechanical computing. He showed that the **Halting Problem** is not effectively computable [Tur37] *i.e.* the function H defined as follows does not have any algorithm that produces an output in finite time for all inputs.

$$H(P, i) = \begin{cases} 1 & \text{if the program } P \text{ halts on input } i; \\ 0 & \text{otherwise.} \end{cases}$$

Self-reference is at the core of this result: a program H that can potentially accept itself as input. This, parallelly with Church's similar result [Chu36a, Chu36b] gave birth to the final major branch of logic, computability theory.

In summary, self-reference has been instrumental to the development of all the four major branches of logic: set theory, proof theory, model theory, and computability theory.

Hilbert's 24th problem

Hilbert had a twenty-fourth problem [Thi03] that was not published as a part of his list of twenty-three problems. The problem asks for a criterion of simplicity in mathematical proofs and the development of a proof theory with the power to prove that a given proof is the simplest possible.

In his Ph.D. thesis, Brandes [Bra08] defined simplicity as follows. Suppose we want to prove a theorem φ in an axiom system Σ with n axioms $\varphi_1, \dots, \varphi_n$. Then, a proof of φ is simple with respect to φ_i if it uses φ_i m times and all other proofs of φ uses φ_i more than m times.

This is the first instance of resource-consciousness in logic [Pam04]: caring not just about what axioms are used in a proof but also how many times. Modern proof-theorists stop short of formalising the notion of simplicity: rather they ask a fundamental question that is already implicit in Brandes' formalisation:

When are two proofs the same?

A lot of it has got to do with design choices. Going back as far as Frege, logicians have tried to formulate several methods to write formal proofs. In modern logic, there are several ways of writing formal proofs *viz.* the Hilbert-Frege style, Gentzen-style and so on. Comparing proofs is thus challenging. Sometimes, the choice of the formal system has a deep impact on the proofs; at other times, the differences are more superficial. Analogously, it is somewhat like comparing two programs written in two different programming languages. In fact, in Frege's lifetime, his formal systems were dismissed as a reformulation of Boole's algebraic account of logic "in the Japanese custom of writing vertically" [Mel09].

In the system of formal proofs designed by Gentzen, two proofs [Cur52, Kle51] were deemed to be equal if they were equivalent up to the order in which independent components were presented. Exploiting this notion of equality has been helpful to understand the relationship between Gentzen-style proofs and the λ -calculus, a model of computation [Her95].

Linear logic [Gir87a] is the quintessential resource-conscious logic and in the theory of linear logic, there is an interesting formulation of the aforementioned equivalence of proofs. Linear logic proofs can be presented geometrically using objects called proof-nets. The equality of proof-nets is trivial from their structure. Moreover, they characterise exactly the same equivalence as that in Gentzen-style proofs *i.e.* two proofs are equivalent by the aforementioned notion if and only if their corresponding proof-nets are the same.

⁵See [Gir11, VH67] for a perspective on the historic dispute between proof theory and model theory.

This work

This thesis aims at studying infinite, circular, and inductive proofs in the context of linear logic. We build toward the semantic meaning of the truth of formulas provable by infinite/circular proofs. We show that infinite proofs are strictly stronger than circular proofs. Finally, we develop the theory of proof-nets in the context of infinite linear logic proofs.

Contents

1	Introduction	19
2	Preliminaries	27
3	Background on proof-theory	35
3.1	A logic independent introduction to proof-theory	35
3.2	The syntax of linear logic	39
3.3	Phase semantics	41
3.4	Properties of proofs	44
3.5	Proof-nets	48
	Notes	51
4	Linear logic with fixed points	52
4.1	Linear logic with fixed points	52
4.2	Explicit vs. implicit induction	55
4.3	Expressiveness	60
4.4	Focussing	65
4.5	Cut-elimination	66
	Notes	68
I	Provability of the sequent calculus	69
5	Phase semantics of μMALL systems	71
5.1	μ MALL ^{ind} phase semantics	71
5.2	A semantics inspired system: μ_{α} MALL ⁺	80
5.3	Properties of μ_{ω} MALL ⁺	87
5.4	Towards phase semantics of non-wellfounded systems	90
6	Complexity of μMALL systems	91
6.1	Counter machines and linear logic	91
6.2	The complexity of μ MALL*	95
6.3	The complexity of μ MALL [∞]	103
6.4	The regularisation problem	108
II	Infinets: the parallel syntax	111
7	Proof-nets for systems of μMLL	113
7.1	MLL proof-nets: a closer look	114
7.2	μ MLL* proof-nets	115
7.3	μ MLL ^{ind} and μ MLL [∪] proof-nets	127
7.4	Towards infinets	131

8 The genesis of inġnets	135
8.1 Trips and simple proofs	135
8.2 Simple non-wellfounded proof structures	138
8.3 Correctness criterion	140
8.4 Canonicity	145
8.5 Regular inġnets	146
8.6 Generalising to non-simple proofs	149
9 The dynamics of inġnets	153
9.1 Cut-elimination in inġnets without axioms	153
9.2 μMALL^∞ cut-elimination revisited	161
9.3 Cut-elimination in inġnets with axioms	166
10 Conclusion	171
Bibliography	174

Résumé Détaillé de la thèse:

La logique linéaire avec les plus petits et les plus grands points fixes: la sémantique de vérité, la complexité, et une syntaxe parallèle

Cette thèse présente les éléments suivants.

La logique linéaire à points fixes possède une relation de prouvabilité complexe et l'étude de son contenu calculatoire peut être faite de manière beaucoup plus systématique dans le cadre des réseaux de preuves.

La logique linéaire est une logique sous-structurale (c'est-à-dire que l'utilisation de règles structurales telles que l'affaiblissement et la contraction y est restreinte) inventée par Girard [Gir87a, Gir95] lors de l'étude de la sémantique cohérence du Système F. Le cadre de cette thèse est la logique μMALL , l'extension de la logique linéaire (additive multiplicative) par les opérateurs de plus petit et de plus grand points fixes. La théorie des points fixes est omniprésente en informatique, notamment en logique. La conception de systèmes déductifs pour les logiques à points fixes est une tâche difficile mais gratifiante. Dans la littérature, trois systèmes ont été étudiés pour μMALL : $\mu\text{MALL}^{\text{ind}}$ (basé sur une (co)induction explicite), μMALL^{∞} (basé sur un raisonnement non bien fondé) et $\mu\text{MALL}^{\circlearrowleft}$ (basé sur un raisonnement circulaire). Cependant, une particularité de la théorie de la preuve non bien fondée est que lorsque l'on considère toutes les dérivations possibles, le système résultant est incohérent. En particulier, on peut dériver le séquent vide de la manière suivante:

$$\frac{\frac{\vdots}{\vdash \mu x.x}(\mu) \quad \frac{\vdots}{\vdash \nu x.x}(\nu)}{\vdash \mu x.x}(\mu) \quad \frac{\frac{\vdots}{\vdash \nu x.x}(\nu)}{\vdash \nu x.x}(\nu)}{\vdash}(\text{cut})$$

Par conséquent, un critère de progrès global est imposé pour séparer les preuves logiquement valides de celles qui ne le sont pas. Typiquement, ce critère exige que chaque branche infinie contienne un fil traçant une formule d'une manière ascendante et témoignant d'une infinité de points de progrès d'une propriété coinductive. De plus, dans ce cadre non bien fondé, la terminaison de la procédure d'élimination des coupures est remplacée par la *productivité i.e.* que des préfixes arbitrairement grands du résultat puissent être calculés en un nombre fini d'étapes. La condition de progrès susmentionnée est une condition suffisante, mais non nécessaire, pour la productivité de la procédure de d'élimination des coupures.

Le système bien fondé a été introduit par Baelde et Miller dans [BM07] et étudié plus en détail dans [BM07, Bae08, Bae12]. Santocanale [San02] a fourni la sémantique catégorique du fragment additif de $\mu\text{MALL}^{\circlearrowleft}$ et plus tard, avec Fortier [FS13, For14], il a prouvé un résultat d'élimination des coupures pour le même système. Le système circulaire et non bien fondé pour le langage complet a été introduit dans [BDS16] qui a également prouvé un résultat non trivial d'élimination de coupure pour le système non bien fondé. Assez récemment, la sémantique cohérence a été étudiée pour $\mu\text{MALL}^{\text{ind}}$ [EJ21] (des résultats préliminaires [EJS21] sur $\mu\text{MALL}^{\circlearrowleft}$ ont également été obtenus).

Brotherston et Simpson ont conjecturé que (dans le cadre des définitions inductives de Martin Löf) les preuves circulaires dérivent les mêmes énoncés que les preuves finitaires avec induction explicite.

La conjecture dite *Brotherston-Simpson* est restée ouverte pendant une dizaine d'années jusqu'à ce que Berardi et Tatsuta [BT17a, BT19] y répondent négativement pour le cas général. D'autre part, si la logique contient l'arithmétique, la conjecture est connue pour être vraie, ce qui a été prouvé indépendamment par Simpson [Sim17] et Berardi et Tatsuta [BT17b]. La conjecture dépend fortement de la logique de base puisque la disponibilité de règles structurelles ou de constructions modales induit des différences subtiles.

Pour μMALL , le problème est difficile. L'utilisation très restreinte des règles structurelles dans le cadre linéaire induit une relation de prouvabilité beaucoup plus raffinée. Le seul travail dans cette direction a été [NST18] qui a montré qu'un fragment de μMALL° avec une comptabilisation explicite des dépliages de points fixes est équivalent à $\mu\text{MALL}^{\text{ind}}$. L'étude de la relation de prouvabilité des systèmes μMALL n'est pas seulement un défi mathématique ; elle a des ramifications et conséquences profondes. $\mu\text{MALL}^{\text{ind}}$ a été établi comme une fondation pour le model checking et les situations classiques de "model checking" sont réduites à la preuve des séquents de $\mu\text{MALL}^{\text{ind}}$ [HM19]. D'autre part, comme nous le verrons dans la section suivante, les formules μMALL constituent un système de types naturel pour les programmes (co)récurifs. Un important problème de décision en théorie des types est le problème d'habitation du types qui demande, étant donné un type τ et un environnement de typage Γ , s'il existe un programme M tel que M est de type τ par rapport à Γ . Ceci est exactement équivalent à la prouvabilité de la séquence $\Gamma \vdash \tau$. Par conséquent, la décidabilité et la complexité de la prouvabilité des systèmes μMALL sont des questions importantes du fait de leurs conséquences.

Il existe une différence flagrante entre la sémantique dénotationnelle des systèmes μMALL (telle que la sémantique cohérence) et la sémantique de vérité (telle que l'algèbre de Boole pour LK et l'algèbre de Heyting pour LJ). La sémantique dénotationnelle interprète les formules ainsi que leurs preuves, préservant ainsi leur contenu calculatoire. En revanche, la sémantique de vérité est une interprétation plus grossière qui met sur un pied d'égalité toutes les preuves d'une même formule et n'interprète que les formules. La sémantique des phases est une sémantique de vérité pour la logique linéaire qui permet d'exprimer des invariants forts de la prouvabilité de la logique linéaire et a notamment été utilisée pour prouver des résultats de décidabilité [Laf97, DM04] et des résultats d'admissibilité de coupures [Oka96, Oka99]. On peut imaginer approcher la conjecture de Brotherston-Simpson dans le cas de μMALL sémantiquement *i.e.* trouver des modèles de preuves μMALL° qui ne sont les interprétations d'aucune preuve $\mu\text{MALL}^{\text{ind}}$.

Pour résumer, la conjecture de Brotherston-Simpson est une question profonde qui se trouve au cœur de la théorie des preuves non bien fondées et, dans le cas de μMALL , particulièrement difficile. Afin de la prouver ou de la réfuter, on peut utiliser des techniques de complexité et de sémantique. Ainsi, la complexité de la prouvabilité et la sémantique de vérité de divers systèmes de μMALL est une étape importante vers la résolution de cette conjecture.

Une idée centrale en logique est l'isomorphisme de Curry-Howard qui établit une correspondance à trois niveaux entre la logique et la programmation fonctionnelle :

- les formules \leftrightarrow les types.
- les preuves \leftrightarrow les programmes
- normalisation/des coupures-élimination \leftrightarrow calcul.

Une des utilisations significatives de cette correspondance est que l'on peut extraire un système de preuve isomorphe à un langage de programmation typé et raisonner sur ce système plutôt que directement sur les programmes. Typiquement, on peut exprimer logiquement des conditions telles que la terminaison, l'absence d'impasse, la sécurité, etc. Inversement, on peut partir d'un système de preuve connu et faire de la rétro-ingénierie sur un langage de programmation isomorphe avec de bonnes garanties [Har16].

En programmation fonctionnelle, le calcul sur des structures de données définies inductivement se fait généralement par récursion. Les programmes corécursifs sont une généralisation de cette approche pour tenir compte des structures de données infinies, du calcul paresseux, des prédicats communicants concurrents, du calcul sur des flux de données, etc. L'état de l'art actuel est que plusieurs assistants de preuve tels que **Agda** et **Coq** ont commencé à supporter les programmes coinductifs. Pour les programmes récursifs, la *terminaison* est garantie par le fait que le programme est typable. Pour les programmes corécursifs, la terminaison est remplacée par la *productivité* : bien que la terminaison du calcul ne soit pas garantie, des préfixes arbitrairement grands du résultat peuvent

néanmoins être calculés en un nombre fini d'étapes. Les langages supportant la corécursion utilisent traditionnellement un vérificateur de type strict qui vérifie une condition de garde syntactique. La productivité étant indécidable, les conditions de garde décidables sont toujours une sous-approximation de l'ensemble de tous les programmes productifs. Par conséquent, la garde est généralement une condition suffisante mais non nécessaire pour assurer la productivité. Cependant, les conditions de garde actuelles sont trop restrictives - elles rejettent trop de programmes productifs - et trop rigides - il peut être non trivial de reformuler des programmes productifs non gardés sous forme gardée. Naturellement, la conception de conditions de garde réalisables qui peuvent accepter de plus en plus de programmes productifs est un domaine de recherche important. Grâce à l'isomorphisme de Curry-Howard, on peut imaginer utiliser la théorie de la preuve de μMALL pour y parvenir, la condition de progression (ou de validité) correspondant à la condition de garde. La condition de progrès peut en effet être étendue pour prendre en compte de plus grandes classes de preuves productives (par conséquent, de programmes) en utilisant des fils rebondissants [BDKS22]. Cependant, une condition de progression formulée à l'aide de fils rebondissants n'est pas robuste sous une permutation triviale des règles d'inférence.

Les réseaux de preuves [Gir87a] sont un formalisme de preuve qui quotiente cette équivalence exacte. Un proof-net peut être vu comme un graphe dont les nœuds sont des règles d'inférence, qui ne sont donc pas ordonnées, et par conséquent moins séquentielles que les preuves du calcul de séquent. Comme ils sont canoniques, les réseaux de preuves sont bien adaptés pour représenter le calcul. Par conséquent, nous pensons que les réseaux de preuves sont le cadre approprié pour traiter la condition de progression des fils rebondissants. Comprendre l'impact de ces permutations et la façon de les quotienter correctement est une motivation profonde pour notre étude des réseaux de preuves pour μMALL^∞ : nous voulons profiter de la canonicité des réseaux de preuves pour améliorer la dynamique des dérivations non bien fondée quant à l'élimination des coupures.

Pour résumer, les conditions de garde pour assurer la productivité peuvent être améliorées en prenant la motivation des conditions de progrès relaxées dans la théorie non bien fondée. Cependant, ces conditions de progrès relaxées ne sont pas robustes en cas de permutation des règles d'inférence et bénéficieront donc de l'étude de la dynamique des réseaux de preuves non bien fondés.

Cette thèse est divisée en trois parties. La première partie est une revue du contexte et de la littérature pertinente. Les deuxième et troisième parties contiennent les contributions originales de l'auteur.

- Chapitre 2 : nous exposons les outils techniques et conceptuels qui seront très utiles tout au long de la thèse.
- Chapitre 3 : nous introduisons le sujet de la théorie de la preuve et discutons de divers aspects de la logique linéaire. Dans les Section 3.2 et Section 3.3, nous discutons respectivement le calcul de séquents et la sémantique de vérité de la logique linéaire. Dans la Section 3.4, nous discutons de diverses propriétés des preuves en logique linéaire telles que l'élimination des coupures et la focalisation. Enfin, dans la Section 3.5, nous discutons de la syntaxe parallèle de la logique linéaire dans le fragment multiplicatif sans unité.
- Chapitre 4 : il sert d'introduction formelle à la logique linéaire des points fixes et donne un aperçu de certains résultats récents (et moins récents). Dans la Section 4.1, nous établissons la syntaxe de μMALL et la notion particulière de sous-formules dans ce contexte. Dans la Section 4.2, nous présentons les trois systèmes de preuve pour μMALL et comparons leur expressivité relative dans la Section 4.3. Nous discutons de la propriété de focalisation de ces systèmes dans la Section 4.4. Enfin, dans la Section 4.5, nous discutons brièvement des résultats d'élimination des coupures pour ces systèmes et de leurs conséquences.

Dans la deuxième partie, notre objectif est d'étudier la relation de prouvabilité complexe susmentionnée de divers systèmes de μMALL . En particulier, nous étudions la sémantique de vérité de ces systèmes et la complexité du problème de décision "Cette formule est-elle prouvable?". Cette partie se compose de deux chapitres (essentiellement indépendants).

- Chapitre 5 : dans ce chapitre, nous nous consacrons à la sémantique des phases de μMALL . Dans Section 5.1, nous élaborons une sémantique de phase correcte et complète pour $\mu\text{MALL}^{\text{ind}}$. Comme d'habitude, cela nous donne une admissibilité de la règle de coupure (non effective)

par une technique due à [Oka96, Oka99]. Par conséquent, ceci sert de preuve alternative de l'admissibilité de la coupure $\mu\text{MALL}^{\text{ind}}$. Dans la Section 5.2, nous introduisons une famille de calculs infiniment ramifiés bien fondés pour μMALL qui bénéficient d'une sémantique de phase très naturelle. Ceci sert de pont vers l'exploration de la sémantique de phase des calculs circulaires et non bien fondés que nous discutons dans la Section 5.4.

- Chapitre 6 : Dans ce chapitre, nous explorons les problèmes de décision sur les différents systèmes de μMALL c'est à dire le problème de décider si une formule donnée (ou, de manière équivalente, un sequent) est prouvable. Nous réduisons le problème de l'atteignabilité dans diverses machines à compteur à ces questions. Par conséquent, nous pouvons calculer la complexité précise de ces questions. Les résultats sont techniquement intéressants car ils s'appuient sur des applications non triviales de la focalisation. Ils ont également des implications profondes : ils permettent de séparer les systèmes comme des ensembles de théorèmes. Dans la Section 6.1, nous présentons les machines à compteur pertinentes et explorons leurs liens avec la logique linéaire. Dans la Section 6.2, nous montrons que μMALL^* est indécidable (par conséquent, $\mu\text{MALL}^{\text{ind}}$ et μMALL^∞ le sont aussi) et que le problème de prouvabilité pour le fragment sans connecteur $\&$ de μMALL^* est équivalent le problème de prouvabilité pour MELL . Dans la Section 6.3, nous obtenons des bornes inférieures sur la prouvabilité de μMALL^∞ , ce qui nous aide finalement à montrer que $\mu\text{MALL}^{\text{ind}}$ prouve un ensemble de théorèmes strictement plus grand que μMALL° . Nous montrons cela et construisons l'argument dans la Section 6.4. De nombreuses parties de ce chapitre sont basées sur la publication [DDS22].

Le but de la troisième et dernière partie est de développer un formalisme de proof-net (*i.e.* une syntaxe parallèle) pour le calcul non-wellfounded de μMALL . Alors que les réseaux de preuves ont une théorie satisfaisante pour le fragment multiplicatif sans unité de la logique linéaire, leur extension aux additifs [JDNM11, Gir96], aux unités multiplicatives [HH16] et aux exponentielles [dC18] est plus compliquée. Par conséquent, nous nous concentrerons sur le fragment multiplicatif sans unités *viz.* μMLL^∞ . Nous développons progressivement la théorie des réseaux μMLL^∞ (ou **infinets**). Une composante importante des infinets sont les "axiomes infinis" : de même que les axiomes habituels contiennent l'information de savoir quelles formules aboutissent dans quelle feuille de l'arbre de preuve, les axiomes infinis contiennent l'information de savoir quelles formules aboutissent dans quelle branche infinie de l'arbre de preuve non bien fondé. Nous avons choisi de travailler avec une présentation algébrique due à Curien [Cur06] au lieu de la présentation graphique habituelle des réseaux de preuves. Bien que ce choix de conception puisse sembler insignifiant, il est crucial lorsqu'il s'agit de modéliser des axiomes infinis. Cette partie se compose de trois chapitres basés sur [DS19, DPS21] et développe de nouveaux matériaux qui sont jusqu'ici inédits :

- Chapitre 7 : dans ce premier chapitre, nous rappelons d'abord les réseaux de preuves MLL via une présentation algébrique due à Curien [Cur05], dans la Section 7.1. Dans les réseaux de preuves non bien fondés, il est nécessaire de connecter les nœuds par des chemins infiniment longs. Pour formaliser de tels concepts dans la théorie des graphes infinitaires, une machinerie topologique lourde est nécessaire. Nous sacrifions la clarté visuelle des graphes pour considérer les réseaux de preuves non bien fondés (ou *infinets*) dans la présentation algébrique. Dans la Section 7.2, nous améliorons directement cette présentation pour développer des réseaux de preuves pour le fragment finitaire de μMLL^∞ (*viz.* μMLL^*). Nous revenons brièvement à la présentation graphique des réseaux de preuves dans la Section 7.3 pour discuter des réseaux de preuves pour $\mu\text{MLL}^{\text{ind}}$ et μMLL° . Dans la Section 7.4, nous discutons de manière semi-informelle des différents pièges de l'adaptation des réseaux au cadre non bien fondé et des diverses constructions apparaissant dans le prochain chapitre.
- Chapitre 8 : dans ce chapitre, nous décrivons la première classe véritablement infinie de réseaux de preuves μMLL . Nous considérons un fragment de μMLL^∞ *viz.* celui qui n'a pas *trips*. Dans la Section 8.1, nous formalisons ce fragment de μMLL^∞ . Dans la Section 8.2, nous définissons la notion appropriée de réseaux de preuves pour ce fragment en généralisant les réseaux de preuves du chapitre précédent. Le caractère non bien fondé pose plusieurs problèmes, dont l'un est une condition de correction plus complexe. Cette condition de correction est introduite, et on montre qu'elle est complète par rapport à la séquentialisation dans la Section 8.3. Dans la Section 8.4, nous montrons que les objets que nous définissons sont effectivement canoniques. Enfin, dans les Section 8.5 et Section 8.6, nous restreignons et généralisons respectivement cette classe de réseaux de preuves non bien fondés. La Section 8.5 considère un fragment finiment présenté et prouve certains résultats de décidabilité et des connexions avec les preuves

circulaires. Dans la Section 8.6, nous introduisons des réseaux de preuves généraux non bien fondés qui contiennent potentiellement des objets correspondant à des voyages.

- Chapitre 9 : Dans ce chapitre, nous étudions la dynamique des réseaux infinis. Nous avons étudié la dynamique des μMLL^* réseaux de preuves dans la Section 7.2. Puisque les règles d'inférence sont les mêmes pour μMLL^* et μMLL^∞ , les règles de réduction pour les infinis sont un sur-ensemble des règles de réduction μMLL^* réseaux de preuves. Dans la Section 9.1, nous traitons les réseaux infinis comme un système de réécriture métrique. Nous devinons d'abord la forme normale (*big step*) et montrons ensuite qu'une séquence de réduction infinie de petits pas converge vers le résultat à grand pas à la limite. Pour deviner la limite, on doit sacrifier une certaine structure *viz.* η -expanser tous les axiomes rendant le calcul sans atomes. Dans la Section 9.1, nous traitons les infinis simples en toute généralité. Cependant, notre preuve n'est pas complètement indépendante du calcul des séquents. En effet, pour obtenir les limites des séquences de réductions infinies, nous passons par un résultat d'élimination des coupures en calcul des séquents que nous prouvons dans la Section 9.2.

Enfin, nous concluons en indiquant des directions futures dans le Chapitre 10.

Chapter 1

Introduction

Fixed points and linear logic

The fixed point of a function F is a value x such that $F(x) = x$. Theorems implying that certain kinds of functions have at least one fixed point have far-reaching implications in computer science. We provide two illustrative examples. The existence of a solution of a non-cooperative game involving two or more players (or *Nash equilibrium* [Nas50]) relies on fixed point theorems involving functions over convex compact subsets of \mathbb{R}^n . In programming language theory, the semantics of recursive function relies on fixed point theorems involving functions over lattices [Kil73, Sco70].

In logic, fixed points were first introduced to capture inductive definitions [Acz77] which predates its first application in computer science as an expressive database query language [AU79]. In order to define the language of a fixed point logic, one introduces explicit fixed point construct(s) and takes the closure under these construct(s) thus obtaining a richer language. For example, a popular choice is two operators μ and ν which are duals of each other and depict the least and greatest fixed points respectively. Over the years, fixed point logics have been studied from various motivations:

1. Perhaps, the most well-known is the (multi)modal μ -calculus [Koz88, Wal94, Niw97] (the extension of basic modal logic K with least and greatest fixed point operators). Introduced by Scott and Bakker in an unpublished manuscript, the logic has been historically studied in formal methods and verification community [CGP99].
2. First order logic extended with various fixed point operators has been extensively explored in finite model theory [Lib04]. In particular, they seem to recur in descriptive complexity, a seminal result being that the properties that can be expressed in first-order logic with a least fixed point operator are exactly those which can be checked PTIME.
3. Another relevant case study is that of Kleene Algebra (and its extensions) where fragments of the Lambek calculus are extended by a ‘Kleene star’ modelling iteration. Such theories have received axiomatisations that have been proved complete (over relational and language models) [Koz91, Kro91, Bru20, DP18].
4. Finally, intensional modelling of inductive and coinductive reasoning has been studied using various fixed point logics [Pau97, San02, BM07]. These works provide an alternate paradigm to Martin-Löf’s inductive predicates [ML75, Bro05, BS10] for similar pursuits.

Proof theory is one of the main branches of mathematical logic. Its original purpose was to secure the consistency of mathematics by finitary methods. This was part of Hilbert’s program; the main goal was to show the correctness of mathematics using formal deductibility by means of a consistency proof. Subsequently, it has broadened into the study of formal deduction systems in general from various other motivations. Naturally, proof theory is syntactic in nature, in contrast to model theory, which is semantic in nature. Our deductive system of choice is the sequent calculus.

In order to design sequent calculi for fixed point logics, there are some fundamental design choices to be made. For instance, one can employ inference rules that explicitly express the (co)induction invariant (*cf.* Figure 1.1).

$$\frac{F(S) \vdash S}{\mu F \vdash S} (\mu) \quad ; \quad \frac{S \vdash F(S)}{S \vdash \nu F} (\nu)$$

Figure 1.1: The Park (co)induction scheme [Par69]

The cut-elimination theorem is the backbone of modern proof theory. Its central position is illustrated by the fact that three fundamental properties of formal logic follow quite directly from this single theorem:

- Subformula property: this means that deductions are modular in the sense that every provable formula φ can be established by a proof in which only subformulas of φ appear.
- Consistency of the logic: this means that the logic is meaningful in the sense that it does not prove the falsity or equivalently does not prove a formula and its negation.
- Completeness theorem: a formula can either be proved or refuted.

However, sequent calculi with explicit (co)induction do not have the subformula property in spite of enjoying cut-elimination. In fact, it is generally accepted that we do not have *true* cut elimination for any logic equipped with a theory of inductive definitions [ML71]. This poses a major challenge when it comes to proof search since one has to essentially guess induction invariants. A more robust and natural alternative formalisation of inductive reasoning is *implicit* induction, which avoids the need for explicitly specifying (co)induction invariants. This formalism generally recovers true cut elimination but at the cost of infinitary axiomatisation of the fixed points.

There are two approaches to implicit (co)induction. The first approach is to consider a Tait-style system *i.e.* infinitary wellfounded derivations which use a so-called ω -rule (*cf.* Figure 1.2) with infinitely many premises of finite approximations of a fixed point. Such rules arise in various areas of logic, notably as Carnap's rule [Car37] in arithmetic. A complete Tait-style system has been proposed for fixed point logics *viz.* for the μ -calculus [Koz88] and star-continuous action lattices [Pal07] (where the ω -rule construes the Kleene star as an ω -iteration of finite concatenations).

$$\frac{\vdash P(0) \quad \vdash P(1) \quad \vdash P(2) \quad \dots}{\vdash \forall n P(n)} (\forall) \quad ; \quad \frac{\vdash \top \quad \vdash F(\top) \quad \vdash F^2(\top) \quad \dots}{\vdash \nu F} (\nu)$$

Figure 1.2: Carnap's rule in arithmetic and the ω -rule in fixed point logic.

The second approach is to define a *non-wellfounded* and/or a *circular* proof system with finitely branching inferences [NW96, San02, DP17]. Such systems potentially admit greater proof-theoretic expressivity while, at the same time, reinforcing connections between these logics and automata theory. Moreover, as explained in later sections they are instrumental in checking the correctness of (co)inductive programs. However, when considering all possible non-wellfounded derivations (aka *pre-proofs*), the resulting system is inconsistent. In particular one can derive the empty sequent.

$$\frac{\vdots}{\vdash \mu x.x} (\mu) \quad \frac{\vdots}{\vdash \nu x.x} (\nu)$$

$$\frac{\vdash \mu x.x \quad \vdash \nu x.x}{\vdash} (\text{cut})$$

Therefore, a global progress criterion is imposed to sieve the logically valid proofs from the unsound ones. Typically, it requires that every infinite branch is supported by some thread tracing some formula in a bottom-up manner and witnessing infinitely many progress points of a coinductive property. Furthermore, in this non-wellfounded setting, termination of the cut-elimination procedure shall be replaced by *productivity* *i.e.* that arbitrarily large prefixes of the result can be computed in a finite number of steps. The aforementioned progress condition is a sufficient, but non-necessary, condition for the productivity of cut-elimination. The cut-elimination dynamics for least/greatest fixed point rules is much simpler in the non-wellfounded setting and it restores the subformula property, making non-wellfounded proofs more suitable to automated proof search.

On the other hand, on account of their infinitude, non-wellfounded proofs have two major drawbacks. Firstly, they cannot be communicated or checked in finite time; and, secondly, in order to be used in an automated theorem prover or a proof assistant, we need finitely representable proof objects. Consequently, we consider a fragment of non-wellfounded derivations *viz.* that of derivation trees with finitely many distinct subtrees, known as *circular*, or cyclic, derivations. Therefore, instead of giving an infinite proof, we give a finite description of an infinite proof, formulated in the meta-theory. Note that computationally this is a strict fragment since because there is an uncountable number of infinite proofs and any system of finite representation is countable.

The setting of this thesis is μ MALL, the extension of (multiplicative additive) linear logic by least and greatest fixed point operators. Coming back to μ MALL, three systems have been studied in the literature: μ MALL^{ind} (based on explicit (co)induction), μ MALL[∞] (based on non-wellfounded reasoning) and μ MALL[○] (based on circular reasoning). Girard [Gir87a, Gir95] reverse-engineered linear logic from the coherence space semantics of System F. In that sense, it is a proto-categorical logic. On the structural proof-theory side, it is (i) a substructural logic (*i.e.* the usage of structural rules such as weakening and contraction is restricted) and consequently (ii) a resource conscious logic (*i.e.* one is concerned about how many times an axiom is being used in a proof). In other words, the sequents $\vdash a$ and $\vdash a, a$ are indeed different, giving the logic the ability to count. Interestingly, Lambek [Lam58] was already using linear fragments of logic as early as 1958 in order to parse sentences in natural languages; a formal connection between Lambek calculus and linear logic was observed in [Abr90]. Resource consciousness provides a unique expressiveness to fixed point logics but also gives rise to unique challenges to their study.

Fixed point theory is omnipresent in computer science notably in logic. Designing deductive systems for fixed point logics is a difficult but rewarding task.

Expressivity of the various systems

Brotherston and Simpson conjectured that (in the setting of Martin L of’s inductive definitions) circular proofs derive the same statements as finitary proofs with explicit induction. The so-called *Brotherston-Simpson conjecture* remained open for about a decade until Berardi and Tatsuta [BT17a, BT19] answered it negatively for the general case. On the other hand, if the logic contains arithmetic, the conjecture is known to be true; proved independently by [Sim17] and [BT17b].

Note that the Brotherston-Simpson conjecture is heavily dependent on the base logic since the availability of structural rules or modal constructs induce subtle differences. For instance, the modal μ -calculus coincides on all systems. On the other hand, in Kleene Algebras, which is a substructural logic, the wellfounded, circular, and Tait-style systems are indeed different [Bus06, DP17, Kuz18]. In terms of expressivity, μ MALL can be seen as an amalgamation of the properties of μ -calculus and Kleene Algebras. Like Kleene Algebras, μ MALL is also ‘resource-conscious’ (indeed, Kleene Algebra and extensions are just fragments of a non-commutative μ MALL); and like the μ -calculus, μ MALL also allows for unrestricted interleaving of fixed points.

Consequently, for μ MALL, the problem is rather difficult. The very restricted use of structural rules in the linear setting induces a much more refined provability relation. The only work in this direction has been [NST18] which showed that a fragment of μ MALL[○] with explicit book-keeping of fixed point unfoldings is equivalent to μ MALL^{ind}. Studying the provability relation of μ MALL systems is not just a mathematical challenge; it has deep ramifications. μ MALL^{ind} has been established as a foundation for model checking and classical model checking situations are reduced to proving μ MALL^{ind} sequents [HM19]. On the other hand, as we will see in the next section μ MALL formulas are a natural type system for (co)recursive programs. An important decision problem in type theory is *type inhabitation* which asks, given a type τ and a typing environment Γ , does there exist a program M such that M is of type τ with respect to Γ ? This is exactly equivalent to the provability of the sequent $\Gamma \vdash \tau$. Therefore, the decidability and complexity of the provability of μ MALL systems are important questions.

Investigating the semantics of fixed point logics has been incredibly fruitful: in particular, [AL13], building on [NW96], in the case of the μ -calculus, [DBHS16], building on [DHL06], in the case of temporal logic, and [DDP18], building on [DP17], in the case of Kleene algebra. We note that this is the style of truth semantics which, as opposed to denotational semantics, equates all proofs of

the same formula; hence only interprets formulas. Categorical semantics of the additive fragment of μMALL° have been studied in [San02]. More recently, coherence space semantics have been studied for $\mu\text{MALL}^{\text{ind}}$ [EJ21] (preliminary results [EJS21] on μMALL° have also been obtained). There is a stark difference between these denotational semantics for the μMALL systems and the aforementioned truth semantics of μ -calculus, temporal logic, and Kleene Algebra. Denotational semantics interpret formulas as well as their proofs thereby preserving their computational content. On the other hand, truth semantics is a coarser interpretation that equates all proofs of the same formula and only interprets formulas. Phase semantics is a truth semantic for linear logic that allows for expressing strong invariant of linear logic provability and has been notably used to prove decidability results [Laf97, DM04] and cut admissibility results [Oka96, Oka99]. One can imagine approaching the Brotherston-Simpson conjecture in the case of μMALL semantically *i.e.* coming up with models of μMALL° proofs that are not the interpretations of any $\mu\text{MALL}^{\text{ind}}$ proofs.

Studying the provability relation *viz.* the decidability and truth semantics of fixed point logics is an extremely important topic.

Proof-nets for corecursive programs

A central idea in logic is the Curry-Howard isomorphism that establishes a three level correspondence between logic and functional programming:

- formulas \leftrightarrow types
- proofs \leftrightarrow programs
- normalisation/cut-elimination \leftrightarrow computation.

One of the significant uses of this correspondence is that one can extract an isomorphic proof system from a typed programming language and reason on that system instead of directly on programs. Typically one can logically express conditions such as termination, deadlock-freedom, security, and so on. Conversely, one could start from a known proof system and reverse-engineer an isomorphic programming language with good safety nets [Har16].

In functional programming, computation over inductively defined data structures is usually done by recursion. Corecursive programs are a generalisation of this approach to account for infinite data-structures, lazy computation, concurrent communicating predicates, computation on streams of data, *etc.* Dating back to at least Bird [Bir84] (who credits John Hughes and Philip Wadler), corecursion was developed in the concurrency and functional programming communities throughout the late '80s and '90s [Hag87, All89, MT91]. This sparked interest in the metatheory of corecursion with several foundational works formalising coinduction [BM96, MD97, Geu92, RT92]. In particular, infinite data structures were investigated in proof assistants [LPM94, Coq94, Fro95] and coinductive proofs were mechanised in higher-order logic (HOL) [Pau97]. The current state-of-the-art is that several proof assistants such as **Agda** and **Coq** have started supporting coinductive programs.

For recursive programs, *termination* is guaranteed by the fact that the program is typable. For corecursive programs, termination is replaced by *productivity*: while the computation is not guaranteed to terminate, arbitrarily large prefixes of the result can nonetheless be computed in a finite number of steps. Languages supporting corecursion traditionally employ a strict type checker that checks for a guard condition. Productivity is undecidable hence decidable guard conditions are always an under-approximation of the set of all productive programs. Therefore, guardedness is usually a sufficient but not necessary condition to ensure productivity. However, current guard conditions are too restrictive (*i.e.* they reject too many productive programs) and are too rigid (*i.e.* it can be non-trivial to reshape unguarded productive programs in guarded form). Naturally, designing tractable guard conditions that can accept more and more productive programs is an important area of research.

In Figure 1.3, we consider several **Coq** coinductive definitions of functions from natural numbers to infinite lists of natural numbers (*a.k.a* streams), which have seemingly minuscule syntactic differences, nevertheless have wildly varying behaviour. We recall that for coinductive types in **Coq**, the syntactic form of definitions is similar to inductive types with just the keyword **CoFixpoint** instead of **Fixpoint** to trigger the correct guard condition [Coq].


```

CoInductive Stream := Cons : nat → Stream → Stream.
CoFixpoint f0(n : nat) : Stream := Cons n (f0 (n+1)).
CoFixpoint f1(n : nat) : Stream := let s := f1 (n+1) in
  Cons n (match s with Cons h t ⇒ Cons h t end).
CoFixpoint f2(n : nat) : Stream := let s := f2 (n+1) in
  (match s with Cons h t ⇒ Cons n (Cons h t) end).
CoFixpoint f3(n : nat) : Stream := let s := f3 (n+1) in
  (match s with Cons h t ⇒ Cons h (Cons n t) end).

```

Figure 1.3: Some productive and non-productive definitions

- `f0` is the only valid Coq coinductive definition; `(f0 n)` computes the stream $n :: n+1 :: n+2 :: \dots$.
- `f1` is a productive term, even though it is rejected by Coq type-checker as it fails to pass its guard condition. It computes the same stream as `f0`.
- `f2` is not productive, but one could introduce a commutation rule: $\text{match } e_1 \text{ with } p \Rightarrow \text{Cons } (h, t) \rightsquigarrow \text{Cons } (h, \text{match } e_1 \text{ with } p \Rightarrow t)$ (if pattern p does not occur free in h and symmetrically with t) to make it so; it is then equivalent to `f1`.
- `f3` is not productive: producing the first element of `(f3 n)` requires to already have produced the first element of each stream `(f3 k)` for $k > n$.

Therefore the broad goal is as follows.

- extend the guard condition so that more programs are accepted;
- provide a more canonical representation of programs so that productivity is more robust.

Proof theory of fixed point logics can tell us about the computational behaviour of these programs: following the guiding principles of the Curry-Howard correspondence, (co)inductive types can be encoded as μMALL formulas, and (co)recursive programs as μMALL° circular proofs. In the context of circular sequent proofs, μMALL is the only logic which has enjoyed an intensional investigation (extensional studies of circular proofs are more traditional since their inception [SS02, SD03, Bro05]). Another natural candidate would be intuitionistic natural deduction with fixed points, but the advantage of μMALL is that it is rich enough to encode many types purely logically. For example, inductive types such as natural numbers and lists can be encoded with a least fixed point.

$$\mathbb{N} = \mu x. \mathbf{1} \oplus x \quad ; \quad \mathbb{L}_{\mathbb{N}} = \mu x. \mathbf{1} \oplus (\mathbb{N} \otimes x)$$

Every proof π_n of \mathbb{N} represents a natural number n , where the least fixed point is unfolded $n + 1$ times in π_n . Analogously, a Church numeral $\lambda f. \lambda x. f^n(x)$ represents n , the number of applications of f . On the other hand, coinductive types such as streams of natural numbers can be encoded with a greatest fixed point $\mathbb{S} = \nu y. \mathbb{N} \otimes y$. Indeed, one can encode the coinductive programs in Figure 1.3 as μMALL° proofs of $\mathbb{N} \vdash \mathbb{S}$, as shown in Figure 1.4: Φ_0, Φ_1, Φ_2 , and Φ_3 represent `f0, f1, f2` and `f3` respectively. To compute the value of `fi (n)` one would need to consider the proof obtained by cutting Φ_i with π_n for $i \in \{1, 2, 3, 4\}$.

$$\Pi_i^n = \frac{\frac{\pi_n}{\vdash \mathbb{N}} \quad \frac{\Phi_i}{\mathbb{N} \vdash \mathbb{S}}}{\vdash \mathbb{S}} (\text{cut})$$

This induces an infinite cut-reduction sequence. For the productivity of the cut-elimination procedure, the progress condition on Φ_i plays a crucial role. It acts as a sort of guard condition for productivity. Naturally, Φ_0 is progressing but the rest are not and indeed, by the cut-elimination result of [BDS16], Π_0^n converges to a proof of \mathbb{S} for all n . Naturally, the condition is sufficient but not necessary for productivity. In fact, it can be relaxed to account for more proofs: [BDKS22] defines a bouncing thread progress condition that generalises the usual progress condition but ensures productivity of cut-elimination. In particular, Φ_1 satisfies the bouncing thread progress condition and Π_1 converges to a proof of \mathbb{S} .

$$\begin{aligned}
\Phi_0 &= \frac{\frac{\frac{\pi_{\text{dup}}}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}}{\mathbb{N} \vdash \mathbb{N}} \text{ (id)} \quad \frac{\frac{\frac{\pi_{\text{succ}}}{\mathbb{N} \vdash \mathbb{N}} \quad \mathbb{N} \vdash \mathbb{S}}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N} \vdash \mathbb{S}} \text{ } (\nu_r), (\otimes_r)}{\mathbb{N} \otimes \mathbb{N} \vdash \mathbb{S}} \text{ } (\otimes_\ell)}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)} \\
\Phi_1 &= \frac{\frac{\frac{\pi_{\text{dup}}}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}}{\mathbb{N} \vdash \mathbb{N}} \text{ (id)} \quad \frac{\frac{\frac{\pi_{\text{succ}}}{\mathbb{N} \vdash \mathbb{N}} \quad \mathbb{N} \vdash \mathbb{S}}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ } (\nu_r), (\otimes_r)}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N} \otimes \mathbb{N} \vdash \mathbb{S}} \text{ } (\otimes_\ell)}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)} \\
\Phi_2 &= \frac{\frac{\frac{\pi_{\text{dup}}}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}}{\mathbb{N} \vdash \mathbb{N}} \text{ (id)} \quad \frac{\frac{\frac{\pi_{\text{succ}}}{\mathbb{N} \vdash \mathbb{N}} \quad \mathbb{N} \vdash \mathbb{S}}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ } (\nu_r), (\otimes_r)}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N}, \mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ } (\nu_\ell), (\otimes_\ell)}{\mathbb{N} \otimes \mathbb{N} \vdash \mathbb{S}} \text{ (cut)} \\
\Phi_3 &= \frac{\frac{\frac{\pi_{\text{dup}}}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}}{\mathbb{N} \vdash \mathbb{N}} \text{ (id)} \quad \frac{\frac{\frac{\pi_{\text{succ}}}{\mathbb{N} \vdash \mathbb{N}} \quad \mathbb{N} \vdash \mathbb{S}}{\mathbb{N} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ } (\nu_r), (\otimes_r)}{\mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ (cut)}}{\mathbb{N}, \mathbb{N}, \mathbb{S} \vdash \mathbb{S}} \text{ } (\nu_\ell), (\otimes_\ell)}{\mathbb{N} \otimes \mathbb{N} \vdash \mathbb{S}} \text{ (cut)}
\end{aligned}$$

Figure 1.4: μMALL^∞ encodings of $\mathbf{f0}$, $\mathbf{f1}$, $\mathbf{f2}$ and $\mathbf{f3}$ from Figure 1.3. For ease of readability, some cut formulas are depicted in blue. The subproofs π_{dup} and π_{cons} are simple circular proofs that duplicates a natural number and appends a natural number to the head of a stream respectively.

On the other hand, every derivation that is reached by reduction sequence from Π_2^n will have a cut as its last inference. Hence, cut cannot be eliminated from Π_2^n , it is a non-productive computation. Interestingly, the difference between Φ_1 and Φ_2 is limited to the relative order of the $(\nu_r)(\otimes_r)$ inferences and the $(\nu_\ell)(\otimes_\ell)$ inferences in Figure 1.4 but this subtle difference is profound enough for cut-elimination in Π_1^n to be productive and in Π_2^n to be non-productive.

This phenomenon is related to the fact that the sequent calculus for LL is non-canonical: a LL proof may be reduced to two cut-free proofs π and π' which are different but guaranteed to be equal up to irrelevant permutations of inference rules. Normalisation for LL sits thus in the middle between classical sequent calculus LK — in which a proof (Lafont's critical pair) can be reduced to any two proofs of the same sequent and natural deduction [Pra65, Gir87a] or λ -calculus [CR36] normalisation which are confluent. In other words, the permutations are denotationally trivial *i.e.* $\llbracket \pi_1 \rrbracket = \llbracket \pi_2 \rrbracket$ in any semantics. The non-canonicity of sequent calculus manifests itself more critically in μMALL^∞ : as discussed above productivity of cut-elimination is not preserved by infinite permutative equivalence [BDKS22]. Therefore the desideratum is a proof paradigm \mathcal{P} such that:

- The Curry-Howard correspondence is preserved. In particular, there is a map Rep that takes $\mathbf{f1}, \mathbf{f2}, \mathbf{n}$ to objects in \mathcal{P} .
- $\text{Rep}(\mathbf{f1}) = \text{Rep}(\mathbf{f2}) = K$.
- Cut elimination productive in K cut against $\text{Rep}(\mathbf{n})$ for all n

Semantically, this desideratum is justified: since denotation is preserved by normalisation we have $\llbracket \Phi_1 \rrbracket = \llbracket \Pi_1 \rrbracket$. Therefore, $\llbracket \Pi_1 \rrbracket = \llbracket \Phi_2 \rrbracket$ hence asserting that K cut against $Rep(\mathbf{n})$ is productive is justified.

Proof-nets [Gir87a] were devised to achieve exactly this. A proof-net can be seen as a graph whose nodes are inference rules, which are thus not ordered, and consequently less sequential than sequent calculus proofs. As they are canonical, proof-nets are well-suited to represent computation. Consequently, we believe that proof-nets are the proper framework for dealing with the bouncing thread progress condition. Understanding the impact of those permutations and how to quotient them properly is a deep motivation for our investigation of proof-nets for μMALL^∞ : we aim at benefiting from the canonicity of proof-nets to improve the dynamics of non-wellfounded derivations *wrt.* cut-elimination.

To sum up, guard conditions for ensuring productivity can be improved by taking motivation from relaxed progress conditions in non-wellfounded proof theory. However, these relaxed progress conditions are not robust under permutation of irrelevant inference rules and will therefore benefit from the study of the dynamics of non-wellfounded proof-nets.

Contribution of the thesis

The elevator pitch of this thesis is as follows.

Linear logic with fixed points has an intricate provability relation and the study of its computational content can be done much more systematically in the framework of proof-nets.

Naturally, this thesis is split into two parts. In the first part, we study the aforementioned intricate provability relation of various systems of μMALL . In the second part, we develop the theory of non-wellfounded proof-nets and go on to prove cut-elimination in that framework. Before developing these two parts containing our main contributions, we start with an introductory part which consists of three chapters:

- Chapter 2: we expose the technical and conceptual tools that will be very useful in the thesis.
- Chapter 3: a case is made for proof theory in logic and then we breeze through the most important ideas of linear logic.
- Chapter 4: serves as a formal introduction to linear logic with fixed points and a methodical survey of some recent (and not-so-recent) results.

Part I

Our goal is to study the provability of the various systems of μMALL . In particular, we study the truth semantics of these systems and the complexity of the decision problem, “Is this formula provable?” This part consists of two (essentially independent) chapters.

- Chapter 5: we devise the phase semantics for $\mu\text{MALL}^{\text{ind}}$ and introduce a family of infinitely branching systems that enjoy cut admissibility. We conclude by discussing ideas to approach the phase semantics of μMALL^∞ via these systems. This chapter is based on the publication [DJS22].
- Chapter 6: we show that the provability of μMALL° and μMALL^∞ have different complexities; hence μMALL^∞ is not conservative over μMALL° . We also identify a fragment of $\mu\text{MALL}^{\text{ind}}$ that is provably equivalent to MELL whose decidability is open. This chapter is based on the publication [DDS22].




Figure 1.5: A roadmap through the chapters of this thesis depicting their dependencies. Partial dependency is depicted by dashed edges.

Part II

The goal of this part is to develop a proof-net formalism (*i.e.* a parallel syntax) for the non-wellfounded calculus of μMALL . While proof-nets have a satisfying theory for the unit-free multiplicative fragment of linear logic, their extension to additives [JDNM11, Gir96], multiplicative units [HH16] and exponentials [dC18] are more complicated. Consequently, we will concentrate on the multiplicative fragment without units *viz.* μMLL^∞ . We incrementally develop the theory of μMLL^∞ proof-nets (or **infinets**). An important component of infinets are “infinite axioms”: Just as usual axioms encapsulate the information of which formulas end up in which leaf of the proof tree, infinite axioms encapsulate the information of which formulas end up in which infinite branch of the non-wellfounded proof tree. We choose to work within an algebraic presentation due to Curien [Cur06] instead of the usual graphical presentation of proof-nets. Although this design choice could seem insignificant, it is crucial when it comes to modelling infinite axioms. This part consists of three chapters based on [DS19, DPS21] and develops new material which is hitherto unpublished:

- Chapter 7: we introduce the proof-nets for ν -free μMLL in the aforementioned algebraic presentation. Then, we revert back to the graphical presentation to develop proof-nets for $\mu\text{MLL}^{\text{ind}}$ and μMLL° and discuss ways to extend them to μMLL^∞ .
- Chapter 8: we develop μMLL^∞ proof-nets as an extension of the nets introduced in the previous chapter and study their properties such as correctness and canonicity. We consider a finitely presentable fragment and discuss some extensions.
- Chapter 9: is the culmination of our development of the theory of non-wellfounded proof-nets where we obtain cut-elimination on these objects using techniques from the world of infinitary rewriting theory.

Chapter 2

Preliminaries

In this chapter, we will introduce some standard notions and notations that will be used throughout this thesis.

Formal language theory

In formal language theory, a language is a set of words and a regular language is a language that can be defined by a *regular expression*. Alternatively, a regular language can be defined as a language recognised by a *finite state automaton*. The equivalence of regular expressions and finite state automata is known as **Kleene's theorem**. Fix a set Σ called the alphabet.

Definition 2.0.1. *The set of regular languages over Σ is defined recursively as follows:*

- *The empty language \emptyset is a regular language.*
- *For each $a \in \Sigma$, the singleton language $\{a\}$ is a regular language.*
- *If A and B are regular languages, then $A \cup B$ and $A \cdot B = \{ww' \mid w \in A, w' \in B\}$ are regular languages.*
- *If A is a regular language, then $A^* = \{w_1 \dots w_n \mid \forall i \in [n], w_i \in A\}$ (Kleene star) is also a regular language.*

Definition 2.0.2. *A finite state automaton \mathcal{A} is a 5-tuple $(Q, \Sigma, \Delta, q_0, F)$ where:*

- *Q is a finite set of states.*
- *Σ is a finite alphabet.*
- *$\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation.*
- *$q_0 \in Q$ is the initial state.*
- *$F \subseteq Q$ is the set of final states.*

We do not define the semantics of regular expressions and finite state automata. We denote the set of words accepted by a finite state automaton \mathcal{A} by $\mathcal{L}(\mathcal{A})$. Clearly, $\mathcal{L}(\mathcal{A}) \subseteq \Sigma^*$, is the set of all finite words over Σ . We note that the notion of a finite state automaton is robust insofar as the expressiveness of finite state automata does not change with structural tweaks such as multiple initial states, epsilon transitions, non-determinism and so on. Similarly, regular languages are a robust notion insofar as they are closed under various operations such as intersection, complementation, *prefix*, and *division*. Let A be a regular language and w be a word. Then, $\bar{A} = \{u \mid \exists v. uv \in A\}$ and $w^{-1}A = \{u \mid wu \in A\}$ are called the **prefix-closure** of A and **division** of A by w respectively are also regular.

Regular languages can be extended to infinite words. The infinite counterpart of the Kleene star is the \bullet^ω operation: $A^\omega = \{w_1 w_2 w_3 \dots \mid \forall i. w_i \in A\}$. The set of all infinite words over Σ is therefore Σ^ω .

Definition 2.0.3. *An ω -language L is ω -regular if it has one of the following forms:*

- *A^ω where A is a regular language not containing the empty string.*




Figure 2.1: A comb with black *teeth* and *spine* $x_1x_2x_3\dots$

- $A \cdot B$ where A is a regular language and B is an ω -regular language.
- $A \cup B$ where A and B are ω -regular languages

Automata over finite words can similarly be extended to infinite words. A **Büchi automaton** is defined exactly like a finite state automaton except it has different semantics. In particular, accepting runs are exactly those in which at least one of the infinitely often occurring states is final. **Büchi's theorem** states that an ω -language is recognised by a Büchi automaton iff it is ω -regular. Note that although deterministic and non-deterministic finite state automata are equally expressive that is not the case for Büchi automata: here, the deterministic is strictly less expressive than non-deterministic.

We denote Σ^∞ as the set of all finite and infinite words over Σ i.e. $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$. Finally, a non-standard notation that will be used throughout this thesis: two words w, w' are said to be **disjoint** if neither w is a prefix of w' nor w' is a prefix of w .

(Infinite) Graph theory

A **graph** is a pair $G = (V, E)$ of sets such that $E \subseteq V \times V$. To avoid notational ambiguities, we shall always assume tacitly that $V \cap E = \emptyset$. The elements of V are called the **vertices** or **nodes** of G and the elements of E are its **edges**. If we allow E to be a multiset, then there can be more than one edge between two vertices. Such graphs are called **multigraphs**. We will now introduce some notions of infinite graph theory that will be useful in this thesis.

An infinite graph (V, E) such that $V = \{x_i\}_{i \in \mathbb{N}}$ and $E = \{(x_i, x_{i+1}) \mid i \in \mathbb{N}\}$ is called a **ray**, and a **double ray** is an infinite graph (V, E) such that $V = \{x_i\}_{i \in \mathbb{Z}}$ and $E = \{(x_i, x_{i+1}) \mid i \in \mathbb{Z}\}$. Thus, up to isomorphism, there is only one ray and one double ray. Note that in the context of infinite graphs, finite paths, rays and double rays are all called paths.

Definition 2.0.4. *The subrays of a ray or double ray are said to be its tails.*

Every ray has infinitely many tails, but any two of them differ only by a finite prefix. An interesting infinite graph is a **comb** which has one ray and infinitely many maximal finite paths. Figure 2.1 is a typical example.

An important concept in infinite graph theory is that of an *end*, which has no finite counterpart.

Definition 2.0.5. *Let $G = (V, E)$ be an infinite graph. Two rays are considered equivalent if, for every finite set $S \subseteq V$, both have a tail in the same component of $G - S$. An end of a graph $G = (V, E)$ is an equivalence class of rays in G .*

The ends of a tree are particularly simple: two rays in a tree are equivalent iff they share a tail. Therefore, the infinite complete binary tree has continuum many ends. On the other hand, although the comb in Figure 2.1 has infinitely many distinct rays, it has exactly one end.

Definition 2.0.6. *A graph is said to be locally finite if all its vertices have finite degrees.*

Proposition 2.0.1. *A connected infinite graph contains a ray or is not locally finite.*

Induction and Coinduction

A recursive datatype is the smallest set containing some *founders* and closed under certain operations, called *constructors*. \mathbb{N} is the simplest recursive type: it has one founder *viz.* 0 and one unary constructor *viz.* `succ`. Another example is \mathbb{L}_A , the set of finite list over the type A : it has one founder *viz.* the empty list $[]$ and one binary constructor *viz.* `::` such that if a is of type A and l is a list of type \mathbb{L}_A , then $a :: l$ is also a type of \mathbb{L}_A .

One can do induction over recursive types which is called *structural induction*. The reason one can do induction over recursive types is that the relation which relates a recursively defined object x with those objects of which x was constructed, is wellfounded. Note that structural induction over \mathbb{N} is the usual notion of mathematical induction.

From a programming perspective, recursion is a technique to define a function over recursive datatypes by possibly invoking itself on the components of the constructors used to build data values. Most programming languages support recursion by allowing a function to call itself from within its own code. It has been proved in computability theory that recursion is expressive enough to write all programs that be can be written using constructs such as `while` and `for`.

Coming back to the example of lists, \mathbb{L}_A is the least fixed point of the function f where f is defined as $f(X) = [] + A \times X$. What if we consider the greatest fixed point? In that case, we will obtain the set of finite and infinite lists over A . Suppose we want to define the set of infinite lists over A . It is not difficult to guess that it is the greatest fixed point of the function f where f is defined as $f(X) = A \times X$. This is the informal idea behind *corecursive datatypes*, which can be construed as a dual of recursive types. In other words, a corecursive datatype is the greatest set closed under certain operations, called *destructors*. For example, \mathbb{L}_A^ω , the set of infinite lists over the type A has two destructors `hd` and `tl` such that if l is an infinite list of type \mathbb{L}_A^ω , then `hd`(l) of type of A and `tl`(l) is an infinite list of type \mathbb{L}_A^ω .

The mathematical dual of structural induction is *coinduction*. Instead of giving a formal definition, we will give an example of a proof by coinduction. We come back to our example of infinite lists. Assume that (A, \leq) is a partially ordered set. Define the ordering $<_{lex}$ on objects of \mathbb{L}_A^ω as the maximum relation $R \subseteq \mathbb{L}_A^\omega \times \mathbb{L}_A^\omega$ satisfying the following property: if $\ell R \ell'$, then

1. `hd`(ℓ) \leq `hd`(ℓ'), and
2. if `hd`(ℓ) = `hd`(ℓ'), then `tl`(ℓ) R `tl`(ℓ').

We will show that $<_{lex}$ is transitive by coinduction on $\mathbb{L}_A^\omega \times \mathbb{L}_A^\omega$.

Theorem 2.0.1. *If $\ell <_{lex} \ell'$ and $\ell' <_{lex} \ell''$ then $\ell <_{lex} \ell''$.*

Proof. By property 1 of R ,

$$\mathbf{hd}(\ell) \leq \mathbf{hd}(\ell') \leq \mathbf{hd}(\ell''). \quad (2.1)$$

By the the transitivity of \leq on A , $\mathbf{hd}(\ell) \leq \mathbf{hd}(\ell'')$. Thus, property 1 holds for ℓ and ℓ'' . If $\mathbf{hd}(\ell) = \mathbf{hd}(\ell'')$, then $\mathbf{hd}(\ell) = \mathbf{hd}(\ell') = \mathbf{hd}(\ell'')$ by Equation (2.1) and the antisymmetry of \leq on A . By the assumption and property 2, $\mathbf{tl}(\ell) <_{lex} \mathbf{tl}(\ell')$ and $\mathbf{tl}(\ell') <_{lex} \mathbf{tl}(\ell'')$. By the *coinduction hypothesis*, $\mathbf{tl}(\ell) <_{lex} \mathbf{tl}(\ell'')$. This establishes property 2 for ℓ, ℓ'' . Since $<_{lex}$ is the maximal relation satisfying that properties 1 and 2, we are done. \square

The magical part is obviously the coinduction hypothesis which seems like induction on non-wellfounded objects. Actually, one can formally show that the argument is sound. Intuitively, one can appeal to the coinductive hypothesis as long as one has *productivity i.e.* there has been progress in observing the elements of the infinite list (*guardedness*) and there is no further analysis of the tails (*opacity*). We summarise the duality of recursion and corecursion in the following table.

Recursion	Corecursion
<p>Recursively defined datatypes are finite objects.</p> <p>Recursive definitions (usually) come with founders and constructors.</p> <p>Recursive definitions compute a least fixed point.</p> <p>We reason over recursive datatypes by induction.</p> <p>Recursive programs come with a guarantee of termination.</p>	<p>Corecursively defined datatypes are potentially infinite objects.</p> <p>Corecursive definitions (usually) come with destructors.</p> <p>Corecursive definitions compute a greatest fixed point.</p> <p>We reason over corecursive datatypes by coinduction.</p> <p>Corecursive programs come with a guarantee of productivity.</p>

Ordinals

Invented by Cantor in 1883, ordinals are a generalisation of ordinal numerals (first, second, n^{th} , etc.) aimed to extend enumeration to infinite sets.

Definition 2.0.7. *A set S is an ordinal if every element of S is also a subset of S and it is strictly well-ordered with respect to membership.*

Observe that defined this way, ordinals are the generalisation of Von Neumann's definition of natural numbers as the following set.

$$\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots\}$$

In fact, the above set *is* an ordinal and denoted by ω . Any ordinal different from \emptyset has the minimum element \emptyset (simply called zero). However, ordinals do not necessarily have a maximum. For example, the finite ordinal 42 has maximum 41 whereas ω does not have a maximum (since there is no largest natural number).

Definition 2.0.8. *If an ordinal has a maximum α , then it is called a **successor ordinal**, written $\alpha + 1$. A non-zero ordinal that is not a successor is called a **limit ordinal**.*

Just like natural numbers, every ordinal α has a successor *viz.* $\alpha \cup \{\alpha\}$. However, contrary to natural number, the class of all ordinals is *not* a set. This is known as the *Burali-Forti paradox*. The class of all ordinals is denoted by **Ord**. One can now state the principle of **transfinite induction** *viz.* a property $P(\alpha)$ is true for all ordinals $\alpha \in \mathbf{Ord}$ if $P(\beta)$ is true for all $\beta < \alpha$, then $P(\alpha)$ is also true. This can be proved in ZFC. Usually a proof by transfinite induction is broken down into three cases:

- **(Base case)** Prove that $P(0)$ is true.
- **(Successor case)** Prove that for any successor ordinal $\alpha + 1$, $P(\alpha + 1)$ follows from $P(\alpha)$.
- **(Limit case)** Prove that for any limit ordinal λ , $P(\lambda)$ follows from $P(\beta)$ for all $\beta < \lambda$.

It can be shown by transfinite induction that every well-ordered set is order-isomorphic to exactly one ordinal. Hence, transfinite induction holds in any well-ordered set. Furthermore, by **Zermelo's theorem** every set can be well-ordered (one of the several statements equivalent to choice). Therefore, in principle, one can induct on any set, provided one is privy to the recipe to well-order it.

Infinite trees

In this section, we will discuss infinite trees in the various ways they can be viewed and the insights each of them provide. There are two distinct sources of infinitude for a tree: infinite branching (*i.e.* a node may have infinitely many children) or non-wellfoundedness (*i.e.* there is an infinite path from the root). In this thesis, we will only talk about finitely branching non-wellfounded trees and infinitely branching wellfounded trees¹.

Non-wellfounded trees can be defined as corecursive datatypes. To do this in full generality requires a lot of work; we do this for complete binary trees.

¹Hence the all discussion pertaining to non-wellfounded trees in this subsection implicitly assumes that they are finitely branching.

Definition 2.0.9. \mathbb{T}_A^ω , the set of infinite complete binary trees over the type A , is a corecursive type given by three destructors, hd , lft , and rgt such that if t is an infinite complete binary tree, then $\mathit{hd}(t)$ is type A , and $\mathit{lft}(t)$ and $\mathit{rgt}(t)$ are of type \mathbb{T}_A^ω .

One can define distance over infinite trees. Let t, t' be two infinite trees. Define $d(t, t') = \frac{1}{2^d}$ where d is the depth of the nodes at which they differ which are nearest from their respective roots. The set of infinite trees is a complete metric space with respect to d .

Recall that a tree is essentially a special type of graph. Therefore, a non-wellfounded tree is a special type of infinite graph. Consequently, one can define the subtree of a non-wellfounded tree as is usual in graph theory. By Proposition 2.0.1, every non-wellfounded tree has an infinite branch. This is known as **König's Lemma** (which can be proved within ZF for countably infinite trees).

Definition 2.0.10. An infinite tree is said to be **regular** if it has finitely many distinct subtrees.

A regular tree may be transformed into a finite graph by “merging” all the nodes from which the same subtrees start. These graphs can be **unfolded** into an infinite trees. Unfolding can be defined as a corecursive process that produces an infinite tree. Note that Figure 2.1 is an example of an infinite tree that is *not* regular.

With the huge success of automata theory leading up to **Rabin's basis theorem** [Rab69, Rab72], it is sometimes overlooked that the infinite trees that first appeared in the context of logic were (potentially) infinitely branching and wellfounded. A good graph-theoretic way to think of the difference between the two sorts of infinite trees is that for wellfounded infinitely branching tree, depth-first search is productive while breadth-first search is not whereas for non-wellfounded finitely branching trees, breadth-first is productive while depth-first search is not.

Wellfounded trees are inductively defined and hence one can induct on them. An important notion in such infinite trees is that the *rank* of a tree which essentially measures how “long” that induction is.

Definition 2.0.11. The **rank** of a wellfounded tree t over A , denoted $\mathit{rk}(t)$, is defined inductively as follows.

- $\mathit{rk}(t) = 0$, if t is tree consisting of just the root.
- $\mathit{rk}(t)$ is the successor of the supremum of the ranks of the immediate subtrees of t , otherwise.

Proposition 2.0.2. The supremum of the ranks of wellfounded trees over A is the cardinality of the set of all subsets of A , and this supremum is not achieved.

In particular, the supremum of the ranks of well-founded trees on ω is ω_1 .

Fixed point theorems

In this section, we will recall some background on the fundamental fixed point theorems of lattice theory. Not only will we use them several times in our technical proofs, but also, they provide intuition about the design of proof systems with fixed point rules and their corresponding semantics. We first recall Tarski's theorem, a lattice-theoretic generalisation of Knaster-Tarski's fixed point theorem on sets.

For the rest of this section, let $(S, \leq_S, \wedge, \vee)$ be a complete lattice with the least element \perp and the greatest element \top .

Theorem 2.0.2 (Tarski fixed point theorem). *Let $f : S \rightarrow S$ be a monotonic function. The set of fixed points of f is non-empty and equipped with \leq_S forms a complete lattice.*

Definition 2.0.12. Let $(T, \leq_T, \wedge, \vee)$ be a directed complete partial order. Let $f : T \rightarrow T$ be a monotonic function. f is said to be **Scott-continuous** if for each directed subset T' we have $f(\bigvee_{T_i \in T'} T_i) = \bigvee_{T_i \in T'} f(T_i)$.

Theorem 2.0.3 (Kleene fixed point theorem). *Every Scott-continuous function f has the least fixed point $\bigvee_{n \in \omega} f^n(\perp)$.*

Observe that this is a constructive formulation of a fixed point. Cousot and Cousot [CC79] proved a constructive version of Tarski's theorem essentially showing that the set of fixed points of f is the image of preclosure operations on S which is defined as limits of stationary transfinite iteration sequences.

Definition 2.0.13. Let $f : S \rightarrow S$ be a monotonic function. The **upper iteration sequence** starting from $x \in S$ is the sequence $\{\mathcal{U}_\alpha \mid \alpha \in \text{Ord}\}$ of elements of S defined by transfinite induction as follows:

$$\begin{aligned} \mathcal{U}_0 &= x; \\ \mathcal{U}_{\alpha+1} &= f(\mathcal{U}_\alpha); \\ \mathcal{U}_\lambda &= \bigwedge_{\alpha < \lambda} \mathcal{U}_\alpha. \end{aligned} \quad [\lambda \text{ is a limit ordinal}]$$

Dually the **lower iteration sequence** starting from $x \in S$ is the sequence $\{\mathcal{U}_\alpha \mid \alpha \in \text{Ord}\}$ of elements of S defined by transfinite induction as follows:

$$\begin{aligned} \mathcal{U}_0 &= x; \\ \mathcal{U}_{\alpha+1} &= f(\mathcal{U}_\alpha); \\ \mathcal{U}_\lambda &= \bigvee_{\alpha < \lambda} \mathcal{U}_\alpha. \end{aligned} \quad [\lambda \text{ is a limit ordinal}]$$

Theorem 2.0.4. Let $f : S \rightarrow S$ be a monotonic function. The lower iteration sequence starting from \perp is increasing and there exists an ordinal θ , called the **closure ordinal** of f , such that $\mathcal{U}_\theta = \mathcal{U}_{\theta+1}$. Moreover, \mathcal{U}_θ is the least fixed point of f . Dually the upper iteration sequence starting from \top is decreasing and its limit is the greatest fixed point of f .

Remark 2.0.1. The closure ordinal of a Scott-continuous function is at most ω .

Recursion theory

Computability theory or recursion theory is concerned with the study of (un)computable functions and their degrees of uncomputability. Informally, an algorithm (a function on ω) is a finite set of instructions, which given x , after finite steps of *computation* outputs $y = f(x)$. A function which is defined on all arguments and can be specified by an algorithm is *computable* or *recursive*. In the following, we will give a precise mathematical formulation of computable functions.

Definition 2.0.14. The set of **primitive recursive functions** \mathcal{C} is the smallest set of functions of the form $\mathbb{N}^k \rightarrow \mathbb{N}$ such that

- **(Constant functions)** for all $m, n \in \mathbb{N}$, $C_m^n \in \mathcal{C}$ where $C_m^n(x_1, \dots, x_n) = m$;
- **(Successor)** $\text{succ} \in \mathcal{C}$ where $\text{succ}(x) = x + 1$;
- **(Projection)** for all $i, n \in \mathbb{N}$, $\Pi_i^n \in \mathcal{C}$ where $\Pi_i^n(x_1, \dots, x_n) = x_i$;
- **(Composition)** If $g_1, g_2, \dots, g_m, h \in \mathcal{C}$, then

$$f(x_1, \dots, x_n) = h(g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n))$$

is in \mathcal{C} where g_1, \dots, g_m are n -ary functions and h is a m -ary function.

- **(Primitive recursion)** If $g, h \in \mathcal{C}$ and then $f \in \mathcal{C}$ where

$$\begin{aligned} f(0, x_2, \dots, x_n) &= g(x_2, \dots, x_n) \\ f(x_1 + 1, x_2, \dots, x_n) &= h(x_1, f(x_1, x_2, \dots, x_n), x_2, \dots, x_n) \end{aligned}$$

assuming g and h are functions of arity $n - 1$ and $n + 1$ respectively.

Note that this is a recursively defined set where $\{C_m^n\}_{m,n \in \mathbb{N}}$, succ , $\{\Pi_i^n\}_{i \leq n \in \mathbb{N}}$ are founders, and composition and primitive recursion are constructors. Although primitive recursive functions include all the usual functions of elementary number theory, it fails to capture all computable functions, a notable example being the **Ackermann function**. In order to characterise all computable function one needs to generalise to partial functions over \mathbb{N}^k .

Definition 2.0.15. *The set of general recursive functions is the smallest set containing constant functions, successor, projections and closed under composition, primitive recursion and minimisation operation μ which is defined as follows. Given a $(k+1)$ -ary function f , the k -ary function $\mu(f)$ is defined by*

$$\mu(f)(x_1, \dots, x_k) = z \stackrel{\text{def}}{\iff} \begin{array}{l} f(i, x_1, \dots, x_k) > 0 \quad \text{for } i = 0, \dots, z-1 \quad \text{and} \\ f(z, x_1, \dots, x_k) = 0 \end{array}$$

A general recursive function is **total** if it is defined on all arguments. Intuitively, since minimisation does an unbounded search from 0 there is a possibility that the search never terminates and the value is undefined. A function is said to be *effectively computable* or *recursive* if it is total.

By the **Church-Turing thesis**, general recursive functions are precisely the functions that can be computed by Turing machines and the one that can be encoded in untyped λ -calculus. Totality corresponds to halting in Turing machines and termination of β -reduction in untyped λ -calculus. In 1936, Church and Turing independently demonstrated that the **Entscheidungsproblem** is not effectively decidable. Consequently, there is no algorithmic procedure that can correctly decide whether a given general recursive function is total or not.

A relation $R \subseteq \mathbb{N}^k$, $k \geq 1$, is recursive if its characteristic function χ_R is recursive where $\chi_R(x_1, \dots, x_n) = 1$ if $(x_1, \dots, x_n) \in R$ and 0 otherwise. Note that a set $B \subseteq \mathbb{N}$ corresponds to the case $k = 1$ so we have the definition of a set being recursive. We are now ready to define degrees of uncomputability.

Definition 2.0.16. 1. *A set B is in Σ_0^0 ($= \Pi_0^0$) iff B is recursive.*

2. *A set B is in Σ_n^0 if there is a recursive relation $R \subseteq \mathbb{N}^{n+1}$ such that $x \in B$ iff*

$$\exists y_1. \forall y_2 \dots Q y_n R(x, y_1, \dots, y_n),$$

where Q is \exists if n is odd, and \forall if n is even.

3. *Likewise, B is in Π_n^0 if $x \in B$ iff*

$$\forall y_1. \exists y_2 \dots Q y_n R(x, y_1, \dots, y_n),$$

where Q is \exists or \forall depending on the parity of n .

4. *$B \in \Delta_n^0$ if $B \in \Sigma_n^0 \cap \Pi_n^0$*

5. *B is arithmetical if $B \in \bigcup_{n \in \omega} (\Sigma_n^0 \cup \Pi_n^0)$.*

An important notion in computability theory is that of Turing-reduction, which allows one to define the relative computability of functions. In particular, an important concept is that of *hardness*, a set B is said to be Σ_n^0 -**hard** if for all $B' \in \Sigma_n^0$ there is a Turing reduction from B' to B . Furthermore, B is said to be Σ_n^0 -**complete** if $B \in \Sigma_n^0$ and B is Σ_n^0 -hard. Likewise, one can define Π_n^0 -hardness.

The arithmetical hierarchy assigns classifications to the formulas in the language of first-order arithmetic. The **analytical hierarchy** is an extension of the arithmetical hierarchy that assigns classifications to the formulas in the language of second-order arithmetic where one can quantify over both natural numbers and the set of natural numbers.


(Infinitary) rewriting theory

Rewriting theory is to the λ -calculus what automata theory is to Turing machines. Rewrite systems, since the λ -calculi, have been insightful formal models for computations. An abstract rewrite system \mathcal{A} consists of a set of rules Φ defined on a particular set of objects A , which in most cases consists of a language of terms. The rules of the system determine how an object a can be rewritten into b (denoted $a \rightarrow b$). Rewriting theory comes with its own set of bespoke terminology.

Definition 2.0.17. *Let $\mathcal{A} = (A, \Phi)$ be an abstract rewrite system.*

1. *Every element $a \in A$ is called a **normal form** of A if there is no $b \in A$ such that $a \rightarrow b$.*

2. *\mathcal{A} has the **diamond property** (DP) if $\leftarrow \circ \rightarrow \subseteq \rightarrow \circ \leftarrow$ i.e. for all $a, b, c \in A$ there exists $d \in A$ such that the following holds.*



3. \mathcal{A} is **confluent** (CR) if $*\leftarrow \circ \rightarrow * \subseteq \rightarrow * \circ * \leftarrow$.
4. \mathcal{A} is **terminating** or **(strongly) normalising** (SN) if there is no infinite reduction sequences.
5. \mathcal{A} is **(weakly) normalising** (WN) if every element in A reduces to a normal form.

Rewrite systems can be generalised to account for infinitary reduction sequences. To this end, a theoretical tool is needed to formalise the intuition of the limit of such sequences. One way to do that is to assign metric spaces as a basis for transfinite reductions.

Definition 2.0.18. A metric rewrite system is a tuple $\mathcal{M} = (A, \Phi, d, h)$ such that:

- (A, Φ) is an abstract rewrite system;
- $d : A \times A \rightarrow \mathbb{R}_{\geq 0}$ is a function such that (A, d) is a metric space;
- $h : \Phi \rightarrow \mathbb{R}_{\geq 0}$ is a function such that if $\varphi \in \Phi : a \rightarrow b$, then $d(a, b) \leq h(\varphi)$.

The metric and the height, are needed to define the limit behaviour of transfinite reduction sequences *viz.* *continuity* and *convergence*, and to distinguish weak and strong variants thereof, respectively.

Definition 2.0.19. Let $S = \{a_i \rightarrow_{h_i} a_{i+1}\}_{i < \alpha}$ be a reduction sequence in a metric rewrite system $\mathcal{M} = (A, \Phi, d, h)$. Then,

1. S is called **weakly continuous** if the sequence $\{a_i\}_{i < \alpha}$ is **continuous** in the metric space (A, d) . If, additionally, $\lim_{i \rightarrow \lambda} h_i = 0$ for each limit ordinal $\lambda < \alpha$ then the sequence is called **strongly continuous**.
2. S is called **weakly convergent** if it is weakly continuous and the sequence $\{a_i\}_{i < \alpha}$ converges, say to some element $a \in A$.
3. S is called **strongly convergent** if it is weakly convergent and $\lim_{i \rightarrow \alpha} h_i = 0$ in case α is a limit ordinal.

Finally, \mathcal{M} is said to be WN^∞ if for every $a \in A$, there is a strongly convergent reduction sequence $\{a_i \rightarrow a_{i+1}\}_{i < \alpha}$ such that $a_0 = a$ and its limit is in the normal form.

Theorem 2.0.5. A strongly convergent reduction sequence has countable length.

Lemma 2.0.1 (Compression Lemma). For every reduction sequence $\{a_i \rightarrow a_{i+1}\}_{i < \alpha}$ strongly converging to a , there is a reduction sequence $\{a'_i \rightarrow a'_{i+1}\}_{i < \beta}$ such that:

- $a'_0 = a_0$,
- $\beta \leq \omega$, and
- it strongly converges to a .

Notes

See [Koz97] for a succinct introduction to formal language theory and [GTW03], specifically for automata on infinite words. See [Die08] for a comprehensive introduction to infinite graph theory. The coinductive proof of Theorem 2.0.1 was adapted from [KS17]. For the history of coinduction in computer science, see [San09].

Tarski's theorem first appeared in [Tar55] which was a lattice-theoretic generalisation of Knaster-Tarski's fixpoint theorem on sets [KT27]. Infinite trees naturally arose in programming language theory and were first methodically studied in [Cou82, Col82]. For ordinal arithmetic and recursion theory, one can look at any of the standard texts [For03, Soa14]. For a quick survey of results in infinitary rewriting theory, see [KSSdV05].

Chapter 3

Background on proof-theory

We used to think that if we knew one, we knew two, because one and one are two. We are finding that we must learn a great deal more about ‘and’.

Arthur Eddington

In this chapter, we establish some background relevant to this thesis. In Section 3.1 we discuss the philosophy behind the sequent calculus and introduce some standard terminology in a logic independent way. In Section 3.2 we introduce our base logic, propositional linear logic. In Section 3.2 we recall the truth semantics of linear logic. In Section 3.4 we discuss some proof-theoretic results pertinent to linear logic, some of them being salient features of its resource consciousness like Curry-Howard correspondence with π -calculus and focussing. Finally, we discuss a pearl of the linear logic community, proof-nets, in Section 3.5.

3.1 A logic independent introduction to proof-theory

Proof theory is one of main branches of logic that studies *proofs* of logical formulas as independent mathematical objects. Therefore, the actual “truth” of a logical formula is not *essential* to study its proof.

Formally, proofs are usually presented as inductive datatypes like lists or trees comprised of *axioms* that *represent* truth and *inference rules* that *preserve* truth. The trio of a *language* (the universe of all logical formulas), inference rules, and axioms is called a *proof calculus*.

Definition 3.1.1. A signature \mathcal{L} is a triple $(S, \text{ar}, \mathcal{A})$ of a finite set of symbols S , a function $\text{ar} : S \rightarrow \alpha$ for some $\alpha \in \text{Ord}$ that assigns to every symbol an ordinal less than α called its **arity**, and a (possibly infinite) set of atoms. The **language**¹, denoted $\mathcal{F}_{\mathcal{L}}$, is the set of **formulas** (φ, ψ, \dots) over the signature \mathcal{L} is defined inductively as follows:

- $a \in \mathcal{F}_{\mathcal{L}}$ for all $a \in \mathcal{A}$;
- $s \in \mathcal{F}_{\mathcal{L}}$ for all $s \in S$ such that $\text{ar}(s) = 0$;
- $\varphi = s(\langle \varphi_i \rangle_{i \in \alpha})$ such that $\text{ar}(s) = \alpha$ and $\varphi_i \in \mathcal{F}_{\mathcal{L}}$ for all $i \in \alpha$.

Hilbert’s announcement of twenty-three open problems at the 1900 Paris ICM² was an important mathematical event. The second problem of the list was what later came to be known as **Hilbert’s program**. Stated anachronistically in the terminology above, it was the search for a proof calculus such that:

- The language could express all mathematical statements.
- The set of axioms would be finite and any mathematical truth could be proved.
- The calculus would be consistent *i.e.* falsity could not be proved.

¹This definition does not cover several types of logical languages such as first-order structures. However, it covers all the logical languages that will be used in this thesis such as propositional fixed point logics and propositional second-order logic.

²International Congress of Mathematics

In 1931, Gödel's incompleteness theorem showed that Hilbert's program was unattainable for something as simple as Peano arithmetic. In 1936, Gentzen established the consistency of Peano arithmetic by a purely combinatorial argument on the structure of proofs. This does not contradict Gödel's theorem since Gentzen's proof uses transfinite induction up to Cantor's ordinal ε_0 which is outside the perview of Peano arithmetic. What matters today is not the consistency result in of itself, but rather Gentzen's formal innovation: the **sequent calculus** and the **cut-elimination theorem**. This framework improved in many ways the proof calculi previously developed by Frege, Russell, and Hilbert and essentially gave birth to modern proof theory. Consequently, we divide proof calculi into two groups:

Frege-Hilbert style. The global set of hypotheses is immutable.

Gentzen style. The global set of hypotheses can be modified at every step of the proof.

We view this distinction not as a binary but as a spectrum where on one extreme, the Frege-Hilbert style proof calculi have many axioms and few inference rules, and on the other extreme, the Gentzen style proof calculi have few axioms and many inference rules. In summary, from the ashes of the Hilbert's program, rose the phoenix of proof theory.

In this thesis, we will mainly deal with the **sequent calculi**, a Gentzen style proof calculi. In sequent calculi, the smallest unit of a proof is a syntactic object of the form

$$\varphi_1, \dots, \varphi_m \vdash \psi_1, \dots, \psi_n \quad (3.1)$$

called a **sequent**. The formulas on the left-hand side of the turnstile are called the **antecedent**, and the formulas on the right-hand side are called the **succedent** or **consequent**. The left-hand side or the right-hand side (or neither or both) may be empty. The sequent in Equation (3.1) should be informally understood as the statement that the conjunction of all the formulas φ_1 through φ_m implies the disjunction of all the formulas ψ_1 through ψ_n .

A **sequent proof** is a tree whose nodes are sequents such that the child relation respects the inference rules and the leaves are axioms. A formula φ is said to be a **theorem** if $\vdash \varphi$ has a proof. A proof calculi is said to be decidable if the following problem is decidable.

Given φ , is φ a theorem?

Proof theorists are generally reluctant to justify the meaning of their sequents in some model. As we mentioned before, the "actual truth" of a formula is inconsequential to the study of proof. Yet, as one will see in Section 3.3, having a semantic notion of truth is actually quite useful for example to get a handle on the above mentioned problem.

Gentzen classifies inference rules as follows.

- **Axiom rule.** The axiom rule is a sequent of the following shape somehow expressing the simple tautology that φ implies φ .

$$\frac{}{\Gamma, \varphi \vdash \Delta, \varphi} \text{(ax)}$$

- **Logical rules.** A logical rule is a sequence of sequents of the following form where $\Gamma_i \subseteq \Gamma$, $\Delta_i \subseteq \Delta$, and φ_i, ψ'_i are subformulas of φ (for some suitable notion of subformulas).

$$\frac{\{\Gamma_i, \psi_i \vdash \Delta_i, \psi'_i\}_{i \in I}}{\Gamma \vdash \Delta, \varphi} (r)$$

The sequent(s) in a rule displayed above the line are **premiss(e)s** and the unique sequent below the line is the **conclusion**. The **principal formula** is the distinguished formula in its conclusion. **Auxiliary formulas** are the formula occurrences distinguished in the premiss(e)s. An **active formula** is either principal or auxiliary. Other formula occurrences in logical or fixed point rules are **side formulas**.

- **Structural rules.** The structural rules manipulate the formulas of the sequent, but do not alter them. We will discuss them in detail next.

The cut rule is arguably the most fundamental inference rule. It reflects the most famous deduction principle of logic: **Modus Ponens**.

$$\frac{\Gamma_1 \vdash \Delta_1, \varphi \quad \Gamma_2, \varphi \vdash \Delta_2}{\Gamma_1, \Gamma_2 \vdash \Delta_1, \Delta_2} (\text{cut})$$

The informal intuition is that in order to prove the conclusion, assume a lemma φ on the right-hand premiss and then prove the lemma φ on the left. Cuts have had a central role in sequent calculi since their advent in Gentzen's works. A sequent calculi is said to be *analytic* if it has the property that a sequent has a proof iff it has a proof without cuts. A cut-elimination theorem or *Hauptsatz* basically grafts the proof of a lemma where-ever it is used to produce a large proof without any extra lemmas. Gentzen proved the consistency of Peano arithmetic as a corollary of the analyticity of his sequent calculi for first-order logic.

Finally, we first introduce three pairs of structural rules that respectively express the possibility to forget, repeat, and use in any order the hypotheses during an argument. They are called *weakening*, *contraction*, and *exchange*.

$$\begin{array}{l} \frac{\Gamma \vdash \Delta}{\Gamma, \varphi \vdash \Delta} (w_l) \\ \frac{\Gamma, \varphi, \varphi \vdash \Delta}{\Gamma, \varphi \vdash \Delta} (c_l) \\ \frac{\Gamma_1, \varphi, \psi, \Gamma_2 \vdash \Delta}{\Gamma_1, \psi, \varphi, \Gamma_2 \vdash \Delta} (ex_l) \end{array} \qquad \begin{array}{l} \frac{\Gamma \vdash \Delta}{\Gamma \vdash \varphi, \Delta} (w_r) \\ \frac{\Gamma \vdash \varphi, \varphi, \Delta}{\Gamma \vdash \varphi, \Delta} (c_r) \\ \frac{\Gamma \vdash \Delta_1, \varphi, \psi, \Delta_2}{\Gamma \vdash \Delta_1, \psi, \varphi, \Delta_2} (ex_r) \end{array}$$

Structurally speaking, these rules are about the structure or the shape of the sequents. Note that no formula in the conclusion of a structural rule is principal. This gives us a few design choices for sequents:

1. **Sequents as lists of formulas.** The original definition of Gentzen, this is the most explicit presentation of sequents. However, it is cumbersome to explicitly use the exchange rule all the time.
2. **Sequents as (multi)sets of formulas.** The rules ex_l, ex_r tell us that the antecedents and consequents of the sequents can be treated as multisets. The rules w_l, w_r, c_l, c_r tell us that the multisets can simply be treated as sets (still requiring explicit weakening). Hence sequents can be construed as (multi)sets of formulas if the proof calculi have the requisite structural rules. Another downside of using (multi)sets is that they are not sufficient to exploit intensional behaviour of proofs via the Curry-Howard isomorphism. For instance, the following two proofs correspond respectively to the λ -terms $\lambda xy.x$ and $\lambda xy.y$ but they will be identified in the sequents as multi-sets presentation.

$$\frac{\frac{}{\varphi \vdash \varphi} (ax)}{\varphi, \varphi \vdash \varphi} (w_l) \quad ; \quad \frac{\frac{}{\varphi \vdash \varphi} (ax)}{\varphi, \varphi \vdash \varphi} (w_r) \quad ; \quad \frac{\frac{}{\varphi \vdash \varphi} (ax)}{\varphi, \varphi \vdash \varphi} (ex_l)$$

Part I of this thesis is about the extensional behaviour of proofs. Therefore, we will use this presentation.

3. **Sequents as sets of formula occurrences.** In this presentation we will distinguish between two occurrences of the same formula by giving them distinct names. This will be our choice for Part II which explores the intensional behaviour of proofs. We detail this next.

Definition 3.1.2. An **address** is a word over the alphabet $\{a_i\}_{i \in \lambda}$ where λ is the supremum of the arities of all connectives in the signature. Two addresses are said to be **disjoint** if neither are prefixes of the other³. A **formula occurrence** (or simply, *occurrence*) is given by a formula φ and an address α , and written φ_α .

³Identical to the definition of disjointness of words in Chapter 2

In this thesis we will have a signature with a binary symbol and a unary symbol. We denote the alphabet set of addresses as $\{l, r, i\}$ standing for “left”, “right”, and “inside” respectively. Whenever we will decompose an occurrence in a logical rule, the address of each of its subformulae will be extended by r if it is the right subformula, by l if it is the left subformula and by i if the connective is unary.

Remark 3.1.1. *We use Greek letters φ, ψ, \dots for formulas and Latin letters A, B, \dots for formula occurrences.*

We need to extend operations from formulas to occurrences. Define $l^\perp = r$, $r^\perp = l$, and $i^\perp = i$. Consequently, this extends to the negation of an address α . Finally, we define $(\varphi_\alpha)^\perp = \varphi_{\alpha^\perp}$. Connectives are extended to occurrences as follows:

- For any binary symbol \odot , we set $A \odot B = (\varphi \odot \psi)_\alpha$ where $A = \varphi_{\alpha l}$ and $B = \varphi_{\alpha r}$.
- For any unary symbol η , we set $\eta F = \eta \varphi_\alpha$ if $F = \varphi_{\alpha i}$.

Finally, on the subject of the shape of sequents, if the logic has *De Morgan duality* (like **LK**), we only need to consider formulas in negation normal form and can use the one-sided sequent presentation. Modern research in structural proof theory teems with rival proof calculi: depending on one’s need, one can either dial up the meta syntax in sequent calculi coming up with *hypersequents* and *nested sequents* or one can tone down the meta syntax with some graphical syntax like *proof-nets*.

Structural rules	$\frac{}{\vdash \varphi, \varphi^\perp} \text{(id)}$	$\frac{\vdash \Gamma_1, \varphi \quad \vdash \Gamma_2, \varphi^\perp}{\vdash \Gamma_1, \Gamma_2} \text{(cut)}$
Logical rules (multiplicative)	$\frac{\vdash \Gamma, \varphi_1, \varphi_2}{\vdash \Gamma, \varphi_1 \wp \varphi_2} \text{(\wp)}$	$\frac{\vdash \Gamma_1, \varphi_1 \quad \vdash \Gamma_2, \varphi_2}{\vdash \Gamma_1, \Gamma_2, \varphi_1 \otimes \varphi_2} \text{(\otimes)}$
Logical rules (additive)	$\frac{\vdash \Gamma, \varphi_i}{\vdash \Gamma, \varphi_1 \oplus \varphi_2} \text{(\oplus}_i\text{)}$	$\frac{\vdash \Gamma, \varphi_1 \quad \vdash \Gamma, \varphi_2}{\vdash \Gamma, \varphi_1 \& \varphi_2} \text{(\&)}$
Unit rules (multiplicative)	$\frac{}{\vdash \mathbf{1}} \text{(\mathbf{1})}$	$\frac{\vdash \Gamma}{\vdash \Gamma, \perp} \text{(\perp)}$
Unit rules (additive)	$\frac{}{\vdash \Gamma, \top} \text{(\top)}$	No rule for $\mathbf{0}$

Figure 3.1: Inference rules for MALL. Here $i \in \{1, 2\}$.

3.2 The syntax of linear logic

Substructural logics are logics lacking at least one of the usual structural rules (or disallowing their unrestricted usage). Two of the most significant substructural logics are *relevance logic* and *linear logic*. In linear logic, the use of structural rules like contraction and weakening is carefully controlled (available only to formulas of a certain form). Connectives of propositional logic each have two versions in linear logic: multiplicative and additive. Consequently, the units have multiplicative and additive versions as well.

	conjunction	disjunction	“true”	“false”
multiplicative	\otimes	\wp	$\mathbf{1}$	\perp
additive	$\&$	\oplus	\top	$\mathbf{0}$

The symbols \otimes , \wp , \oplus , and $\&$ are read as “tensor”, “par”, “plus”, and “with”. More formally, let \mathcal{A} be a countable set of propositional constants $\{a, b, \dots\}$.

Definition 3.2.1. MALL formulas are given by the following grammar:

$$\varphi, \psi ::= \mathbf{0} \mid \top \mid \perp \mid \mathbf{1} \mid a \mid a^\perp \mid \varphi \wp \psi \mid \varphi \otimes \psi \mid \varphi \oplus \psi \mid \varphi \& \psi$$

where $a \in \mathcal{A}$.

Negation, $(\bullet)^\perp$, is not part of the syntax and is defined as a meta-operation on formulas. φ^\perp is often read as “ φ perp”.

Definition 3.2.2. Negation of a MALL formula is defined inductively as follows.

$$\begin{array}{ll} \mathbf{0}^\perp = \top & \top^\perp = \mathbf{0} \\ \perp^\perp = \mathbf{1} & \mathbf{1}^\perp = \perp \\ (a)^\perp = a^\perp & a^{\perp\perp} = a \\ (\varphi \wp \psi)^\perp = \psi^\perp \otimes \varphi^\perp & (\varphi \otimes \psi)^\perp = \psi^\perp \wp \varphi^\perp \\ (\varphi \oplus \psi)^\perp = \psi^\perp \& \varphi^\perp & (\varphi \& \psi)^\perp = \psi^\perp \oplus \varphi^\perp \end{array}$$

Linear implication can be defined as a macro as follows $\varphi \multimap \psi := \varphi^\perp \wp \psi$. $\varphi \multimap \psi$ is colloquially read as “ φ loli ψ ”. We also denote linear equivalence $\varphi \multimap \psi$ as $(\varphi \multimap \psi) \otimes (\psi \multimap \varphi)$.

The logical system thus obtained is called multiplicative-additive linear logic and abbreviated as MALL. The inference rules of MALL are depicted in Figure 3.1 (sequents being construed as finite multisets). Observe that in MALL neither do we have contraction and weakening nor are they derivable. Operationally, this means that logical deduction is no longer merely about an ever-expanding collection of persistent “truths”, but also a way of manipulating resources that cannot always be duplicated or thrown away at will. In order to get back a correspondence with λ -calculus, one reintroduces structural rules, but only on modal formulas of the form $?\varphi$ (and consequently on its dual $!\varphi$). These are called exponential formulas. $!\varphi$ is read as “bang φ ” or “of course φ ”. $?\varphi$ is read as “question mark φ ” or “why not φ ”. MALL with exponentials is called full linear logic, or simply, LL.

Definition 3.2.3. LL formulas are given by the following grammar:

$$\varphi, \psi ::= \mathbf{0} \mid \top \mid \perp \mid \mathbf{1} \mid a \mid a^\perp \mid \varphi \wp \psi \mid \varphi \otimes \psi \mid \varphi \oplus \psi \mid \varphi \& \psi \mid ?\varphi \mid !\varphi$$

where $a \in \mathcal{A}$.

Negation of LL formulas can be defined by extending the definition of negation of MALL formulas such that ! and ? are duals of each. They are called exponentials since they transform multiplicatives into additives *i.e.* because the following formulas are provable:

$$!(\varphi \& \psi) \multimap !\varphi \otimes !\psi \quad ?(\varphi \oplus \psi) \multimap ?\varphi \wp ?\psi$$

As to why the binary connectives are called multiplicative and additive, the reason is syntactically opaque. An observation from coherence semantics is that multiplicatives correspond to Cartesian *product* and additives correspond to direct *sum*.

Finally, the inference rules for the exponentials are as follows:

$$\begin{array}{c} \frac{\vdash \Gamma, ?\varphi, ?\varphi}{\vdash \Gamma, ?\varphi} \text{ (c)} \\ \frac{\vdash \Gamma, \varphi}{\vdash \Gamma, ?\varphi} \text{ (d)} \end{array} \quad \begin{array}{c} \frac{\vdash \Gamma}{\vdash \Gamma, ?\varphi} \text{ (w)} \\ \frac{\vdash ?\Gamma, \varphi}{\vdash ?\Gamma, !\varphi} \text{ (p)} \end{array}$$

The rules are called *contraction*, *weakening*, *dereliction*, and *promotion* respectively. Observe that promotion is a non-local rule *i.e.* the rule depends on the shape of the context.

Definition 3.2.4. Let φ be a formula. The set of **subformulas** of φ , denoted $\text{SF}(\varphi)$, is defined as the smallest set such that:

- $\varphi \in \text{SF}(\varphi)$
- $\psi \odot \psi' \in \text{SF}(\varphi) \implies \{\psi, \psi'\} \subseteq \text{SF}(\varphi)$ for $\odot \in \{\otimes, \wp, \&, \oplus\}$.
- $\Box\psi \in \text{SF}(\varphi) \implies \psi \in \text{SF}(\varphi)$ for $\Box \in \{!, ?\}$.

This induces a natural ordering \preceq called the **subformula ordering** on the set of formulas *viz.* $\psi \preceq \varphi$ iff $\psi \in \text{SF}(\varphi)$. Moreover, ψ is called an **immediate subformula** of φ if $\psi \neq \varphi$ and for all $\psi' \preceq \varphi$ such that $\psi \preceq \psi'$ we have $\psi = \psi'$.

Definition 3.2.5. The **syntax tree** of a formula φ , denoted $\mathfrak{T}(\varphi)$, is the tree whose nodes are $\text{SF}(\varphi)$ and $\psi \rightarrow \psi'$ if ψ' is an immediate subformula of ψ .

The syntax tree is a graphical representation of the subformula partial order. Note that a syntax tree $\mathfrak{T}(\varphi)$ of an LL formula φ induces a prefix closed language $\mathcal{L}_\varphi \subset \{l, r, i\}^*$ such that there is a natural bijection between the words in \mathcal{L} and the set of all simple paths starting from the root of the syntax tree.

A logic is said to have the **subformula property** if for all formulas φ , every sequent of every cut-free proof of φ consists of subformulas of φ . At first glance, the subformula property implies the finitude of the proof-search space. However, such finitude is a rather peculiar property in structural proof theory at large and in general, a cut-free LL proof can have sequents of unbounded⁴ size.

Theorem 3.2.1. LL is undecidable.

The fragment of LL with just the multiplicative (respectively, additive) connectives, and multiplicative (respectively, additive) units is called multiplicative linear logic or, MLL (respectively, additive linear logic or, ALL). The fragment with multiplicative connectives, exponential modalities, and multiplicative units is called multiplicative exponential linear logic or, MELL.

Theorem 3.2.2. MALL is PSPACE-complete. MLL is NP-complete.

Open Question

Is MELL decidable?

⁴unbounded in the size of the conclusion

3.3 Phase semantics

Truth semantics interprets logical formulas in a mathematical structure. For instance, LK is interpreted in Boolean algebras, LJ is interpreted in Heyting algebras, the modal logic S4 is interpreted in interior algebras, and so on. Similar to LJ and S4, linear logic is not based on an a priori existence of truth values, but it has a truth value semantics, which is given by *phase spaces*.

In linear logic, denotational semantics is perhaps more popular which interprets not just formulas but also proofs. While truth semantics asks the question “What does it mean for φ to be true?”, denotational semantics asks “What does it mean that φ has a proof π ?”. Phase semantics, at the very least, establish safeguards against dubious category-theoretic isomorphisms like $\mathbf{0} \cong \top$ which to a non-expert may seem to suggest that the logic is inconsistent.

To get a flavour of *phase semantics*, the truth semantics of linear logic, it is informative to characterise the set of lists Γ of formulas that make a formula φ provable.

Definition 3.3.1. For a formula φ , define $\text{Pr}(\varphi) = \{\Gamma \mid \text{MALL} \vdash \Gamma, \varphi\}$ and $\text{Pr}_{cf}(\varphi) = \{\Gamma \mid \text{MALL} \vdash_{cf} \Gamma, \varphi\}$ where $\text{MALL} \vdash \Gamma, \varphi$ means that $\vdash \Gamma, \varphi$ is provable in MALL and $\text{MALL} \vdash_{cf} \Gamma, \varphi$ means that $\vdash_{cf} \Gamma, \varphi$ is cut-free provable in MALL.

Let us examine some properties of $\text{Pr}(\varphi)$. First, notice that the axiom rule ensures that for any φ , $\varphi^\perp \in \text{Pr}(\varphi)$. Invertibility of the (\perp) rule gives us that $\text{Pr}(\perp)$ is the set of all provable sequents. Similar observations on the invertibility of the $(\&)$ rule inform that $\text{Pr}(\varphi \& \psi) = \text{Pr}(\varphi) \cap \text{Pr}(\psi)$. For the (non-invertible) connectives \otimes and \oplus , we only have $\text{Pr}(\varphi \otimes \psi) \supseteq \text{Pr}(\varphi) \cdot \text{Pr}(\psi) = \{\Gamma, \Delta \mid \Gamma \in \text{Pr}(\varphi), \Delta \in \text{Pr}(\psi)\}$ and $\text{Pr}(\varphi \oplus \psi) \supseteq \text{Pr}(\varphi) \cup \text{Pr}(\psi)$. This suggests that the algebraic model for linear logic should simultaneously be a monoid and lattice *i.e.* a *residuated lattice*.

$\text{Pr}(\perp)$ plays a major role in this approach, especially when considering it together with the cut inference. Indeed, for any φ , one has that $\text{Pr}(\varphi^\perp) = \{\Gamma \mid \forall \Delta \in \text{Pr}(\varphi), \Gamma, \Delta \in \text{Pr}(\perp)\}$. This naturally suggests to consider the operation $S^\perp = \{\Gamma \mid \forall \Delta \in S, \Gamma \cdot \Delta \in \text{Pr}(\perp)\}$ which induces a closure operator $(\bullet)^{\perp\perp}$ on the set of multisets of linear formulas. As we will soon see, $\text{Pr}(\varphi)$ is closed under the double negation operation for any φ .

These are the basic design principles of phase semantics: interpreting linear formulas as closed subsets of a monoid for the closure operation induced by the orthogonality relation *w.r.t.* a specific subset \perp of the monoid which is an *abstraction of the set of all provable sequents*.

Definition 3.3.2. A **phase space** is a 4-tuple $\mathcal{M} = (M, 1, \cdot, \perp)$ where $(M, 1, \cdot)$ is a commutative monoid and $\perp \subseteq M$. For $X, Y \subseteq M$, define the following operations.

$$XY := \{xy \mid x \in X, y \in Y\}$$

$$X^\perp := \{y \mid \forall x \in X, xy \in \perp\}$$

A **fact** is defined as $X \subseteq M$ such that $X = X^{\perp\perp}$. Equivalently, $X = Y^\perp$ for some $Y \subseteq M$. Given a phase space \mathcal{M} we define $\mathcal{X}_{\mathcal{M}}$ as the set of facts.

Example 3.3.1. Consider the additive monoid $(\mathbb{Z}, 0, +)$ and let $\perp = \{0\}$. For any set $S \subseteq \mathbb{Z}$, $S^\perp = \{y \mid \forall x \in S, x + y = 0\}$. Therefore, if S is not singleton then $S^\perp = \emptyset$; so, $S^{\perp\perp} = \mathbb{Z}$. On the other hand, $\{x\}^\perp = \{-x\}$. The facts of this phase space are \emptyset , singleton sets, and \mathbb{Z} .

Proposition 3.3.1. The following properties hold:

1. $X \subseteq Y^\perp \iff XY \subseteq \perp$
2. $XX^\perp \subseteq \perp$
3. $X \subseteq Y \implies Y^\perp \subseteq X^\perp$
4. $X \subseteq X^{\perp\perp}$
5. $X^{\perp\perp\perp} = X^\perp$
6. $(X \cup Y)^\perp = X^\perp \cap Y^\perp$

We define the following operations on facts. Let X, Y be facts in the following.

$$X \otimes Y := (XY)^{\perp\perp}$$

$$X \wp Y := (X^\perp Y^\perp)^\perp$$

$$X \& Y := X \cap Y$$

$$X \oplus Y := (X \cup Y)^{\perp\perp}$$

Proposition 3.3.2. *Let X, Y be facts. Then, $1 \in X \wp Y \iff X^\perp \subseteq Y$.*

Fix a phase space \mathcal{M} and let $\mathcal{X}_{\mathcal{M}}$ be its set of facts. Fix $V : \mathcal{A} \rightarrow \mathcal{X}_{\mathcal{M}}$ where \mathcal{A} is the set of atoms. A phase space along with a valuation is called a **phase model**. The semantics $\llbracket \varphi \rrbracket$ of a MALL formula φ is parameterised by a valuation (suppose, V) which we will denote by $\llbracket \varphi \rrbracket^V$. We are now ready to define the semantics which is defined inductively as follows:

$$\begin{array}{lll} \llbracket a \rrbracket^V = V(a) & \llbracket a^\perp \rrbracket^V = \llbracket a \rrbracket^{V^\perp} & [a \in \mathcal{A}] \\ \llbracket \mathbf{1} \rrbracket^V = \{1\}^{\perp\perp} & \llbracket \perp \rrbracket^V = \perp & \\ \llbracket \mathbf{0} \rrbracket^V = \{\emptyset\}^{\perp\perp} & \llbracket \top \rrbracket^V = M & \\ \llbracket A \odot B \rrbracket^V = \llbracket A \rrbracket^V \odot \llbracket B \rrbracket^V & & [\odot \in \{\otimes, \wp, \&, \oplus\}] \end{array}$$

When V is clear from the context, we shall simply drop it, writing $\llbracket A \rrbracket$. Finally, we generalise the definition to define the semantics of a finite multiset $\Gamma = \varphi_1, \dots, \varphi_n$ as $\llbracket \Gamma \rrbracket = \llbracket \varphi_1 \wp \varphi_2 \wp \dots \wp \varphi_n \rrbracket$.

Theorem 3.3.1 (Soundness for MALL). *If the sequent Γ is provable in MALL then for all phase models $(M, 1, \cdot, \perp, V)$, $1 \in \llbracket \Gamma \rrbracket^V$.*

Example 3.3.2. *To illustrate the utility of the phase semantic, we show that in any provable multiplicative formula φ (i.e. a MALL formula with only multiplicative connectives), an atom occurs exactly as many times as its negation. Fix an arbitrary atom a occurring in φ . Let $\llbracket \varphi \rrbracket^V$ be the interpretation of φ in the phase space in Example 3.3.1 w.r.t. the valuation V that maps the atom a to $\{1\}$ and every other atom to $\{0\}$. In this phase space, it is easy to see that $X \otimes Y = X \wp Y = \{x + y \mid x \in X, y \in Y\}$. By Theorem 3.3.1, if φ is provable, $0 \in \llbracket \varphi \rrbracket^V$ hence the number of occurrences of a in φ is equal to the number of occurrences of a^\perp .*

Note that a syntactic proof would require the heavy tool of MALL cut-admissibility.

Definition 3.3.3. *The syntactical model, denoted $(\text{MALL}^\bullet, \emptyset, \cdot, \perp, V)$, is a phase model defined as follows:*

- $(\text{MALL}^\bullet, \emptyset, \cdot)$, called **syntactic monoid**, is the free commutative monoid generated by all formulas. In other words, MALL^\bullet is the set of all sequents construed as finite multisets, the empty multiset \emptyset is the monoid identity, and the multiset union is the monoid operation.
- $\perp = \text{Pr}(\perp)$ i.e. \perp is set of all provable sequents.
- $V(a) = \text{Pr}(a)$ for all atoms $a \in \mathcal{A}$.

For the syntactic model to be well-defined note that one needs to show that $\text{Pr}(a)$ and $\text{Pr}(\perp)$ are facts in the phase space $(\text{MALL}^\bullet, \emptyset, \cdot, \perp)$. We will prove something more general.

Proposition 3.3.3. *For any formula φ , $\text{Pr}(\varphi)$ is a fact in the syntactic model.*

Proof. We will first show the following.

$$\text{Pr}(\varphi^\perp) \subseteq \text{Pr}(\varphi)^\perp \quad (3.2)$$

Let $\Gamma \in \text{Pr}(\varphi^\perp)$. Then, there is a proof π of $\vdash \Gamma, \varphi^\perp$. In order, to show that $\Gamma \in \text{Pr}(\varphi)^\perp$, we need to show that for all $\Delta \in \text{Pr}(\varphi)$, $\Gamma, \Delta \in \perp$ i.e. for all Δ such that there is a proof π' of $\vdash \Delta, \varphi$, we have that Γ, Δ is provable. This can be obtained by a cut:

$$\frac{\begin{array}{c} \pi \\ \nabla \\ \vdash \Gamma, \varphi^\perp \end{array} \quad \begin{array}{c} \pi' \\ \nabla \\ \vdash \Delta, \varphi \end{array}}{\vdash \Gamma, \Delta} \text{ (cut)}$$

Now we will show the following.

$$\text{Pr}(\varphi^\perp)^\perp \subseteq \text{Pr}(\varphi) \quad (3.3)$$

Let $\Gamma \in \text{Pr}(\varphi^\perp)^\perp$. Then, for all Δ such that $\vdash \Delta$, φ^\perp is provable, we have $\vdash \Gamma, \Delta$ is provable. Plugging $\Delta = \varphi$, we have $\vdash \Gamma, \varphi$ is provable *i.e.* $\Gamma \in \text{Pr}(\varphi)$. We combine these two results in the following way.

$$\begin{aligned} \text{Pr}(\varphi^\perp) &\subseteq \text{Pr}(\varphi)^\perp && \text{[Equation (3.2)]} \\ \Rightarrow \text{Pr}(\varphi)^{\perp\perp} &\subseteq \text{Pr}(\varphi^\perp)^\perp && \text{[Proposition 3.3.1]} \\ \Rightarrow \text{Pr}(\varphi)^{\perp\perp} &\subseteq \text{Pr}(\varphi) && \text{[Equation (3.3)]} \end{aligned}$$

□

Lemma 3.3.1 (Adequation Lemma for MALL). *For all formulas φ , $\llbracket \varphi \rrbracket^V \subseteq \text{Pr}(\varphi)$.*

Theorem 3.3.2 (Completeness for MALL). *If for any phase model (\mathcal{M}, V) , $1 \in \llbracket \Gamma \rrbracket^V$ then $\vdash \Gamma$.*

Proof. Suppose for any phase model (\mathcal{M}, V) , $1 \in \llbracket \Gamma \rrbracket^V$. In particular, this holds for the syntactic model. By Lemma 3.3.1, $\llbracket \Gamma \rrbracket^V \subseteq \text{Pr}(\Gamma)$ (construing Γ as a par formula). Therefore, $\emptyset \in \text{Pr}(\Gamma)$. (Recall \emptyset is the unit of syntactic monoid.) Hence, $\vdash \Gamma$. □

Lemma 3.3.1 can be strengthened. Using the exact same proof, one can in fact prove that for all formulas φ , $\llbracket \varphi \rrbracket^V \subseteq \text{Pr}_{cf}(\varphi)$. This gives cut-free completeness:

Theorem 3.3.3 (Cut-free completeness for MALL). *If for any phase model (\mathcal{M}, V) , $1 \in \llbracket \Gamma \rrbracket^V$ then $\vdash_{cf} \Gamma$.*

As a direct corollary of Theorem 3.3.3, we have the cut-admissibility of MALL:

Corollary 3.3.3.1. *MALL admits cuts.*

Finally, we note that this truth semantics can be extended to LL by suitably enriching the phase model. A phase model can be constructed out of any monoid $(M, 1, \cdot)$. To model exponentials we need more structure on it *viz.* we designate a submonoid J satisfying $\forall x \in J, \{x\}^{\perp\perp} = \{xx\}^{\perp\perp}$. One can think of this as a sort of idempotence property. Note that any monoid $(M, 1, \cdot)$ contains at least one such that J *viz.* $J = \{1\}$. Let $I = \perp^\perp \cap J$. Then, we interpret the exponentials as follows.

$$\begin{aligned} \llbracket !\varphi \rrbracket^V &= (I \cap \llbracket \varphi \rrbracket^V)^{\perp\perp} \\ \llbracket ?\varphi \rrbracket^V &= (I \cap \llbracket \varphi \rrbracket^{V\perp})^\perp \end{aligned}$$

$$\begin{array}{c}
\frac{\frac{\pi}{\vdash \Gamma, \varphi} \quad \frac{}{\vdash \varphi^\perp, \varphi} \text{ ax}}{\Gamma, \varphi} \text{ (cut)} \rightarrow_{\text{LL}} \frac{\pi}{\vdash \Gamma, \varphi} \\
\\
\frac{}{\vdash \mathbf{1}} \text{ (1)} \quad \frac{\frac{\pi}{\vdash \Gamma}}{\vdash \perp, \Gamma} (\perp) \text{ (cut)} \rightarrow_{\text{LL}} \frac{\pi}{\vdash \Gamma} \\
\\
\frac{\frac{\pi_1}{\vdash \Gamma_1, \varphi} \quad \frac{\pi_2}{\vdash \Gamma_2, \psi}}{\vdash \Gamma_1, \Gamma_2, \varphi \otimes \psi} (\otimes) \quad \frac{\frac{\pi_3}{\vdash \varphi^\perp, \psi^\perp, \Delta}}{\vdash \varphi^\perp \wp \psi^\perp, \Delta} (\wp) \text{ (cut)} \rightarrow_{\text{LL}} \frac{\frac{\pi_2}{\vdash \Gamma_2, \psi} \quad \frac{\frac{\pi_1}{\vdash \Gamma_1, \varphi} \quad \frac{\pi_3}{\vdash \varphi^\perp, \psi^\perp, \Delta}}{\vdash \psi^\perp, \Gamma_1, \Delta} \text{ (cut)}}{\vdash \Gamma_1, \Gamma_2, \Delta} \text{ (cut)} \\
\\
\frac{\frac{\pi_i}{\vdash \Gamma, \varphi_i} \text{ (}\oplus_i\text{)}}{\vdash \Gamma, \varphi_1 \oplus \varphi_2} \quad \frac{\frac{\pi'_1}{\vdash \varphi_1^\perp, \Delta} \quad \frac{\pi'_2}{\vdash \varphi_2^\perp, \Delta}}{\vdash \varphi_1^\perp \& \varphi_2^\perp, \Delta} (\&) \text{ (cut)} \rightarrow_{\text{LL}} \frac{\frac{\pi_i}{\vdash \Gamma, \varphi_i} \quad \frac{\pi'_i}{\vdash \varphi_i^\perp, \Delta}}{\vdash \Gamma, \Delta} \text{ (cut)} \\
\\
\frac{\frac{\pi_1}{\vdash ?\Gamma, \varphi} \text{ (p)} \quad \frac{\frac{\pi_2}{\vdash \varphi^\perp, \Delta}}{\vdash ?\varphi^\perp, \Delta} \text{ (d)}}{\vdash ?\Gamma, \Delta} \text{ (cut)} \rightarrow_{\text{LL}} \frac{\frac{\pi_1}{\vdash ?\Gamma, \varphi} \quad \frac{\pi_2}{\vdash \varphi^\perp, \Delta}}{\vdash ?\Gamma, \Delta} \text{ (cut)} \\
\\
\frac{\frac{\pi_1}{\vdash ?\Gamma, \varphi} \text{ (p)} \quad \frac{\frac{\pi_2}{\vdash ?\varphi^\perp, ?\varphi^\perp, \Delta}}{\vdash ?\varphi^\perp, \Delta} \text{ (c)}}{\vdash ?\Gamma, \Delta} \text{ (cut)} \rightarrow_{\text{LL}} \frac{\frac{\pi_1}{\vdash ?\Gamma, \varphi} \text{ (p)} \quad \frac{\frac{\pi_2}{\vdash ?\Gamma, ?\varphi^\perp, \Delta}}{\vdash ?\varphi^\perp, ?\Gamma, \Delta} \text{ (c)}}{\vdash ?\Gamma, \Delta} \text{ (c)} \\
\\
\frac{\frac{\pi_1}{\vdash ?\Gamma, \varphi} \text{ (p)} \quad \frac{\frac{\pi_2}{\vdash \Delta}}{\vdash ?\varphi^\perp, \Delta} \text{ (w)}}{\vdash ?\Gamma, \Delta} \text{ (cut)} \rightarrow_{\text{LL}} \frac{\frac{\pi_2}{\vdash \Delta}}{\vdash ?\Gamma, \Delta} \text{ (w)}
\end{array}$$

Figure 3.2: Key cases of cut-elimination in LL

3.4 Properties of proofs

3.4.1 Cut elimination

The cut-admissibility result of Theorem 3.3.3 shows that if Γ is provable then *there exists* a cut-free proof of Γ . This existential is strongly non-constructive *i.e.* it does not shed any light on the cut-elimination *procedure*. In particular, it is not discernible whether the cut-elimination equivalence equates all proofs of a particular sequent. This is not desirable from a Curry-Howard perspective since cut-elimination corresponds to computation (in an informal sense for now) and different cut-free proofs of the same sequent potentially correspond to different programs.

Definition 3.4.1. *The cut elimination relation \rightarrow_{LL} is the binary relation over proofs generated by the key rules and commutation rules in Figure 3.2 and Figure 3.3 respectively. We denote the reflexive transitive closure of \rightarrow_{LL} by $\rightarrow_{\text{LL}}^*$.*

Note that if $\pi \rightarrow_{\text{LL}} \pi'$ then π, π' have the same conclusion. Moreover, if π is a proof that contains an instance of the cut rule, then there exists π' such that $\pi \rightarrow_{\text{LL}} \pi'$. Construing the \rightarrow_{LL} relation as the reduction relation of a rewriting system over proofs, the set of normal forms is exactly the set of cut-free proofs. Therefore, all it remains to show is that \rightarrow_{LL} is normalising. We can in fact prove some stronger *viz.* \rightarrow_{LL} is in fact terminating (modulo certain conditions on commutation) by showing a bespoke termination measure to be wellfounded. Consequently, we have the following.

$$\begin{array}{c}
\frac{\frac{\pi_1}{\vdash \varphi, \Gamma} \quad \frac{\pi_2}{\vdash \psi, \Gamma', \xi}}{\vdash \varphi \otimes \psi, \Gamma, \Gamma', \xi} (\otimes) \quad \frac{\pi_3}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \varphi, \Gamma} \quad \frac{\pi_2}{\vdash \psi, \Gamma', \xi} \quad \frac{\pi_3}{\xi^\perp, \Delta}}{\vdash \varphi \otimes \psi, \Gamma, \Gamma', \Delta} (\text{cut}) \\
\frac{\vdash \varphi \otimes \psi, \Gamma, \Gamma', \Delta}{} \\
\frac{\frac{\pi_1}{\vdash \varphi_1, \varphi_2, \Gamma, \xi} \quad \frac{\pi_2}{\vdash \varphi_1 \wp \varphi_2, \Gamma, \xi}}{\vdash \varphi_1 \wp \varphi_2, \Gamma, \Delta} (\wp) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \varphi_1, \varphi_2, \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash \varphi_1, \varphi_2, \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \varphi_1, \varphi_2, \Gamma, \Delta}{} \\
\frac{\frac{\pi_1}{\vdash \varphi, \Gamma, \xi} \quad \frac{\pi_2}{\vdash \psi, \Gamma, \xi}}{\vdash \varphi \& \psi, \Gamma, \xi} (\&) \quad \frac{\pi_3}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \varphi, \Gamma, \xi} \quad \frac{\pi_3}{\xi^\perp, \Delta}}{\vdash \varphi, \Gamma, \Delta} (\text{cut}) \quad \frac{\frac{\pi_2}{\vdash \psi, \Gamma, \xi} \quad \frac{\pi_3}{\xi^\perp, \Delta}}{\vdash \psi, \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \varphi \& \psi, \Gamma, \Delta}{} \\
\frac{\frac{\pi_i}{\vdash \varphi_i, \Gamma, \xi}}{\vdash \varphi_1 \oplus \varphi_2, \Gamma, \xi} (\oplus_i) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_i}{\vdash \varphi_i, \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash \varphi_i, \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \varphi_i, \Gamma, \Delta}{} \\
\frac{\vdash \top, \Gamma, \xi}{} (\top) \quad \frac{\pi}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{}{\vdash \top, \Gamma, \Delta} (\top) \\
\frac{\frac{\pi_1}{\vdash \Gamma, \xi}}{\vdash \perp, \Gamma, \xi} (\perp) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \Gamma, \Delta}{} \\
\frac{\frac{\pi_1}{\vdash \varphi, \Gamma, \xi}}{\vdash ?\varphi, \Gamma, \xi} (d) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \varphi, \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash \varphi, \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \varphi, \Gamma, \Delta}{} \\
\frac{\frac{\pi_1}{\vdash ?\varphi, ?\varphi, \Gamma, \xi}}{\vdash ?\varphi, \Gamma, \xi} (c) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash ?\varphi, ?\varphi, \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash ?\varphi, ?\varphi, \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash ?\varphi, ?\varphi, \Gamma, \Delta}{} \\
\frac{\frac{\pi_1}{\vdash \Gamma, \xi}}{\vdash ?\varphi, \Gamma, \xi} (w) \quad \frac{\pi_2}{\xi^\perp, \Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \Gamma, \xi} \quad \frac{\pi_2}{\xi^\perp, \Delta}}{\vdash \Gamma, \Delta} (\text{cut}) \\
\frac{\vdash \Gamma, \Delta}{} \\
\frac{\frac{\pi_1}{\vdash \varphi, ?\Gamma, ?\xi}}{\vdash !\varphi, ?\Gamma, ?\xi} (p) \quad \frac{\pi_2}{!\xi^\perp, ?\Delta} \quad \rightarrow_{\text{LL}} \quad \frac{\frac{\pi_1}{\vdash \varphi, ?\Gamma, ?\xi} \quad \frac{\pi_2}{!\xi^\perp, ?\Delta}}{\vdash \varphi, ?\Gamma, ?\Delta} (\text{cut}) \\
\frac{\vdash \varphi, ?\Gamma, ?\Delta}{} \\
\frac{\vdash !\varphi, ?\Gamma, ?\Delta}{} (p)
\end{array}$$

Figure 3.3: Commutation cases of cut-elimination in LL

Theorem 3.4.1. *Let π be an LL proof. Then there exists π' such that $\pi \rightarrow_{\text{LL}}^* \pi'$ such that π' is cut-free.*

Corollary 3.4.1.1. *LL has the subformula property.*

Corollary 3.4.1.2. *The empty sequent \vdash is not provable. Subsequently, it is impossible to prove both a formula φ and its negation φ^\perp ; it is impossible to prove $\mathbf{0}$ or \perp .*

However, note that \rightarrow_{LL} is not confluent. This is not a unique symptom of linear logic. In fact, in LK, proof may be reduced to two completely different proofs along two different reduction sequences. An example of such a situation is given by the following derivation, called **Lafont's critical pair**. This proof will reduce either to π_1 or π_2 depending on the direction we choose to reduce cuts.

$$\begin{array}{ccc}
 \begin{array}{c} \pi_1 \\ \nabla \\ \hline \vdash \Gamma \\ \hline \vdash \Gamma, \Gamma \text{ (w)} \\ \hline \vdash \Gamma \text{ (c)} \end{array} & \begin{array}{c} \pi_1 \qquad \pi_2 \\ \nabla \qquad \nabla \\ \hline \vdash \Gamma \qquad \vdash \Gamma \\ \hline \vdash \Gamma, \varphi \text{ (w)} \quad \vdash \Gamma, \varphi^\perp \text{ (w)} \\ \hline \vdash \Gamma, \Gamma \text{ (cut)} \\ \hline \vdash \Gamma \text{ (c)} \end{array} & \begin{array}{c} \pi_2 \\ \nabla \\ \hline \vdash \Gamma \\ \hline \vdash \Gamma, \Gamma \text{ (w)} \\ \hline \vdash \Gamma \text{ (c)} \end{array} \\
 \xrightarrow{\text{LL}^*} & & \xrightarrow{\text{LK}^*}
 \end{array}$$

Besides, in LL, the non-confluence is less critical since one can show that if a proof reduces to two different proofs then they are equivalent up to trivial commutation of inference rules (consequently, being denotationally equivalent). Confluence is recovered in the proof formalism called proof-nets.

3.4.2 Axiom expansion

In type theory, η -expansion refers to rule $M \rightsquigarrow \lambda x.Mx$ (where $x \notin \text{fv}(M)$). This rule asserts that every term is a function and in fact, two functions are equivalent iff they evaluate to the same term on all possible arguments. Via the Curry-Howard correspondence, just like there is a counterpart to β -reduction in logic (*viz.* cut-elimination), there is a counterpart to η -expansion in logic (*viz.* axiom expansion). Note that in this thesis which only talks about logic, η -expansion and axiom expansion will be used interchangeably.

Proposition 3.4.1 (η -expansion). *For every proof π of $\vdash \Gamma$, there is a proof π' of $\vdash \Gamma$ in which the axiom rule is only used with atomic formulas. Moreover, if π is cut-free, then so is π' .*

Proof. It suffices to prove that for every formula φ , the sequent $\vdash \varphi, \varphi^\perp$ has a cut-free proof in which the axiom rule is used only on atomic formulas. We prove this by induction on φ .

If φ is atomic, then $\vdash \varphi, \varphi^\perp$ is an instance of the atomic axiom rule. If $\varphi = \psi_1 \otimes \psi_2$ then we have

$$\begin{array}{c}
 \begin{array}{c} \pi_2 \qquad \pi_2 \\ \nabla \qquad \nabla \\ \hline \vdash \psi_1, \psi_1^\perp \quad \vdash \psi_2, \psi_2^\perp \\ \hline \vdash \psi_1 \otimes \psi_2, \psi_1^\perp, \psi_2^\perp \text{ (}\otimes\text{)} \\ \hline \vdash \psi_1 \otimes \psi_2, \psi_1^\perp \wp \psi_2^\perp \text{ (}\wp\text{)} \end{array}
 \end{array}$$

where π_1 and π_2 are cut-free proofs in which the axiom rule is used only on atomic formulas by induction hypothesis. Other connectives follow similarly. \square

Proposition 3.4.1 allows us to assume wlog that every subformula is principle in a logical rule (except in the case where there is an occurrence of the (w) or (\top)-rule). Such a 'locality' feature is more intrinsic to proof formalisms like deep inference where even structural rules such as contraction and cuts can be turned into their atomic versions.

3.4.3 Invertibility of inference rules and focussing

In structural proof theory, focussed proofs are a family of proofs that are more structural than usual sequent calculus proofs. They arise through goal-directed proof search where the search space is vastly reduced for focussed proofs. A sequent calculus is said to have the **focussing property** when focussed proofs are complete with respect to provability.

The starting point of focusing is the classification of the inference rules of linear logic into two categories: *invertible* and *non-invertible*. The conclusion of an invertible inference rule is provable iff its premises are provable. An inference rule is non-invertible if it is not invertible. For example, the (\wp) -rule is invertible since we can derive its premisses from its conclusion using the following derivation:

$$\frac{\frac{\frac{}{\vdash \varphi^\perp, \varphi}(\text{id}) \quad \frac{}{\vdash \psi^\perp, \psi}(\text{id})}{\vdash \varphi^\perp \otimes \psi^\perp, \varphi, \psi}(\otimes)}{\vdash \Gamma, \varphi \wp \psi}(\text{cut})}{\vdash \Gamma, \varphi, \psi}$$

Proposition 3.4.2. *The $\wp, \&, \text{ and } \perp$ rules of MALL are invertible. The \otimes and \oplus rules of MALL are non-invertible.*

Therefore, one can apply invertible rules without losing provability. Note that one can also apply trivially apply the (\top) -rule without losing provability. The invertible rules along with (\top) -rule are called the *negative* rules.

Let us now consider an invertible rule. Imagine we have the sequent $\vdash a \wp b, a^\perp \otimes b^\perp$. If we apply the tensor rule immediately, we lose provability. Moreover, after the application of the par rule, the tensor rule has to be of a certain shape: if the premisses are $\vdash b, a^\perp$ and $\vdash a, b^\perp$ then again we lose provability. So, one cannot apply the tensor rule context-freely. Similarly, observe that one also cannot apply the $(\mathbf{1})$ -rule whenever one wants and needs to make sure if there are no side formulas. The non-invertible rules along with the $(\mathbf{1})$ -rule and $(\mathbf{0})$ (vacuously since there is no rule for $\mathbf{0}$) are called the *positive* rules. Hence, the intuition is that applying negative rules preserves provability whereas applying positive rules may potentially lead to a loss of provability. By assigning arbitrary polarities to atomic variables one can extend the notion of polarities to formulas.⁵

The crux of focusing is the following proof search strategy, called the *focusing discipline*:

- **Negative phase:** If the sequent contains a negative formula N , then decompose N .
- **Positive phase:** If the sequent contains only positive formulas, then some formula can be chosen as a *focus*. Recursively select its positive subformulas as principal formulas until a negative subformula is reached.

Theorem 3.4.2 (Focussing Theorem). *MALL has the focusing property i.e. $\vdash \Gamma$ has a proof iff $\vdash \Gamma$ has a focussed proof.*

This theorem ensures that the focussing discipline is a complete proof-search strategy. Note that in this subsection we implicitly discuss only cut-free proofs since it is motivated by proof search and we know that if there is a proof of $\vdash \Gamma$ there is a cut-free proof of Γ .

⁵Historically positive and negative rules were called *synchronous* and *asynchronous* rules respectively.

3.5 Proof-nets

In Section 3.4.1, we noted that although an LL proof may be reduced to two different cut-free proofs we are guaranteed that they are equal up to irrelevant permutations of inference rules. In particular, we would like to devise a proof formalism where the following two proofs π_1 and π_2 are the same.

$$\frac{\frac{\frac{\pi'}{\triangle}}{\vdash \varphi_1, \varphi_2, \psi_1, \psi_2} (\wp)}{\vdash \varphi_1, \varphi_2, \psi_1 \wp \psi_2} (\wp)}{\vdash \varphi_1 \wp \varphi_2, \psi_1 \wp \psi_2} (\wp) \quad ; \quad \frac{\frac{\frac{\pi'}{\triangle}}{\vdash \varphi_1, \varphi_2, \psi_1, \psi_2} (\wp)}{\vdash \varphi_1 \wp \varphi_2, \psi_1, \psi_2} (\wp)}{\vdash \varphi_1 \wp \varphi_2, \psi_1 \wp \psi_2} (\wp)$$

However, one must be careful that one does not quotient more than necessary *i.e.* equating two proofs with distinct computational content. In particular, we would like to differentiate the following two proofs in our formalism. (There are exactly two cut-free proofs of this sequent, and hence it is an alternative encoding of booleans in LL without using units. Also note, that the exchange rule is crucial; otherwise only one of the proofs would exist.)

$$\frac{\frac{\frac{\frac{\text{id}}{\vdash a, a^\perp} (\text{id})}{\vdash a, a^\perp} (\otimes)}{\vdash a, a, a^\perp \otimes a^\perp} (\wp)}{\vdash a \wp a, a^\perp \otimes a^\perp} (\wp)}{\vdash (a \wp a) \wp (a^\perp \otimes a^\perp)} (\wp) \quad ; \quad \frac{\frac{\frac{\frac{\text{id}}{\vdash a, a^\perp} (\text{id})}{\vdash a, a^\perp} (\otimes)}{\vdash a, a, a^\perp \otimes a^\perp} (\wp)}{\vdash a \wp a, a^\perp \otimes a^\perp} (\wp)}{\vdash (a \wp a) \wp (a^\perp \otimes a^\perp)} (\wp)$$

Proof-nets are a geometrical method of representing proofs that were invented to eliminate such kind of syntactic *bureaucracy*. A proof-net can be seen as a graph whose nodes are inference rules, which are thus not ordered, and consequently less sequential than sequent calculus proofs. In particular, proof-nets recover the confluence of cut-elimination.

In this subsection, we will recall proof-nets for MLL without units. Proof-nets for MLL with units or MELL is not fully canonical and proof-nets for MALL are quite cumbersome and out of the scope of this thesis.


Definition 3.5.1. An MLL proof-structure is a vertex-labelled and edge-labelled directed multi-graph where the nodes are labelled by rules $\{\mathbf{ax}, \mathbf{cut}, \otimes, \wp, \mathbf{c}\}$ and the edges are labelled by formulas such that:

- Nodes labelled \mathbf{ax} have two incoming edges labelled φ and φ^\perp for some formula φ and no outgoing edges.
- Nodes labelled \mathbf{cut} have two outgoing edges labelled φ and φ^\perp for some formula φ and no incoming edges.
- Nodes labelled \otimes (respectively, \wp) have two outgoing edges labelled φ and ψ (from left to right) and one incoming edge labelled $\varphi \otimes \psi$ (respectively, $\varphi \wp \psi$), for some formula φ and ψ .
- Nodes labelled \mathbf{c} have exactly one outgoing edge and no incoming edges.


A sequent proof, π , in MLL can be translated into a proof-structure $\mathbf{dsq}(\pi)$ such that there is a bijection between the internal nodes of $\mathbf{dsq}(\pi)$ and the inference rules of π .

Definition 3.5.2. Let π be a MLL proof. **Desquentialisation** of π , denoted $\mathbf{dsq}(\pi)$, is defined by induction on the structure of π as follows. There are several cases based on the root node of π .


The root node is (id). The proof-structure corresponding to $\frac{\text{id}}{\vdash \varphi, \varphi^\perp} (\text{id})$ is the graph containing a single node \mathbf{ax} with incoming edges labelled φ and φ^\perp which have \mathbf{c} nodes as targets.




The root node is (cut). Let π be of the form $\frac{\vdash \Gamma, \varphi \quad \vdash \Delta, \varphi^\perp}{\vdash \Gamma, \Delta}$ (cut) such that \mathcal{R}_1 and \mathcal{R}_2 are the desequentialisation of the subproofs rooted at $\vdash \Gamma, \varphi$ and $\vdash \Delta, \varphi^\perp$ respectively. Then, $\text{dsq}(\pi)$ is obtained by removing the c nodes of \mathcal{R}_1 and \mathcal{R}_2 with outgoing edges e_1, e_2 labelled φ and φ^\perp respectively, and by introducing a new cut node with outgoing edges e_1 and e_2 .



The root node is (\otimes). Let π be of the form $\frac{\vdash \Gamma, \varphi \quad \vdash \Delta, \psi}{\vdash \Gamma, \Delta, \varphi \otimes \psi}$ (\otimes) such that \mathcal{R}_1 and \mathcal{R}_2 are the desequentialisation of the subproofs rooted at $\vdash \Gamma, \varphi$ and $\vdash \Delta, \psi$ respectively. Then, $\text{dsq}(\pi)$ is obtained by removing the c nodes of \mathcal{R}_1 and \mathcal{R}_2 with outgoing edges e_1, e_2 labelled φ and ψ respectively, and by introducing a new \otimes node with outgoing edges e_1 and e_2 and an incoming edge labelled $\varphi \otimes \psi$ whose source is a new c node.




The root node is (\wp) rule. Let π be of the form $\frac{\vdash \Gamma, \varphi, \psi}{\vdash \Gamma, \varphi \wp \psi}$ (\wp) such that \mathcal{R} is the desequentialisation of the subproof rooted at $\vdash \Gamma, \varphi, \psi$. Then, $\text{dsq}(\pi)$ is obtained by removing the c nodes of \mathcal{R} with outgoing edges e_1, e_2 labelled φ and ψ respectively, and by introducing a new \wp node with outgoing edges e_1 and e_2 and an incoming edge labelled $\varphi \wp \psi$ whose source is a new c node.



Definition 3.5.3. A proof net is a proof-structure that is the desequentialisation of some proof.


Example 3.5.1. The desequentialisation of the proofs π_1 and π_2 discussed at the beginning of this subsection are indeed the same.



The desequentialisation of the booleans are indeed different proof-nets \mathcal{R}_1 and \mathcal{R}_2 . Note that non-commutativity (i.e. absence of the exchange rule) corresponds to the planarity of the proof-nets.



Proof-structures, therefore, allow to present sequent proofs in a non-sequential way but the objects are not inductively presented anymore which makes their logical correctness non-trivial. Observe that not every proof-structure represents (or is the desequentialisation of) an MLL proof. For example,




Hence one imposes a correctness criterion to delineate a subset of “correct” proof-structures which belong to the image of the translation $\mathbf{dsq}(\bullet)$. There are several correctness criteria known in the literature. We will discuss the most well-known criterion, the so-called *Danos-Regnier criterion*.

Definition 3.5.4. A **switching**, sw , of a proof-structure \mathcal{R} is a function from the set of \wp nodes of \mathcal{R} to $\{\text{left, right}\}$. Given a switching sw of \mathcal{R} , the **correction graph** is defined as the undirected graph \mathcal{R} where for each \wp node p , its $sw(p)$ outgoing edge has been deleted.

Definition 3.5.5. A proof-structure is said to be **DR-correct** if for every switching, the correction graph is acyclic and connected.

Example 3.5.2. We will show that \mathcal{R}_1 from Example 3.5.1 is DR-correct. There are four switch-




ings viz.

Note that each of these graphs is acyclic and connected. Therefore, \mathcal{R}_1 is DR-correct.

Theorem 3.5.1. A proof-structure is a proof net iff it is DR-correct.

Definition 3.5.6. The **cut elimination relation** \rightarrow_{MLL} is the binary relation over proof-nets generated by the following set of graph rewrite rules.



Lemma 3.5.1 (Correctness preservation). *Let \mathcal{R} be a proof-net and $R \rightarrow_{\text{MLL}} \mathcal{R}'$. Then \mathcal{R}' is a proof-net.*

Theorem 3.5.2. *\rightarrow_{MLL} is confluent and terminating.*

Notes

Linear logic was invented by Girard in [Gir87a] motivated comes from a semantical analysis of the models of System F. Many of the results and concepts stated here were born in that seminal paper.

For a survey of decidability results in linear logic, see [Lin95]. Many of the results were obtained by Lincoln during his PhD [Lin92]. Interestingly, a decidability proof for MELL was published but later the status of the problem was reverted when a crucial bug [Str19] was noticed. Girard published a book providing many insights into linear logic [Gir11]. A general technical survey is due to Curien [Cur05].

See [Avr88] for an alternative to phase semantics with regards to truth semantics of linear logic. See [AD93] for Kripke models of linear logic.

Focussing was invented and studied by Andreoli in his PhD thesis [And92]. Saurin [Sau08] observed that every proof could be focussed by simply permuting inference rules.

The correctness condition presented here was introduced in [DR89] (thus the “DR” in our presentation stands for “Danor-Regnier”). The original correctness condition of Girard is called the long-trip condition.

The Curry-Howard aspect of linear logic has not been mentioned in this brief summary. Natural deduction of linear logic was studied [Tro95, BACJ98] with the motivation to interpret λ -terms. Herbelin [Her95] provided a proofs vs programs correspondence directly in sequent proofs. A recent breakthrough has been interpreting sequent proofs as processes in π -calculus [Mil93, CP10].

Chapter 4

Linear logic with fixed points

What is Matter? - Never mind.
What is Mind? - No matter.

Punch, or The London Charivari (Vol. 29)

This chapter serves as a survey on μ MALL, the extension of propositional linear logic with the least and greatest fixed points. In Section 4.1 we will establish the syntax of μ MALL and the peculiar notion of subformulas in this context. There are several sequent calculi for μ MALL. The first one is a wellfounded system with the so-called *Park's rules*, which are rules directly inspired by the (co)induction principle. The second approach is inspired by the infinite descent proof technique that yields a non-wellfounded system and a circular system. We introduce these proof systems in Section 4.2 and compare the relative expressiveness in Section 4.3. Finally, we discuss the focussing property of these systems in Section 4.4. In Section 4.5 we briefly discuss the cut-elimination results for these systems and their consequences.

4.1 Linear logic with fixed points

In order to define a fixed point logic, one introduces explicit fixed point construct(s) and closes the language under these construct(s) thus obtaining a (potentially) richer logic. In this subsection, we will introduce μ MALL, the extension of MALL with the least and greatest fixed point operators.

Definition 4.1.1. Fix a countable set of propositional constants $\mathcal{A} = \{a, b, \dots\}$ and variables $\mathcal{V} = \{x, y, \dots\}$ such that $\mathcal{A} \cap \mathcal{V} = \emptyset$. μ MALL **pre-formulas** are given by the following grammar:

$$\varphi, \psi ::= \mathbf{0} \mid \top \mid \perp \mid \mathbf{1} \mid a \mid a^\perp \mid x \mid \varphi \wp \psi \mid \varphi \otimes \psi \mid \varphi \oplus \psi \mid \varphi \& \psi \mid \mu x. \varphi \mid \nu x. \varphi$$

where $a \in \mathcal{A}$, $x \in \mathcal{V}$, and μ, ν bind the variable x in φ . Free and bound variables, and capture-avoiding substitution are defined as usual. The subformula ordering is denoted \preceq and $\text{fv}(\bullet)$ denotes free variables. When a pre-formula is closed (i.e. no free variables), we simply call it a **formula**.

Negation, $(\bullet)^\perp$, defined as a meta-operation on pre-formulas, will be used only on formulas. As expected, the least and greatest fixed points are the dual of each other.

Definition 4.1.2. Negation of a pre-formula is defined inductively as follows.

$$\begin{array}{ll} \mathbf{0}^\perp = \top; & \top^\perp = \mathbf{0}; \\ \perp^\perp = \mathbf{1}; & \mathbf{1}^\perp = \perp; \\ (a)^\perp = a^\perp; & a^{\perp\perp} = a; \quad [a \in \mathcal{A}] \\ x^\perp = x; & [x \in \mathcal{V}] \\ (\varphi \wp \psi)^\perp = \psi^\perp \otimes \varphi^\perp; & (\varphi \otimes \psi)^\perp = \psi^\perp \wp \varphi^\perp; \\ (\varphi \oplus \psi)^\perp = \psi^\perp \& \varphi^\perp; & (\varphi \& \psi)^\perp = \psi^\perp \oplus \varphi^\perp; \\ (\mu x. \varphi)^\perp = \nu x. \varphi^\perp; & (\nu x. \varphi)^\perp = \mu x. \varphi^\perp. \end{array}$$

One might be surprised to see that variables are self-dual *i.e.* $x = x^\perp$. One can choose an alternate presentation: in lieu of such a condition, one imposes a positivity condition *viz.* fixed point operators are only applied on a positive¹ formula *wrt.* to the binder where positive formulas are defined as follows.

Definition 4.1.3. *The occurrence of a variable is positive if it occurs within an even number of $(\bullet)^\perp$ nestings. A formula φ is said to be **positive** wrt. a variable x if every occurrence of x is positive in φ .*

The former presentation syntactically prevents us from writing formulas such as $\mu x.x^\perp$. Note that exponentials and non-monotonic definitions combine to yield inconsistency: for example, the formula $\mu x.x^\perp$ does not lead to an inconsistency, whereas $\mu x.?x^\perp$ does. However, in MALL, it is easy to check that the positivity condition and self-duality of variables are equivalent notions.

By an easy induction of the structure of formulas, we have that for all formulas φ , $\varphi^{\perp\perp} = \varphi$. Write $\varphi[\psi/x]$ to denote the substitution of every occurrence of x in φ by ψ . We have $(\varphi[\psi/x])^\perp = \varphi^\perp[\psi^\perp/x]$.

Example 4.1.1. *Let $\psi = \nu x.x \otimes (a \wp a^\perp)$. Then,*

$$\begin{aligned} \psi^\perp &= \mu x.(x \otimes (a \wp a^\perp))^\perp \\ &= \mu x.x^\perp \wp (a \wp a^\perp)^\perp \\ &= \mu x.x \wp (a^\perp \otimes a^{\perp\perp}) \\ &= \mu x.x \wp (a^\perp \otimes a) \end{aligned}$$

In order to work with fixed points logics, it is often necessary to generalise the notion of a subformula. In the following, we will introduce Fischer–Ladner subformulas and discuss their properties. Note that these results are logic-independent and hold for other fixed point logics such as the modal μ -calculus.

Definition 4.1.4. *Let \rightarrow_{FL} be the binary relation between μ MALL formulas given by the following:*

- for all μ MALL formulas $\varphi_0, \varphi_1, \varphi_0 \odot \varphi_1 \rightarrow_{\text{FL}} \varphi_i$ where $i \in \{0, 1\}$ and $\odot \in \{\wp, \otimes, \oplus, \&\}$, and
- for all μ MALL formulas $\varphi, \eta x.\varphi \rightarrow_{\text{FL}} \varphi[\eta x.\varphi/x]$ for $\eta \in \{\mu, \nu\}$.

Let $\rightarrow_{\text{FL}}^*$ be the reflexive and transitive closure of \rightarrow_{FL} . The **Fischer–Ladner closure** of a formula φ , denoted $\text{FL}(\varphi)$, is defined as the set $\{\psi \mid \varphi \rightarrow_{\text{FL}}^* \psi\}$. The **FL-graph** of a formula φ , denoted $\mathfrak{G}(\varphi)$, is the directed graph $(\text{FL}(\varphi), \rightarrow_{\text{FL}})$.

It may not be immediately clear that the Fischer–Ladner closure is always finite since one cannot argue that the membership of one formula in this set only entails the membership of strictly smaller formulas, as it can be done for the usual notion of subformulas. We first give an example of a formula ψ that has a finite Fischer–Ladner closure and invite the reader to observe that its FL-graph is, in fact, a regular tree rooted at ψ .

Example 4.1.2. *The FL-graph of the formula ψ of Example 4.1.1:*

$$\begin{array}{c} \mathfrak{G}(\psi) = \nu x.x \otimes (a \wp a^\perp) \\ \begin{array}{ccc} & a & a^\perp \\ & \swarrow & \searrow \\ \psi \otimes (a \wp a^\perp) & \longrightarrow & a \wp a^\perp \end{array} \end{array}$$

Lemma 4.1.1. *Let φ, φ_1 , and φ_2 be arbitrary formulas. We have the following.*

1. $\text{FL}(\varphi_1 \odot \varphi_2) = \{\varphi_1 \odot \varphi_2\} \cup \text{FL}(\varphi_1) \cup \text{FL}(\varphi_2)$ for $\odot \in \{\wp, \otimes, \oplus, \&\}$.
2. $\text{FL}(\varphi_1[\varphi_2/x]) = \{\psi[\varphi_2/x] \mid \psi \in \text{FL}(\varphi_1)\} \cup \text{FL}(\varphi_2)$ where $x \in \text{fv}(\varphi_1)$.
3. $\text{FL}(\eta x.\varphi) = \{\eta x.\varphi\} \cup \{\psi[\eta x.\varphi/x] \mid \psi \in \text{FL}(\varphi)\}$ for $\eta \in \{\mu, \nu\}$.

¹Not to be confused with positive formulas in the context of focussing.

Proof. Item (1) is immediate from the definition.

For item (2), we first observe that $\text{FL}(\varphi_1[\varphi_2/x]) \subseteq \{\psi[\varphi_2/x] \mid \psi \in \text{FL}(\varphi_1)\} \cup \text{FL}(\varphi_2)$ can be proved by showing that the RHS has the required closure properties. In the other direction, we divide the proposition into two smaller claims: Equations (4.1) and (4.2).

$$\text{FL}(\varphi_1[\varphi_2/x]) \supseteq \{\psi[\varphi_2/x] \mid \psi \in \text{FL}(\varphi_1)\}. \quad (4.1)$$

$$\text{FL}(\varphi_1[\varphi_2/x]) \supseteq \text{FL}(\varphi_2). \quad (4.2)$$

To prove Equation (4.1), suppose we have $\varphi_1 = \psi_0 \rightarrow_{\text{FL}} \psi_1 \rightarrow_{\text{FL}} \dots \rightarrow_{\text{FL}} \psi_n = \psi$. We will induct on n . The base case is $n = 0$. So, $\psi = \varphi_1$. Since $\varphi_1[\varphi_2/x] \in \text{FL}(\varphi_1[\varphi_2/x])$ we are done. The induction case is $n > 0$. Consider ψ_{n-1} . Note that it cannot be an atom or a variable since there exists ψ_n such that $\psi_{n-1} \rightarrow_{\text{FL}} \psi_n$. Suppose $\psi_{n-1} = \xi_1 \odot \xi_2$ where $\odot \in \{\wp, \otimes, \oplus, \&\}$. By the induction hypothesis, $\xi_1 \odot \xi_2[\varphi_2/x] \in \text{FL}(\varphi_1[\varphi_2/x])$ which implies $\xi_1[\varphi_2/x] \text{FL}(\varphi_1[\varphi_2/x])$ and $\xi_2[\varphi_2/x] \text{FL}(\varphi_1[\varphi_2/x])$. Since $\psi_{n-1} \rightarrow_{\text{FL}} \psi_n$, either $\xi_1[\varphi_2/x] = \psi_n$ or $\xi_2[\varphi_2/x] = \psi_n$. Hence done. Now suppose $\psi_{n-1} = \eta y. \xi$ for $\eta \in \{\mu, \nu\}$. Then, by hypothesis, $\eta y. \xi[\varphi_2/x] \in \text{FL}(\varphi_1[\varphi_2/x])$ which implies $\xi[\varphi_2/x, \xi/y] \in \text{FL}(\varphi_1[\varphi_2/x])$. Since $\psi_{n-1} \rightarrow_{\text{FL}} \psi_n$, $\xi[\varphi_2/x, \xi/y] = \psi_n$ and hence we are done.

To prove Equation (4.2), we will induct on the depth of x in $\mathfrak{T}(\varphi_1)$. Suppose we have that $x = \psi_0 \preceq \psi_1 \preceq \dots \preceq \psi_n = \varphi_1$. The base case is $n = 0$. So, $\varphi_1 = x$ (since $x \in \text{fv}(\varphi_1)$, it cannot be an atom or a different variable). Therefore, this case is trivial. The induction case is $n > 0$. Note that $x \in \text{fv}(\psi_i)$ for all i . Consider ψ_{n-1} . By the induction hypothesis, $\text{FL}(\varphi_2) \subseteq \text{FL}(\psi_{n-1}[\varphi_2])$. Note that $\psi_{n-1} \preceq \varphi_1$ implies that $\varphi_1 \rightarrow_{\text{FL}} \psi_{n-1}$. One can easily show that $\text{FL}(\psi_{n-1}[\varphi_2]) \subseteq \text{FL}(\varphi_1[\varphi_2])$. Hence done.

For item (3), we have two cases, either $x \in \text{fv}(\varphi)$ or not. In the first case, we have:

$$\begin{aligned} \text{FL}(\eta x. \varphi) &= \{\eta x. \varphi\} \cup \text{FL}(\varphi[\eta x. \varphi]) && \text{[By definition]} \\ &= \{\eta x. \varphi\} \cup \{\psi[\eta x. \varphi/x] \mid \psi \in \text{FL}(\varphi)\} \cup \text{FL}(\eta x. \varphi) && \text{[By item (2)]} \end{aligned}$$

Therefore, we have $\text{FL}(\eta x. \varphi) \supseteq \{\eta x. \varphi\} \cup \{\psi[\eta x. \varphi/x] \mid \psi \in \text{FL}(\varphi)\}$. The opposite direction is trivial. In the case $x \notin \text{fv}(\varphi)$, we have $x \notin \text{fv}(\psi)$ for all $\psi \in \text{FL}(\varphi)$ and hence $\{\psi[\eta x. \varphi/x] \mid \psi \in \text{FL}(\varphi)\} = \text{FL}(\varphi)$. Hence, we are done. \square

Theorem 4.1.1. *For any formula φ , $\text{FL}(\varphi)$ is a finite set.*

Proof. By induction on φ . The base case is when $\varphi \in \mathcal{A} \cup \mathcal{V}$ or φ is a unit. Then $\text{FL}(\varphi) = \{\varphi\}$ and we are done. For the induction case, we have two subcases.

Case 1. $\varphi = \varphi_1 \odot \varphi_2$ for $\odot \in \{\wp, \otimes, \oplus, \&\}$.

By Lemma 4.1.1, $\text{FL}(\varphi) = \{\varphi\} \cup \text{FL}(\varphi_1) \cup \text{FL}(\varphi_2)$. Since $\text{FL}(\varphi_1)$ and $\text{FL}(\varphi_2)$ are finite by induction hypothesis, we are done.

Case 2. $\varphi = \eta x. \varphi'$ for $\eta \in \{\mu, \nu\}$.

By Lemma 4.1.1, $\text{FL}(\varphi) = \{\varphi\} \cup \{\psi[\varphi/x] \mid \psi \in \text{FL}(\varphi')\}$. Since $\text{FL}(\varphi')$ is finite by induction hypothesis, we are done. \square

Definition 4.1.5. *A syntax tree of a formula φ , denoted $\mathfrak{T}(\varphi)$ is the (possibly infinite) unfolding tree of its Fischer-Ladner graph $\mathfrak{G}(\varphi)$.*

Recall that the *syntax tree* $\mathfrak{T}(\varphi)$ of an LL formula φ induces a prefix closed language, $\mathcal{L}_\varphi \subseteq \{l, r, i\}^*$ such that there is a natural bijection between \mathcal{L}_φ and the branches of the tree. In the case of μMALL , the syntax tree $\mathfrak{T}(\varphi)$ could potentially be infinite, hence the induced language can also contain infinite words *i.e.* $\mathcal{L}_\varphi \subseteq \{l, r, i\}^\infty$.

One can argue that since the \wp s are invertible rules, one can have the following without having to allow for arbitrary invariants.

$$\frac{\vdash \Gamma, \varphi \left[\left(\bigotimes_{i \in [n]} \psi_n^\perp \right) / x \right]}{\vdash \Gamma, \nu x. \varphi} (\nu)$$

where $\Gamma = \psi_1, \dots, \psi_n$ and if $\Gamma = \emptyset$ choose $\mathbf{1}$ as the invariant (by invertibility of the (\perp) rule, one can assume $\Gamma = \perp$). Note that although this rule and previous one are equivalent up to provability, the proof of equivalence goes through a cut. Therefore, the inadmissibility of cuts in a system with the previous rule does not automatically imply the inadmissibility of cuts in a system with the above rule. Unfortunately, this new system also does not admit cuts as exemplified by the following proof of $\vdash b, \nu x. a \wp x$ which the readers can verify has no cut-free proof (a, b are arbitrary atoms here).

$$\frac{\frac{\frac{\frac{}{\vdash \mathbf{0}, a, \top} (\top)}{\vdash \mathbf{0}, a \wp \top} (\wp)}{\vdash b, \top} (\top)}{\vdash \mathbf{0}, \nu x. a \wp x} (\nu)}{\vdash b, \nu x. a \wp x} (\text{cut})$$

Note that, essentially, we had to change the invariant with help of a cut.

Definition 4.2.1. *The system $\mu\text{MALL}^{\text{ind}}$ is generated from the inference rules of MALL given in Figure 3.1 and the following rules for the fixed points.*

$$\frac{\vdash \Gamma, \varphi(\mu x. \varphi/x)}{\vdash \Gamma, \mu x. \varphi} (\mu) \quad ; \quad \frac{\vdash \Gamma, \psi \quad \vdash \psi^\perp, \varphi[\psi/x]}{\vdash \Gamma, \nu x. \varphi} (\nu)$$

Example 4.2.1. *Consider ψ from Example 4.1.1. We have the following proof of the sequent $\vdash \psi$.*

$$\frac{\frac{\frac{\frac{}{\vdash \mathbf{1}} (\mathbf{1})}{\vdash \mathbf{1} \otimes (a \wp a^\perp)} (\otimes)}{\vdash \perp, \mathbf{1} \otimes (a \wp a^\perp)} (\perp)}{\vdash \mathbf{1}} (\mathbf{1})}{\vdash \psi} (\nu)$$

The choice of a ψ in the (ν) rule is akin to choosing the correct induction hypothesis. However, this rule does not preserve subformulas. In fact, one can also think of choosing ψ as choosing a cut formula. Therefore, analyticity does not guarantee the subformula property in $\mu\text{MALL}^{\text{ind}}$. Moreover, this still poses a major automation challenge.

An alternative is to consider systems with implicit (co)induction. In such systems, the fixed point rules are fold and unfold.

$$\frac{\Gamma, \varphi[\mu x. \varphi/x] \vdash \Delta}{\Gamma, \mu x. \varphi \vdash \Delta} (\mu_l) \quad ; \quad \frac{\Gamma \vdash \varphi[\mu x. \varphi/x], \Delta}{\Gamma \vdash \mu x. \varphi, \Delta} (\mu_r)$$

$$\frac{\Gamma, \varphi[\nu x. \varphi/x] \vdash \Delta}{\Gamma, \nu x. \varphi \vdash \Delta} (\nu_l) \quad ; \quad \frac{\Gamma \vdash \varphi[\nu x. \varphi/x], \Delta}{\Gamma \vdash \nu x. \varphi, \Delta} (\nu_r)$$

As before, construing logical entailment as the underlying partial order, we have that $\varphi(\mu\varphi) \leq \Delta \implies \mu\varphi \leq \Delta$ and $\Gamma \leq \varphi(\mu\varphi) \implies \Gamma \leq \mu\varphi$. Plugging $\Delta = \Gamma = \varphi(\mu\varphi)$, we have that $\mu\varphi$ is a pre-fixed point and a post-fixed point. Therefore, $\mu\varphi$ is indeed a fixed-point (but not necessarily the least).

Observe that the rules of μ and ν are identical. Therefore, one can do the same argument for the ν rules and conclude that it is also a fixed point. So, one cannot differentiate between $\mu\varphi$ and $\nu\varphi$. This is problematic logically since both $\mu x. x$ and $\nu x. x$ are not provable using these rules but semantically and from a Curry-Howard perspective, it makes sense for the latter to be provable. Regardless, this is

a useful logic in our investigation, and we call the system μMALL^* . Note that μMALL^* is equivalent to the logic with a unique self-dual fixed point operator with the above rules.

In order to make a distinction between μ and ν , the first step is to move to non-wellfounded proofs. We will now describe the one-sided non-wellfounded and circular systems of μMALL .

Definition 4.2.2. A **pre-proof** of μMALL^∞ is a possibly infinite tree generated from the inference rules of MALL (see Figure 3.1) and the one-sided version of the above rules viz.

$$\frac{\vdash \Gamma, \varphi[\mu x. \varphi/x]}{\vdash \Gamma, \mu x. \varphi} (\mu) \quad ; \quad \frac{\vdash \Gamma, \varphi[\nu x. \varphi/x]}{\vdash \Gamma, \nu x. \varphi} (\nu)$$

One of the key caveats of non-wellfounded proof theory is that, unconstrained, they admit inconsistencies: it is possible to derive any sequent:

$$\frac{\begin{array}{c} \vdots \\ \vdash \mu x. x \\ \vdash \mu x. x \end{array} (\mu) \quad \begin{array}{c} \vdots \\ \vdash \Gamma, \nu x. x \\ \vdash \Gamma, \nu x. x \end{array} (\nu)}{\vdash \Gamma} (\text{cut})$$

For this reason, we impose a global criterion on pre-proofs. This criterion also helps us establish the necessitous distinction between μ and ν .

Definition 4.2.3. Given a pre-proof π , for all inference rules r occurring in π , we define the **immediate ancestor relation** $\text{IA}(r)$ on formulas of r by: $(\varphi, \psi) \in \text{IA}(r)$ if either one of the following holds

- φ is principal and ψ is auxiliary; or,
- φ is a side formula in the conclusion and ψ is a corresponding side formula in a premiss; or,
- r is structural and φ is a formula in the conclusion and ψ is a corresponding formula in a premiss.

For the ease of the reader, we explicitly exhibit the $\text{IA}(\bullet)$ for the inference rules of μMALL^∞ which have non-empty premisses.

$$\begin{array}{ccc} \frac{\vdash \Gamma_1, \varphi \quad \vdash \Gamma_2, \varphi^\perp}{\vdash \Gamma_1, \Gamma_2} (\text{cut}) & \frac{\vdash \Gamma, \varphi_1, \varphi_2}{\vdash \Gamma, \varphi_1 \wp \varphi_2} (\wp) & \frac{\vdash \Gamma_1, \varphi_1 \quad \vdash \Gamma_2, \varphi_2}{\vdash \Gamma_1, \Gamma_2, \varphi_1 \otimes \varphi_2} (\otimes) \\ \\ \frac{\vdash \Gamma, \varphi_i}{\vdash \Gamma, \varphi_1 \oplus \varphi_2} (\oplus_i) & \frac{\vdash \Gamma, \varphi_1 \quad \vdash \Gamma, \varphi_2}{\vdash \Gamma, \varphi_1 \& \varphi_2} (\&) & \frac{\vdash \Gamma}{\vdash \Gamma, \perp} \perp \\ \\ \frac{\vdash \Gamma, \varphi[\mu x. \varphi/x]}{\vdash \Gamma, \mu x. \varphi} (\mu) & \frac{\vdash \Gamma, \varphi[\nu x. \varphi/x]}{\vdash \Gamma, \nu x. \varphi} (\nu) & \end{array}$$

Definition 4.2.4. Let $\beta = (\Gamma_i)_{i < \omega}$ be an infinite branch of a μMALL^∞ pre-proof π and let r_i be the rule with conclusion Γ_i . A **thread** of β is given by $k \in \mathbb{N}$ and a sequence of formulas $\{\varphi_i\}_{k \leq i < \omega}$ such that, for $k < i < \omega$, we have $(\varphi_i, \varphi_{i+1}) \in \text{IA}(r_i)$. The thread is said to be **stationary** if it is not infinitely often principal i.e. there exists $k' > k$ such that for all $i \geq k'$, $\varphi_i = \varphi_{k'}$ and φ_i is not principal.

Proposition 4.2.1. Let τ be a thread. The set of formulas occurring infinitely often in τ , denoted $\text{Inf}(\tau)$, is non-empty and admits a minimum with respect to the \preceq (subformula) ordering.

Proof. Let $\tau = \{\varphi_i\}_{k \leq i < \omega}$. By definition of immediate ancestors, $\varphi_{i+1} \in \text{FL}(\varphi_i)$ for all $k \leq i < \omega$. Therefore, $\varphi_i \in \text{FL}(\varphi_k)$ for all $k \leq i < \omega$. By Theorem 4.1.1, $\text{FL}(\varphi_k)$ is finite. Therefore, $\text{Inf}(\tau)$ is non-empty. We will now show that it admits a minimum with respect to the \preceq ordering.

Suppose τ is stationary i.e. there exists $k' > k$ such that for all $i \geq k'$, $\varphi_i = \varphi_{k'}$; therefore, $\text{Inf}(\tau) = \{\varphi_{k'}\}$ and we are done.

Suppose τ is not stationary. Then, $\{\varphi_i\}_{k \leq i < \omega}$ traces an infinite path in the unfolding of $\mathfrak{G}(\varphi_k)$. Therefore, $\text{Inf}(\tau)$ is isomorphic to a strongly connected subgraph S of $\mathfrak{G}(\varphi_k)$. Observe that among the

nodes of S , there is a unique node n which is nearest to the root φ_k . If it were the case that there were two nodes equidistant to the root then there would be a third node, nearer to the root contradicting the minimality of the distance between n and φ_k . The formula corresponding to the node n is the minimum with respect to the \preceq ordering. \square

Corollary 4.2.0.1. *Let τ be a non-stationary thread. Then, the minimum formula in $\text{Inf}(\tau)$ is a fixed point formula.*

The result is subtle. The set of recurring formulas has a minimum in the \preceq ordering, but it is not totally ordered. Consider the following proof of the formula $\varphi = \nu x.(x \oplus a) \otimes (x \oplus b)$.

$$\frac{\frac{\frac{\vdash \varphi}{\vdash \varphi \oplus a} (\oplus_1) \quad \frac{\vdash \varphi}{\vdash \varphi \oplus b} (\oplus_1)}{\vdash (\varphi \oplus a) \otimes (\varphi \oplus b)} (\otimes)}{\vdash \varphi} (\nu)$$

Each branch of this proof has exactly one thread. The set of formulas occurring infinitely often along a thread in any branch that is not ultimately left-leaning or right-leaning is $\{\varphi, (\varphi \oplus a) \otimes (\varphi \oplus b), \varphi \oplus a, \varphi \oplus b\}$. This set has minimum *viz.* φ but it is not totally ordered. In particular, $\varphi \oplus a$ and $\varphi \oplus b$ are incomparable.

Let π be a μMALL^∞ pre-proof. A branch that has at least one thread which is not ultimately stationary is called a **real branch**. If a branch is not real, it is called a **virtual branch**.

Definition 4.2.5. *A thread τ is **progressing** if it is not stationary and the outermost connective of the smallest formula occurring infinitely often in τ is ν .*

Definition 4.2.6. *A μMALL^∞ pre-proof is called a **proof** if every infinite branch has a progressing thread.*

Example 4.2.2. *Coming back to our inconsistent example, note that while the right infinite branch has a progressing thread along $\nu x.x$ (indicated red), the left branch has no progressing thread, so the pre-proof is not a proof.*

$$\frac{\frac{\frac{\vdots}{\mu x.x} (\mu) \quad \frac{\vdots}{\nu x.x, \Gamma} (\nu)}{\mu x.x} (\mu) \quad \frac{\frac{\vdots}{\nu x.x, \Gamma} (\nu)}{\nu x.x, \Gamma} (\nu)}{\Gamma} (\text{cut})$$

Now consider the formula ψ in Example 4.1.1. The readers are encouraged to convince themselves that the following is a proof of the sequent $\vdash \psi$.

$$\frac{\frac{\frac{\frac{\vdash a, a^\perp}{\vdash a, a^\perp} (\text{id})}{\vdash \psi} (\wp) \quad \frac{\vdash a \wp a^\perp}{\vdash \psi \otimes (a \wp a^\perp)} (\otimes)}{\vdash \psi} (\nu)}$$

Definition 4.2.7. *A μMALL^∞ pre-proof is said to be **circular** (a.k.a. **regular**) if it has finitely many distinct sub-trees. The class of circular proofs is denoted by μMALL° .*

Theorem 4.2.1. *Given a regular pre-proof π , checking whether it is a proof is decidable. Moreover, the problem is PSPACE-complete.*

We end this section by making a few passing observations that reinforce the robustness of the systems introduced.

Canonicity of fixed-point operators. An intrinsic property of logical operators is their *canonicity* (or lack thereof) *i.e.* if one adds duplicates of the same type to the language and copy-cat inference rules, then the duplicates are equivalent to their corresponding originals. An interesting question about a logic, and more precisely about its connectives, is whether a connective is equivalent to all of its duplicates. For example, **MALL** connectives are canonical whereas atoms and exponentials are non-canonical. It turns out that fixed point operators are also canonical.

Canonicity of the progress condition. The progress condition is conjectured to be maximal for cut-free proofs in the following sense: the set of valid cut-free proofs is the largest set of consistent cut-free pre-proofs.

Progress condition vs. parity condition. Moreover, the progress condition is very natural. Let $c : \mathbb{F}_\tau \rightarrow \mathbb{N}$ where \mathbb{F}_τ is the set of formulas occurring in a thread τ such that $c(\varphi) \leq c(\psi)$ iff $\varphi \preceq \psi$ and ν -formulas are assigned even numbers. Then the progress condition is exactly the parity condition of combinatorial games. In fact, this connection is exploited to prove the decidability of the progress condition for regular pre-proofs (cf. Theorem 4.2.1) by reduction to the universality of non-deterministic parity ω -word automata.

4.3 Expressiveness

In this subsection, we will explore the relative provabilities of the proof systems that we have introduced. In order to do so, we will first prove, one important property of these logics, *viz.* functoriality.

Definition 4.3.1. *A logic is said to have the **functoriality** property if the following rule is derivable where φ is a formula such that $x \in \text{fv}(\varphi)$.*

$$\frac{\vdash \psi^\perp, \psi'}{\vdash \varphi^\perp[\psi^\perp/x], \varphi[\psi'/x]} \text{ (func)}$$

Proposition 4.3.1. $\mu\text{MALL}^{\text{ind}}$, μMALL° , and μMALL^∞ have the functoriality property.

Proof. We will prove by induction on the maximum depth of x in φ . Therefore the base case is $\varphi = x$ (it cannot be an atom or another variable distinct from x since $x \in \text{fv}(\varphi)$). In this case **func** is a trivial rule with identical premisses and conclusion. There are several subcases for the induction step. The subcases for the multiplicative additive connectives follow from Proposition 3.4.1. We will exhibit the subcases when the outermost operator of φ is a fixed point.

Suppose $\varphi = \nu y.\xi$. We assume that $x \neq y$ and so $\{x, y\} \subseteq \text{fv}(\xi)$. We will first exhibit for $\mu\text{MALL}^{\text{ind}}$. Let $\rho = \varphi^\perp[\psi^\perp/x]$.

$$\frac{\frac{\frac{\text{IH}}{\triangle} \quad \frac{\vdash \xi^\perp[\psi^\perp/x, \rho/y], \xi[\psi'/x, \rho^\perp/y]}{\vdash \rho, \rho^\perp} \text{ (id)} \quad \frac{\vdash \xi^\perp[\psi^\perp/x, \rho/y], \xi[\psi'/x, \rho^\perp/y]}{\vdash \rho, \xi[\psi'/x, \rho^\perp/y]} \text{ (\mu)}}{\vdash \rho, \nu y.\xi[\psi'/x]} \text{ (\nu)}}{\vdash \rho, \nu y.\xi[\psi'/x]} \text{ (\nu)}$$

Note that although the sizes of the formulas blows up, the maximum depth of x strictly decreases from ρ to $\xi^\perp[\psi^\perp/x, \rho/y]$. Therefore, induction hypothesis can be applied. The case when $\varphi = \mu y.\xi$ follows exactly similarly. We will now exhibit for μMALL^∞ and μMALL° (since the proof is finite, they can be tackled at one go). Let $\rho' = \varphi[\psi'/x]$.

$$\frac{\frac{\text{IH}}{\triangle} \quad \frac{\vdash \xi^\perp[\psi^\perp/x, \rho/y], \xi[\psi'/x, \rho'/y]}{\vdash \rho, \xi[\psi'/x, \rho'/y]} \text{ (\mu)}}{\vdash \rho, \rho'} \text{ (\nu)}$$

□

There are several ways to understand the functoriality property. Firstly, note that the proof indicates that it is a generalisation of the axiom expansion. Indeed the proof of Proposition 4.3.1 appeals to the axiom expansion of **MALL** and the case for the fixed points simplifies two fixed point formulas that are (almost) the duals of each other.

Consider the two-sided version of the rule:

$$\frac{\psi \vdash \psi'}{\varphi[\psi/x] \vdash \varphi[\psi'/x]} \text{ (func)}$$

Construing logical entailment as a morphism we have that φ respects the morphism $\psi \rightarrow \psi'$ which is the defining feature of a functor. Operationally, this corresponds to the map operator in functional programming. Imagine ψ is the type of integers and ψ' is the type of booleans and the proof of $\psi \vdash \psi'$ corresponds to the parity function. Now suppose φ is the polymorphic type of lists *i.e.* $\varphi[\psi]$ is type of integer lists and $\varphi[\psi']$ is the type of boolean lists. Then, functoriality is an operator that takes the parity function and a list of integers and returns a list of booleans which is the result of applying the parity function to each element of the list.

Finally, functoriality can also be seen as a *deep inference property*. In deep inference, not only can one apply inference rules to the outermost connectives but also to ‘deeper’ connectives. Functoriality somehow provides a similar power: one can bypass the connectives of φ and apply the inference rules on ψ or ψ' without losing provability.

4.3.1 Relative expressiveness of the different systems

In the following, we construe proof systems as a set of provable sequents. In this sense, if \mathcal{S} and \mathcal{T} are two proof systems of the same logical language, then $\mathcal{S} \subseteq \mathcal{T}$ means that if a sequent is provable in \mathcal{S} then it is also provable in \mathcal{T} .

We first note that $\mu\text{MALL}^\circ \subseteq \mu\text{MALL}^\infty$ since a circular proof can be simply unfolded into a non-wellfounded proof.

Open Question

(Regularisation) Is $\mu\text{MALL}^\infty \subseteq \mu\text{MALL}^\circ$?

Cuts play a crucial role in regularisation. Although cuts are admissible in μMALL^∞ , regularity is not preserved by cut elimination (more details in Section 4.5). Therefore in the case of circular proofs, the calculus without cuts is strictly weaker than the calculus with cuts. In fact, we can show that the cut-free μMALL^∞ and the cut-free μMALL° are not equiprovable since $\varphi = \nu x.x\wp x$ has a cut-free μMALL^∞ proof but no corresponding cut-free μMALL° proof.

$$\begin{array}{c} \vdots \\ \frac{\vdash \varphi, \varphi, \varphi}{\vdash \varphi \wp \varphi, \varphi} (\wp) \\ \frac{\vdash \varphi, \varphi}{\vdash \varphi \wp \varphi} (\nu) \\ \frac{\vdash \varphi, \varphi}{\vdash \varphi \wp \varphi} (\wp) \\ \frac{\vdash \varphi \wp \varphi}{\vdash \varphi} (\nu) \end{array}$$

However, there is indeed a regular proof with cuts of the aforementioned theorem.

$$\frac{\frac{\frac{\frac{\vdash \varphi, \mu x.x}{\vdash \varphi, \varphi, \mu x.x} (\wp)}{\vdash \varphi \wp \varphi, \mu x.x} (\nu)}{\vdash \varphi, \mu x.x} (\text{cut})}{\vdash \varphi} (\text{cut})}{\vdash \varphi} (\text{cut})$$

Therefore, the non-trivial question of regularisation is: Is $\mu\text{MALL}^\infty \subseteq \mu\text{MALL}^\circ$ (possibly with cuts)? Our next observation is that $\mu\text{MALL}^{\text{ind}} \subseteq \mu\text{MALL}^\circ$. It is enough to show that the Park coinduction rule can be simulated using circular proofs.

$$\frac{\frac{\frac{\frac{\vdash \psi^\perp, \varphi[\psi/x]}{\vdash \psi^\perp, \varphi[\psi/x]} (\mu)}{\vdash \psi^\perp, \varphi[\psi/x]} (\text{cut})}{\vdash \Gamma, \psi} (\text{cut})}{\vdash \Gamma, \nu x.\varphi} (\text{cut})$$

Open Question

(Brotherston-Simpson conjecture) Is $\mu\text{MALL}^\circ \subseteq \mu\text{MALL}^{\text{ind}}$?

Note that the unfolding rule for ν can be derived using the Park coinduction rule as follows.

$$\frac{\frac{\frac{\frac{\vdash \varphi^\perp[\mu x.\varphi^\perp/x], \varphi[\nu x.\varphi/x]}{\vdash \mu x.\varphi^\perp, \varphi[\nu x.\varphi/x]} (\mu)}{\vdash \Gamma, \varphi[\nu x.\varphi/x]} (\text{func})}{\vdash \Gamma, \nu x.\varphi} (\nu)$$

This cannot be directly used to finitise regular proofs since a naive transformation will not guarantee the wellfoundedness of the translation. However, this indeed shows that $\mu\text{MALL}^* \subseteq \mu\text{MALL}^{\text{ind}}$. Since $\nu x.x$ is not provable in μMALL^* , the inclusion is strict. To summarise, we have the following.

$$\mu\text{MALL}^* \subsetneq \mu\text{MALL}^{\text{ind}} \subseteq \mu\text{MALL}^{\circ} \subseteq \mu\text{MALL}^{\infty}$$

4.3.2 Expressiveness in terms of other logics

Fixed points can be encoded in second-order linear logic (LL^2) vis-à-vis the following translation $[\bullet] : \mu\text{MALL} \rightarrow \text{LL}^2$.

$$\begin{aligned} [\varphi] &= \varphi & [\varphi \in \mathcal{A} \cup \mathcal{V} \cup \{\perp, \mathbf{1}, \mathbf{0}, \top\}] \\ [\varphi \odot \psi] &= [\varphi] \odot [\psi] & [\odot \in \{\wp, \otimes, \oplus, \&\}] \\ [\mu x.\varphi] &= \forall \psi. ?([\varphi](\psi^\perp) \otimes \psi) \wp \psi^\perp \\ [\nu x.\varphi] &= \exists \psi. !(\psi^\perp \wp [\varphi](\psi)) \otimes \psi \end{aligned}$$

Lemma 4.3.1. *Let φ be a μMALL formula. If φ is provable in $\mu\text{MALL}^{\text{ind}}$ then $[\varphi]$ is provable in LL^2 .*

Lemma 4.3.2. *Let φ be a LL formula such that $\varphi = [\psi]$ for some μMALL formula ψ . If φ is provable in LL^2 then ψ is provable in $\mu\text{MALL}^{\text{ind}}$.*

From Lemma 4.3.1 and Lemma 4.3.2, we get the following.

Theorem 4.3.1. *$\mu\text{MALL}^{\text{ind}}$ provability can be encoded in LL^2 .*

Therefore, fixed points in a sense are special cases of second-order quantification. Note that we crucially needed exponentials in order to simulate them. But in the propositional fragment, one can simulate the exponentials using fixed points.

Definition 4.3.2. *The translation $[\bullet] : \text{LL} \rightarrow \mu\text{MALL}$ is defined as follows (we reuse the same notation from LL^2 encoding).*

$$\begin{aligned} [\varphi] &= \varphi & [\varphi \in \mathcal{A} \cup \mathcal{V} \cup \{\perp, \mathbf{1}, \mathbf{0}, \top\}] \\ [\varphi \odot \psi] &= [\varphi] \odot [\psi] & [\odot \in \{\wp, \otimes, \oplus, \&\}] \\ [?\varphi] &= \mu x. \perp \oplus [\varphi] \oplus (x \wp x) \\ [!\varphi] &= \nu x. \mathbf{1} \& [\varphi] \& (x \otimes x) \end{aligned}$$

Lemma 4.3.3. *Let φ be a LL formula. If φ is provable in LL then $[\varphi]$ is provable in $\mu\text{MALL}^{\text{ind}}$.*

Open Question

Let φ be a μMALL formula such that $\varphi = [\psi]$ for some LL formula ψ . If φ is provable in $\mu\text{MALL}^{\text{ind}}$ then is ψ provable in LL ?

Therefore, in terms of expressiveness, $\mu\text{MALL}^{\text{ind}}$ lies between LL and LL^2 .

4.3.3 Expressiveness in terms of computational content

μMALL is expressive enough to encode several (co)inductive data-types like lists and streams. In particular, natural numbers can be expressed. Recall that in MALL one can already define the type $\text{Bool} := \mathbf{1} \oplus \mathbf{1}$ as follows:

$$\text{true} := \frac{\frac{}{\vdash \mathbf{1}} (\mathbf{1})}{\vdash \text{Bool}} (\oplus_1) \quad \text{false} := \frac{\frac{}{\vdash \mathbf{1}} (\mathbf{1})}{\vdash \text{Bool}} (\oplus_2)$$

We now define the type $\mathbb{N} := \mu x. \mathbf{1} \oplus x$. The set of natural numbers is defined by the isomorphism $n \mapsto \pi_n$ where π_n is defined inductively as follows.

$$\pi_0 := \frac{\frac{\frac{\overline{\vdash \mathbf{1}} \text{ (1)}}{\vdash \mathbf{1} \oplus \mathbb{N}} \text{ } (\oplus_1)}}{\vdash \mathbb{N}} \text{ } (\mu)}{\vdash \mathbb{N}} \text{ } (\mu)$$

$$\pi_n := \frac{\frac{\frac{\frac{\pi_{n-1}}{\vdash \mathbb{N}} \text{ } (\oplus_2)}}{\vdash \mathbf{1} \otimes \mathbb{N}} \text{ } (\oplus_2)}}{\vdash \mathbb{N}} \text{ } (\mu)}$$

Notably one can encode the set of primitive recursive functions over natural numbers in $\mu\text{MALL}^{\text{ind}}$. In order to maintain a natural distinction between input and output, we switch to the two-sided system for the rest of this subsection. For example, the successor function will be encoded as a proof of the sequent $\mathbb{N} \vdash \mathbb{N}$:

$$\text{succ} := \frac{\frac{\frac{\overline{\mathbb{N} \vdash \mathbb{N}} \text{ (id)}}{\mathbb{N} \vdash \mathbf{1} \oplus \mathbb{N}} \text{ } (\oplus)}{\mathbb{N} \vdash \mathbb{N}} \text{ } (\mu)}$$

We will first show that natural numbers can be contracted. Since there is only one formula in the context which works as the invariant, we use the one sequent (μ_ℓ) rule.

$$\text{dup} := \frac{\frac{\frac{\frac{\frac{\frac{\pi_0}{\vdash \mathbb{N}} \text{ } (\oplus_r)}{\vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\oplus_r)}{\mathbf{1} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mathbf{1}_\ell)}{\mathbf{1} \oplus (\mathbb{N} \otimes \mathbb{N}) \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\oplus_\ell^1)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}$$

In the case of circular proofs, the invariant $\mathbb{N} \otimes \mathbb{N}$ which occurs naturally in the $\mu\text{MALL}^{\text{ind}}$ proof has to be reinstalled using a cut.

$$\frac{\frac{\frac{\frac{\frac{\frac{\pi_0}{\vdash \mathbb{N}} \text{ } (\oplus_r)}{\vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\oplus_r)}{\mathbf{1} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mathbf{1}_\ell)}{\mathbf{1} \oplus \mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\oplus_\ell)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}{\mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N}} \text{ } (\mu_\ell)}$$

Encoding projection and composition is relatively straightforward. We will show that one can encode primitive recursion over natural numbers. Let $h : \mathbb{N}^{k+1} \rightarrow \mathbb{N}$, $f : \mathbb{N}^k \rightarrow \mathbb{N}$, and $g : \mathbb{N}^{k+2} \rightarrow \mathbb{N}$ such that

$$h(0, \bar{x}) = f(\bar{x})$$

$$h(y + 1, \bar{x}) = g(y, h(y, \bar{x}), \bar{x})$$

Let π_f and π_g be the proofs corresponding to f and g respectively. Let $\psi = (\mathbb{N}^k \multimap \mathbb{N}) \otimes \mathbb{N}$ be the

4.4 Focussing

The focussing result of MALL can be extended to proof systems of μ MALL. The first step is to classify the fixed point operators as positive and negative. Since the dual of a positive formula is a negative formula (and vice versa) there are basically two choices, either μ is positive and ν is negative, or, μ is negative and ν is positive.

At this point, we observe that for μ MALL^{ind} the μ rule is invertible whereas the ν rule is not in general. If one chooses a clever hypothesis (for example, the unfolding always works) then it is reversible.

$$\frac{\frac{\frac{}{\vdash \varphi^\perp[\nu x.\varphi^\perp/x], \varphi[\mu x.\varphi/x]}{\vdash \nu x.\varphi^\perp, \varphi[\mu x.\varphi/x]} (\nu)}{\vdash \Gamma, \mu x.\varphi} (\text{id})}{\vdash \Gamma, \varphi[\mu x.\varphi/x]} (\text{cut})$$

By conventional reasoning, this goes on to indicate that μ is negative and ν is positive. However, consider the sequent $\vdash \top \otimes \mathbf{1}, \mu x.x$. The focussing discipline forces one to apply the μ rule and as a result, one is perpetually stuck in a loop. Note that, by classifying μ as negative, one still conserves the property that provability is invariant under the application of negative rules. However, there is no clear syntactic rationale to declare μ to be positive (and dually ν to be negative).

There is a denotational intuition for the polarity of fixed points. In denotational interpretations of polarised linear logic, positive formulas are interpreted as an object of the *Eilenberg-Moore category* of the !-comonad of any categorical model of LL. On the other hand, it is natural to interpret $\nu x.\varphi$ as a final coalgebra. Therefore, $\nu x.\varphi$ must be negative and by duality $\mu x.\varphi$ is positive. It should be further noted that obtaining the polarity of fixed point formulas from their LL² encoding is erroneous since the encoding does not preserve computational content.

At this juncture, we mention that the focussing property for MALL can be refined. We can, in fact, prove that π is a proof of Γ then either π is a focussed proof or one can permute inference rules of π to get a focussed proof. As an aside, this shows that focussing does not constrain the computational meaning of proofs. Back to permutations, this is a more robust characterisation of polarity: negative rules can be permuted with any other rule and positive rules can only be permuted with other positive rules. With this in mind, consider the following proof.

$$\frac{\frac{\frac{}{\vdash \mu x.x, \nu x.x} (\mu)}{\vdash \mu x.x, \nu x.x} (\nu)}{\vdash \mu x.x, \nu x.x}$$

Note that if we permute down infinitely many μ rules, we have a pre-proof that is not a proof whereas if we permute down infinitely many ν rules, we indeed have a proof.

Theorem 4.4.1 (Focussing Theorem). μ MALL^{ind}, μ MALL[◊], and μ MALL[∞] have the focussing property i.e. $\vdash \Gamma$ has a proof iff $\vdash \Gamma$ has a focussed proof.

4.5 Cut-elimination

4.5.1 Cut-elimination of $\mu\text{MALL}^{\text{ind}}$

We first note that in Lemma 4.3.2, if the LL^2 proof is cut-free then so is the $\mu\text{MALL}^{\text{ind}}$ proof. Therefore one has the cut-admissibility of $\mu\text{MALL}^{\text{ind}}$ for free. Suppose we are given a $\mu\text{MALL}^{\text{ind}}$ proof of a sequent $\vdash \Gamma$ possibly with cuts. By Lemma 4.3.1, there is a LL^2 proof of $\vdash [\Gamma]$. By the cut-elimination result of LL^2 , there is a cut-free proof of $\vdash [\Gamma]$. By Lemma 4.3.2, there is a cut-free proof of $\vdash \Gamma$. However, this uses the cut-elimination result of LL^2 as a black box and the use of the encoding prevents one from refining the cut-admissibility into a cut-elimination theorem.

It is indeed possible to get a properly constructive cut-elimination result by considering the following reduction rule along with the key cases of cut-elimination in MALL .

$$\begin{array}{c}
 \begin{array}{c} \pi_1 \\ \nabla \\ \hline \vdash \Gamma, \varphi[\mu x.\varphi] \end{array} \quad (\mu) \quad \frac{\begin{array}{c} \pi_2 \\ \nabla \\ \hline \vdash \Delta, \psi \end{array} \quad \begin{array}{c} \pi_3 \\ \nabla \\ \hline \vdash \psi^\perp, \varphi^\perp[\psi] \end{array} (\nu)}{\vdash \Delta, \nu x.\varphi^\perp} \text{(cut)} \\
 \hline
 \vdash \Gamma, \mu x.\varphi \quad \vdash \Delta, \nu x.\varphi^\perp \\
 \hline
 \vdash \Gamma, \Delta \text{ (cut)}
 \end{array}
 \rightarrow_{\mu\text{MALL}^{\text{ind}}}
 \begin{array}{c}
 \begin{array}{c} \pi_1 \\ \nabla \\ \hline \vdash \Gamma, \varphi[\mu x.\varphi] \end{array} \quad \frac{\begin{array}{c} \pi_2 \\ \nabla \\ \hline \vdash \Delta, \psi \end{array} \quad \frac{\begin{array}{c} \pi_3 \\ \nabla \\ \hline \vdash \psi^\perp, \varphi^\perp[\psi] \end{array} (\nu) \quad \frac{\vdash \psi^\perp, \psi \quad \vdash \psi^\perp, \varphi^\perp[\psi]}{\vdash \psi^\perp, \nu x.\varphi^\perp} \text{(func)}}{\vdash \psi^\perp, \varphi^\perp[\nu x.\varphi^\perp]} \text{(cut)}}{\vdash \Delta, \varphi^\perp[\nu x.\varphi^\perp]} \text{(cut)} \\
 \hline
 \vdash \Gamma, \varphi[\mu x.\varphi] \quad \vdash \Delta, \varphi^\perp[\nu x.\varphi^\perp] \\
 \hline
 \vdash \Gamma, \Delta \text{ (cut)}
 \end{array}
 \end{array}$$

Obviously one also needs to add appropriate commutation rules with fixed-point formulas to commutation rules of cut-elimination in MALL . We do not explicitly write them here.

Theorem 4.5.1. *Let π be a $\mu\text{MALL}^{\text{ind}}$ proof. Then there exists π' such that $\pi \rightarrow_{\mu\text{MALL}^{\text{ind}}}^* \pi'$ and π' is cut-free.*

4.5.2 Cut-elimination of μMALL^∞

In finitary proof theory, a successful cut-elimination procedure guarantees that any proof can be reduced to a cut-free proof after a finite number of steps. Via the Curry-Howard correspondence, if reduced properly, this implies the termination of the program that the proof corresponds to. In the non-wellfounded setting, what we need to establish is not the termination of the cut-elimination procedure, but rather its *productivity* i.e. every finite prefix of the result can be computed in a finite number of steps. In particular, consider the following function $f : \mathbb{N} \rightarrow \text{Stream}_{\mathbb{N}}$ that takes a natural number n and produces the stream $n :: n + 1 :: n + 2 :: \dots$

$$f := \frac{\begin{array}{c} \text{dup} \\ \nabla \\ \hline \mathbb{N} \vdash \mathbb{N} \otimes \mathbb{N} \end{array} \quad \frac{\begin{array}{c} \text{succ} \\ \nabla \\ \hline \mathbb{N} \vdash \mathbb{N} \quad \mathbb{N} \vdash \text{Stream}_{\mathbb{N}} \end{array} \quad \frac{\mathbb{N}, \mathbb{N} \vdash \mathbb{N} \otimes \text{Stream}_{\mathbb{N}} \quad \mathbb{N}, \mathbb{N} \vdash \text{Stream}_{\mathbb{N}}}{\mathbb{N}, \mathbb{N} \vdash \text{Stream}_{\mathbb{N}}} (\otimes_r) \quad \frac{\mathbb{N}, \mathbb{N} \vdash \text{Stream}_{\mathbb{N}} \quad \mathbb{N} \otimes \mathbb{N} \vdash \text{Stream}_{\mathbb{N}}}{\mathbb{N} \otimes \mathbb{N} \vdash \text{Stream}_{\mathbb{N}}} (\otimes_\ell)}{\mathbb{N} \otimes \mathbb{N} \vdash \text{Stream}_{\mathbb{N}}} (\text{cut})}$$

A successful cut-elimination procedure guarantees that for all k , one can inspect $n :: n + 1 :: n + 2 :: \dots :: n + k$ in a finite time. Note that this also shows that circular proofs are not closed under cut-elimination. Therefore, we will only consider μMALL^∞ .




Figure 4.1: A discrepancy due to non-canonicity of sequent proofs

In finitary proof theory, cut elimination may proceed by reducing topmost cuts but there is no such thing, in general, as a topmost cut in non-wellfounded proof theory. Instead one relies on the reduction of bottom-most cuts using a generalized cut-rule, the *multicut rule*. A multicut is a rule with arbitrary number of premisses where pairs of premisses constitute a cut rule. One can visualise this as the flattening of a tree of cuts. The following is an instance of a multicut rule. The blue lines indicate two cuts that have been flattened.

$$\frac{\vdash \Gamma, \varphi \quad \vdash \varphi^\perp, \psi \quad \vdash \psi^\perp, \Delta}{\vdash \Gamma, \Delta} (\text{mcut})$$

A μMALL^∞ proof (possibly with cuts) can be transformed into a μMALL_m^∞ proof by first assigning explicit addresses to formulas, then permuting cuts so that they are consecutive, and then applying the following rule called **merge** in a bottom-up manner such that there is at most one **mcut** rule on every branch.

$$\frac{\mathcal{C} \quad \frac{\frac{\vdash \Delta, F \quad \vdash \Gamma, F^\perp}{\vdash \Delta, \Gamma} (\text{cut})}{\vdash \Sigma} (\text{cut})}{\vdash \Sigma} \rightarrow_{\text{merge}} \frac{\mathcal{C} \quad \vdash \Delta, F \quad \vdash \Gamma, F^\perp}{\vdash \Sigma} (\text{mcut})$$

We do not mention other multicut reduction rules since that is not the focus of this thesis. We denote them by \rightarrow_r where r is the label indicating the type of the reduction and the pair of cut occurrences (again, working with occurrences is crucial here) it is acting on. Let \rightarrow be the union of all such \rightarrow_r .

There is a final piece to the puzzle. In order to obtain productivity of the cut-elimination procedure, one needs to restrict the set of reduction sequences and start from pre-proofs that satisfy the progress condition. A priori, productivity only guarantees that we obtain a pre-proof at the limit of the cut-elimination procedure. One also needs to ensure that the limit is indeed a proof typically while considering (the productivity of) higher-order functions, where the result of a cut-elimination may itself be used as a function.

Definition 4.5.1. Let $(\pi_i \rightarrow \pi_{i+1})_{i < \omega}$ be a multicut reduction sequence. The sequence is said to be **fair** if for all $i \in \omega$ such that $\pi_i \rightarrow_r \pi'$ there is some $j \in \omega$ such that $j \geq i$ and $\pi_j \rightarrow_r \pi_{j+1}$.

Theorem 4.5.2. Let $(\pi_i \rightarrow \pi_{i+1})_{i < \omega}$ be a fair reduction sequence such that π_0 is a proof. Then there exists a cut-free proof $\pi \in \mu\text{MALL}^\infty$ such that $\lim_{n \rightarrow \infty} d(\pi, \pi_n) = 0$.

The proof of Theorem 4.5.2 is especially intricate. There are two things to be shown: (i) fair multicut reductions are strongly convergent (ii) the limit is progressing. To obtain both of these, it is crucial that π_0 satisfies the progress condition. Both are proofs by contradiction. For (i), assuming that there exists an infinite fair sequence of multicut commutation steps, one can derive the proof of the empty sequent in a suitably defined proof-system \mathcal{S} . For (ii), assuming that the limit is not progressing, one gets a proof of a sequent with a list of **0**s in the aforementioned proof system \mathcal{S} . In both cases, the contradiction is dependent on the soundness of \mathcal{S} which is established by a semantic argument.

Before concluding, we mention that there have been some very recent advances in generalising Theorem 4.5.2. First, note that although the progress condition is crucial for proving Theorem 4.5.2 but it is not necessary for the productivity of cut-elimination. For example, the pre-proof in Figure 4.1b although not progressing can be productively reduced. One can generalise the progress condition to the so-called *bouncing-thread progress condition* in order to capture some (but not all) such proofs.

Theorem 4.5.3. *Let $(\pi_i \rightarrow \pi_{i+1})_{i < \omega}$ be a fair reduction sequence such that π_0 is a pre-proof satisfying the bouncing-thread progress condition. Then there exists a cut-free proof $\pi \in \mu\text{MALL}^\infty$ such that $\lim_{n \rightarrow \infty} d(\pi, \pi_n) = 0$.*

The bouncing-thread progress condition, however, is not robust under the permutation of inference rules. For example, the pre-proof in Figure 4.1a does not satisfy the bouncing thread progress condition although it is permutatively equivalent to the one in Figure 4.1b.

Notes

Fischer-Ladner closure was first introduced in the context of propositional dynamic logic [FL79]. Early approaches [Gir92] to adding fixed points to linear logic considered a sort of a Y -combinator that could not discriminate between a least and greatest fixed point.

$\mu\text{MALL}^{\text{ind}}$ with first order predicates and equality was first introduced by Baelde and Miller [BM07, Bae08]. Baelde established its focussing property, expressivity, and cut-elimination (initially an indirect proof via LL^2 and then a direct proof [Bae12] using *candidates of reducibility*). Subsequently, ludics [BDS15] and coherence space semantics [EJ21], on the semantic side, and deep inference systems [CG14], on the syntactic side, have been studied for $\mu\text{MALL}^{\text{ind}}$. On a related note, game semantics [Cla09a, Cla09b, Cla10] has been studied for a wellfounded calculus of μLJ .

Santocanale [San02] introduced a circular proof system for the additive fragment of μMALL and showed a correspondence to simple computations on (co)inductive data. Along with Fortier [FS13, For14], he subsequently established cut-elimination for that fragment. Non-wellfounded and circular proof systems for μMALL were introduced in [BDS16] which studied μMALL^∞ focussing and cut-elimination, the latter being especially challenging since the combinatorial and topological techniques of [FS13] do not scale to the multiplicatives. Bouncing threads were introduced in [BDKS22].

The decidability of the progress condition of μMALL° was shown in [Dou17] and the hardness result was proved in [NST19]. For an abstract categorical treatment of the progress condition in circular proofs, see [Weh21].

On the Curry-Howard side, a corollary of the result in [KPP21] is that μMALL° is at least as expressive as Gödel's System T and [DP22, CP22] explores fragments of μMALL° and μMALL^∞ as session-typed processes.

Finally, we remark that non-wellfounded proof theory is not a peculiarity of fixed point logics or of logics that express some form of inductive reasoning. A notable example is a provability logic called *Gödel-Löb logic* [Sha14].

Part I

Provability of the sequent calculus

Chapter 5

Phase semantics of μ MALL systems

This part, consisting of Chapter 5 and Chapter 6, is dedicated to studying the provability of the different proof systems of μ MALL introduced in Chapter 4. In Chapter 5, we will explore the truth semantics and in Chapter 6, we will explore the complexity of the decision problem “Given Γ , is Γ provable?” for the various systems. The results in this part indicate both familiar similarities and striking differences between μ MALL and fixed-point logics in other settings.

In this chapter, we dedicate ourselves to the phase semantics of μ MALL. In Section 5.1 we devise sound and complete phase semantics for μ MALL^{ind}. As usual, this gives us a (non-effective) cut-admissibility by a technique due to [Oka96, Oka99]. Therefore this serves as an alternate proof of μ MALL^{ind} cut-admissibility. In Section 5.2, we introduce a family of wellfounded infinitely branching calculi for μ MALL that enjoy very natural phase semantics. This serves as a bridge towards exploring the phase semantics of the circular and non-wellfounded calculi which we discuss in Section 5.4.

5.1 μ MALL^{ind} phase semantics

We start by recalling Theorem 4.3.1 *i.e.* one can faithfully encode μ MALL^{ind} in LL^2 . Therefore, one can use LL^2 phase semantics [Oka96, Oka99] to define the phase semantics of μ MALL^{ind}. We will briefly discuss the phase semantics of LL^2 now.

Fix a phase space $\mathcal{M} = (M, \perp, \top)$. We denote the set of facts by $\mathcal{X}_{\mathcal{M}}$. We enrich the phase model by a particular subset of $\mathcal{X}_{\mathcal{M}}$ which we will denote by \mathcal{D} such that the interpretation of LL^2 formulas are elements of \mathcal{D} . We will now set up the definitions required to specify \mathcal{D} .

Definition 5.1.1. For any set of facts $\mathcal{D} \subseteq \mathcal{X}_{\mathcal{M}}$, the set of contexts $\mathbb{C}_{\mathcal{D}}$ is given by the following grammar

$$f, g ::= [] \mid X \in \mathcal{D} \mid f \odot g$$

where $\odot \in \{\otimes, \wp, \&, \oplus\}$. Substitution of contexts is defined as expected: for any $X \in \mathcal{D}$, $f[X] \in \mathbb{C}_{\mathcal{D}}$ is f where every occurrence of $[]$ has been replaced by X .

Definition 5.1.2. $\mathcal{D} \subseteq \mathcal{X}_{\mathcal{M}}$ is said to be LL^2 -closed if

- $\{\perp, \perp^\perp, M, M^\perp\} \subseteq \mathcal{D}$;
- \mathcal{D} is closed under the operations $(\bullet)^\perp, \otimes, \wp, \&$, and \oplus ; and
- for all $f \in \mathbb{C}_{\mathcal{D}}$, $\bigcap_{X \in \mathcal{D}} f[X] \in \mathcal{D}$ and $(\bigcup_{X \in \mathcal{D}} f[X] \in \mathcal{D})^{\perp\perp}$.

The phase model of LL^2 is a 3-tuple $(\mathcal{M}, \mathcal{D}, V)$ where \mathcal{M} is a phase space, \mathcal{D} is an LL^2 -closed set of facts, and V is a \mathcal{D} -valuation *i.e.* a map of the form $V : \mathcal{A} \cup \mathcal{V} \rightarrow \mathcal{D}$. Given a \mathcal{D} -valuation V , define the enrichment $V[x \mapsto X]$ of V as follows:

$$V[x \mapsto X](y) := \begin{cases} V(y) & \text{if } x \neq y; \\ X & \text{otherwise.} \end{cases}$$

Note that $V[x_1 \mapsto X_1] \dots [x_n \mapsto X_n] = V[x_{\pi(1)} \mapsto X_{\pi(1)}] \dots [x_{\pi(n)} \mapsto X_{\pi(n)}]$ for any permutation π of $[n]$. Therefore the sequence of enrichments can be treated as a set. We write it as $V[x_1 \mapsto X_1, \dots, x_n \mapsto X_n]$.

The interpretation is now easy to define. The interpretation of MALL operators and units are standard. Quantifier formulas are interpreted as follows.

$$\begin{aligned} \llbracket \exists^2 x. \varphi(x) \rrbracket^V &:= \left(\bigcup_{X \in \mathcal{D}} \llbracket \varphi \rrbracket^V [x \mapsto X] \right)^{\perp\perp} \\ \llbracket \forall^2 x. \varphi(x) \rrbracket^V &:= \bigcap_{X \in \mathcal{D}} \llbracket \varphi \rrbracket^V [x \mapsto X] \end{aligned}$$

Note that the closure properties of \mathcal{D} ensure that the codomain of the interpretation is \mathcal{D} .

Theorem 5.1.1. *Let φ be an LL^2 formula. Then, φ is provable iff for all phase models $(M, \perp, \mathcal{D}, V)$, we have $1 \in \llbracket \varphi \rrbracket^V$ where 1 is the unit of the monoid.*

Let φ be a μMALL formula. Then, by Theorem 4.3.1 and Theorem 5.1.1, φ is provable in $\mu\text{MALL}^{\text{ind}}$ \iff $\llbracket \varphi \rrbracket$ is provable in LL^2 \iff $1 \in \llbracket \llbracket \varphi \rrbracket \rrbracket$. Therefore, $\llbracket \bullet \rrbracket$ is a sound and complete interpretation for $\mu\text{MALL}^{\text{ind}}$.

5.1.1 Direct interpretation of μMALL formulas

The above trick is barely insightful since it relies on the phase semantics of LL^2 as a black box. In this section, we essentially peek into this black box and develop a direct interpretation of μMALL with minimal requirements. We begin by observing that the set of all facts of a phase space induces a complete lattice.

Let $\mathcal{M} = (M, 1, \perp)$ be a phase space and $\mathcal{X}_{\mathcal{M}}$ be the set of all facts.

Proposition 5.1.1. *$\mathcal{X}_{\mathcal{M}}$ is closed under arbitrary $\&$ and \oplus operations.*

Proof. Let $\{X_i\}_{i \in I}$ be a set of facts. Let $U = \bigoplus_{i \in I} \{X_i\}$. Then

$$\begin{aligned} U^\perp &= \left(\bigcup_{i \in I} X_i \right)^{\perp\perp\perp} \\ &= \left(\bigcup_{i \in I} X_i \right)^\perp \quad \text{[By Proposition 3.3.1.5]} \end{aligned}$$

Therefore, $U^{\perp\perp} = U$. So, U is a fact. Let $V = \bigcap_{i \in I} X_i$. Then,

$$\begin{aligned} V &= \bigcap_{i \in I} X_i^{\perp\perp} \quad \text{[}X_i\text{s are facts]} \\ &= \left(\bigcup_{i \in I} X_i^\perp \right)^\perp \quad \text{[By Proposition 3.3.1.6]} \end{aligned}$$

Therefore, V is a fact (since it is of the form Y^\perp for some Y). □

Therefore, all subsets of $\mathcal{X}_{\mathcal{M}}$ have both a supremum and an infimum in $\mathcal{X}_{\mathcal{M}}$. We have the following.

Corollary 5.1.1.1. *$(\mathcal{X}_{\mathcal{M}}, \subseteq, \oplus, \&, \emptyset^{\perp\perp}, M)$ is a complete lattice.*

Therefore, by Tarski's theorem (Theorem 2.0.2), any monotonic function $\xi : \mathcal{X}_{\mathcal{M}} \rightarrow \mathcal{X}_{\mathcal{M}}$ has a fixed point. Furthermore, the least fixed point $\mu\xi$ (respectively, the greatest fixed point $\nu\xi$ by duality) is given by

$$\mu\xi = \bigcap_{X \in \mathcal{X}_{\mathcal{M}}} \{X \mid \xi(X) \subseteq X\} \quad ; \quad \nu\xi = \left(\bigcup_{X \in \mathcal{X}_{\mathcal{M}}} \{X \mid X \subseteq \xi(X)\} \right)^{\perp\perp}.$$

In order to extend the phase semantics of MALL to $\mu\text{MALL}^{\text{ind}}$ we extend valuations to variables *i.e.* for any valuation V , $\text{dom}(V) = \mathcal{A} \cup \mathcal{V}$ (as we already did for LL^2) and define $\llbracket F \rrbracket^V$ by induction on F

with the usual interpretation of atoms, units, and multiplicative-additive connectives and as follows for fixed points formulas:

$$\begin{aligned} \llbracket x \rrbracket^V &= V(x) & x \in \mathcal{V} \\ \llbracket \mu x. \varphi \rrbracket^V &= \bigcap_{X \in \mathcal{X}_{\mathcal{M}}} \{X \mid \llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq X\} \\ \llbracket \nu x. \varphi \rrbracket^V &= \left(\bigcup_{X \in \mathcal{X}_{\mathcal{M}}} \{X \mid X \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto X]}\} \right)^{\perp\perp} \end{aligned}$$

This interpretation is sound but the completeness is not clear. Indeed, not all facts necessarily have a pre-image, therefore $\llbracket F \rrbracket^{V[x \mapsto X]}$ does not exactly correspond to syntactic substitution and the syntactic model is not a phase model. We need to allow strict subsets of \mathcal{X} for building fixed points. As above in the case of \mathbf{LL}^2 , one cannot consider any subset of \mathcal{X} for this purpose and we shall require that they satisfy some closure properties. We essentially restrict the codomain of $\llbracket \bullet \rrbracket^V$ to subspaces of $\mathcal{X}_{\mathcal{M}}$ closed under μ MALL operations using the same technique as \mathbf{LL}^2 .

Recall $\mathbb{C}_{\mathcal{D}}$ is the set of contexts. For $f \in \mathbb{C}_{\mathcal{D}}$, define $\mu f = \bigcap_{X \in \mathcal{D}} \{X \mid f(X) \subseteq X\}$ and $\nu f = \left(\bigcup_{X \in \mathcal{D}} \{X \mid X \subseteq f(X)\} \right)^{\perp\perp}$.

Definition 5.1.3. $\mathcal{D} \subseteq \mathcal{X}_{\mathcal{M}}$ is said to be μ -closed if

- $\{\perp, \perp^{\perp}, M, M^{\perp}\} \subseteq \mathcal{D}$;
- \mathcal{D} is closed under the operations $(\bullet)^{\perp}, \otimes, \wp, \&, \text{ and } \oplus$; and
- for all $f \in \mathbb{C}_{\mathcal{D}}$, $\mu f \in \mathcal{D}$ and $\nu f \in \mathcal{D}$.

A phase space \mathcal{M} equipped with a μ -closed set of facts \mathcal{D} is called a μ -phase space. A μ -phase space along with a \mathcal{D} -valuation is called a μ -phase model. The μ -phase semantics $\llbracket \bullet \rrbracket$ is a function that takes a μ MALL pre-formula φ and returns a fact in \mathcal{D} . Note that $\llbracket \mu x. \varphi \rrbracket^V$ and $\llbracket \nu x. \varphi \rrbracket^V$ are defined as before except X ranges over \mathcal{D} . A priori, the semantics of pre-formula is only an element of $\mathcal{X}_{\mathcal{M}}$. The closure properties of \mathcal{D} ensure $\llbracket F \rrbracket^V \in \mathcal{D}$ for every formula F and \mathcal{D} -valuation V .

Proposition 5.1.2. For all μ MALL pre-formulas φ , all phase models $(\mathcal{M}, \mathcal{D}, V)$, $\llbracket \varphi \rrbracket^V \in \mathcal{D}$.

Proof. By simple induction on preformulas. The base case holds since the codomain of V is \mathcal{D} and Definition 5.1.3 ensures that the interpretations of the units are in \mathcal{D} . The induction case has two subcases.

- Suppose $\varphi = \psi \odot \psi'$ for some MALL operator \odot . By induction hypothesis $\llbracket \psi \rrbracket^V$ and $\llbracket \psi' \rrbracket^V$ are in \mathcal{D} and by the closure properties of \mathcal{D} , $\llbracket \psi \rrbracket^V \odot \llbracket \psi' \rrbracket^V \in \mathcal{D}$.
- Suppose $\varphi = \mu x. \psi$. Note that $\llbracket \varphi \rrbracket^V = \mu f$ where $f = \llbracket \psi \rrbracket^{V[x \mapsto \square]}$. By induction hypothesis $\llbracket \psi \rrbracket^{V[x \mapsto \square]}$ is an element of $\mathbb{C}_{\mathcal{D}}$. Therefore, by the closure property of \mathcal{D} , $\mu f \in \mathcal{D}$.

□

We will prove the monotonicity of μ MALL operations. For that, we need to prove a basic property of facts.

Proposition 5.1.3. If S, T, U, V are subsets of M such that $S \subseteq T$ and $U \subseteq V$ then $SU \subseteq TV$.

Proof. Suppose $x = su \in SU$ such that $s \in S$ and $u \in U$. Then, $s \in T$ and $u \in V$. Hence $x = st \in TV$. □

Lemma 5.1.1 (Monotonicity). Let F be a μ MALL pre-formula with at most one free variable x . If $X \subseteq Y$ then $\llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}$.

Proof. By induction on φ . The base case is when φ is an atom, a variable, or a unit which are trivial.

- Suppose $\varphi = p \in \mathcal{A} \cup \{\perp, \mathbf{1}, \mathbf{0}, \top\}$. Then, $\llbracket \varphi \rrbracket^{V[x \mapsto X]} = \llbracket \varphi \rrbracket^{V[x \mapsto Y]} = V(p)$.
- Suppose $\varphi = y \in \mathcal{V}$. There are two cases. If $y \neq x$, then $\llbracket \varphi \rrbracket^{V[x \mapsto X]} = \llbracket \varphi \rrbracket^{V[x \mapsto Y]} = V(y)$. Otherwise, we have $\llbracket \varphi \rrbracket^{V[x \mapsto X]} = X \subseteq Y = \llbracket \varphi \rrbracket^{V[x \mapsto Y]}$.

There are several subcases for the induction case.

- Suppose $\varphi = \psi \otimes \psi'$.

$$\begin{aligned}
& \llbracket \psi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]}; \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi' \rrbracket^{V[x \mapsto Y]} && \text{[By IH]} \\
& \Rightarrow \llbracket \psi \rrbracket^{V[x \mapsto X]} \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]} \llbracket \psi' \rrbracket^{V[x \mapsto Y]} && \text{[By Proposition 5.1.3]} \\
& \Rightarrow (\llbracket \psi \rrbracket^{V[x \mapsto X]} \llbracket \psi' \rrbracket^{V[x \mapsto X]})^{\perp\perp} \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto Y]} \llbracket \psi' \rrbracket^{V[x \mapsto Y]})^{\perp\perp} && \text{[By Proposition 3.3.1]} \\
& \Rightarrow \llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}
\end{aligned}$$

- Suppose $\varphi = \psi \wp \psi'$.

$$\begin{aligned}
& \llbracket \psi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]}; \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi' \rrbracket^{V[x \mapsto Y]} && \text{[By IH]} \\
& \Rightarrow (\llbracket \psi \rrbracket^{V[x \mapsto Y]})^{\perp} \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto X]})^{\perp}; (\llbracket \psi' \rrbracket^{V[x \mapsto Y]})^{\perp} \subseteq (\llbracket \psi' \rrbracket^{V[x \mapsto X]})^{\perp} && \text{[By Proposition 3.3.1]} \\
& \Rightarrow (\llbracket \psi \rrbracket^{V[x \mapsto Y]})^{\perp} (\llbracket \psi' \rrbracket^{V[x \mapsto Y]})^{\perp} \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto X]})^{\perp} (\llbracket \psi' \rrbracket^{V[x \mapsto X]})^{\perp} && \text{[By Proposition 5.1.3]} \\
& \Rightarrow \left((\llbracket \psi \rrbracket^{V[x \mapsto X]})^{\perp} (\llbracket \psi' \rrbracket^{V[x \mapsto X]})^{\perp} \right)^{\perp} \subseteq \left((\llbracket \psi \rrbracket^{V[x \mapsto Y]})^{\perp} (\llbracket \psi' \rrbracket^{V[x \mapsto Y]})^{\perp} \right)^{\perp} && \text{[By Proposition 3.3.1]} \\
& \Rightarrow \llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}
\end{aligned}$$

- Suppose $\varphi = \psi \& \psi'$.

$$\begin{aligned}
& \llbracket \psi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]}; \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi' \rrbracket^{V[x \mapsto Y]} && \text{[By IH]} \\
& \Rightarrow \llbracket \psi \rrbracket^{V[x \mapsto X]} \cap \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]} \cap \llbracket \psi' \rrbracket^{V[x \mapsto Y]} \\
& \Rightarrow \llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}
\end{aligned}$$

- Suppose $\varphi = \psi \oplus \psi'$.

$$\begin{aligned}
& \llbracket \psi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]}; \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi' \rrbracket^{V[x \mapsto Y]} && \text{[By IH]} \\
& \Rightarrow \llbracket \psi \rrbracket^{V[x \mapsto X]} \cup \llbracket \psi' \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y]} \cup \llbracket \psi' \rrbracket^{V[x \mapsto Y]} \\
& \Rightarrow (\llbracket \psi \rrbracket^{V[x \mapsto X]} \cup \llbracket \psi' \rrbracket^{V[x \mapsto X]})^{\perp\perp} \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto Y]} \cup \llbracket \psi' \rrbracket^{V[x \mapsto Y]})^{\perp\perp} && \text{[By Proposition 3.3.1]} \\
& \Rightarrow \llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}
\end{aligned}$$

- Suppose $\varphi = \mu y. \psi$. Observe that $y \neq x$ since we assumed that x is not a bound variable in φ . Now by hypothesis, for any fact Z , $\llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]}$. Therefore for every Z such that $\llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]} \subseteq Z$ we have $\llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]} \subseteq Z$. Therefore, $\{Z \mid \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]} \subseteq Z\} \subseteq \{Z \mid \llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]} \subseteq Z\}$. Hence, $\bigcap_{\llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]} \subseteq Z} \{Z\} \subseteq \bigcap_{\llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]} \subseteq Z} \{Z\}$. We conclude $\llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}$.
- Suppose $\varphi = \nu y. \psi$. As before we comment that $y \neq x$ and therefore by hypothesis, for any fact Z , $\llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]} \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]}$. Therefore for every Z such that $Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]}$ we have $Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]}$. Therefore, $\{Z \mid Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]}\} \subseteq \{Z \mid Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]}\}$. Hence, $\bigcup_{Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto X, y \mapsto Z]}} \{Z\} \subseteq \bigcup_{Z \subseteq \llbracket \psi \rrbracket^{V[x \mapsto Y, y \mapsto Z]}} \{Z\}$. Applying Proposition 3.3.1 twice, we conclude $\llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto Y]}$.

□

An application of monotonicity is showing that the interpretation of the fixed point operators are indeed fixed points in the mathematical sense.

Proposition 5.1.4. *Let \mathcal{D} be μ -closed and $f \in \mathbb{F}_{\mathcal{D}}$. Then μf and νf are the least and greatest fixed point of f in \mathcal{D} .*

Proof. We show it for μf . First of all, $\{X \mid f(X) \subseteq X\}$ is non-empty since $M^\perp \in \mathcal{D}$ (also by Proposition 5.1.1). First, we show that it is indeed a fixed point. Observe that $\mu_f \subseteq X$ for any $X \in \mathcal{D}$ which is a pre-fixed point of f . By Lemma 5.1.1, we apply f on both sides. So $f(\mu_f) \subseteq f(X)$ for all $X \in \mathcal{D}$ satisfying $f(X) \subseteq X$ and therefore $f(\mu_f) \subseteq \bigcap_{X \in \mathcal{D}} \{X \mid f(X) \subseteq X\} = \mu f$. So μf is a prefixed point.

But then, since $\mu_f \in \mathcal{D}$, and thanks to the closure properties of \mathcal{D} , so is $f(\mu_f)$. By monotonicity of f , one gets that $f(f(\mu_f)) \subseteq f(\mu_f)$, ensuring that $f(\mu_f)$ is a prefixed point of f . But μf is the least prefixed point; so, we conclude that $\mu_f \subseteq f(\mu_f)$. Therefore, $\mu_f = f(\mu_f)$. Finally, recall $\mu f \subseteq X$ for any pre-fixed point X in \mathcal{D} , so it is the least fixed point in \mathcal{D} . \square

Note that Proposition 5.1.4 cannot be proved directly by Theorem 2.0.2 since \mathcal{D} is not necessarily a complete lattice. Moreover, it does not also imply that \mathcal{D} is a complete lattice by the converse of Theorem 2.0.2 since we show that it has fixed points of a particular kind of monotonic function, not any arbitrary monotonic function.

Given a valuation V , define

$$V^\perp(p) = \begin{cases} V(p) & \text{if } p \in \mathcal{A}; \\ V(p)^\perp & \text{if } p \in \mathcal{V}. \end{cases}$$

Note that we have $V^{\perp\perp} = V$ and $V[x \mapsto X]^\perp = V^\perp[x \mapsto X^\perp]$.

Lemma 5.1.2 (Duality preservation). *Let φ be a μ MALL preformula. Then, $\llbracket \varphi^\perp \rrbracket^{V^\perp} = (\llbracket \varphi \rrbracket^V)^\perp$.*

Proof. By induction on φ . The base case is when φ is an atom, a variable or a unit.

- Suppose $\varphi = a \in \mathcal{A}$. Then, $\llbracket \varphi^\perp \rrbracket^{V^\perp} = V^\perp(a^\perp) = V(a^\perp) = V(a)^\perp = (\llbracket \varphi \rrbracket^V)^\perp$.
- Suppose $\varphi = x \in \mathcal{V}$. Then, $\llbracket \varphi^\perp \rrbracket^{V^\perp} = V^\perp(x^\perp) = V^\perp(x) = V(x)^\perp = (\llbracket \varphi \rrbracket^V)^\perp$.
- The case for the units is easy.

There are several subcases for the induction case.

- Suppose $\varphi = \psi \otimes \psi'$.

$$\begin{aligned} \llbracket (\psi \otimes \psi')^\perp \rrbracket^{V^\perp} &= \llbracket \psi^\perp \wp \psi'^\perp \rrbracket^{V^\perp} \\ &= ((\llbracket \psi^\perp \rrbracket^{V^\perp})^\perp (\llbracket \psi'^\perp \rrbracket^{V^\perp})^\perp)^\perp \\ &= (\llbracket \psi^{\perp\perp} \rrbracket^V \llbracket \psi'^{\perp\perp} \rrbracket^V)^\perp && \text{[By IH]} \\ &= (\llbracket \psi \rrbracket^V \llbracket \psi' \rrbracket^V)^\perp \\ &= (\llbracket \psi \rrbracket^V \llbracket \psi' \rrbracket^V)^{\perp\perp\perp} && \text{[By Proposition 3.3.1]} \\ &= (\llbracket (\psi \otimes \psi') \rrbracket^V)^\perp \end{aligned}$$

Negating both sides we have, $((\llbracket (\psi \otimes \psi')^\perp \rrbracket^{V^\perp})^\perp)^\perp = (\llbracket (\psi \otimes \psi') \rrbracket^V)^{\perp\perp}$. Hence, $(\llbracket \psi^\perp \wp \psi'^\perp \rrbracket^{V^\perp})^\perp = \llbracket (\psi^\perp \wp \psi'^\perp)^\perp \rrbracket^V$. This takes care of the case when the outermost connective of φ is a \wp .

- Suppose $F = \psi' \oplus \psi'$.

$$\begin{aligned} \llbracket (\psi' \oplus \psi')^\perp \rrbracket^{V^\perp} &= \llbracket \psi^\perp \& \psi'^\perp \rrbracket^{V^\perp} \\ &= \llbracket \psi^\perp \rrbracket^{V^\perp} \cap \llbracket \psi'^\perp \rrbracket^{V^\perp} \\ &= (\llbracket \psi \rrbracket^V)^\perp \cap (\llbracket \psi' \rrbracket^V)^\perp && \text{[By IH]} \\ &= (\llbracket \psi \rrbracket^V \cup \llbracket \psi' \rrbracket^V)^\perp \\ &= (\llbracket \psi \rrbracket^V \cup \llbracket \psi' \rrbracket^V)^{\perp\perp\perp} && \text{[By Proposition 3.3.1]} \\ &= (\llbracket \psi \oplus \psi' \rrbracket^V)^\perp \end{aligned}$$

As in the previous case, negating both sides, we derive the case when the outermost connective of φ is a $\&$.

- Suppose $\varphi = \mu x.\psi$.

$$\begin{aligned}
\llbracket (\mu x.\psi)^\perp \rrbracket^{V^\perp} &= \llbracket \nu x.\psi^\perp \rrbracket^{V^\perp} \\
&= \left(\bigcup_{X \in \mathcal{D}} \{X \mid X \subseteq \llbracket \psi^\perp \rrbracket^{V^\perp[x \mapsto X]}\} \right)^{\perp\perp} \\
&= \left(\bigcup_{X \in \mathcal{D}} \{X \mid X \subseteq \llbracket \psi^\perp \rrbracket^{(V[x \mapsto X^\perp])^\perp}\} \right)^{\perp\perp} \\
&= \left(\bigcup_{X \in \mathcal{D}} \{X \mid X \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto X^\perp]})^\perp\} \right)^{\perp\perp} && \text{[By IH]} \\
&= \left(\bigcap_{X \in \mathcal{D}} \{X^\perp \mid X \subseteq (\llbracket \psi \rrbracket^{V[x \mapsto X^\perp]})^\perp\} \right)^{\perp} && \text{[By Proposition 3.3.1]} \\
&= \left(\bigcap_{X \in \mathcal{D}} \{X^\perp \mid \llbracket \psi \rrbracket^{V[x \mapsto X^\perp]} \subseteq X^\perp\} \right)^{\perp} && \text{[By Proposition 3.3.1]} \\
&= \left(\bigcap_{X \in \mathcal{D}} \{X \mid \llbracket \psi \rrbracket^{V[x \mapsto X]} \subseteq X\} \right)^{\perp} && \text{[Closure property of } \mathcal{D} \text{]} \\
&= (\llbracket \mu x.\psi \rrbracket^V)^\perp
\end{aligned}$$

As in the previous case, negating both sides, we derive the case when the outermost operator of φ is a ν .

□

5.1.2 Soundness and completeness

In this subsection, we will prove that with regards to $\mu\text{MALL}^{\text{ind}}$, the interpretation of μMALL given above is sound and complete.

Soundness for wellfounded systems is easy to prove since one can rely on induction on wellfounded trees. The interesting cases are when the inference rule applied on the root sequent is a fixed point rule.

Lemma 5.1.3 (Soundness for $\mu\text{MALL}^{\text{ind}}$). *If $\vdash \Gamma$ then for all μ -phase models $(\mathcal{M}, \mathcal{D}, V)$, $1 \in \llbracket \Gamma \rrbracket^V$.*

Proof. Fix an arbitrary μ -phase model $(\mathcal{M}, \mathcal{D}, V)$. Given a proof π of $\vdash \Gamma$ we will induct on π . The case when the last rule is a **MALL** connective follows from the proof of Theorem 3.3.1. In this proof, we only detail the fixed point cases.

Suppose it is a μ rule. We have that $\Gamma = \Gamma', \mu x.\varphi$.

$$\frac{\vdash \Gamma', \varphi[\mu x.\varphi/x]}{\vdash \Gamma', \mu x.\varphi} (\mu)$$

Assume that we have proved $\llbracket \varphi[\mu x.\varphi/x] \rrbracket^V \subseteq \llbracket \mu x.\varphi \rrbracket^V$. We have the following:

$$\begin{aligned}
(\llbracket \mu x.\varphi \rrbracket^V)^\perp &\subseteq (\llbracket \varphi[\mu x.\varphi/x] \rrbracket^V)^\perp && \text{[Proposition 3.3.1.3]} \\
\Rightarrow (\llbracket \Gamma' \rrbracket^V)^\perp \cdot (\llbracket \mu x.\varphi \rrbracket^V)^\perp &\subseteq (\llbracket \Gamma' \rrbracket^V)^\perp \cdot (\llbracket \varphi[\mu x.\varphi/x] \rrbracket^V)^\perp \\
\Rightarrow \left((\llbracket \Gamma' \rrbracket^V)^\perp \cdot \llbracket \varphi[\mu x.\varphi/x] \rrbracket^{V^\perp} \right)^\perp &\subseteq \left((\llbracket \Gamma' \rrbracket^V)^\perp \cdot \llbracket \mu x.\varphi \rrbracket^{V^\perp} \right)^\perp && \text{[Proposition 3.3.1.3]} \\
\Leftrightarrow \llbracket \Gamma' \wp \varphi[\mu x.\varphi/x] \rrbracket^V &\subseteq \llbracket \Gamma' \wp \mu x.\varphi \rrbracket^V \\
\Rightarrow 1 \in \llbracket \Gamma' \wp \mu x.\varphi \rrbracket^V &&& \text{[IH]}
\end{aligned}$$

Therefore, it suffices to prove $\llbracket \varphi[\mu x.\varphi/x] \rrbracket^V \subseteq \llbracket \mu x.\varphi \rrbracket^V$. Observe that $\llbracket \varphi[\mu x.\varphi/x] \rrbracket^V = \llbracket \varphi \rrbracket^{V[x \mapsto \llbracket \mu x.\varphi \rrbracket^V]}$. Let $X \in \mathcal{D}$ such that $\llbracket \varphi \rrbracket^{V[x \mapsto X]} \subseteq X$ (we thus have $\llbracket \mu x.\varphi \rrbracket^V \subseteq X$). We need to show that $\llbracket \varphi \rrbracket^{V[x \mapsto \llbracket \mu x.\varphi \rrbracket^V]} \subseteq X$. It suffices to show that $\llbracket \varphi \rrbracket^{V[x \mapsto \llbracket \mu x.\varphi \rrbracket^V]} \subseteq \llbracket \varphi \rrbracket^{V[x \mapsto X]}$ which is true by Lemma 5.1.1.

Now suppose the last rule is a ν rule i.e. $\Gamma = \Gamma', \nu x.\varphi$.

$$\frac{\vdash \Gamma', \psi \quad \vdash \psi^\perp, \varphi[\psi/x]}{\vdash \Gamma', \nu x.\varphi} (\nu)$$

We need to show that $1 \in \llbracket \Gamma' \wp \nu x.\varphi \rrbracket^V$ which by Proposition 3.3.2 is equivalent to showing $\llbracket \Gamma' \rrbracket^{V^\perp} \subseteq \llbracket \nu x.\varphi \rrbracket^V$. By hypothesis, we have that $1 \in \llbracket \Gamma' \wp \psi \rrbracket^V$ which is similarly equivalent to $\llbracket \Gamma' \rrbracket^{V^\perp} \subseteq \llbracket \psi \rrbracket^V$. Therefore it suffices to show that $\llbracket \psi \rrbracket^V \subseteq \llbracket \nu x.\varphi \rrbracket^V$. Wlog assume that ψ is a closed formula. We have the following:

$$\begin{aligned} 1 &\in \llbracket \psi^\perp \wp \varphi[\psi/x] \rrbracket^V && \text{[IH]} \\ \Leftrightarrow (\llbracket \psi^\perp \rrbracket^V)^\perp &\subseteq \llbracket \varphi[\psi/x] \rrbracket^V && \text{[Proposition 3.3.2]} \\ \Leftrightarrow \llbracket \psi \rrbracket^V &\subseteq \llbracket \varphi[\psi/x] \rrbracket^V && \text{[Lemma 5.1.2]} \\ \Leftrightarrow \llbracket \psi \rrbracket^V &\subseteq \llbracket \varphi \rrbracket^{V[x \mapsto \llbracket \psi \rrbracket^V]} \\ \Rightarrow \llbracket \psi \rrbracket^V &\subseteq \llbracket \nu x.\varphi \rrbracket^V \end{aligned}$$

□

Completeness for fixed point logics is generally quite difficult since analyticity does not guarantee a subformula property. One is faced with a similar *cul de sac* in proving the cut-elimination of $\mu\text{MALL}^{\text{ind}}$ since it is not straightforward to define the notion of the complexity of the cut formula which reduces with each step of cut-elimination. This problem is solved in [Bae12] by invoking a technique similar to the Tait-Girard reducibility candidates (originally formulated to establish certain properties of various typed lambda calculi [Tai75, Gir71]). It is not surprising that in order to prove completeness one needs to invoke reducibility candidates since the completeness result will give cut admissibility as a corollary.

Recall that $\text{Pr}(\varphi)$ is the set of all sequents Γ such that $\vdash \Gamma, \varphi$ is cut-free provable.

Definition 5.1.4. Let (\mathcal{M}, V) be the syntactic model. Given a μMALL formula φ , the **reducibility candidates** of φ , denoted $\langle \varphi \rangle$, are given by

$$\{X \in \mathcal{X}_{\mathcal{M}} \mid \varphi^\perp \in X \subseteq \text{Pr}(\varphi)\}.$$

Proposition 5.1.5. $X \in \langle \varphi \rangle \iff X^\perp \in \langle \varphi^\perp \rangle$

Proof. Let $X \in \langle \varphi \rangle$. Then $\{\varphi^\perp\} \subseteq X \implies X^\perp \subseteq \{\varphi^\perp\}^\perp = \text{Pr}(\varphi^\perp)$. Also, $X \subseteq \text{Pr}(\varphi) \implies \text{Pr}(\varphi)^\perp = \text{Pr}(\text{Pr}(\varphi)) \subseteq X$. But $\varphi \in \text{Pr}(\text{Pr}(\varphi))$. Hence done. □

We are now ready to define the μ -syntactic model.

Definition 5.1.5. The μ -syntactic model, denoted $(\mu\text{MALL}^\bullet, \emptyset, \cdot, \perp, V)$, is defined as:

- $(\mu\text{MALL}^\bullet, \emptyset, \cdot)$ is the free commutative monoid generated by all formulas.
- $\perp = \text{Pr}(\perp)$.
- $V(p) = \text{Pr}(p)$ for all $p \in \mathcal{A} \cup \mathcal{V}$.
- $\mathcal{D} = \bigcup_{\varphi \in \text{Form}} \langle \varphi \rangle$ where *Form* is the set of all μMALL formulas.

Observe that $\perp = \text{Pr}(\perp) \in \mathcal{D}$ and that \mathcal{D} indeed contains \perp^\perp , μMALL^\bullet and $\mu\text{MALL}^{\bullet\perp}$. It is not a priori obvious that the μ -syntactical model is indeed a μ -phase model. The following two lemmas not only help to establish that sanity check but do much more: they also prove completeness.

Let φ be a preformula with m free variables $\bar{x} = (x_1, \dots, x_m)$. Let $\bar{\psi} = (\psi_1, \dots, \psi_m)$ be an m -tuple of formulas and let $\bar{X} = (X_1 \dots X_m)$ be an m -tuple of facts such that $X_i \in \langle \psi_i \rangle$.

Lemma 5.1.4 (Adequation Lemma for $\mu\text{MALL}^{\text{ind}}$). $\llbracket \varphi \rrbracket^{V[\bar{x} \mapsto \bar{X}]} \subseteq \text{Pr}(\varphi[\bar{\psi}/\bar{x}])$.

Proof. By induction on φ . The proof is similar to the proof of Lemma 3.3.1 except for the following fixed point cases.

Case 1. Suppose $\varphi = \mu y.\varphi'$. Let $\xi = \mu y.\varphi'[\bar{\psi}/\bar{x}]$. In the proof of Lemma 5.1.3, we showed that for any formula φ_0 , $\Pr(\varphi_0[\mu x.\varphi_0/x]) \subseteq \Pr(\mu x.\varphi_0)$. Therefore, $\Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y]) \subseteq \Pr(\xi)$. Let $Y^* = \llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto \Pr(\xi)]}$. By induction hypothesis and the fact that $\Pr(\varphi_0) \in \langle \varphi_0 \rangle$ for all formulas φ_0 , we have the following.

$$Y^* \subseteq \Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y]) \quad (5.1)$$

Therefore, it is enough to show that $\llbracket \mu y.\varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}]} \subseteq Y^*$. Take $\Gamma \in \llbracket \mu y.\varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}]}$. For any fact $Y \in \mathcal{D}$, to show $\Gamma \in Y$ it is enough to show $\llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto Y]} \subseteq Y$. Therefore we need to check that $\llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto Y^*]} \subseteq Y^* = \llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto \Pr(\xi)]}$. This follows by Lemma 5.1.1 from Equation (5.1).

Case 2. Suppose $\varphi = \nu y.\varphi'$. For any fact Y , define $Z_Y = \llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto Y]}$. Let $\Gamma \in \bigcup \{Y \in \mathcal{D} \mid Y \subseteq Z_Y\}$. Therefore, there exists, $Y^* \in \mathcal{D}$ such that $\Gamma \in Y^* \subseteq Z_{Y^*}$. Since $Y^* \in \mathcal{D}$, $Y^* \in \langle \xi \rangle$ for some formula ξ . By induction hypothesis, $\llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto Y^*]} \subseteq \Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y])$.

We will now show that $\Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y]) \subseteq \Pr(\varphi[\bar{\psi}/\bar{x}])$. Let $\Delta \in \Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y])$. If we show that $\xi^\perp \in \Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y])$, we have the following.

$$\frac{\frac{\vdash \xi^\perp, \varphi'[\bar{\psi}/\bar{x}, \xi/y]}{\vdash \Delta, \varphi'[\bar{\psi}/\bar{x}, \xi/y]} \text{ (func)}}{\vdash \Delta, \nu x.\varphi'[\bar{\psi}/\bar{x}]} \text{ (\nu)}$$

In order to show $\xi^\perp \in \Pr(\varphi'[\bar{\psi}/\bar{x}, \xi/y])$, we use the induction hypothesis to reduce the problem to showing $\xi^\perp \in \llbracket \varphi' \rrbracket^{V[\bar{x} \mapsto \bar{X}, y \mapsto Y^*]}$ which is true since $Y^* \in \langle \xi \rangle$. Therefore we have,

$$\begin{aligned} & \Gamma \in \Pr(\nu y.\varphi'[\bar{\psi}/\bar{x}]) \\ & \Rightarrow \bigcup \{Y \in \mathcal{D} \mid Y \subseteq Z_Y\} \subseteq \Pr(\nu y.\varphi'[\bar{\psi}/\bar{x}]) \\ & \Rightarrow \left(\bigcup \{Y \in \mathcal{D} \mid Y \subseteq Z_Y\} \right)^{\perp\perp} \subseteq \Pr(\nu y.\varphi'[\bar{\psi}/\bar{x}]) \end{aligned}$$

This concludes our proof. \square

Lemma 5.1.5. $\varphi^\perp(G/x) \in \llbracket \varphi \rrbracket^{V[x \mapsto X]}$.

Proof. Observe that Lemma 5.1.4 and Proposition 5.1.5 imply $\llbracket \varphi^\perp \rrbracket^{V[\bar{x} \mapsto \bar{X}^\perp]} \subseteq \Pr(\varphi^\perp(\bar{G}^\perp/\bar{x}))$. Therefore, $\{\varphi^\perp(\bar{G}^\perp/\bar{x})\} \llbracket \varphi^\perp \rrbracket^{V[\bar{x} \mapsto \bar{X}^\perp]} \subseteq \{\varphi^\perp(\bar{G}^\perp/\bar{x})\}.\Pr(\varphi^\perp(\bar{G}^\perp/\bar{x})) \subseteq \Pr(\perp) = \perp$. By Proposition 3.3.1.1, $\varphi^\perp(\bar{G}^\perp/\bar{x}) \subseteq \left(\llbracket \varphi^\perp \rrbracket^{V[\bar{x} \mapsto \bar{X}^\perp]} \right)^\perp$ which is $\llbracket \varphi \rrbracket^{V[\bar{x} \mapsto \bar{X}]}$ by Lemma 5.1.2. \square

Observe that by Lemma 5.1.4 and Lemma 5.1.5 we have $\llbracket \varphi \rrbracket^{V[x \mapsto X]} \in \langle \varphi(\bar{G}/\bar{x}) \rangle$. This ensures that \mathcal{D} is μ -closed. Consequently, $(\mu\text{MALL}^\bullet, \emptyset, \cdot, \perp, V)$ is a μ -phase space model.

Theorem 5.1.2 (Cut-free completeness for $\mu\text{MALL}^{\text{ind}}$). *If for any μ -phase model $(\mathcal{M}, \mathcal{D}, V)$, $1 \in \llbracket \Gamma \rrbracket^V$ then $\vdash_{cf} \Gamma$.*

Corollary 5.1.2.1. $\mu\text{MALL}^{\text{ind}}$ admits cuts.

Note that this is an alternate proof of Theorem 4.5.1. To exhibit the power of the phase semantics, we prove Lemma 4.3.3 (which shows that exponentials can be encoded with fixed points) semantically.

Proposition 5.1.6. *Let (M, \perp, V) be a phase model. Extend it to an LL phase model $\mathcal{M} = (M, \perp, V, J)$ and a μ -phase model $\mathcal{M}' = (M, \perp, V, \mathcal{D})$ such that $J \in \mathcal{D}$. Then, for any LL formula φ , $\llbracket \varphi \rrbracket^V \subseteq \llbracket \varphi \rrbracket^V$.*

Proof. First note that since $J \in \mathcal{D}$, $\llbracket \varphi \rrbracket^V \in \mathcal{D}$ for LL formula φ . Now we will show our result. We will induct on φ . The base cases and the cases where $\varphi = \psi \odot \psi'$ are trivial. Suppose $\varphi = !\psi$. It suffices to show

$$I \cap \llbracket \psi \rrbracket^{V^{\perp\perp}} \subseteq \bigcup \left\{ X \in \mathcal{D} \mid X \subseteq \llbracket \mathbf{1} \& (x \otimes x) \& [\psi] \rrbracket^{V[x \mapsto X]} \right\}.$$

Now $\llbracket \mathbf{1} \& (x \otimes x) \& [\psi] \rrbracket^{V[x \mapsto X]} = \perp\perp \cap (XX)^{\perp\perp} \cap \llbracket [\psi] \rrbracket^V$. Therefore, we need to show that $P^{\perp\perp} \subseteq \perp\perp \cap (PP)^{\perp\perp} \cap \llbracket [\psi] \rrbracket^V$ where $P = I \cap \llbracket \psi \rrbracket^V$. Let $p \in P$. Since $p \in J$, $\{p\}^\perp \subseteq \{pp\}^\perp$. Hence $P^{\perp\perp} \subseteq (PP)^{\perp\perp}$. Now, we are done by hypothesis. By duality this also covers the case $\varphi = ?\psi$. Hence done. \square

By completeness of the phase semantics of LL and $\mu\text{MALL}^{\text{ind}}$, we have Lemma 4.3.3.

5.1.3 Closure ordinals

Closure ordinals are a standard measure of the complexity of any (class of) monotone functions. The closure ordinal of a fixed point formula is essentially the closure ordinal of the corresponding monotone function in the truth semantics. In [DG02], the closure ordinal is construed as a function of the size of the finite model. The study of closure ordinals of modal logic formulas is a young and exciting area of research [Cza10, AL13, GS17]. It departs from the previous notion of closure ordinals in its model-independence. In this case, closure ordinal really serves as a measure of the complexity of a formula.

Definition 5.1.6. *Let φ be a preformula such that $x \in \text{fv}(\varphi)$. Fix a μ -phase model \mathcal{M} . We define the **closure ordinal** of φ with respect to x and the μ -phase model \mathcal{M} , denoted $\mathcal{O}_{\mathcal{M}}(\varphi)$, as the closure ordinal of $\lambda X. \llbracket \varphi \rrbracket^{V[x \mapsto X]}$. The closure ordinal of φ with respect to x (across all models) is defined as $\mathcal{O}(\varphi) := \sup_{\mathcal{M}} \{\mathcal{O}_{\mathcal{M}}(\varphi)\}$. Finally, φ is said to be **constructive** if $\mathcal{O}(\varphi) \leq \omega$.*

Spelt out more explicitly, define $\Theta_{\alpha}^{\varphi}$ for all ordinals α as follows.

$$\begin{aligned} \Theta_0^{\varphi} &= \emptyset^{\perp\perp} \\ \Theta_{\alpha+1}^{\varphi} &= \llbracket \varphi \rrbracket^{V[x \mapsto \Theta_{\alpha}^{\varphi}]} \\ \Theta_{\alpha}^{\varphi} &= \left(\bigcup_{\beta} \Theta_{\beta}^{\varphi} \right)^{\perp\perp} \quad \text{if } \alpha \text{ is a limit ordinal.} \end{aligned}$$

Then, the closure ordinal of φ with respect to x is the smallest ordinal such that $\Theta_{\alpha}^{\varphi} = \Theta_{\alpha+1}^{\varphi}$. Following the proof of Theorem 2.0.4, we can show that for any pre-formula φ , the sequence $\{\Theta_{\alpha}^{\varphi}\}_{\alpha \in \text{Ord}}$ is ultimately stationary. Consequently, for all μ -phase models \mathbb{M} , $\mathcal{O}_{\mathbb{M}}(\varphi)$ exists and $\llbracket \mu x. \varphi \rrbracket = \Theta_{\mathcal{O}_{\mathbb{M}}(\varphi)}^{\varphi}$.

Example 5.1.1. *Observe that*

$$\begin{aligned} \Theta_1^{a \&x} &= \llbracket a \&x \rrbracket^{V[x \mapsto \Theta_0^{a \&x}]} \\ &= V(a) \cap \Theta_0^{a \&x} \\ &= V(a) \cap \llbracket a \&x \rrbracket^{V[x \mapsto \Theta_0^{a \&x}]} \\ &= V(a) \cap (V(a) \cap \emptyset^{\perp\perp}) \\ &= \emptyset^{\perp\perp} \\ &= \Theta_0^{a \&x} \end{aligned}$$

Therefore, $\mathcal{O}(a \&x) = 0$.

In the tradition of μ -calculus, the name ‘constructive’ is used loosely, motivated by the observation that if $\mathcal{O}(\varphi)$ is a finite ordinal (*i.e.* strictly below ω) for any pre-formula φ , then $\mu x. \varphi$ is provably equivalent to $\varphi^{\mathcal{O}(\varphi)}(0)$. Therefore, the class of μMALL formulas with closure ordinal strictly less than ω can be embedded in MALL and enjoys several good properties like finite model property and decidability. Observe that if the interpretation of a formula in any μ -phase model is Scott-continuous, then one can show that it is constructive by mimicking the proof of Theorem 2.0.3. The converse does not hold in general. In the following section, we consider a proof system of μMALL where fixed points are approximated by their α^{th} approximation for some ordinal α .

5.2 A semantics inspired system: $\mu_\alpha \text{MALL}^+$

In this section, we devise an infinite family of wellfounded but infinitely branching calculi inspired by the phase semantics devised in the previous section.

Let φ be a preformula with $x \in \text{fv}(\varphi)$. Then it is reasonable to consider $\llbracket \mu x.\varphi \rrbracket = \Theta_{\mathcal{O}(\varphi)}^\varphi$. But, computing closure ordinals is difficult. Instead, what we can assert is that for some fixed α and for all φ , $\llbracket \mu x.\varphi \rrbracket = \Theta_\alpha^\varphi$. What we are essentially doing is approximating fixed points. If $\alpha < \mathcal{O}(\varphi)$ then we are underapproximating otherwise it is exactly the fixed point. What is the syntactic counterpart to this? Explicit approximants of fixed points in the language of the logic.

5.2.1 Setting up $\mu_\alpha \text{MALL}$

Definition 5.2.1. *Given an ordinal $\alpha \in \text{Ord}$, the set of the $\mu_\alpha \text{MALL}$ formulas is given by the following grammar.*

$$\varphi, \psi := a \mid a^\perp \mid x \mid \varphi \otimes \psi \mid \varphi \wp \psi \mid \varphi \& \psi \mid \varphi \oplus \psi \mid \mu^\beta x.\varphi \mid \nu^\beta x.\varphi \mid \mu x.\psi \mid \nu x.\psi$$

where $a \in \mathcal{A}$, $x \in \mathcal{V}$, and $\beta < \alpha$. As usual $\mathcal{A} \cap \mathcal{V} = \emptyset$.

As sets of formulas, we have that $\mu \text{MALL} \subset \mu_0 \text{MALL} \subset \mu_1 \text{MALL} \subset \dots$ and we denote the supremum of this infinite ascending chain as $\mu_\approx \text{MALL}$. Now we devise a family of calculi $\{\mu_\alpha \text{MALL}^+\}_{\alpha \in \text{Ord}}$ such that $\mu_\alpha \text{MALL}^+$ is the calculus corresponding to $\mu_\alpha \text{MALL}$ and the fixed point rule in $\mu_\alpha \text{MALL}$ is

$$\frac{\vdash \mu^\alpha x.\varphi}{\vdash \mu x.\varphi} (\mu) \quad ; \quad \frac{\vdash \nu^\alpha x.\varphi}{\vdash \nu x.\varphi} (\nu)$$

We still need to specify the inference rules for the fixed point approximants. We essentially reverse-engineer from the intuition that $\llbracket \mu^\beta x.\varphi \rrbracket = \Theta_\beta^\varphi$. Therefore, we have,

$$\frac{\vdash \Gamma, \mathbf{0}}{\vdash \Gamma, \mu^0 x.\varphi} (\mu_0) \quad ; \quad \frac{\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x]}{\vdash \Gamma, \mu^{\beta+1} x.\varphi} (\mu_{\beta+1})$$

Note that when β is a limit ordinal, Θ_β^φ is an infinitary supremum in the space of facts. Syntactically, this is analogous to an infinitary \oplus . In the following $\beta \leq \alpha$ is a limit ordinal and $\beta' < \beta$.

$$\frac{\vdash \Gamma, \mu^{\beta'} x.\varphi}{\vdash \Gamma, \mu^\beta x.\varphi} (\mu_{\beta'}^{\beta'})$$

Therefore, this is an infinite collection of inference rules. Dually, for the greatest fixed point approximants,

$$\frac{\vdash \Gamma, \top}{\vdash \Gamma, \nu^0 x.\varphi} (\nu_0) \quad ; \quad \frac{\vdash \Gamma, \varphi[\nu^\beta x.\varphi/x]}{\vdash \Gamma, \nu^{\beta+1} x.\varphi} (\nu_{\beta+1}) \quad ; \quad \frac{\{\vdash \Gamma, \nu^{\beta'} x.\varphi\}_{\beta' < \beta}}{\vdash \Gamma, \nu^\beta x.\varphi} (\nu_\beta)$$

Note that the wellfoundedness of ordinals ensures that $\mu_\alpha \text{MALL}^+$ is wellfounded for every α . However, for all $\alpha \geq \omega$, they are infinitely branching. Such infinitary systems (called Tait-style systems) where proof trees are wellfounded with possible infinite branching are well-studied in various areas of logic *viz.* arithmetic [Car37, Min78] and fixed point logics [Koz88, JKS08].

Example 5.2.1. *Let $\psi = (\mu x.a^\perp \wp x)^\perp \wp (a^{\perp p} \wp \mathbf{0})$ and $\varphi = a \otimes a \otimes y$ for some $p \in \omega$ and $\Gamma = \mu x.a^\perp \wp x, a \otimes (\nu y.\varphi)$. We will show that $\vdash \Gamma, \psi^\perp$ is provable in $\mu_\omega \text{MALL}^+$.*

$$\frac{\frac{\frac{\frac{\frac{\frac{\vdash \nu^n y.\varphi, \mu^{2n} x.a^\perp \wp x}_{n \in \omega}}{\vdash \nu^n y.\varphi, \mu x.a^\perp \wp x}_{n \in \omega}} (\mu, \mu_\omega^{2n})}{\vdash \nu y.\varphi, \mu x.a^\perp \wp x} (\nu, \nu^\omega)}{\vdash \mu^p x.a^\perp \wp x, (a^{\perp p} \wp \mathbf{0})^\perp} (\otimes)}{\vdash \mu^p x.a^\perp \wp x, \nu y.\varphi, \psi^\perp} (\otimes)}{\vdash a^\perp, a} (\text{id}) \quad \frac{\frac{\frac{\frac{\vdash a^\perp, \mu^p x.a^\perp \wp x, a \otimes (\nu y.\varphi), \psi^\perp}{\vdash a^\perp \wp (\mu^p x.a^\perp \wp x), a \otimes (\nu y.\varphi), \psi^\perp} (\wp)}{\vdash \mu^{p+1} x.a^\perp \wp x, a \otimes (\nu y.\varphi), \psi^\perp} (\mu_{p+1})}{\vdash \mu^{p+1} x.a^\perp \wp x, a \otimes (\nu y.\varphi), \psi^\perp} (\mu, \mu_\omega^{p+1})}{\vdash \Gamma, \psi^\perp} (\otimes)$$

It is easy to show that for all $p, n \in \omega$, $\vdash \mu^p x. a^\perp \wp x, (a^{\perp p} \wp \mathbf{0})^\perp$ and $\vdash \nu^n y. \varphi, \mu^{2n} x. a^\perp \wp x$ are provable by induction on p and n respectively. We will show this for the first sequent. The second one is very similar. The base case is $p = 0$ in which case $a^{\perp p} \wp \mathbf{0}$ is simply $\mathbf{0}$. Therefore, we have to prove the sequent $\vdash \mu^0 x. a^\perp \wp x, \top$ which is provable by a (\top) rule. For the induction case assume $p = q + 1$ and we have the following.

$$\frac{\frac{\frac{\overline{\vdash a, a^\perp} \text{ (id)}}{\vdash \mu^q x. a^\perp \wp x, (a^{\perp q} \wp \mathbf{0})^\perp} (\otimes)}{\vdash a^\perp, \mu^q x. a^\perp \wp x, (a^{\perp q+1} \wp \mathbf{0})^\perp} (\wp)}{\vdash a^\perp \wp (\mu^q x. a^\perp \wp x), (a^{\perp q+1} \wp \mathbf{0})^\perp} (\mu_{q+1})}{\mu^{q+1} x. a^\perp \wp x, (a^{\perp q+1} \wp \mathbf{0})^\perp} (\mu_{q+1})$$

5.2.2 The rank of a formula

It is quite tricky to define a proper notion of the complexity of a fixed point formula in such settings. Closely based on [JKS08], our notion of the rank of a μ_α MALL formula is a finite sequence of ordinals.

First we will set up some notation. If $\alpha_1, \dots, \alpha_n$ are ordinals, we write $\langle \alpha_1, \dots, \alpha_n \rangle$ for the sequence σ whose length $|\sigma|$ is n and whose i th component σ_i is the ordinal α_i . Let $<_{lex}$ be the strict lexicographical ordering of finite sequences of ordinals and \leq_{lex} its reflexive closure. Note that $<_{lex}$ is a well-ordering on any set of sequences of bounded lengths but not a well-ordering in general. In particular, $\langle 1 \rangle, \langle 0, 1 \rangle, \langle 0, 0, 1 \rangle, \dots$ is an infinite descending chain in $<_{lex}$. Given two finite sequences of ordinals σ, τ , we define the component-wise ordering \leq as σ, τ iff $|\sigma| \leq |\tau|$ and $(\sigma)_i \leq (\tau)_i$ for all $1 \leq i \leq |\sigma|$. Clearly, the relation \leq is transitive. We denote the standard concatenation of sequences by $*$. Finally, we define a component-wise maximum operation \sqcup by setting: (i) $\sigma \sqcup \langle \rangle := \langle \rangle \sqcup \sigma := \langle \rangle$; (ii) if $\sigma = \langle b_1, \dots, b_m \rangle$ and $\tau = \langle b'_1, \dots, b'_n \rangle$, then

$$\sigma \sqcup \tau = \begin{cases} \langle \max(b_1, b'_1), \dots, \max(b_m, b'_m), b'_{m+1}, \dots, b'_n \rangle & \text{if } m \leq n; \\ \langle \max(b_1, b'_1), \dots, \max(b_n, b'_n), b_{n+1}, \dots, b_m \rangle & \text{otherwise.} \end{cases}$$

Proposition 5.2.1. *For all sequences σ, σ_1 , and σ_2 , the following holds.*

$$(\sigma * \sigma_1) \sqcup (\sigma * \sigma_2) = \sigma * (\sigma_1 \sqcup \sigma_2)$$

Proof. Let $\sigma = \langle a_1, \dots, a_n \rangle$. Then,

$$\begin{aligned} (\sigma * \sigma_1) \sqcup (\sigma * \sigma_2) &= \langle \max(a_1, a_1), \dots, \max(a_n, a_n) \rangle * (\sigma_1 \sqcup \sigma_2) \\ &= \langle a_1, \dots, a_n \rangle * (\sigma_1 \sqcup \sigma_2) \\ &= \sigma * (\sigma_1 \sqcup \sigma_2) \end{aligned}$$

□

Now we are ready to define the rank of a μ_α MALL formula. The rank of every μ_α MALL formula will be a finite sequence of ordinals less than or equal to $\alpha + 1$.

Definition 5.2.2. *The rank of a μ_α MALL formula φ , denoted $\text{rk}(\varphi)$, is defined by induction on φ as follows:*

- if φ is an atom, a variable, or a unit, then $\text{rk}(\varphi) = \langle 0 \rangle$;
- if $\varphi = \psi \odot \psi'$, then $\text{rk}(\varphi) = (\text{rk}(\psi) \sqcup \text{rk}(\psi')) * \langle 0 \rangle$;
- if $\varphi = \eta^\beta x. \psi$, then $\text{rk}(\varphi) = \text{rk}(\psi) * \langle \beta \rangle$;
- if $\varphi = \eta x. \psi$, then $\text{rk}(\varphi) = \text{rk}(\psi) * \langle \alpha + 1 \rangle$.

where $\odot \in \{\wp, \otimes, \oplus, \&\}$ and $\eta \in \{\mu, \nu\}$.

Remark 5.2.1. *For all formulas φ , $|\text{rk}(\varphi)|$ is the length of the longest path from the root to a leaf in the syntax tree of φ .*

Example 5.2.2. Let $\psi = (\nu x.a^\perp \otimes x) \wp (a^p \wp \mathbf{0})$ be a μ_ω MALL formula.

$$\begin{aligned}
\text{rk}(H) &= \text{rk}(\nu x.a^\perp \otimes x) \sqcup \text{rk}(a^p \wp \mathbf{0}) * \langle 0 \rangle \\
&= \left((\text{rk}(a^\perp \otimes x) * \langle \omega \rangle) \sqcup \overbrace{\langle 0, \dots, 0 \rangle}^{p+1} \right) * \langle 0 \rangle \\
&= \left(\langle 0, 0, \omega \rangle \sqcup \overbrace{\langle 0, \dots, 0 \rangle}^{p+1} \right) * \langle 0 \rangle \\
&= \langle 0, 0, \omega, \overbrace{0, \dots, 0}^{\max(1, p-1)} \rangle
\end{aligned}$$

Lemma 5.2.1. Let φ be a μ_α MALL preformula such that $x \in \text{fv}(\varphi)$. Let ξ be a preformula such that $\text{rk}(\varphi) \leq \text{rk}(\xi)$. Then, there exists a finite (possibly empty) sequence of ordinals σ such that $\text{rk}(\varphi[\xi/x]) = \text{rk}(\xi) * \sigma$.

Proof. We induct on $|\text{rk}(\varphi)|$. The base case is when $|\text{rk}(\varphi)| = 1$. Since $x \in \text{fv}(\varphi)$, φ cannot be an atom or a unit. Therefore, $\varphi = x$. Plugging $\sigma = \langle \rangle$, we are done. The induction case has several subcases.

- Suppose $\varphi = \psi \odot \psi'$ where $\odot \in \{\otimes, \wp, \&, \oplus\}$. We have two cases. Either $x \in \text{fv}(\psi) \cap \text{fv}(\psi')$ or x is free in only one of them. Note that $\text{rk}(\psi), \text{rk}(\psi') \leq \text{rk}(\varphi)$. Therefore if x is free in them, the induction hypothesis can be fired. In the first case, we have $\text{rk}(\psi[\xi/x]) = \text{rk}(\xi) * \sigma_1$ and $\text{rk}(\psi'[\xi/x]) = \text{rk}(\xi) * \sigma_2$. Therefore,

$$\begin{aligned}
\text{rk}(\varphi[\xi/x]) &= \text{rk}(\psi[\xi/x] \odot \text{rk}(\psi'[\xi/x])) \\
&= (\text{rk}(\psi[\xi/x]) \sqcup \text{rk}(\psi'[\xi/x])) * \langle 0 \rangle \\
&= ((\text{rk}(\xi) * \sigma_1) \sqcup (\text{rk}(\xi) * \sigma_2)) * \langle 0 \rangle && \text{[By Proposition 5.2.1]} \\
&= \text{rk}(\xi) * (\sigma_1 \sqcup \sigma_2) * \langle 0 \rangle
\end{aligned}$$

Therefore by plugging $\sigma = (\sigma_1 \sqcup \sigma_2) * \langle 0 \rangle$, we are done. In the other case, wlog assume $x \notin \text{fv}(\psi')$. Therefore, we have $\psi'[\xi/x] = \psi'$. Firing the induction hypothesis for ψ , we have $\text{rk}(\psi[\xi/x]) = \text{rk}(\xi) * \sigma_1$ as before. Therefore,

$$\begin{aligned}
\text{rk}(\varphi[\xi/x]) &= \text{rk}(\psi[\xi/x] \odot \text{rk}(\psi'[\xi/x])) \\
&= (\text{rk}(\psi[\xi/x]) \sqcup \text{rk}(\psi')) * \langle 0 \rangle \\
&= ((\text{rk}(\xi) * \sigma_1) \sqcup \text{rk}(\psi')) * \langle 0 \rangle \\
&= \text{rk}(\xi) * \sigma_1 * \langle 0 \rangle && \text{[Since } \text{rk}(\psi') \leq \text{rk}(\xi)\text{]}
\end{aligned}$$

Therefore by plugging $\sigma = \sigma_1 * \langle 0 \rangle$, we are done.

- Suppose $\varphi = \eta^\beta y. \psi$ where $\eta \in \{\mu, \nu\}$, $\beta < \alpha$, and $y \neq x$. Clearly, $x \in \text{fv}(\psi)$ and $\text{rk}(\psi) \leq \text{rk}(\varphi)$. Therefore, by hypothesis, $\text{rk}(\psi[\xi/x]) = \text{rk}(\xi) * \sigma'$.

$$\begin{aligned}
\text{rk}(\varphi[\xi/x]) &= \text{rk}(\eta^\beta y. \psi[\xi/x]) \\
&= \text{rk}(\psi[\xi/x]) * \langle \beta \rangle \\
&= (\text{rk}(\xi) * \sigma') * \langle \beta \rangle
\end{aligned}$$

By plugging $\sigma = \sigma' * \langle \beta \rangle$, we are done. The case when $\varphi = \eta y. \psi$ goes exactly similarly. □

Theorem 5.2.1. The following hold for any μ_α MALL formulas φ :

1. $\text{rk}(\varphi) <_{lex} \text{rk}(\varphi \odot \psi)$ and $\text{rk}(\psi) <_{lex} \text{rk}(\varphi \odot \psi)$;
2. $\text{rk}(\mathbf{0}) <_{lex} \text{rk}(\mu^0 x. \varphi)$, $\text{rk}(\top) <_{lex} \text{rk}(\nu^0 x. \varphi)$;

3. $\text{rk}(\varphi[\eta^\beta x.\varphi/x]) <_{lex} \text{rk}(\eta^{\beta+1}x.\varphi)$ for all $\beta < \alpha$;
4. $\text{rk}(\eta^{\beta'}x.\varphi) <_{lex} \text{rk}(\eta^\beta x.\varphi)$ for all $\beta' < \beta \leq \alpha$;
5. $\text{rk}(\eta^\alpha x.\varphi) <_{lex} \text{rk}(\eta x.\varphi)$.

Proof. The first, second, fourth, and fifth assertions are immediate from Definition 5.2.2. For the third one, we have two cases.

Case 1: $x \notin \text{fv}(\varphi)$. In this case, $\varphi[\eta^\beta x.\varphi/x] = \varphi$. Then the result follows by definition of rank.

Case 2: $x \in \text{fv}(\varphi)$. By definition, $\text{rk}(\varphi) \trianglelefteq \text{rk}(\eta^\beta x.\varphi)$. By Lemma 5.2.1, $\text{rk}(\varphi[\eta^\beta x.\varphi/x]) = \text{rk}(\eta^\beta x.\varphi) * \sigma = \text{rk}(\varphi) * \langle \beta \rangle * \sigma$ for some σ . But $\text{rk}(\eta^{\beta+1}x.\varphi) = \text{rk}(\varphi) * \langle \beta + 1 \rangle$. Hence we are done. \square

Definition 5.2.3. The strong closure $\text{SC}(\varphi)$ of a μ_α MALL formula φ is the least set such that:

- $\varphi \in \text{SC}(\varphi)$;
- $\psi \odot \psi' \in \text{SC}(\varphi) \implies \{\psi, \psi'\} \subset \text{SC}(\varphi)$ where $\odot \in \{\wp, \otimes, \oplus, \&\}$;
- $\mu^0 x.\psi \in \text{SC}(\varphi) \implies \mathbf{0} \in \text{SC}(\varphi)$;
- $\nu^0 x.\psi \in \text{SC}(\varphi) \implies \mathbf{T} \in \text{SC}(\varphi)$;
- $\eta^{n+1}x.\psi \in \text{SC}(\varphi) \implies \psi(\eta^n x.\psi/x) \in \text{SC}(\varphi)$ for all $n \in \omega$ and $\eta \in \{\mu, \nu\}$;
- $\eta x.\psi \in \text{SC}(\varphi) \implies \eta^\beta x.\psi \in \text{SC}(\varphi)$ for all $\beta \leq \alpha$ and $\eta \in \{\mu, \nu\}$.

Define φ^- to be image of φ under the forgetful functor that erases the explicit approximations occurring in φ . For example, $(a \otimes \mu^6 x.\nu y.x \oplus y)^- = a \otimes \mu x.\nu y.x \oplus y$.

Theorem 5.2.2. For any formula φ , the set $\{\text{rk}(\psi) \mid \psi \in \text{SC}(\varphi)\}$ is a well-order with respect to the $<_{lex}$ ordering.

Proof. By contradiction. Note that $|\text{rk}(\psi)| = |\text{rk}(\psi^-)|$. Furthermore, if $\psi \in \text{SC}(\varphi)$ then $\psi^- \in \text{FL}(\varphi)$. But $\text{FL}(\varphi)$ is a finite set, therefore the set $\{|\text{rk}(\psi)| \mid \psi \in \text{SC}(\varphi)\}$ is finite.

Assume there exists $\{\sigma_i\}_{i \in I}$ an infinite descending chain in $\{|\text{rk}(\psi)| \mid \psi \in \text{SC}(\varphi)\}$. By the Infinite Ramsey Theorem, there is an infinite subsequence $\{\sigma_{i_j}\}_{j \in \omega}$ such that for all $j, i_j \in I$ and for all j, j' , $|\sigma_{i_j}| = |\sigma_{i_{j'}}| = n$. Then, there exist $k \leq n$ and $N \in \mathbb{N}$ such that $\{(\sigma_{i_j})_k\}_{j > N}$ is a descending chain. This contradicts the wellfoundedness of natural numbers. \square

We conclude this subsection by exhibiting the use of ranks. Namely, we prove the functoriality property in μ_α MALL⁺.

Theorem 5.2.3. The following rule is derivable in μ_α MALL⁺ when $x \in \text{fv}(\varphi)$.

$$\frac{\vdash \psi^\perp, \psi'}{\vdash \varphi^\perp[\psi^\perp/x], \varphi[\psi'/x]} \text{ (func)}$$

Proof. We will induct on the rank of φ (this is wellfounded by Theorem 5.2.2). The base case is $\varphi = x$. In this case **func** is a trivial rule with identical premisses and conclusion. There are several subcases for the induction step. The subcases for the multiplicative additive connectives follow from Proposition 3.4.1. We will exhibit the subcase when the outermost operator of φ is a fixed point. Suppose $\varphi = \mu^\beta y.\xi$. Clearly $x \neq y$. There are two cases.

Case 1: β is successor ordinal i.e. $\beta = \gamma + 1$. We have

$$\begin{array}{c} \text{IH} \\ \nabla \\ \frac{\vdash \xi^\perp[\psi^\perp/x, \nu^\gamma y.\xi^\perp/y], \xi[\psi'/x, \mu^\gamma y.\xi^\perp/y]}{\vdash \xi^\perp[\psi^\perp/x, \nu^\gamma y.\xi^\perp/y], \mu^{\gamma+1} y.\xi[\psi'/x]} (\mu_{\gamma+1}) \\ \frac{\vdash \xi[\psi^\perp/x, \nu^\gamma y.\xi^\perp/y], \mu^{\gamma+1} y.\xi[\psi'/x]}{\vdash \nu^{\gamma+1} y.\xi^\perp[\psi^\perp/x], \mu^{\gamma+1} y.\xi[\psi'/x]} (\nu_{\gamma+1}) \end{array}$$

The induction hypothesis can be applied since we are ensured the rank decreases using Theorem 5.2.1.2.

Case 2: β is a limit ordinal. We have

$$\frac{\text{IH} \quad \frac{\frac{\{\vdash \nu^{\beta'} y. \xi^\perp[\psi^\perp/x], \nu^{\beta'} y. \xi^\perp[\psi'/x]\}_{\beta' < \beta}}{\vdash \nu^{\beta'} y. \xi^\perp[\psi^\perp/x], \mu^{\beta'} y. \xi[\psi'/x]_{\beta' < \beta}} (\mu_{\beta'}^{\beta'})}{\vdash \nu^\gamma y. \xi^\perp[\psi^\perp/x], \mu^\beta y. \xi[\psi'/x]} (\nu_\beta)}$$

The case when $\varphi = \nu^\beta y. \xi$ is symmetric and the case when $\varphi = \eta x. \xi$ is trivial for $\eta = \{\mu, \nu\}$. This concludes the proof. \square

5.2.3 Soundness and completeness of $\mu_\alpha \text{MALL}^+$

The phase semantics for $\mu_\alpha \text{MALL}^+$ is much simpler to define. Like MALL the semantics can be defined given a phase space and a valuation (without the extra structure over the set of facts and extension of valuations to variables as in $\mu \text{MALL}^{\text{ind}}$).

We recall that given a phase space \mathcal{M} , its set of facts is denoted by $\mathcal{X}_\mathcal{M}$. Fix a valuation $V : \mathcal{A} \rightarrow \mathcal{X}_\mathcal{M}$. The interpretation of $\mu_\alpha \text{MALL}^+$ is an extension of the interpretation of MALL . Hence we only need to specify the interpretation of the fixed point operators.

$$\begin{aligned} \llbracket \mu^0 x. \varphi \rrbracket^V &= \emptyset^{\perp\perp} \\ \llbracket \nu^0 x. \varphi \rrbracket^V &= \top^{\perp\perp} \\ \llbracket \eta^{\beta+1} x. \varphi \rrbracket^V &= \llbracket \varphi[\eta^\beta x. \varphi/x] \rrbracket^V && [\eta \in \{\mu, \nu\}] \\ \llbracket \mu^\beta x. \varphi \rrbracket^V &= \left(\bigcup_{\beta' < \beta} \llbracket \mu^{\beta'} x. \varphi \rrbracket^V \right)^{\perp\perp} && [\beta \text{ is a limit ordinal.}] \\ \llbracket \nu^\beta x. \varphi \rrbracket^V &= \left(\bigcap_{\beta' < \beta} \llbracket \nu^{\beta'} x. \varphi \rrbracket^V \right) && [\beta \text{ is a limit ordinal.}] \\ \llbracket \eta x. \varphi \rrbracket^V &= \llbracket \eta^\alpha x. \varphi \rrbracket^V && [\eta \in \{\mu, \nu\}] \end{aligned}$$

We will now prove the soundness and completeness of this interpretation. Soundness is obtained by a straightforward induction on the structure of the proof. Completeness depends on the wellfoundedness of rank.

Theorem 5.2.4 (Soundness for $\mu_\alpha \text{MALL}^+$). *If $\vdash \Gamma$ then for all phase models (\mathcal{M}, V) , $1 \in \llbracket \Gamma \rrbracket^V$.*

Proof. Let π be a proof of $\vdash \Gamma$ as usual in LL. We will induct on the structure of π . The base case (when the proof is just an application of the (id), (1), or (T) rule) is easily taken care of. For the induction case, consider the rule applied at the root of π . If it is a MALL operator, then the proof follows exactly like that of Theorem 3.3.1. Therefore, we describe only fixed point cases.

- Suppose $\Gamma = \Gamma', \varphi$ such that φ is principal, and either $\varphi = \eta x. \psi$ or $\varphi = \eta^0 x. \psi$ or $\varphi = \eta^{\beta+1} x. \psi$ for $\eta = \{\mu, \nu\}$. Then, the premisses are of the form $\vdash \Gamma', \varphi'$ where $\llbracket \varphi \rrbracket^V = \llbracket \varphi' \rrbracket^V$. By hypothesis, $1 \in \llbracket \Gamma', \varphi' \rrbracket^V$ which implies $1 \in \llbracket \Gamma', \varphi \rrbracket^V$.
- Suppose $\Gamma = \Gamma', \mu^\lambda x. \varphi$ such that $\mu^\lambda x. \varphi$ is principal, and λ is a limit ordinal. Suppose the rule applied is (μ_λ^β) for some $\beta < \lambda$. Then the premisses are $\vdash \Gamma', \mu^{\beta'} x. \varphi$ and by hypothesis, $1 \in \llbracket \Gamma', \mu^{\beta'} x. \varphi \rrbracket^V$. By Proposition 3.3.2, this is equivalent to $(\llbracket \Gamma \rrbracket^V)^\perp \subseteq \llbracket \mu^{\beta'} x. \varphi \rrbracket^V$. Now, $\llbracket \mu^{\beta'} x. \varphi \rrbracket^V \subseteq \bigcup_{\beta' < \lambda} \llbracket \mu^{\beta'} x. \varphi \rrbracket^V$. Taking double negations on both sides by Proposition 3.3.1.3, we have, $(\llbracket \Gamma \rrbracket^V)^\perp \subseteq \llbracket \mu^\lambda x. \varphi \rrbracket^V$. By Proposition 3.3.2, we are done.
- Suppose $\Gamma = \Gamma', \nu^\lambda x. \varphi$ such that $\nu^\lambda x. \varphi$ is principal, and λ is a limit ordinal. For all premisses $\vdash \Gamma', \nu^{\beta'} x. \varphi$ apply hypothesis to get $1 \in \llbracket \Gamma', \nu^{\beta'} x. \varphi \rrbracket^V$. By Proposition 3.3.2, this is equivalent to $(\llbracket \Gamma \rrbracket^V)^\perp \subseteq \llbracket \nu^{\beta'} x. \varphi \rrbracket^V$. Therefore, $(\llbracket \Gamma \rrbracket^V)^\perp \subseteq \llbracket \nu^\lambda x. \varphi \rrbracket^V$. By Proposition 3.3.2, we are done.

\square

Lemma 5.2.2 (Adequation Lemma for $\mu_\alpha\text{MALL}^+$). *For all formulas φ , $\llbracket\varphi\rrbracket^V \subseteq \text{Pr}(\varphi)$.*

Proof. By induction on $\text{rk}(\varphi)$. The base case is when φ is an atom or a unit in which case by definition $\llbracket\varphi\rrbracket^V = \text{Pr}(\varphi)$. There are several subcases for the induction case. For the **MALL**, one can reuse the proof of Lemma 3.3.1. The only thing that one needs to observe is that the measure of the complexity of a formula has changed but thanks to Theorem 5.2.1 we are safe. In this proof, we only tackle the fixed point cases.

- It is trivial for formulas of the form $\eta^0 x.\psi$ where $\eta \in \{\mu, \nu\}$.
- If $\varphi = \eta^{\beta+1} x.\psi$, then

$$\begin{aligned} \llbracket\varphi\rrbracket^V &= \llbracket\psi[\eta^\beta x.\psi/x]\rrbracket^V \\ &\subseteq \text{Pr}(\psi[\eta^\beta x.\psi/x]) \quad [\text{IH since } \text{rk}(\psi[\eta^\beta x.\psi/x]) <_{lex} \text{rk}(\varphi) \text{ by Theorem 5.2.1}] \\ &\subseteq \text{Pr}(\varphi) \end{aligned}$$

- If $\varphi = \mu^\beta x.\psi$ such that β is a limit ordinal, then

$$\begin{aligned} \llbracket\varphi\rrbracket^V &= \left(\bigcup_{\beta' < \beta} \llbracket\mu^{\beta'} x.\varphi\rrbracket^V \right)^{\perp\perp} \\ &\subseteq \left(\bigcup_{\beta' < \beta} \text{Pr}(\psi[\eta^{\beta'} x.\psi/x]) \right)^{\perp\perp} \quad [\text{IH since } \text{rk}(\psi[\eta^{\beta'} x.\psi/x]) <_{lex} \text{rk}(\varphi) \text{ by Theorem 5.2.1}] \end{aligned}$$

Now, $\text{Pr}(\psi[\eta^{\beta'} x.\psi/x]) \subseteq \text{Pr}(\varphi)$ for all $\beta' < \beta$. Therefore, $\bigcup_{\beta' < \beta} \text{Pr}(\psi[\eta^{\beta'} x.\psi/x]) \subseteq \text{Pr}(\varphi)$. By applying Proposition 3.3.1, we have $\left(\bigcup_{\beta' < \beta} \text{Pr}(\psi[\eta^{\beta'} x.\psi/x])\right)^{\perp\perp} \subseteq \text{Pr}(\varphi)$.

- The case when $\varphi = \nu^\beta x.\psi$ and β is a limit ordinal goes similarly as above. The case when $\varphi = \eta x.\psi$ is trivial where $\eta \in \{\mu, \nu\}$.

This concludes the proof. \square

As usual we get the completeness from the adequation lemma exactly in the same way as Theorem 3.3.3.

Theorem 5.2.5 (Completeness for $\mu_\alpha\text{MALL}^+$). *If for any phase model (\mathcal{M}, V) , $1 \in \llbracket\Gamma\rrbracket^V$ then $\vdash \Gamma$.*

Theorem 5.2.5 gives cut admissibility for free as usual in phase semantics [Oka96, Oka99].

Corollary 5.2.5.1. *For all α , $\mu_\alpha\text{MALL}^+$ admits cuts.*

However, note that Corollary 5.2.5.1 does not shed any light on the cut-elimination procedure. In particular, it is not discernible whether the cut-elimination equivalence equates all proofs of a particular sequent. One can consider more constructive cut-elimination with explicit reduction sequences. The new reduction rules are quite straightforward (β being a limit ordinal in the latter):

$$\begin{array}{c} \frac{\frac{\frac{\pi_\beta}{\triangle} \quad \frac{\pi'_\beta}{\triangle}}{\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x]} (\mu_{\beta+1}) \quad \frac{\frac{\pi'_\beta}{\triangle}}{\vdash \Delta, \varphi^\perp[\nu^\beta x.\varphi^\perp/x]} (\nu_{\beta+1})}{\vdash \Gamma, \Delta} (\text{cut})}{\vdash \Gamma, \Delta} (\text{cut}) \quad \longrightarrow \quad \frac{\frac{\pi_\beta}{\triangle} \quad \frac{\pi'_\beta}{\triangle}}{\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x] \quad \vdash \Delta, \varphi^\perp[\nu^\beta x.\varphi^\perp/x]} (\text{cut})}{\vdash \Gamma, \Delta} (\text{cut}) \\ \\ \frac{\frac{\frac{\pi_\gamma}{\triangle} \quad \frac{\pi'_\gamma}{\triangle}}{\vdash \Gamma, \mu^\gamma x.\varphi} (\mu_\beta^\gamma) \quad \frac{\{\vdash \Delta, \nu^\gamma x.\varphi^\perp\}_{\gamma < \beta}}{\vdash \Delta, \nu^\beta x.\varphi^\perp} (\nu_\beta)}{\vdash \Gamma, \Delta} (\text{cut})}{\vdash \Gamma, \Delta} (\text{cut}) \quad \longrightarrow \quad \frac{\frac{\pi_\gamma}{\triangle} \quad \frac{\pi'_\gamma}{\triangle}}{\vdash \Gamma, \mu^\gamma x.\varphi \quad \vdash \Delta, \nu^\gamma x.\varphi^\perp} (\text{cut})}{\vdash \Gamma, \Delta} (\text{cut}) \end{array}$$

However, standard cut-elimination techniques for finitary proof theory (say, for example, notions of cut rank based on the depth of the topmost cut) do not work for Tait-like systems since there could be infinitely many cuts in a proof each at a height higher than the next. The reduction sequences shall be potentially infinite and fair reduction sequences shall ensure wellfoundedness of the limit (analogous to ensuring productivity in non-wellfounded settings).

There are however well-known techniques to get cut admissibility in Tait-like systems, the first being due to Schütte [Sch68] where proofs are assigned a cut rank. One shows that if there is a proof π of a sequent $\vdash \Gamma$ with cut-rank $\text{rk}(\pi) > 0$ then there is a proof π' of $\vdash \Gamma$ such that $\text{rk}(\pi') = 0$ (possibly incurring a blowup in the size of the proof). Cut-admissibility has been proved previously for Tait-style systems of fixed point logics in [Pal07, BS12] using this technique. Semantic proofs of cut-admissibility have been explored in various logics [Sch60, Mae91, Avi01] but to our knowledge, this is the first semantic proof of cut-admissibility in a Tait-style system.

5.3 Properties of $\mu_\omega\mathbf{MALL}^+$

We will now explore the provability of $\mu_\alpha\mathbf{MALL}^+$ with special attention to $\mu_\omega\mathbf{MALL}^+$. Note that for any α, β such that $\alpha \leq \beta$, we have that $\llbracket \mu^\alpha x.\varphi \rrbracket \subseteq \llbracket \mu^\beta x.\varphi \rrbracket$ and $\llbracket \nu^\beta x.\varphi \rrbracket \subseteq \llbracket \nu^\alpha x.\varphi \rrbracket$. This intuition can be exploited to prove that if $\alpha \neq \beta$, then $\mu_\alpha\mathbf{MALL}^+ \neq \mu_\beta\mathbf{MALL}^+$ as sets of theorems. Wlog, assume $\alpha < \beta$. Then, $\vdash \mu x.a^\perp \wp x, \nu^\beta y.a \otimes y$ is provable in $\mu_\beta\mathbf{MALL}^+$ but not in μ_α . Similarly, $\vdash \mu^\alpha x.a^\perp \wp x, \nu y.a \otimes y$ is provable in $\mu_\alpha\mathbf{MALL}^+$ but not in $\mu_\beta\mathbf{MALL}^+$. The proofs in both situations are trivial. But how does one prove that a sequent is *not* provable?

There are a few techniques in the bag of tricks of linear logic. Firstly, one can use semantics means to show that there is a phase model where the interpretation of the sequent does not contain the unit of the monoid. This is difficult in our situation since we have the formulas $\mu x.a^\perp \wp x$ and $\mu^\alpha x.a^\perp \wp x$ in the sequents. Not only are we not able to compute their closure ordinal, we are also unable to find a phase model where the fixed point iteration sequence does not collapse at 0. It is not possible to work with such trivialising examples.

However, all is not lost as the phase semantics does provide cut admissibility. Armed with Corollary 5.2.5.1, one can provide a more syntactic proof-search argument. These techniques will be used several times in the next chapter but we will already get a flavour here. We begin by showing that the (μ^α) is invertible if α is a successor ordinal.

Proposition 5.3.1. *Let $\beta + 1 \leq \alpha$. If $\vdash \Gamma, \mu^{\beta+1}x.\varphi$ is provable in $\mu_\alpha\mathbf{MALL}^+$ then so is $\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x]$.*

Proof. Let π be a proof of $\vdash \Gamma, \mu^{\beta+1}x.\varphi$. We have the following.

$$\frac{\frac{\frac{\pi}{\vdash \Gamma, \mu^{\beta+1}x.\varphi} \quad \frac{\frac{\frac{\overline{\vdash \nu^\beta x.\varphi^\perp, \mu^\beta x.\varphi}}{\vdash \varphi^\perp[\nu^\beta x.\varphi^\perp/x], \varphi[\mu^\beta x.\varphi/x]} \text{(func)}}{\vdash \nu^{\beta+1}x.\varphi^\perp, \varphi[\mu^\beta x.\varphi/x]} \text{(\nu_{\gamma+1})}}{\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x]} \text{(cut)}}{\vdash \Gamma, \varphi[\mu^\beta x.\varphi/x]} \text{(id)}$$

□

Lemma 5.3.1. *Suppose $\alpha < \beta$. Then,*

1. $\vdash \mu x.a^\perp \wp x, \nu^\beta y.a \otimes y$ is not provable in $\mu_\alpha\mathbf{MALL}^+$.
2. $\vdash \mu^\alpha x.a^\perp \wp x, \nu y.a \otimes y$ is not provable in $\mu_\beta\mathbf{MALL}^+$.

Proof. We will only prove (1). (2) follows very similarly. Let (α, β) be the smallest pair of ordinals (in lexicographic ordering) such that $\alpha < \beta$ and $\vdash \mu x.a^\perp \wp x, \nu^\beta y.a \otimes y$ is provable in $\mu_\alpha\mathbf{MALL}^+$. Observe that there is no proof where only subformulas of $\mu x.a^\perp \wp x$ is active. Therefore, the (ν^β) rule is applied at some point. Since, (\wp) is invertible, wlog, one can assume that the conclusion of the (ν^β) rule is of the form $\vdash a^{\perp n}, \mu^{\alpha'} x.a^\perp \wp x, \nu^\beta y.a \otimes y$ for some $\alpha' < \alpha$. If β is a limit ordinal, then consider any premiss $\vdash \Delta$ where $\Delta = a^{\perp n}, \mu^{\alpha'} x.a^\perp \wp x, \nu^{\beta'} y.a \otimes y$ and $\beta' < \beta$. By Proposition 5.3.1, if Δ is provable, so is $\mu^{\alpha'+n} x.a^\perp \wp x, \nu^{\beta'} y.a \otimes y$. Since $(\alpha' + n, \beta') <_{lex} (\alpha, \beta)$, we have a contradiction! □

Consequently we have the following theorem.

Theorem 5.3.1. *For any two distinct $\alpha, \beta \in \text{Ord}$, $\mu_\alpha\mathbf{MALL}^+$ and $\mu_\beta\mathbf{MALL}^+$ are orthogonal systems i.e. as sets of theorems neither $\mu_\alpha\mathbf{MALL}^+ \subseteq \mu_\beta\mathbf{MALL}^+$ nor $\mu_\alpha\mathbf{MALL}^+ \supseteq \mu_\beta\mathbf{MALL}^+$.*

For the rest of this section, let $\Gamma = \mu x.a^\perp \wp x, a \otimes (\nu y.a \otimes a \otimes y)$ and $\psi = (\mu x.a^\perp \wp x)^\perp \wp (a^{\perp p} \wp \mathbf{0})$.

Lemma 5.3.2. $\vdash \Gamma$ is not provable in $\mu_\omega\mathbf{MALL}^+$.

Proof. Suppose there is a proof. By Corollary 5.2.5.1, we can assume that this proof is cut-free. Therefore, the only possibilities for the first rule are (\otimes) or (μ) followed by one of $\{(\mu_\omega^n)\}_{n \in \omega}$. If it is the former, then the left premiss has to contain $\mu x.a^\perp \wp x$ since $\vdash a$ cannot be proved. Consequently, the right premiss is $\vdash \nu y.\varphi$. If it were provable, so would be $\vdash \nu^n x.\varphi$ for all $n \in \omega$. It is easy to observe that $\vdash \nu^n y.\varphi$ is not provable for all $n > 0$.

Now suppose the first rule is (μ) followed by (μ_ω^n) for some $n \in \omega$. Since the (\wp) rule and, by Proposition 5.3.1, the (μ^n) rule is invertible for all $n \in \omega$, it suffices to show that $\vdash (a^\perp)^n, a \otimes \nu y.\varphi$

is not provable. The only possible rule here is the (\otimes) . The only splitting of the context which renders the left premiss provable is one where the right premiss is $\vdash \overbrace{a^\perp, \dots, a^\perp}^{n-1}, \nu y. \varphi$. The only possible rule that can be applied here is the (ν^ω) rule. Consider any premiss other than the $\lfloor \frac{n-1}{2} \rfloor$ th premiss. Since the number of a and a^\perp are different in it, it is not provable by Example 3.3.2. \square

By Example 4.2.1 $\vdash \Gamma$ is provable in $\mu\text{MALL}^{\text{ind}}$ and by Lemma 5.3.2, it is not provable in $\mu_\omega\text{MALL}^+$. Consequently, we have the following.

Theorem 5.3.2. $\mu_\omega\text{MALL}^+$ does not prove the same theorems as $\mu\text{MALL}^{\text{ind}}$.

The opposite direction *i.e.* the (non)existence of a sequent that is provable in $\mu_\omega\text{MALL}^+$ but is not provable in $\mu\text{MALL}^{\text{ind}}$ is more difficult to establish. In fact, the Park's coinduction can be simulated in the (ν_ω) rule.

$$\frac{\{\vdash \Gamma, \nu^n x. \varphi\}_{n \in \omega}}{\vdash \Gamma, \nu^\omega x. \varphi} (\nu_\omega) \quad ; \quad \frac{\vdash \Gamma, \psi \quad \frac{\vdash \psi^\perp, \varphi[\psi/x] \quad \vdash \varphi^\perp[\psi^\perp/x], \nu^n x. \varphi}{\vdash \psi^\perp, \nu^n x. \varphi} (\text{cut})}{\vdash \Gamma, \nu^n x. \varphi} (\text{cut})$$

Now we will show that $\vdash \varphi^\perp[\psi^\perp/x], \nu^n x. \varphi$ is provable by induction on n . The base case is trivial. Applying (ν^0) , we get a premiss with \top and hence done. Otherwise,

$$\begin{array}{c} \text{IH} \\ \nabla \\ \frac{\vdash \varphi^\perp, \varphi[\psi/x] \quad \vdash \varphi^\perp[\psi^\perp/x], \nu^{n-1} x. \varphi}{\vdash \psi^\perp, \nu^{n-1} x. \varphi} (\text{cut}) \\ \frac{\vdash \psi^\perp, \nu^{n-1} x. \varphi}{\vdash \varphi^\perp[\psi^\perp/x], \varphi[\nu^{n-1} x. \varphi/x]} (\text{func}) \\ \frac{\vdash \varphi^\perp[\psi^\perp/x], \varphi[\nu^{n-1} x. \varphi/x]}{\vdash \varphi^\perp[\psi^\perp/x], \nu^n x. \varphi} (\nu_n) \end{array}$$

Furthermore, observe that $\vdash \Gamma$ is provable in $\mu_{\omega+1}\text{MALL}^+$. We conjecture the following.

Open Question

If a sequent $\vdash \Delta$ is provable in $\mu\text{MALL}^{\text{ind}}$, then there exists α such that $\vdash \Delta$ is provable in $\mu_\alpha\text{MALL}^+$. In other words, $\mu\text{MALL}^{\text{ind}} \subseteq \mu_{\approx}\text{MALL}^+$.

In fact, Theorem 5.3.2 is revealing. It shows that formulas of the form $\mu x. \varphi(x) \wp \psi(x)$ cannot contract at will in $\mu_\omega\text{MALL}$. The number of times they can be contracted (say, n) is determined the moment the (μ_ω^n) rule is fired. This is reminiscent of the *multiplexing* rule of light linear logics [Gir98, Asp98, AR02, Laf04]. Recall that exponentials are encoded by fixed point formulas as follows.

$$[?\varphi] = \mu x. \perp \oplus [\varphi] \oplus (x \wp x) \quad ; \quad [!\varphi] = \nu x. \mathbf{1} \& [\varphi] \& (x \otimes x)$$

Now the usual rules for exponentials consist of dereliction, contraction, weakening, and promotion. However, it has been observed that, one can replace them with the following set of rules without losing provability.

$$\frac{\vdash \Gamma, \varphi}{\vdash ?\Gamma, !\varphi} (\text{sp}) \quad \frac{\vdash \Gamma \vdash \overbrace{\varphi, \dots, \varphi}^n}{\Gamma \vdash ?\varphi} (\text{mpx}) \quad \frac{\vdash \Gamma, ??\varphi}{\vdash \Gamma, ?\varphi} (\text{dig})$$

These are called *functorial promotion*, *multiplexing*, and *digging*. Without the digging rule, these are called **soft exponentials** and the corresponding logic is called **soft linear logic** (SLL) [Laf04]. It is straightforward to see that the soft promotion and multiplexing rule can be encoded in $\mu_\omega\text{MALL}^+$.

Open Question

Digging with respect to the μMALL encoding of exponentials is not admissible in $\mu_\omega\text{MALL}^+$.

Soft linear logic is one of the several light linear logics that have been designed with implicit complexity motivations. The phase semantics of SLL was described in [DM04]. Following [Laf97], they further showed that SLL does not have finite model property [DM04]. If we could show that $\mu_\omega \text{MALL}^+$ faithfully encodes soft linear logic, then $\mu_\omega \text{MALL}^+$ would inherit SLL's lack of finite model property. However, as with the encoding of LL in $\mu \text{MALL}^{\text{ind}}$, it is difficult to show that the encoding of SLL in $\mu_\omega \text{MALL}^+$ is faithful. However, as it turns out, it is quite simple to give a direct proof exploiting the above (non)provability results.

Theorem 5.3.3. $\mu_\omega \text{MALL}^+$ does not have finite model property.

Proof. The proof goes by contradiction. A finite phase model (\mathcal{M}, V) has finitely many facts. Therefore, by the Pigeonhole Principle, there exists p, q such that $p < q$ and $\llbracket a^p \wp \mathbf{0} \rrbracket^V = \llbracket a^q \wp \mathbf{0} \rrbracket^V$. Therefore, $\llbracket \mu x. a \wp x \rrbracket^V = \llbracket a^p \wp \mathbf{0} \rrbracket^V$. Consequently, $1 \in \llbracket \psi \rrbracket^V$ where $\psi = (\mu x. a \wp x)^\perp \wp (a^p \wp \mathbf{0})$. By Theorem 5.2.5, $\vdash \psi$ is provable in μMALL_ω^+ . In Example 5.2.1, we show that $\vdash \Gamma, \psi^\perp$ is provable in $\mu_\omega \text{MALL}^+$. By an application of the cut-rule, we have $\vdash \Gamma$ is provable in $\mu_\omega \text{MALL}^+$. This is a contradiction by Lemma 5.3.2. \square

Note that for all $n \in \omega$, $\mu_n \text{MALL}^+$ can be embedded in MALL and enjoys several good properties like finite model property and decidability. Therefore, Theorem 5.3.3 shows the ω is the smallest ordinal for which $\mu_\alpha \text{MALL}^+$ is a non-trivial system.

5.4 Towards phase semantics of non-wellfounded systems

Observe that since $\mu\text{MALL}^{\text{ind}} \subseteq \mu\text{MALL}^{\circ} \subseteq \mu\text{MALL}^{\infty}$, μ -phase models are complete for μMALL° and μMALL^{∞} . However, we will see in the next chapter that they prove different set of theorems; hence their interpretations are not sound. Soundness proof in non-wellfounded systems usually goes by contradiction: one assumes that there is a non-wellfounded proof of a sequent Γ but its interpretation is not sound (in our case that would amount to asserting $1 \notin \llbracket \Gamma \rrbracket$). From this one extracts a chain which either contradicts the progress condition of π or some wellfoundedness condition in the model. Neither of these are clear from a phase semantics interpretation.

The first idea is to rehash the constructive soundness proof of [DP17] in our setting. Since the progress condition in the presence of unbounded interleaving of fixed points is quite complicated, the adaptation is not straightforward.

We consider a strictly larger system $\mu_{\approx}\text{MALL}^{\infty}$. The language is $\mu_{\approx}\text{MALL}$. We have the usual (μ) (respectively, (ν)) unfolding rules for formulas of the form $\mu x.\varphi$ (respectively, $\nu x.\varphi$). For fixed point approximants we have the rules as in Section 5.2. This system is non-wellfounded and infinitely branching. Note that μMALL^{∞} is a subset of $\mu_{\approx}\text{MALL}^{\infty}$ as sets of theorems.

Lemma 5.4.1. *If $\vdash \Delta, \nu x.\varphi$ is provable in $\mu_{\approx}\text{MALL}^{\infty}$, then $\vdash \Delta, \nu^{\alpha}x.\varphi$ is also provable in $\mu_{\approx}\text{MALL}^{\infty}$ for any $\alpha \in \text{Ord}$.*

TO prove this the idea is to consider an infinitary η -expansion of the proof of $\vdash \Delta, \nu x.\varphi$ and then construct the proof of $\vdash \Delta, \nu^{\alpha}x.\varphi$ by transfinite induction on α .

The idea is to infer the soundness of μMALL^{∞} from the soundness of $\mu_{\approx}\text{MALL}^{\infty}$. Note that we define the interpretation of $\mu x.\varphi$ (respectively, of $\nu x.\varphi$) as the least (respectively, greatest) fixed point of $\llbracket \varphi \rrbracket$. Now we lift the notion of ranks to $\mu_{\approx}\text{MALL}$ by defining $\text{rk}(\eta x.\varphi) = \text{rk}(\varphi) * \langle \star \rangle$ where $\alpha < \star$ for all $\alpha \in \text{Ord}$. Furthermore, one can define the rank of a sequent as the multiset of the rank of the formulas in the sequent.

Proof idea for soundness of μMALL^{∞} . Let π be a $\mu_{\approx}\text{MALL}^{\infty}$ proof and let us proceed by induction on the rank of the conclusion sequent Γ . For every infinite branch β of π , let $\nu x.\varphi$ be the minimal formula that occurs infinitely often. Consider the rule which introduces $\nu x.\varphi$ for the first time; their conclusions form a bar B through the infinite tree of π . The prefix closure of B must be finite by the progress condition and thus, if each of the sequents of B is valid then so is the conclusion of π by the soundness of well-founded $\mu_{\approx}\text{MALL}^{\infty}$ derivations. The soundness of wellfounded $\mu_{\approx}\text{MALL}^{\infty}$ proofs can be obtained by combining results from Section 5.1 and Section 5.2. Now consider a subproof π' that derives a sequent in B . This sequent must have the form $\vdash \Delta, \nu x.\varphi$ where $\nu x.\varphi$ is principal for the concluding (ν)-rule of π' . Along any branch rank of sequent decreases, so $\text{rk}(\Delta, \nu x.\varphi) \leq \text{rk}(\Gamma)$. Now, by Lemma 5.4.1, π' can be transformed into a proof π'' of $\Delta, \nu^{\alpha}x.\varphi$ where α is the closure ordinal of φ . Now $\text{rk}(\Delta, \nu^{\alpha}x.\varphi) < \text{rk}(\Delta, \nu x.\varphi) \leq \text{rk}(\Gamma)$. By the induction hypothesis, $1 \in \llbracket \Delta, \nu^{\alpha}x.\varphi \rrbracket$. But $\llbracket \Delta, \nu^{\alpha}x.\varphi \rrbracket = \llbracket \Delta, \nu x.\varphi \rrbracket$ since α is the closure ordinal of φ . Hence done. \square

We note that we can define phase semantics of any fragment of μMALL° which can be finitised [Dou17, NST18] *i.e.* any fragment for which the Brotherston-Simpson conjecture holds. In that case, one can have the same interpretation for the wellfounded system and the circular system. Since we can handle the semantics of wellfounded systems more easily, a conjecture stems from the aforementioned proof idea (which essentially reduces the soundness of a non-wellfounded proof to soundness of a wellfounded proof):

Open Question

There exists a wellfounded infinitely branching system provably equivalent to μMALL^{∞} .

Chapter 6

Complexity of μ MALL systems

In this chapter, we explore the decision problems about the various systems of μ MALL *viz.* the problem of deciding if a given formula (equivalently, a sequent) is provable. We reduce various the reachability problem in various counter machines to these questions. Consequently, we can compute the precise complexity of these questions. The results are technically interesting from a proof-theoretic point of view since they involve non-trivial applications of focussing. They also have deep implications *viz.* they can separate systems as sets of theorems.

In Section 6.1, we introduce the relevant counter machines and explore their connections with linear logic. In Section 6.2, we show that μ MALL* is undecidable (consequently, so is μ MALL^{ind} and μ MALL[∞]) and that the provability problem of $\&$ -free fragment of μ MALL* is equivalent to the provability problem of MELL. In Section 6.3, we obtain lower bounds on μ MALL[∞] provability which ultimately helps us show that μ MALL^{ind} proves a strictly larger set of theorems than μ MALL[◊]. We show this and constructivise the argument in Section 6.4.

6.1 Counter machines and linear logic

6.1.1 Petri nets

Petri nets [Pet62] are a model of concurrency. We will explain the basic components of Petri nets and their behaviour by means of an example (Figure 6.1). Consider the graph in Figure 6.1a. There are two types of nodes: **places** and **transitions**. Places are usually depicted by circles (*cf.* the nodes labelled P_1, P_2, P_3) and transitions by rectangles (*cf.* the nodes labelled T_1, T_2). A **Petri net** is given by a bipartite graph $(P, T, E, \mathcal{M}_{in})$ where P is the set of places, T is the set of transitions, and $E \subseteq (P \times T) \cup (T \times P)$ is the set of edges between P and T . Furthermore, each place can hold zero or more **tokens**. A **marking** $\mathcal{M} \in \mathbb{N}^{|P|}$ denotes the number of tokens in each place at a certain time. \mathcal{M}_{in} is the initial marking. In Figure 6.1a, $\mathcal{M} = (0, 1, 0)$ is the marking. **Firing** of a transition constitutes consuming the tokens from its incoming places and producing them in its outgoing places. As different transitions are fired, we obtain different configurations *i.e.* different markings. A transition is **fireable** if every place that points to it has at least one token. In Figure 6.1a, T_2 is fireable while T_1 is not. If we fire a transition t , then one token disappears from every place that points to t and one token is added to every place that emanates from t . On firing T_2 , the token from P_2 disappears and P_1 and P_3 get one token each. Consequently, we have the marking $(1, 0, 1)$ which corresponds to Figure 6.1b. Now T_1 is fireable and if it is fired we obtain $(0, 1, 2)$ which corresponds to Figure 6.1c.

Since its inception, linear logic was advertised as the logic for concurrency [Gir87b] and its relation with Petri nets has been explored from both provability and denotational points of view. On the provability side, [Asp87, Bro89, GG89] have independently observed that places are like formulas in linear logic and transitions like proofs. In Figure 6.1, we model the transition T_1 as the formula $P_1^\perp \otimes (P_2 \wp P_3)$. The sequent $\vdash P_1, P_3$ denotes the marking in Figure 6.1b. We have the following:

$$\frac{\frac{\vdash P_1^\perp, P_1 \quad \frac{\vdash P_2, P_3, P_3}{\vdash P_3, P_2 \wp P_3} (\wp)}{\vdash P_1, P_3, P_1^\perp \otimes (P_2 \wp P_3)} (\otimes)}{\vdash P_1, P_3, P_1^\perp \otimes (P_2 \wp P_3)} (\otimes)$$





Figure 6.1: The behaviour of a Petri net

We obtain the marking in Figure 6.1c. Using this observation, one can devise provability semantics of linear logic by establishing a formula φ is provable iff a certain marking is reachable in a Petri Net N_φ . However, one can make this more fine-grained: by letting Petri nets freely generate a linear category, one can interpret linear logic in that setting. This generates a triangular correspondence [BG90, BG95]:



However, turns out, the connections between Petri nets and linear logic are not as deep as it was once suspected. Linear logic is still a logic for concurrency but in a very different way than Girard envisioned (as types of π -calculus). One of the biggest successes of this line of research has been obtaining the complexity of various fragments of linear logic. We explore these results in the following subsection. Nevertheless, the intuitions from this subsection will be helpful in understanding some constructions in forthcoming sections.

6.1.2 Counter machines

There are several Turing equivalent models of computation studied in theoretical computer science to model computation. While a Turing machine is an abstraction of running a sequential algorithm, a **counter machine** is an abstraction of running a parallel algorithm. A counter machine comprises a finite set of one or more *registers*, each of which can hold a single non-negative integer, and a list of arithmetic and control instructions for the machine to follow. A mutual exclusion principle avoids interlocking *i.e.* the simultaneous writing operation by two (or more) threads to the same *register*. A *Minsky machine* is one of the most well-known counter machines.

Definition 6.1.1. A **Minsky machine** \mathcal{M} is a tuple (Q, r_1, r_2, I) where Q is a finite set of states, r_1, r_2 are two registers, and I is a set of instructions of the form $\text{INC}(\bullet, \bullet, \bullet)$ and $\text{JZDEC}(\bullet, \bullet, \bullet, \bullet)$ that manipulate the current state and the contents of the registers. The operational semantics of \mathcal{M} is given by its configuration graph, the vertices of which are configurations of form $\langle q, a, b \rangle \in Q \times \mathbb{N} \times \mathbb{N}$ and edges are one of the following forms:

$$\begin{array}{ll}
 \langle p, a, b \rangle \xrightarrow{\text{INC}(p, r_1, q)} \langle q, a + 1, b \rangle & \langle p, a, b \rangle \xrightarrow{\text{INC}(p, r_2, q)} \langle q, a, b + 1 \rangle \\
 \langle p, 0, b \rangle \xrightarrow{\text{JZDEC}(p, r_1, q_0, q_1)} \langle q_0, 0, b \rangle & \langle p, a, 0 \rangle \xrightarrow{\text{JZDEC}(p, r_2, q_0, q_1)} \langle q_0, a, 0 \rangle \\
 \langle p, a + 1, b \rangle \xrightarrow{\text{JZDEC}(p, r_1, q_0, q_1)} \langle q_1, a, b \rangle & \langle p, a, b + 1 \rangle \xrightarrow{\text{JZDEC}(p, r_2, q_0, q_1)} \langle q_1, a, b \rangle
 \end{array}$$

Given a state q_s , a **run** of \mathcal{M} is a sequence of configurations $\{s_i\}_{i \in \mathbb{O}}$ ($\mathbb{O} \in \omega + 1$) such that $s_0 = \langle q_s, 0, 0 \rangle$ and for all $i \in \mathbb{O}$ with $i + 1 \in \mathbb{O}$, (s_i, s_{i+1}) is an edge in the configuration graph.

Theorem 6.1.1 ([Min67]). Given a Minsky machine \mathcal{M} and an initial state q_s , checking that it has an infinite run from q_s is Π_1^0 -hard.

Our next example is of vector addition systems. As presented the registers are not explicit. However, it is not difficult to unearth the k registers in the following definition.

Definition 6.1.2. A **vector addition system** \mathcal{V} is a 3-tuple (k, \mathbf{A}, T_u) such that $k \in \mathbb{N}$, and $\mathbf{A} \in \mathbb{N}^k$ and $T_u \subset \mathbb{Z}^k$ are finite sets. The operational semantics of \mathcal{V} is given by its configuration graph (\mathbb{N}^k, E_{T_u}) such that $(\vec{v}, \vec{v}') \in E_{T_u}$ if $\vec{v} - \vec{v}' \in T_u$. A configuration \vec{v} is said to be *reachable* if there is a path in the configuration graph starting from \vec{v} and ending in some $\vec{v}' \in \mathbf{A}$.

Let us come back to Petri nets. Given a Petri net $N = (P, T, E, \mathcal{M}_{in})$, define the **marking graph** as $(\mathbb{N}^{|P|}, E_T)$ such that $\mathcal{M}, \mathcal{M}' \in \mathbb{N}^k$ are construed as markings and $(\mathcal{M}, \mathcal{M}') \in E_T$ if there exists $t \in T$ such that $\mathcal{M} \xrightarrow{t} \mathcal{M}'$. A marking \mathcal{M} is said to be *reachable* if there is a path in the marking graph of N , starting from \mathcal{M}_{in} and ending at \mathcal{M} . It is easy to note that a Petri net N induces a vector addition system \mathcal{V}_N such that the marking graph of N and the configuration graph of \mathcal{V}_N are isomorphic. The vector addition system corresponding to our running example is $(3, (0, 1, 0), \{\vec{t}_1, \vec{t}_2\})$ where $\vec{t}_1 = (1, -1, -1)$ and $(1, -1, 1)$.

Theorem 6.1.2 ([CO21, Ler21]). *Reachability in vector addition systems is Ackermann-complete.*

There are several ways of extending the expressiveness of Petri nets or vector addition systems. For example, one can add weights to edges in order to indicate how many tokens are consumed or produced; one can have coloured tokens to distinguish between them; and so on. The extension which is interesting to our discussion is branching behaviour. A firing sequence is analogous to a run of a deterministic or non-deterministic automaton. Now, the run of an alternating automaton is a tree (instead of a list). The question, therefore, arises: is there a meaningful counterpart in the world of Petri nets?

Definition 6.1.3. A **branching vector addition system with states or ABVASS** is a tuple $\mathcal{B} = (Q, Q_\ell, k, \mathbf{A}, T_u, T_s, T_f)$ such that:

- Q is a finite set of states with $Q_\ell \subseteq Q$;
- $k \in \mathbb{N}$ is called the **dimension**;
- \mathbf{A} is a finite subset of \mathbb{N}^k called the set of **axioms**;
- $T_u \subset Q \times \mathbb{Z}^k \times Q$, $T_s \subset Q^3$, $T_f \subset Q^3$ are finite and called the **unary, split, and fork rules** respectively.

If $T_s = \emptyset$ and $T_f = \emptyset$, then it is called a **VASS**. If $T_s = \emptyset$ then it is called an **AVASS**, and if $T_f = \emptyset$ then it is called a **BVASS**.

Definition 6.1.4. Given an ABVASS $\mathcal{B} = (Q, Q_\ell, k, \mathbf{A}, T_u, T_s, T_f)$, a **configuration** is a pair $(q, \vec{v}) \in Q \times \mathbb{N}^k$ where Q is the set of states of \mathcal{B} and k is its dimension. $(q, \vec{v}) \in Q \times \mathbb{N}^k$ is said to be **reachable** if there is a binary tree labelled by configurations such that:

- The root node is labelled by (q, \vec{v}) .
- If a node (q, \vec{v}) has a unique child (q', \vec{v}') then $(q, \vec{v} - \vec{v}', q') \in T_u$.
- If a node (q, \vec{v}) has children (q', \vec{v}') and (q'', \vec{v}'') then:
 - either, $\vec{v}' + \vec{v}'' = \vec{v}$ and $(q, q', q'') \in T_s$, or,
 - $\vec{v} = \vec{v}' = \vec{v}''$ and $(q, q', q'') \in T_f$.
- The leaves are labelled by elements of $Q_\ell \times \mathbf{A}$.

Such a binary tree is called a **run tree** of the configuration.

Example 6.1.1. Let $\mathcal{B} = (\{q_0, q_1, q_2\}, \{q_1, q_2\}, 1, \{0\}, \{(q_1, 3, q_2), (q_2, 3, q_1)\}, \emptyset, \{(q_0, q_1, q_2)\})$. \mathcal{B} is an AVASS since there are no fork rules. We claim that $(q_0, 5)$ is reachable. The split rule takes $(q_0, 5)$ to $(q_1, 5)$ and $(q_2, 5)$. From $(q_1, 5)$, one can trigger two unary rules to reach $(q_2, 2)$ and then $(q_1, 0)$. Similarly, from $(q_2, 5)$, one can trigger two unary rules to reach $(q_1, 3)$ and then $(q_2, 0)$. Since, $(q_2, 0), (q_1, 0) \in \{q_1, q_2\} \times \{0\}$, we are done¹. On the other hand, observe that $(q, 3)$ and (q_2) are not reachable.

¹It is easy to generalise this to see that a configuration of the form $(q, 5k)$ is reachable for all $k \in \mathbb{N}$.

Now consider $\mathcal{B}' = (\{q_0, q_1, q_2\}, \{q_1, q_2\}, 1, \{0\}, \{(q_1, 3, q_2), (q_2, 3, q_1)\}, \{(q_0, q_1, q_2)\}, \emptyset)$. \mathcal{B}' is a BVASS since there are no split rules. We claim that $(q_0, 5)$ is reachable. The fork rule can be triggered to reach $(q_1, 3)$ and $(q_2, 2)$ from $(q_0, 5)$. Since these two configurations are reachable, we are done. However, note that in contrast to \mathcal{B} , $(q_0, 3)$ is reachable by forking it to $(q_1, 3)$ and $(q_2, 0)$. Similarly, $(q_0, 2)$ is also reachable by forking it to $(q_1, 0)$ and $(q_2, 2)$.

Theorem 6.1.3 ([HP79]). *An n -dimension VASS can be simulated in a $n + 3$ dimension VAS.*

Similarly, one can simulate the branching behaviour without the use of states by having an operational semantics of the form: if a node \vec{v} has children \vec{v}' and \vec{v}'' , then $\vec{v} - \vec{v}' - \vec{v}'' \in T_s$. On the other hand, without states, it is impossible to mimic the alternating behaviour.

Theorem 6.1.4 ([dGGS04]). *Reachability in BVASSes is equivalent to provability in MELL.*

Naturally, the decidability of the reachability problem in BVASSes is an open question (since the decidability of MELL provability is open). However, it inherits the Ackermann-hardness of VASSes since every BVASS is also a VASS *i.e.* we have the following corollary to Theorem 6.1.2.

Corollary 6.1.4.1. *Reachability in BVASSes is Ackermann-hard.*

Propositional linear logic was shown to be undecidable [Lin92, LMAS92] by a reduction from the reachability problem in an *and-branching* two counter machine without zero-test. Such machines are essentially equivalent to AVASSes [Kop01, LS14, LS15] (in particular, the fork rule is exactly the same).

Theorem 6.1.5 ([LMAS92]). *The AVASS reachability problem is Σ_1^0 -complete.*

In the following sections, we reduce reachability in AVASSes to the provability of μMALL^* , reduce reachability in BVASSes to the provability of $\&$ -free μMALL^* , and reduce the non-halting of Minsky machines to the provability of μMALL^∞ .

6.2 The complexity of μMALL^*

6.2.1 Σ_1^0 -completeness via full LL

Recall the encoding of the exponential modalities by fixed point formulas from Definition 4.3.2. Lemma 4.3.3 shows that if φ is provable in LL then $[\varphi]$ is provable in $\mu\text{MALL}^{\text{ind}}$. The converse is not known to be true. However, we prove the converse in a restricted setting which immediately gives us the complexity of μMALL^* .

Theorem 6.2.1. *Let φ be a $!$ -free LL formula. Then, $\vdash \varphi$ is provable in LL iff $\vdash [\varphi]$ is provable in μMALL^* .*

Proof. The only if part follows from Lemma 4.3.3. For the if part, we generalise the statement from a formula φ to a sequent $\Gamma = \varphi_1, \dots, \varphi_n$. Define $[\Gamma]$ as expected as $[\varphi_1], \dots, [\varphi_n]$. Suppose there is a proof π_0 of $\vdash [\Gamma]$. Then, by Theorem 4.4.1 and Theorem 4.5.1, there exists a cut-free focussed proof π of $\vdash [\Gamma]$. We will induct on π .

The base case is simple. If $[\Gamma] = [\varphi], [\varphi^\perp]$ then $\Gamma = \varphi, \varphi^\perp$ therefore if one can apply (id) rule in π , one can do the same for Γ in LL. Similarly, for units $\top, \mathbf{1}$. Now suppose the last rule of π is a multiplicative additive rule (\odot) with principal formula $[\varphi \odot \psi]$ and auxiliary formula(s) one (or both) of $[\varphi]$ and $[\psi]$. Noting the definition of the $[\bullet]$, we see that one can apply the (\odot) on $\varphi \odot \psi$ and auxiliary formula(s) one (or both) of φ and ψ such that the premisses are equal up to $[\bullet]$ to the corresponding premisses in π . Conclude by applying the induction hypothesis to these premisses.

The only remaining case is when the last rule of π is (μ) and the principal formula is $[\varphi]$. (Note that since Γ is $!$ -free, $[\Gamma]$ is ν -free.) Recall that π is a focussed proof. Therefore, (i) $[\Gamma]$ is a positive sequent, and moreover (ii) $[\varphi]$ is the focus. Therefore, the next rule is the ternary (\oplus) rule with principal formula $\perp \oplus ([\varphi] \wp [\varphi]) \oplus [\varphi]$. We will now subject the reader to the obvious case analysis.

Case 1. The next rule (\oplus_1). Then, the premiss is a negative sequent with the only negative formula \perp . Therefore the next rule is necessarily (\perp). In LL, we apply $?_w$ on φ and conclude by induction hypothesis.

Case 2. The next rule (\oplus_2). Again, we end up in a negative sequent with the only negative formula $[\varphi] \wp [\varphi]$. Therefore the next rule is necessarily (\wp). In LL, we apply $?_c$ on φ and conclude by induction hypothesis.

Case 3. The next rule (\oplus_3). The premiss is a sequent of the form $[\Gamma']$ and therefore immediately apply the induction hypothesis. \square

Corollary 6.2.1.1. μMALL^* is Σ_1^0 -complete.

Proof. This follows from Theorem 6.2.1 and the fact that the reduction in Theorem 3.2.1 only uses $!$ -free formulas. \square

Corollary 6.2.1.2. $\mu\text{MALL}^{\text{ind}}$ and μMALL° are Σ_1^0 -complete.

Proof. Σ_1^0 -membership for $\mu\text{MALL}^{\text{ind}}$ since wellfounded proof are finitely presentable and hence recursively checkable. μMALL° pre-proofs are also finitely presentable and since given a circular pre-proof, checking the progress condition is decidable, we have Σ_1^0 -membership. For hardness, note that via Corollary 6.2.1.1 the reduction only uses ν -free formulas and hence the result follows since all systems *i.e.* μMALL° , $\mu\text{MALL}^{\text{ind}}$, and μMALL^* coincide on ν -free formulas. \square

It is folklore that if φ is an LK pre-formula with a free variable x then φ and $\varphi[\varphi/x]$ are equivalent. This immediately gives us a conservative embedding of μLK (note that this is different from μ -calculus since there are no modalities) in LK with a polynomial blowup. In the same vein, [GGS16, GGS19] shows that there is a conservative embedding of μLJ in LJ with an exponential blowup. MALL is known to be PSPACE-complete [LMAS92]. Therefore we have the following corollary.

Corollary 6.2.1.3. *There is no effectively computable reduction from μMALL^* (or $\mu\text{MALL}^{\text{ind}}$, μMALL°) to MALL.*

6.2.2 A potentially decidable fragment

In Section 6.1, we discussed that the provability of MELL is equivalent to the reachability problem in BVASSs, which are important open questions in the respective communities. In the spirit of this line of research, we show that provability of the $\&$ -free fragment of μMALL^* is equivalent to reachability in BVASSs. We denote the $\&$ -free fragment of μMALL^* by μMLL_{\oplus}^* .

We fix $k + 1$ propositional variables, a_1, \dots, a_k, z , and define below an encoding of integer vectors of dimension up to k (the unique vector of dimension 0 is written ε). For the purpose of the encoding vectors will be read from left to right *i.e.* a vector \vec{v} of dimension $l + 1$ will be of the form (n, \vec{u}) for an integer n and a vector \vec{u} of dimension l . For typographic ease, we use a^n to denote $\overbrace{a, \dots, a}^{n \text{ times}}$ in a sequent.

Definition 6.2.1. *The encoding of an integer vector \vec{v} of dimension d , relative to propositional variables b_i, \dots, b_{d+i-1}, z , written $[\vec{v}]_{b_i, \dots, b_{d+i-1}, z}$, is defined inductively as follows:*

$$[\vec{v}]_{b_i, \dots, b_{d+i-1}} \triangleq \begin{cases} z & \text{if } \vec{v} = \varepsilon; \\ b_i \wp [\vec{v}']_{b_i, \dots, b_{d+i-1}, z} & \text{if } \vec{v} = (n, \vec{u}), n \geq 1, \text{ and } \vec{v}' = (n-1, \vec{u}); \\ b_i^\perp \otimes [\vec{v}']_{b_i, \dots, b_{d+i-1}, z} & \text{if } \vec{v} = (n, \vec{u}), n \leq -1, \text{ and } \vec{v}' = (n+1, \vec{u}); \\ [\vec{u}]_{b_{i+1}, \dots, b_{d+i-1}, z} & \text{if } \vec{v} = (0, \vec{u}). \end{cases}$$

We will simply write $[\vec{v}]$ for the encoding of a vector of dimension k relative to a_1, \dots, a_k, z . (We also use this lighter notation for vectors of lower dimension when the dimension and the $\{a_i, \dots, a_k, z\}$ to be used are clear from the context.)

Example 6.2.1. *Consider the encoding of $(-1, 0, 1)$ relative to b_1, b_2, b_3, z .*

$$\begin{aligned} [(-1, 0, 1)]_{b_1, b_2, b_3} &= b_1^\perp \otimes [(0, 0, 1)]_{b_1, b_2, b_3} \\ &= b_1^\perp \otimes [(1)]_{b_3} = b_1^\perp \otimes (b_3 \wp [(0)]_{b_3}) \\ &= b_1^\perp \otimes (b_3 \wp [\varepsilon]_{b_3}) = b_1^\perp \otimes (b_3 \wp z) \end{aligned}$$

Observe that the i^{th} coordinate is represented by the propositional variable b_i . The following lemma shows that the encoding is meaningful with respect to vector equality.

Lemma 6.2.1. *Let \vec{u} and \vec{v} be two vectors. Then $\vdash [\vec{u}]^\perp, [\vec{v}]$ is provable iff $\vec{u} = \vec{v}$.*

Proof. The if direction is a trivial induction on the dimension. For the only if direction, first note that for any vector \vec{u} , $[\vec{u}]$ is a purely MLL formula. Therefore, by Example 3.3.2, for all i , the number of times a_i occurs in $[\vec{v}]$ is equal to the number of times a_i^\perp occurs in $[\vec{u}]^\perp$. This is enough to ensure that $\vec{u} = \vec{v}$. \square

The following technical lemma will allow us to reason by induction on the dimension via the encoding at the provability level, which is crucial to prove our forthcoming theorem.

Lemma 6.2.2. *Let $1 \leq i \leq k$, $m \geq 0$ and let s be an integer such that $m \geq s$. Let \vec{q} be an integer vector of dimension $k - i$. If $\vdash [\vec{q}]^\perp, \Gamma, a_i^m, a_{i+1}^{m+1}, \dots, a_k^{m_k}, z$ is provable, then so is $\vdash [\vec{r}]^\perp, \Gamma, a_i^{m+s}, a_{i+1}^{m+1}, \dots, a_k^{m_k}, z$ where $\vec{r} = (s, \vec{q})$ and $\vec{v} = (m, \vec{u})$.*

Proof. We will induct on $|s|$. The base case is $s = 0$. Then, $[\vec{r}] = [\vec{q}]$ hence this case is trivial. For the induction case, we have two subcases.

Case 1: If s is positive.

$$\frac{\frac{\text{IH}}{\vdash [(s+1, \vec{q})]^\perp, \Gamma, a_i^{m+(s-1)}, a_{i+1}^{m+1}, \dots, a_k^{m_k}, z} \text{(id)} \quad \vdash a_i^\perp, a_i}{\vdash a_i^\perp \otimes [(s+1, \vec{q})]^\perp, \Gamma, a_i^{m+s}, a_{i+1}^{m+1}, \dots, a_k^{m_k}, z} (\otimes)$$

Case 2: If s is negative.

$$\frac{\text{IH} \quad \nabla}{\frac{\vdash [(s-1, \vec{q})]^\perp, \Gamma, a_i^{m+s+1}, a_{i+1}^{m_{i+1}}, \dots, a_k^{m_k}, z}{\vdash a_i \wp [(s-1, \vec{q})]^\perp, \Gamma, a_i^{m+s}, a_{i+1}^{m_{i+1}}, \dots, a_k^{m_k}, z} (\wp)}$$

□

In the following, let S be a finite set and $[\bullet] : S \rightarrow \mu\text{MALL}$. Define CH_S to be the formula that offers a choice of picking the dual of one of the (encoding of) elements of S .

$$\text{CH}_S \triangleq \bigoplus_{s \in S} [s]^\perp$$

When S is a set of instructions we rely on the above encoding, when S is a set of states, we use the identity encoding benefiting from the fact that states are indeed propositional variables. The reader might be surprised by our use of the logical duality here: it is simply because we are working in the one-sided calculus. Furthermore, note that the formula is well-defined only when we have fixed an order on the elements of S ; however, the choice of an order is irrelevant from the provability point of view.

We can now define the encoding of a BVASS $\mathcal{B} = (Q, Q_\ell, k, \mathbf{A}, T_u, T_s)$ with $|Q_\ell \times \mathbf{A}| = \alpha$, $|T_u| = \beta$, and $|T_s| = \gamma$.

- For a unary rule $t \in T_u$ of the form (p, \vec{r}, q) we have $[t] \triangleq p \wp (q^\perp \otimes [\vec{r}])$.
- For a split rule $t \in T_s$ of the form (p, q_1, q_2) we have that $[t] \triangleq p \wp (q_1^\perp \otimes q_2^\perp)$.
- For a final configuration $(q, \vec{v}) \in Q_\ell \times \mathbf{A}$, we have that $[(q, \vec{v})] \triangleq q \wp [\vec{v}]$.
- Finally, \mathcal{B} is encoded as follows.

$$B \triangleq \mu x. \text{CH}_{Q_\ell \times \mathbf{A}} \oplus (\text{CH}_{T_u} \wp (z \otimes x)) \oplus (\text{CH}_{T_s} \wp ((z \wp x) \otimes x)).$$

Lemma 6.2.3. *If the configuration (q, \vec{v}) is reachable in \mathcal{B} then $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}, z$ is provable in μMALL_\oplus^* where $\vec{v} = (v_1, \dots, v_k)$.*

Proof. Since (q, \vec{v}) is reachable, there exists a run tree. We will mimic the run tree to produce a proof tree. The proof goes by induction on the run tree. The base case is that of a one node run tree where the node is labelled by $(p, \vec{v}) \in Q_\ell \times \mathbf{A}$. We have,

$$\frac{\frac{\frac{\text{[Lemma 6.2.1]} \quad \nabla}{\frac{\vdash p^\perp, p}{\vdash p^\perp, p} (\text{id})} \quad \frac{\vdash [\vec{v}]^\perp, a_1^{v_1}, \dots, a_k^{v_k}, z}{\vdash [\vec{v}]^\perp, a_1^{v_1}, \dots, a_k^{v_k}, z} (\otimes)}{\vdash p^\perp \otimes [\vec{v}]^\perp, p, a_1^{v_1}, \dots, a_k^{v_k}, z} (\oplus)}{\vdash \text{CH}_{Q_\ell \times \mathbf{A}}, p, a_1^{v_1}, \dots, a_k^{v_k}, z} (\mu), (\oplus_1)}{\vdash B, p, a_1^{v_1}, \dots, a_k^{v_k}, z}$$

where the α -ary \oplus on $\text{CH}_{Q_\ell \times \mathbf{A}}$ chooses the configuration (p, \vec{v}) . There are two cases for the inductive step depending on whether a unary or split rule is applied at the root node.

Case 1. Suppose the unary rule (p, \vec{r}, q) is applied to the node labelled $(p, \vec{v} + \vec{r})$ where $\vec{r} = (r_1, \dots, r_k)$. We have the following. The β -ary \oplus chooses the (encoding of) rule $(p, \vec{r}, q) \in T_u$.

$$\begin{array}{c}
 \text{IH} \\
 \nabla \\
 \frac{\frac{\frac{\text{---}(\text{id})}{\vdash z^\perp, z} \quad \frac{\text{---}}{\vdash q, B, a_1^{v_1}, \dots, a_k^{v_k}, z}}{\vdash q, z^\perp, z \otimes B, a_1^{v_1}, \dots, a_k^{v_k}, z} (\otimes)}{\vdash q, [\vec{r}]^\perp, z \otimes B, a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z} (\text{Lemma 6.2.2, } k \text{ times})}{\frac{\frac{\text{---}(\text{id})}{\vdash p^\perp, p} \quad \frac{\text{---}(\text{id})}{\vdash q \wp [\vec{r}]^\perp, z \otimes B, a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z} (\wp)}{\vdash p^\perp \otimes (q \wp [\vec{r}]^\perp), z \otimes B, p, a_1^{v_1+r_1}, \dots, a_k^{v_k}, z} (\otimes)}{\frac{\frac{\text{---}}{\vdash \text{CH}_{T_u}, z \otimes B, p, a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z} (\wp)}{\vdash \text{CH}_{T_u} \wp(z \otimes B), p, a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z} (\wp)}{\vdash B, p, a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z} (\oplus)} (\oplus)
 \end{array}$$

Case 2. Suppose the split rule (p, q_1, q_2) is applied to the node labelled $(p, \vec{u} + \vec{v})$ to produce (q_1, \vec{u}) and (q_2, \vec{v}) . The γ -ary \oplus chooses the (encoding of) rule $(p, q_1, q_2) \in T_u$.

$$\begin{array}{c}
 \text{IH} \\
 \nabla \\
 \frac{\frac{\frac{\text{---}(\text{id})}{\vdash p^\perp, p} \quad \frac{\frac{\text{---}}{\vdash q_1, z, B, a_1^{u_1}, \dots, a_k^{u_k}} (\wp)}{\vdash q_1, z \wp B, a_1^{u_1}, \dots, a_k^{u_k}} (\wp)}{\vdash p^\perp \otimes (q_1 \wp q_2), (z \wp B) \otimes B, a_1^{u_1+v_1}, \dots, a_k^{u_k+v_k}, z} (\otimes)}{\frac{\frac{\text{---}}{\vdash q_2, B, a_1^{v_1}, \dots, a_k^{v_k}, z} (\otimes)}{\vdash q_1, q_2, (z \wp B) \otimes B, a_1^{u_1+v_1}, \dots, a_k^{u_k+v_k}, z} (\wp)}{\vdash q_1 \wp q_2, (z \wp B) \otimes B, a_1^{u_1+v_1}, \dots, a_k^{u_k+v_k}, z} (\otimes)}{\frac{\frac{\text{---}}{\vdash \text{CH}_{T_s}, (z \wp B) \otimes B, p, a_1^{u_1+v_1}, \dots, a_k^{u_k+v_k}, z} (\wp)}{\vdash \text{CH}_{T_s} \wp((z \wp B) \otimes B), p, a_1^{u_1+v_1}, \dots, a_k^{u_k+v_k}, z} (\wp)}{\vdash B, p, a_1^{v_1}, \dots, a_k^{v_k}, z} (\oplus)} (\oplus)
 \end{array}$$

□

Before proving the opposite direction, we will prove a technical lemma.

Lemma 6.2.4. *For all v_1, v_2, \dots, v_n , the sequent $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}$ is not provable.*

Proof. We shall reason by contradiction. Assume that there are some $v_1, \dots, v_k \in \mathbb{N}$ such that $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}$ is derivable. *Wlog*, let us assume that v_1, \dots, v_k are chosen such that the sequent can be proved with a cut-free, focussed and minimal proof π *i.e.* no subproof of π is rooted at a sequent of the form $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}$. Note that the first two conditions can be assumed by Theorem 4.4.1 and Theorem 4.5.1 respectively; then it is possible to assert the third condition since μMALL^* is well-founded.

Note that the sequent is positive and B is the only formula on which any rule can be applied. Therefore, it is the focus. After the application of the μ rule (which is the only rule that can be applied at the root of the proof) there are three possibilities, \oplus_1, \oplus_2 or \oplus_3 resulting in either one of the following sequents:

1. $\vdash \text{CH}_{Q_\ell \times A}, q, a_1^{v_1}, \dots, a_k^{v_k}$, or,
2. $\vdash \text{CH}_{T_u} \wp(z \otimes B), q, a_1^{v_1}, \dots, a_k^{v_k}$, or,
3. $\vdash \text{CH}_{T_s} \wp((z \wp B) \otimes B), q, a_1^{v_1}, \dots, a_k^{v_k}$.

Case 1. Assuming that one applies \oplus_1 , we are still in the positive phase and in a purely MLL sequent. Using Example 3.3.2 with respect to z , we have that this is not provable.

Case 2. Assuming that one applies \oplus_2 , we end up in a negative sequent. After the mandatory (\wp) rule, we have a sequent of the form $\vdash \text{CH}_{T_u}, z \otimes B, q, a_1^{v_1}, \dots, a_k^{v_k}$. This proof shall contain a (\otimes) rule of principal formula $z \otimes B$ of conclusion. The left premisses must contain either CH_{T_u} or a subformula of CH_{T_u} of the form $[\vec{r}]^\perp$ (since CH_{T_u} that contains z^\perp as subformula and otherwise we have a contradiction by Example 3.3.2). By a similar argument, we ensure that there exists $q' \in Q$ in the right premisses. Therefore, the right premisses is of the form $\vdash B, q', a_1^{v_1}, \dots, a_k^{v_k}$. This contradicts the minimality of the proof.

Case 3. Assuming that one applies \oplus_3 , we end up in a negative sequent. After the mandatory (\wp) rule, we have a sequent of the form $\vdash \text{CH}_{T_u}, (z \wp B) \otimes B, q, a_1^{v_1}, \dots, a_k^{v_k}$. If $(z \wp B) \otimes B$ is the focus, then the right premisses is of the form $\vdash \text{CH}_{T_s}, B, q, a_1^{v_1}, \dots, a_k^{v_k}$. If q is not present, then no axiom, unary or split rule would be applicable at any point, so the proof would be a non-wellfounded unfolding of B . If CH_{T_s} is not present then we would contradict the minimality of the proof. Therefore, the left premisses is of the form $\vdash z \wp B, a_1^{v_1}, \dots, a_k^{v_k}$. Since q is not present, then no axiom, unary or split rule would be applicable at any point, so the proof would be a non-wellfounded unfolding of B . Hence we have exhausted all three possible cases and can therefore deduce the expected contradiction. \square

Lemma 6.2.5. *If $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}, z$ is provable in μMLL^*_\oplus then the configuration (q, \vec{v}) is reachable in \mathcal{B} where $\vec{v} = (v_1, \dots, v_k)$.*

Proof. Assume that $q \in Q$ and $\{a_1, \dots, a_k\}$ are negative atoms and z is a positive atom. By Theorem 4.4.1 and Theorem 4.5.1, there exists a cut-free focussed proof of $\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}, z$. We will induct on the height of the proof. B is the only formula on which any rule can be applied. Therefore, it is the focus. The proof starts off as follows.

$$\frac{\vdash \text{CH}_{Q_\ell \times \mathbf{A}} \oplus (\text{CH}_{T_u} \wp (z \otimes B)) \oplus (\text{CH}_{T_s} \wp (z \wp B \wp B)), q, a_1^{v_1}, \dots, a_k^{v_k}, z}{\vdash B, q, a_1^{v_1}, \dots, a_k^{v_k}, z} (\mu)$$

Case 1. The next rule is (\oplus_1). The auxiliary formula is positive. Therefore, the next rule is the α -ary (\oplus) that chooses $q'^\perp \otimes [\vec{v}']^\perp$ for some $(q', \vec{v}') \in Q_\ell \times \mathbf{A}$. This is also a positive formula; hence the next rule is (\otimes) and its left premisses is of the form $\vdash q'^\perp, \Gamma$. Since q' is a negative atom, this is still in the positive phase. Therefore, it must be the conclusion of an (id) rule. Therefore, $q' = q, \Gamma = \{q\}$, and the right premisses of the tensor rule is of the form $\vdash [\vec{v}']^\perp, a_1^{v_1}, \dots, a_k^{v_k}, z$. By Lemma 6.2.1, $\vec{v}' = (v_1, \dots, v_k)$. Hence $(q, \vec{v}) \in Q_\ell \times \mathbf{A}$ and the run tree is a single node labelled by (q, \vec{v}) .

Case 2. The next rule is (\oplus_2). The following rule is necessarily (\wp). Now there are two possibilities, either CH_{T_u} is the focus or $z \otimes B$ is the focus. Suppose the latter happens. Then, the left premisses of the tensor rule with principal formula $z \otimes B$ is of the form $\vdash z, \Gamma$. Observe that this is a positive sequent and it must be the conclusion of an (id) rule. This is not possible. Therefore CH_{T_u} is the focus. The next rule is thus a β -ary \oplus and a unary rule of the form $q'^\perp \otimes (q'' \wp [\vec{r}]^\perp)$ is chosen. The next rule is thus a tensor. Since the left premisses (say, $\vdash q'^\perp, \Gamma$) is positive, the only possibility is that $\Gamma = \{q\}$ and $q' = q$. The right premisses is negative and after an application of (\wp) we have a sequent of the form $\vdash [\vec{r}]^\perp, z \otimes B, q'', a_{v_1}, \dots, a_{v_k}, z$.

Observe that if the first coordinate of \vec{r} is negative then the outermost connective of $[\vec{r}]^\perp$ is negative. Suppose the first ℓ coordinates of \vec{r} are negative, then the subsequent inference rules necessarily are several \wp rules until we reach the sequent $\vdash [\vec{q}]^\perp, z \otimes B, a_1^{v_1+r_1}, \dots, a_\ell^{v_\ell+r_\ell}, a_{\ell+1}^{v_{\ell+1}}, \dots, a_k^{v_k}, z$ where $\vec{q} = (0, \dots, 0, r_{\ell+1}, \dots, r_k)$ and $r_{\ell+1} > 0$. This is a positive sequent. There are again two possibilities for the foci: $[\vec{q}]^\perp$ or $z \otimes B$. In the latter case, we will argue as before.

Thus $[\vec{q}]^\perp$ is the focus. Observe that the left premisses of the tensor has to be $\vdash a_{\ell+1}^\perp, \Gamma$. We assumed that a_i is a negative atom for all i . Therefore, this has to be the conclusion of an axiom and hence $\Gamma = \{a_\ell\}$. Continuing like this we reach $\vdash z^\perp, z \otimes B, q'', a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z$. The only possible focus is now $z \otimes B$ and by a similar argument as the above instances the left premisses is necessarily of the form $\vdash z^\perp, z$ and the conclusion of an axiom. The right premisses of the form $\vdash B, q'', a_1^{v_1+r_1}, \dots, a_k^{v_k+r_k}, z$ and we can apply induction hypothesis on this subproof of lower height.

Case 3. The next rule is (\oplus_3) . We have a negative sequent and the following rules is necessarily a (\wp) s. Now there are two possibilities, either CH_{T_s} is the focus or $(z\wp B) \otimes B$ is the focus.

Suppose $(z\wp B) \otimes B$ is the focus. We will consider two subcases. Assume that CH_{T_s} goes to the left premiss. Now, q must be on the right premiss, otherwise on the right premiss, no axiom, unary or binary rule can ever be applied and it will be a nonwellfounded unfolding of B . Finally, by Lemma 6.2.4, z also goes to the right. Therefore, the left premiss is of the form $\vdash z\wp B, a^{v_1}, \dots, a^{v_k}, \text{CH}_{T_s}$. Again, observe that due to the absence of an atom $p \in Q$, no axiom, unary or binary rule can ever be applied and it will be a nonwellfounded unfolding of B . The second subcase when CH_{T_s} goes to the right premiss is symmetric.

If CH_{T_s} is the focus, then a $|\gamma|$ -ary (\oplus) -rule chooses the encoding of a split rule t of the form (p', q_1, q_2) . Since $[t]^\perp$ is a positive formula whence the next rule is (\otimes) with principal formula $[t]^\perp$. Arguing as we have done in several instances, the left premiss is necessarily of the form $\vdash p'^\perp, p$ and $p = p'$. The right premiss is a negative sequent. After an application of the (\wp) rule, we have a sequent of the form $\vdash q_1, q_2, (z\wp B) \otimes B, a^{v_1}, \dots, a^{v_k}, z$. The focus is $(z\wp B) \otimes B$. If q_1 and q_2 , then there is a premiss with no atom from Q , and thus no axiom, unary or split rule can ever be applied on that subproof. If the contexts split in such a way that B occurs in only one of the premisses then consider the premiss where it does *not* occur. Note that it is a purely MLL sequent and applying Example 3.3.2 with respect to z , we have a contradiction. Finally, by Lemma 6.2.4, the only possible splitting of the context has exactly one z and one B on either side, allowing us to fire the induction hypothesis on each premiss. \square

From Lemma 6.2.3 and Lemma 6.2.5, we conclude that:

Theorem 6.2.2. *BVASS reachability reduces to μMLL_{\oplus}^* provability.*

We will now show that μMLL_{\oplus}^* provability reduces to BVASS reachability. It is equivalent to prove a sequent and a formula. Fix a formula φ and order $\text{FL}(\varphi)$ lexicographically as $\{\varphi_1, \dots, \varphi_k\}$. Let Seq be the set of multiset with the φ_i s as elements. Let $[\bullet] : \text{Seq} \rightarrow \mathbb{N}^k$ be an injective function such that the i th coordinate of $[\Gamma]$ is the multiplicity of φ_i in Γ . Let $e_\psi = (v_1, \dots, v_k)$ be the vector such that

$$v_i = \begin{cases} 0 & \text{if } \varphi_i \neq \psi; \\ 1 & \text{otherwise.} \end{cases}$$

We will now define the BVASS that will determine the provability of φ . Consider the following BVASS $\mathcal{B}_\varphi = (Q, Q_\ell, k, \mathbf{A}, T_u, T_s)$ such that:

- $Q = \text{FL}(\varphi) \cup \{\star, \emptyset\}$;
- $Q_\ell = \{\star\}$;
- $\mathbf{A} = \{[\psi, \psi^\perp] \mid \{\psi, \psi^\perp\} \subseteq \text{FL}(F)\} \cup \{[\top], [\mathbf{1}]\}$;
- The unary rules are given in the following way as (we present them as transitions over configurations):

1. $(\psi \oplus \psi', [\Gamma]) \xrightarrow{\vec{0}} (\psi, [\Gamma])$
2. $(\psi \oplus \psi', [\Gamma]) \xrightarrow{\vec{0}} (\psi', [\Gamma])$
3. $(\psi \wp \psi', [\Gamma]) \xrightarrow{e_\psi + e_{\psi'}} (\emptyset, [\Gamma, \psi, \psi'])$
4. $(\mu X.\psi, [\Gamma]) \xrightarrow{\vec{0}} (\psi[\mu X.\psi], [\Gamma])$
5. $(\perp, [\Gamma]) \xrightarrow{\vec{0}} (\emptyset, [\Gamma])$
6. $(\top, [\Gamma]) \xrightarrow{-e_\psi} (\top, [\Gamma \setminus \{\psi\}])$
7. $(\top, [\emptyset]) \xrightarrow{e_\top} (\star, [\top])$
8. $(\mathbf{1}, [\emptyset]) \xrightarrow{e_\mathbf{1}} (\star, [\mathbf{1}])$
9. $(S, [\Gamma]) \xrightarrow{e_S - e_{\psi'}} (\psi', [\Gamma \setminus \{\psi'\}, S])$ where $S \in \{\emptyset, \{\psi\}\}$ and $e_\emptyset = \vec{0}$
10. $(\emptyset, [\Gamma]) \xrightarrow{\vec{0}} (\star, [\Gamma])$

- There is exactly one split rule (as before we present it as a transition over configurations):

$$\begin{array}{c} (\psi \otimes \psi', [\Gamma, \Delta]) \\ \swarrow \quad \searrow \\ (\psi, [\Gamma]) \quad (\psi', [\Delta]) \end{array}$$

Lemma 6.2.6. *If φ is provable in μMLL_{\oplus}^* then $(\emptyset, [\varphi])$ is reachable in \mathcal{B}_{φ} .*

Proof. By Theorem 4.5.1, we can assume that we have a cut-free proof. We will induct on the height of the proof. We will choose a stronger hypothesis *viz.* for any sequent Γ made of formulas from $\text{FL}(\varphi)$, if Γ is provable in μMLL_{\oplus}^* then $(\emptyset, [\Gamma])$ is reachable in \mathcal{B} .

For the the base case there are three subcases:

- The proof of the form

$$\frac{}{\vdash \psi, \psi^{\perp}} (\text{id})$$

In \mathcal{B} , $(\emptyset, [\psi, \psi^{\perp}])$ is indeed reachable via a unary rule to $(\star, [\psi, \psi^{\perp}]) \in Q_{\ell} \times \mathbf{A}$.

- The proof is of the form

$$\frac{}{\vdash \Gamma', \top} (\top)$$

We have $(\emptyset, [\Gamma, \top]) \xrightarrow{-e_{\top}} (\top, [\Gamma])$. A (finite) series of unary rules leads us to $(\top, [\emptyset])$ from where we can reach $(\star, [\top])$.

- The case for the **(1)** rule is trivial.

For the induction case, we have several subcases:

- Suppose $\Gamma = \Gamma', \psi \oplus \psi'$ and the proof is of the form:

$$\frac{\vdash \Gamma', \psi}{\vdash \Gamma', \psi \oplus \psi'} (\oplus_1)$$

We have $(\emptyset, [\Gamma', \psi \oplus \psi']) \xrightarrow{-e_{\psi \oplus \psi'}} (\psi \oplus \psi', [\Gamma']) \xrightarrow{\vec{0}} (\psi, [\Gamma']) \xrightarrow{e_{\psi}} (\emptyset, [\Gamma', \psi])$. We can now apply the induction hypothesis. This goes exactly similarly for the rules (\oplus_2) and (μ) .

- Suppose $\Gamma = \Gamma', \psi \wp \psi'$ and the proof is of the form:

$$\frac{\vdash \Gamma', \psi, \psi'}{\vdash \Gamma', \psi \wp \psi'} (\wp)$$

We have $(\emptyset, [\Gamma', \psi \wp \psi']) \xrightarrow{-e_{\psi \wp \psi'}} (\psi \wp \psi', [\Gamma']) \xrightarrow{e_{\psi} + e_{\psi'}} (\emptyset, [\Gamma', \psi, \psi'])$. We can now apply the induction hypothesis. This idea works for the (\perp) rule as well.

- Suppose $\Gamma = \Gamma', \psi \otimes \psi'$ and the proof is of the form:

$$\frac{\vdash \Delta, \psi \quad \vdash \Delta', \psi'}{\vdash \Gamma', \psi \otimes \psi'} (\otimes)$$

We have,

$$\begin{array}{c} (\emptyset, [\Gamma', \psi \otimes \psi']) \\ \downarrow \\ (\psi \otimes \psi', [\Gamma']) \\ \swarrow \quad \searrow \\ (\psi, [\Delta]) \quad (\psi', [\Delta']) \\ \downarrow e_{\psi} \quad \downarrow e_{\psi'} \\ (\emptyset, [\Delta, \psi]) (\emptyset, [\Delta', \psi']) \end{array}$$

We can now apply the induction hypothesis. □

Lemma 6.2.7. *Let ψ, ψ' be formulas such that $\psi \neq \top$. Let $S \in \{\emptyset, \{\psi\}\}$ and $S' \in \{\emptyset, \{\psi'\}\}$. If in a run of \mathcal{B}_φ , there is a path from $(S, [\Gamma])$ to $(S', [\Gamma'])$ then the following is derivable in μMLL_{\oplus}^* .*

$$\frac{\vdash S', \Gamma'}{\vdash S, \Gamma}$$

Proof. We will prove by induction on the length of the path from $(S, [\Gamma])$ to $(S', [\Gamma'])$. The base case is trivial as $\Gamma' = \Gamma$ and $S = S'$. For the induction case, there are two subcases.

- If $S = \psi \odot \psi'$, $S = \mu x. \psi$, $S = \perp$ then we apply the corresponding rule and then the induction hypothesis.
- Since we assume that $S' \neq \star$, therefore S cannot be $\mathbf{1}$. If $S = \emptyset$, we cannot have $(\emptyset, [\Gamma]) \rightarrow (\star, [\Gamma])$ for the same reason. Now the only rule left to examine is $(\psi, [\Gamma]) \rightarrow (\psi', [\Gamma \setminus \{\psi'\}, \psi])$. Note that as sequents these are identical, hence we are done. □

Lemma 6.2.8. *If $(\emptyset, [\varphi])$ is reachable in \mathcal{B} then φ is provable in μMLL_{\oplus}^* .*

Proof. There is a run in normal form starting from $(\emptyset, [\varphi])$. We use Lemma 6.2.7, to obtain a maximal derivation of φ . The process ends at a configuration of the form $(\top, [\Gamma])$, $(\mathbf{1}, [\emptyset])$, or $(\emptyset, [\Gamma])$. Consequently, the leaves of the derivation are of the form $\vdash \top, \Gamma$, $\vdash \mathbf{1}$, or $[\Gamma]$. In the first and second case, conclude by applying the (\top) and $(\mathbf{1})$ rule respectively. In the final case, note that if the next configuration in the run-tree $(\psi', [\Gamma \setminus \{\psi'\}, \psi])$ then the derivation is not maximal. Thus the next rule is $(\star, [\Gamma])$. Since this is a leaf of the run-tree, it is necessarily a final configuration of the form $(\star, [\psi, \psi^\perp])$. We conclude by applying an (id) rule. □

From Lemma 6.2.6 and Lemma 6.2.8, we conclude:

Theorem 6.2.3. μMLL_{\oplus}^* provability reduces to BVASS reachability.

Finally, from Theorem 6.2.2 and Theorem 6.2.3, we have:

Theorem 6.2.4. μMLL_{\oplus}^* provability is equivalent to BVASS reachability.

Corollary 6.2.4.1. μMLL_{\oplus}^* provability is equivalent to MELL provability. Furthermore, μMLL_{\oplus}^* is at least Ackermann-hard.

6.3 The complexity of μMALL^∞

In this section, we will show that μMALL^∞ provability reduces from the non-halting of Minsky machines. Our reduction is inspired by the one in [Kuz20] for infinitary commutative action logic.

Action logic is the equational theory of non-commutative intuitionistic logic with a unary operator $*$ called the Kleene star. Action logic without left and right linear implication (called *residuals* in the community of algebraic logic) is called the Kleene algebra. An action lattice A is said to be $*$ -continuous, if for any $a \in A$, we have $a^* = \sup\{a^n \mid n \in \omega\}$ where $a^0 = \mathbf{1}$ and $a^n = \overbrace{a \otimes \dots \otimes a}^{n \text{ times}}$ for $n > 0$. Consequently, we have the following rules for the Kleene star.

$$\frac{\Gamma, \mathbf{1} \vdash \psi \quad \Gamma, \varphi \vdash \psi \quad \Gamma, \varphi \otimes \varphi \vdash \psi \quad \dots}{\Gamma, \varphi^* \vdash \psi} (*_\ell) \quad ; \quad \frac{\Gamma_1 \vdash \varphi \quad \dots \quad \Gamma_n \vdash \varphi}{\Gamma_1, \dots, \Gamma_n \vdash \varphi^*} (*_r^n)$$

Kuznetsov [Kuz20] reduces to the non-halting of Minsky machines to provability in infinitary commutative action logic. Note that, we can encode of the Kleene star as $\varphi^* = \mu x.(\mathbf{1} \oplus (\varphi \otimes x))$ and then these rules exactly correspond the fixed point rules on this formula in the $\mu_\omega\text{MALL}^+$. However, there are a couple of issues with directly importing Kuznetsov's result even for $\mu_\omega\text{MALL}^+$. Action logic is intuitionistic, requiring an extension of the conservativity of linear logic over intuitionistic linear logic [Sch91] to $\mu_\omega\text{MALL}$. Strictly speaking, this is not possible since $\mathbf{0}$ is itself encodable as a fixed point *viz.* $\mu x.x$, and it is not obvious what language such a conservativity result might hold over.

An extra issue in the case of μMALL^∞ is the inference rule for the Kleene star is ω -branching. Therefore, one would also need to establish translations from the omega-branching μMALL (say, $\mu_\omega\text{MALL}$ for example) to μMALL^∞ (and vice versa) which seem to be quite non-trivial and require yet further intermediary systems. Therefore, we provide a direct reduction.

6.3.1 The hardness result

Fix a Minsky machine $\mathcal{M} = (Q, r_1, r_2, I)$. We construe $\{a, b, z_a, z_b\} \cup Q$ as a set of propositional variables (assuming $\{a, b, z_a, z_b\} \cap Q = \emptyset$). We use a and z_a (respectively b and z_b) to represent the contents of the register r_1 (respectively r_2). We encode instructions (with any extra 0-ary instruction **zero-check**) as follows:

$$\begin{aligned} [\text{INC}(p, r_1, q)] &\triangleq p \wp (q^\perp \otimes a^\perp) \\ [\text{JZDEC}(p, r_1, q_0, q_1)] &\triangleq (p \wp (q_0^\perp \oplus z_a^\perp)) \& ((p \wp a) \wp q_1^\perp) \\ [\text{zero-check}] &\triangleq (z_a \otimes z_a^\perp) \oplus (z_b \otimes z_b^\perp) \end{aligned}$$

For any formula φ , define $\varphi^* = \mu x.(\mathbf{1} \oplus (\varphi \otimes x))$ and $\varphi^\omega = \nu x.(\perp \& (F \wp x))$. Observe that $(\varphi^*)^\perp = (\varphi^\perp)^\omega$.

Proposition 6.3.1. *For any formula φ and any $n \in \mathbb{N}$, $\vdash \varphi^n, (\varphi^\perp)^*$ is provable in μMALL^∞ .*

Proof. We proceed by induction on n . We call π_φ^n the proof of $\vdash \varphi^n, (\varphi^\perp)^*$.

Base Case. $n = 0$. We have

$$\frac{\frac{\frac{}{\vdash \mathbf{1}} (\mathbf{1})}{\vdash \mathbf{1} \oplus (\varphi^\perp \otimes (\varphi^\perp)^*)} (\oplus_1)}{\vdash (\varphi^\perp)^*} (\mu)}$$

Induction Case. $n = m + 1$. We have

$$\frac{\frac{\frac{}{\vdash \varphi, \varphi^\perp} (\text{id}) \quad \frac{\text{IH} = \pi_\varphi^m}{\vdash \varphi^m, (\varphi^\perp)^*} (\otimes)}{\vdash \varphi^{m+1}, \varphi^\perp \otimes (\varphi^\perp)^*} (\otimes)}{\vdash \varphi^{m+1}, (\varphi^\perp)^*} (\mu), (\oplus_2)$$

□

Finally, we encode the invariant to be maintained by

$$\boxed{\text{Inv} \triangleq ((a^\perp)^* \otimes (b^\perp)^* \otimes \text{CH}_Q) \oplus ((b^\perp)^* \otimes z_a) \oplus ((a^\perp)^* \otimes z_b).$$

It checks one of the three following conditions: (i) the control is at a valid configuration (ii) r_1 is zero (iii) r_2 is zero. Note that $[q] = q$ where the left-hand side is the state q and the right-hand side is the propositional variable q .

Lemma 6.3.1. *For all $m, n \in \mathbb{N}$, $\vdash z_a, \text{CH}_I^m, b^n, \text{Inv}$ and $\vdash z_b, \text{CH}_I^m, a^n, \text{Inv}$ are provable in μMALL^∞ .*

Proof. We will show that $\vdash z_a, \text{CH}_I^m, b^n, \text{Inv}$ is provable for any m, n by induction on m . A symmetric proof will work for $\vdash z_b, \text{CH}_I^m, a^n, \text{Inv}$.

Base Case. $m = 0$. We have

$$\frac{\frac{\frac{\pi_b^n}{\nabla} \quad \frac{}{\vdash z_a, z_a^\perp} (\text{id})}{\vdash z_a, b^n, (b^\perp)^* \otimes z_a^\perp} (\otimes)}{\vdash z_a, b^n, \text{Inv}} (\oplus_2)$$

Induction Case. $m = k + 1$. We have

$$\frac{\frac{\frac{}{\vdash z_a, z_a^\perp} (\text{id}) \quad \frac{\text{IH}}{\nabla} \quad \frac{}{\vdash z_a, \text{CH}_I^k, b^n, \text{Inv}} (\otimes)}{\vdash z_a, z_a \otimes z_a^\perp, \text{CH}_I^k, b^n, \text{Inv}} (\otimes)}{\vdash z_a, \text{CH}_I^{k+1}, b^n, \text{Inv}} (\oplus_1)$$

□

Lemma 6.3.2. *If \mathcal{M} performs k steps from $\langle p, m, n \rangle$, then $\vdash \text{CH}_I^k, p, a^m, b^n, \text{Inv}$ is derivable in μMALL^∞ .*

Proof. We will proceed by induction on k . The base case is $k = 0$ i.e. \mathcal{M} is at $\langle p, m, n \rangle$. We have that

$$\frac{\frac{\frac{\pi_a^m}{\nabla} \quad \frac{\pi_b^n}{\nabla} \quad \frac{}{\vdash p, p^\perp} (\text{id})}{\vdash a^m, (a^\perp)^* \otimes b^n, (b^\perp)^* \otimes p, \text{CH}_Q} (\otimes)}{\vdash p, a^m, b^n, (a^\perp)^* \otimes (b^\perp)^* \otimes \text{CH}_Q} (\oplus_1)$$

For the induction case, assume $k = \ell + 1$. We will examine the first step of the execution. We have three sub-cases: incrementation, decrementation of a non-zero register, and decrementation of a zero-valued register.

Case 1. The first step is $\text{INC}(p, r_1, q)$ ($\text{INC}(p, r_1, q)$ is similar). We have

$$\frac{\frac{\frac{}{\vdash p^\perp, p} (\text{id}) \quad \frac{\text{IH}}{\nabla} \quad \frac{}{\vdash \text{CH}_I^\ell, q, a^{m+1}, b^n, \text{Inv}} (\otimes)}{\vdash q \wp a, \text{CH}_I^\ell, a^m, b^n, \text{Inv}} (\otimes)}{\vdash [\text{INC}(p, r_1, q)]^\perp, \text{CH}_I^\ell, p, a^m, b^n, \text{Inv}} (\otimes)}{\vdash \text{CH}_I^{\ell+1}, p, a^m, b^n, \text{Inv}} (\oplus)$$

where the $|I|$ -ary \oplus chooses the instruction $\text{INC}(p, r_1, q)$.

Case 2. The first step is $\text{JZDEC}(p, r_1, q_0, q_1)$ and $m = 0$ (a first step $\text{JZDEC}(p, r_2, q_0, q_1)$ with $n = 0$ is similar). We have:

$$\frac{\frac{\frac{\frac{\text{IH}}{\triangle}}{\vdash p^\perp, p} \text{(id)} \quad \frac{\vdash q_0, \text{CH}_I^\ell, b^n, \text{Inv}}{\vdash (q_0 \& z_a), \text{CH}_I^\ell, b^n, \text{Inv}} \text{(\&)}}{\vdash p^\perp \otimes (q_0 \& z_a), \text{CH}_I^\ell, p, b^n, \text{Inv}} \text{(\otimes)}}{\vdash [\text{JZDEC}(p, r_1, q_0, q_1)]^\perp, \text{CH}_I^\ell, p, b^n, \text{Inv}} \text{(\oplus_1)}}{\vdash \text{CH}_I^{\ell+1}, p, b^n, \text{Inv}} \text{(\oplus)}$$

where we select the appropriate instruction by applying the corresponding \oplus inference, as in Case 1 and where the $\dot{\vdash}$ part repeats the pattern decreasing the number of CH_I formulas in the sequent.

Case 3. The first step is $\text{JZDEC}(p, r_1, q_0, q_1)$ and $m \neq 0$ (a first step $\text{JZDEC}(p, r_2, q_0, q_1)$ with $n \neq 0$ is similar). We have:

$$\frac{\frac{\frac{\text{IH}}{\triangle}}{\vdash p^\perp, p} \text{(id)} \quad \frac{\vdash a^\perp, a}{\vdash p^\perp \otimes a^\perp, p, a} \text{(\otimes)}}{\vdash (p^\perp \otimes a^\perp) \otimes q_1, \text{CH}_I^\ell, p, a^m, b^n, \text{Inv}} \text{(\otimes)}}{\vdash [\text{JZDEC}(p, r_1, q_0, q_1)]^\perp, \text{CH}_I^\ell, p, a^m, b^n, \text{Inv}} \text{(\oplus_2)}}{\vdash \text{CH}_I^{\ell+1}, p, a^m, b^n, \text{Inv}} \text{(\oplus)}$$

where we select the appropriate instruction by applying the corresponding \oplus inference. \square

We appeal to the focussing property of μMALL^∞ to prove the opposite direction. Therefore, for the rest of this subsection we assign atomic polarities as follows: a, b and q are negative for any state $q \in Q$, z_a, z_b are positive.

Lemma 6.3.3. *In any focused proof of $\vdash \text{CH}_I^k, p, a^m, b^n, \text{Inv}$ where $k \neq 0$, the positive formula Inv is not the focus.*

Proof. Suppose Inv is the focus, aiming at a contradiction. We have three cases depending on the rule applied on Inv .

Case 1. The first rule is (\oplus_1) with principal formula Inv . Then, the auxilliary formula is $(a^\perp)^* \otimes (b^\perp)^* \otimes \text{CH}_Q$. Since the outermost connective is positive, we must immediately apply the ternary tensor rule which has three premisses of the form (i) $\vdash \Gamma, (a^\perp)^*$, (ii) $\vdash \Delta, (b^\perp)^*$, and (iii) $\vdash \Xi, \text{CH}_Q$ with foci $(a^\perp)^*$, $(b^\perp)^*$ and CH_Q respectively. Consider the first premiss. We are still in the positive phase with $(a^\perp)^*$ as focus. The next rules are (μ) and (\oplus) rule respectively. If Γ is non-empty one cannot choose 1 hence the next rule is a \otimes . Therefore we again have premisses of the form $\vdash \Gamma', a^\perp$ and $\vdash \Gamma'', (a^\perp)^*$ where $\Gamma = \Gamma', \Gamma''$. Observe that we can repeat our argument on the second premiss since $(a^\perp)^*$ is still under focus but we cannot continue like that ad infinitum since that would be an infinite branch without any progressing threads. Hence, at some point $\mathbf{1}$ is chosen. Therefore we conclude $\Gamma = \{a^{m'}\}$ where $m' \leq m$. Similarly, $\Delta = \{b^{n'}\}$ where $n' \leq n$. Therefore $\Xi = \text{CH}_I^k, p, a^{m-m'}, b^{n-n'}$. Now consider the last premiss $\vdash \Xi, \text{CH}_Q$ (with focus CH_Q). We are forced to choose q^\perp for some $q \in Q$, which shall be conclusion of an (id) rule and Ξ must be $\{q\}$. This is only possible if p is chosen, $m = m', n = n'$ and $k = 0$. But we assumed $k \neq 0$. Contradiction!

Case 2. The first rule is (\oplus_2) with principal formula Inv . Then, the auxilliary formula is $(b^\perp)^* \otimes z_a$. Since the outermost connective is positive, we must immediately apply the tensor rule. One of the premisses is of the form $\vdash \Delta, z_a$ with z_a as focus and we cannot apply any inference rule. This is because Δ cannot be z_a^\perp so the identity rule is ruled out and z_a is a positive atom. The reasoning is symmetric if the first rule is (\oplus_3) . \square

Lemma 6.3.4. *If $\vdash \text{CH}_I^k, p, a^m, b^n, \text{Inv}$ is derivable in μMALL^∞ , then \mathcal{M} performs k steps starting from $\langle p, m, n \rangle$.*

Proof. By Theorem 4.4.1 and Theorem 4.5.2 there is a cut-free and focussed proof of $\vdash \text{CH}_I^k, p, a^m, b^n, \text{Inv}$. We will induct on the height of this proof. The base case is vacuous. For the induction case, from Lemma 6.3.3, we get that CH_I is the focus. Recall that CH_I is a n -ary \oplus for some n that "chooses" the encoding of an instruction and **zero-check**. Suppose **zero-check** is chosen. Since it is a positive formula, it will be the principal formula again. *Wlog* assume $z_a \otimes z_a^\perp$ is chosen, which being a positive formula will be the focus. The tensor rule has a premiss of the form $\vdash \Delta, z_a$ and z_a being positive, the only applicable rule is (id). This is impossible. So **zero-check** cannot be chosen and some other instruction is chosen. Observe that if it is an incrementation, a focussed proof will follow exactly like the proof fragment exhibited while proving Lemma 6.3.2. If a decrementation is chosen, we need to make sure that the control goes to the appropriate state depending on whether the register in question is zero or not. Observe that an erroneous choice is doomed to fail. We have two cases:

Case 1. Suppose we have $\vdash (p^\perp \otimes a^\perp) \otimes q, \text{CH}_I^\ell, b^n, \text{Inv}$. Here $(p^\perp \otimes a^\perp) \otimes q$ is the focus since in the earlier step $[\text{JZDEC}(p, r_1, q_0, q_1)]^\perp$ was necessarily the focus for some state q_0 . Therefore we have sequent of the form $\vdash \Delta, a^\perp$ where a^\perp is the focus and Δ cannot be $\{a\}$.

Case 2. Suppose we have $\vdash p^\perp \otimes (q_0 \& z_a), \text{CH}_I^\ell, p, a^m, b^n, \text{Inv}$. As before $p^\perp \otimes (q_0 \& z_a)$ is the focus and we have the sequent $\text{CH}_I^\ell, z_a, a^m, b^n, \text{Inv}$. Using the exact same argument as in Lemma 6.3.3, we have that **Inv** cannot be the focus. If an instruction (other than **zero-check**) is chosen which has the state t as the current state then a focused proof leads us to a sequent of the form Δ, t^\perp with t as focus and $\Delta \neq \{t^\perp\}$. Therefore only **zero-check** can be chosen and we end up in the sequent $\vdash z_a, a^m, b^n, \text{Inv}$. It is clear that this does not have a proof.

Therefore after choosing a decrementation one cannot be led astray into the wrong state. Hence the proof follow exactly as in Lemma 6.3.2 and we will end up in a subproof of the shape $\vdash \text{CH}_I^{k-1}, q, a^{m'}, b^{n'}, \text{Inv}$ for some state q and some natural number m', n' . We can then apply the induction hypothesis and get the desired result. \square

From Lemma 6.3.2 and Lemma 6.3.4, we have the following.

Theorem 6.3.1. \mathcal{M} performs n steps starting from $\langle q_s, 0, 0 \rangle$ iff $\vdash \text{CH}_I^n, q_s, \text{Inv}$ is derivable in μMALL^∞ .

Theorem 6.3.2. A Minsky machine \mathcal{M} has an infinite run from the state q_s iff $\text{CH}_I^\omega, q_s, \text{Inv}$ is derivable in μMALL^∞ .

Proof. For the only if part we assume that \mathcal{M} loops. So, \mathcal{M} runs for n steps for all $n \in \mathbb{N}$. Therefore, by Theorem 6.3.1, we have that $\Gamma_n = \text{CH}_I^n, q_s, \text{Inv}$ is derivable for all $n \in \mathbb{N}$. Let us call π_n a proof of Γ_n , for $n \in \mathbb{N}$. We have

$$\frac{\frac{\frac{\pi_0}{\vdash q_s, \text{Inv}}}{\vdash \perp, q_s, \text{Inv}} (\perp) \quad \frac{\frac{\frac{\pi_1}{\vdash \text{CH}_I, q_s, \text{Inv}}}{\vdash \text{CH}_I, \perp, q_s, \text{Inv}} (\perp) \quad \vdash \text{CH}_I, \text{CH}_I^\omega, q_s, \text{Inv}}{\vdash \text{CH}_I \wp \text{CH}_I^\omega, q_s, \text{Inv}} (\wp), (\nu), (\&)}{\vdash \text{CH}_I^\omega, q_s, \text{Inv}} (\nu), (\&)}{\vdash \text{CH}_I^\omega, q_s, \text{Inv}} (\nu), (\&)} \quad \vdots$$

Observe that this pre-proof is indeed a proof as the right-most non-wellfounded branch is validated by a thread on CH_I^ω . For the if direction assume that we have a proof π of $\vdash \text{CH}_I^\omega, q_s, \text{Inv}$. Observe that for all $n \in \mathbb{N}$ we have a proof of $\text{CH}_I^n, q_s, \text{Inv}$:

$$\frac{\frac{\pi_{\text{CH}_I}^n}{\vdash \text{CH}_I^n, (\text{CH}_I^\perp)^*} \quad \frac{\pi}{\vdash \text{CH}_I^\omega, q_s, \text{Inv}}}{\vdash \text{CH}_I^n, q_s, \text{Inv}} (\text{cut})$$

By Theorem 6.3.1, \mathcal{M} runs at least n steps for all $n \in \mathbb{N}$. We collect all these runs and get a finitely branching infinite tree rooted at $\langle q_s, 0, 0 \rangle$. König's lemma ensures that there is an infinite run of \mathcal{M} from q_s . \square

As a direct consequence of Theorem 6.3.2 and Theorem 6.1.1 we have the following:

Corollary 6.3.2.1. μMALL^∞ provability is Π_1^0 -hard.

Corollary 6.3.2.2. μMALL^∞ provability is $(\Sigma_1^0 \cup \Pi_1^0)$ -hard.

Proof. Since $\mu\text{MALL}^* \subseteq \mu\text{MALL}^\infty$, the Σ_1^0 -hardness of Corollary 6.2.1.1 is inherited. From Corollary 6.3.2.1, μMALL^∞ is Π_1^0 -hard. Therefore, it is $(\Sigma_1^0 \cup \Pi_1^0)$ -hard. \square

6.3.2 Towards a tight upper bound

There is a trivial upper bound for μMALL^∞ in the analytical hierarchy. Provability can be trivially encoded as

$$\boxed{\exists \text{ pre-proof } \pi. \forall \text{ branches } \beta \in \pi. \exists \text{ thread } t \in \beta. t \text{ is progressing}}$$

Checking the progress condition is arithmetical. Therefore μMALL^∞ provability is in Σ_3^1 . This leaves a chasm between the $(\Sigma_1^0 \cup \Pi_1^0)$ lower bound and the Σ_3^1 upper bound. We conjecture that both of these can be improved.

Open Question

μMALL^∞ provability is Π_1^1 -complete.

We will sketch a few ideas in both directions. For the lower bound, we note that it is not surprising that Σ_1^0 and Π_1^0 bounds are not tight. The reduction for the Σ_1^0 bound uses only ν -free formulas while the Σ_1^0 uses alternation-free formulas. The full expressiveness of μMALL^∞ is possibly captured only with formulas with alternations of fixed point operators. On a related note, the universal Horn theory of $*$ -continuous Kleene algebras is Π_1^1 -complete [Koz02].

Theorem 6.3.3 ([AH89]). *Given a Minsky machine \mathcal{M} and an initial state q_s , checking whether there exists an infinite run from q_s such that q_s occurs infinitely often is Σ_1^1 -complete.*

Consequently, checking whether there does *not* exist an infinite run satisfying such a Büchi condition is Π_1^1 -complete. We envision a formula of the form $\mu x. \nu y. f(x, y)$ such that:

- The μ unfolding corresponds to the fact q_s is visited.
- One ν unfolding corresponds to one step in the configuration graph.
- A finite proof would correspond to a finite run of \mathcal{M} .
- The only infinite progressing derivations would unfold μ finitely many times, thereby ensuring that all infinite runs of \mathcal{M} visit q_s at most finitely many times.

For the upper bound, similar to ideas in Section 5.4, we conjecture that there is some wellfounded infinitely branching system provably equivalent to μMALL^∞ . Such systems are in Π_1^1 . We end this section by showing the following.

Theorem 6.3.4. *The set of $\mu_\omega\text{MALL}^+$ provable sentences are in Π_1^1 .*

Proof. Consider a function $f : \mu_\omega\text{MALL} \rightarrow \mu_\omega\text{MALL}$ that takes a $\mu_\omega\text{MALL}$ formula and replaces every occurrence of ν by ν^n for some n . Now quantifying over f is morally a quantification over a function from \mathbb{N} to \mathbb{N} (courtesy some Gödel encoding of $\mu_\omega\text{MALL}$ formulas). Now, a $\mu_\omega\text{MALL}$ formula φ is provable iff for all f , $f(\varphi)$ is provable. Any proof of $f(\varphi)$ is finitely branching and wellfounded; hence proof-search is arithmetical. Therefore $\mu_\omega\text{MALL}$ proof-search is in Π_1^1 . \square

6.4 The regularisation problem

6.4.1 Regularisation in other logics

We will start with a brief discussion on regularisation in the modal μ -calculus. The non-wellfounded calculus (with the unfolding rules for the fixed points) admits cuts whence the induced (cut-free) calculus enjoys the subformula property (with respect to Fischer-Ladner subformulas) meaning that only finitely many distinct sequents may occur in a proof.² As a result, once a particular sequent to be proved is fixed, the progress condition becomes an ω -regular property on infinite branches. This allows us to reduce *regular* completeness of the system to *non-wellfounded* completeness of the system, thanks to Rabin's basis theorem [Rab72]. This idea is implicit in Niwinski and Walukiewicz's seminal work [NW96].

They introduced a notion of validity game for the μ -calculus, whose Opponent strategies may be identified with countermodels, and whose Prover strategies may be identified with certain 'non-wellfounded' derivations in a simple cut-free sequent calculus. The determinacy of these games immediately yields completeness of the induced class of non-wellfounded derivations. Furthermore, upon careful inspection of the winning condition, one can deduce that these games are determined by strategies of uniformly bounded finite memory, which in turn yields regular completeness of μ -calculus.

This reduction is, a priori, non-constructive: it asserts the existence of a regular proof but does not tell us how to construct one from a given non-wellfounded one. However, it is possible to define a constructive procedure that 'cuts' branches of an infinite proof tree to transform it into a regular one, using automata-theoretic techniques. In this subsection, we will sketch this idea.

Let Seq be the set of μ -calculus sequents (construed as sets of formulas). The correctness condition can be construed as an infinite word language $\Theta \subseteq \text{Seq}^\omega$. Furthermore, we define a graph $P = (\text{Seq}, E)$ such that (Γ, Γ') if there is an instance of a μ -calculus inference rule with $\vdash \Gamma$ as conclusion and $\vdash \Gamma'$ one of the premisses.

Theorem 6.4.1. *Given a non-wellfounded proof π , there is a regular proof π_{reg} of the same conclusion obtained by only cutting branches of π and adding back edges to descendants i.e. every infinite branch of π_{reg} has the form xy^ω , where x and xy are finite prefixes of some infinite branch of π .*

Proof. We write Σ for the set of sequents occurring in π . We note that Σ is finite and that $\Theta_\Sigma := \Theta \cap \Sigma^\omega$ is ω -regular. We also fix a deterministic parity automaton \mathcal{T}_Σ that recognises Θ_Σ . Let us write $P \times \mathcal{T}_\Sigma$ for the annotation of P by states of \mathcal{T}_Σ obtained from running \mathcal{T}_Σ through the branches of P .

Let $B = (s_i, q_i)_i$ be an infinite branch of $P \times \mathcal{T}_\Sigma$; since $(s_i)_i$ is accepted by \mathcal{T}_Σ , there must be some even colour c that is the least among the infinitely occurring colours of $(q_i)_i$. By the pigeon-hole principle, B must have a prefix xy where $y : (s, q) \rightarrow^+ (s, q)$ such that the least colour in y is even (and $\geq c$). We define π_{reg} by cutting each such infinite branch B at xy and placing a back-edge between the latter and former occurrences of (s, q) in y . (More formally, the least such cuttings will form a bar through $P \times \mathcal{T}_\Sigma$ that can be recursively obtained by blind search from the root). Note that the only simple loops of π_{reg} are the y 's obtained from such cuttings.

For the correctness of π_{reg} , note that any infinite branch must visit some simple loops infinitely often. By construction, the least colour occurring infinitely often will be the least of the even colours associated to such simple loops, and so must be even. \square

6.4.2 Regularisation in μMALL

It is worth pointing out that the argument we mentioned for regularisation in the μ -calculus in Section 6.4.1 can in fact be adapted to certain fragments of μMALL , in particular the additive fragment. Writing μALL^∞ and μALL° for the restriction of μMALL^∞ and μMALL° , respectively, to only additive connectives, we have:

Theorem 6.4.2. *If $\vdash \Gamma$ is provable in μALL^∞ , then it is also (cut-free) provable in μALL° .*

²More precisely, the number of formulas in $\text{FL}(\Gamma)$ is linear in the size of Γ and subsequently there are at most $2^{O(|\Gamma|)}$ such sequents.

Proof. By Theorem 4.5.2, we may assume that $\vdash \Gamma$ has a cut-free μALL^∞ proof π . Note that each (non-cut) rule of μALL preserves, bottom-up, the number of formulas in a sequent. Since there are only finitely many formulas that can occur, π contains at most finitely many distinct sequents. The rest of the proof follows exactly as the proof of Theorem 6.4.1. \square

Note that this also implies the *decidability* of μALL^∞ since, after guessing a (exponential-size) pre-proof of Γ , checking that it is a proof is decidable (in space polynomial in the size of the proof).

Corollary 6.4.2.1. μALL^∞ (equivalently μALL°) is decidable in EXSPACE.

We stop short of attempting to optimise this result since, in particular, it seems sensitive to the precise presentation of μALL . Often $(\mu)\text{ALL}$ is presented with *exactly* two formulas in a sequent, e.g. [San02, FS13], and this invariant is maintained by the rules of $(\mu)\text{ALL}$. In such a presentation, there are only quadratically many distinct sequents in a μALL^∞ proof.

However, the regular and non-wellfounded calculi of μMALL are different, in general. This follows immediately from the complexity results obtained in the previous sections.

Theorem 6.4.3. There are theorems of μMALL^∞ that are not provable in μMALL° .

Proof. By Corollary 6.2.1.2, μMALL° is Σ_1^0 -complete and by Corollary 6.3.2.2, μMALL^∞ is $(\Sigma_1^0 \cup \Pi_1^0)$ -hard. Since $\Sigma_1^0 \subsetneq \Pi_1^0$ and $\mu\text{MALL}^\circ \subseteq \mu\text{MALL}^\infty$, we conclude that actually $\mu\text{MALL}^\circ \subsetneq \mu\text{MALL}^\infty$. \square

Observe that this proof is apparently non-constructive in the sense that we do not explicitly exhibit a sequent in $\mu\text{MALL}^\infty \setminus \mu\text{MALL}^\circ$. Furthermore, note that we could have made our conclusion already from the Π_1^0 -hardness of μMALL^∞ (not requiring the $(\Sigma_1^0 \cup \Pi_1^0)$ -hardness). While it is clear that not all sequents of the form $\vdash \text{CH}_I^\omega, q_s, \text{Inv}$ from Section 6.3 can be derivable in μMALL° , it is not clear which particular Minsky machine \mathcal{M} to choose to witness this underderivability. In fact, the argument can indeed be constructivised using established recursion-theoretic techniques, namely the notion of *productive function* [Soa14].

Definition 6.4.1. Let $\ulcorner \bullet \urcorner$ be the Gödel encoding of recursively enumerable sets. Let W_x be the set S such that $\ulcorner S \urcorner = x$. A set P is called **productive** if there exists an effectively computable partial function f such that if $W_x \subset P$, then $f(x)$ is defined and is an element of $P \setminus W_x$. The function f is called a **productive function**.

Proposition 6.4.1. The set $\overline{K} = \{x \mid x \notin W_x\}$ is productive with respect to the trivial productive function $\text{id}(x) = x$.

Proof. Suppose $W_x \subset \overline{K}$. Then, we need to show that $x \in \overline{K} \setminus W_x$. Suppose $x \in W_x$ then $x \notin \overline{K}$. But since $W_x \subseteq \overline{K}$, $x \in \overline{K}$. So, $x \notin W_x$. Then, $x \in \overline{K}$ by definition. Therefore, $x \in \overline{K} \setminus W_x$. \square

Theorem 6.4.4. If P is productive and it is many-one reducible to P' , then P' is also productive.

Proposition 6.4.2. \overline{K} is a Π_1^0 -complete set.

Proof. Note that $K = \{x \mid x \in W_x\} = \{x \mid \text{TM}_x \text{ halts on } x\}$ is the complement of \overline{K} . Membership in K is in Σ_1^0 since W_x is a recursively enumerable set for all x . We will show that K is not recursive. Define

$$f(x) = \begin{cases} \text{TM}_x(x) + 1 & \text{if } x \in K; \\ 0 & \text{otherwise.} \end{cases}$$

Note that if K is recursive then f is recursive. Suppose $f = \text{TM}_y$ for some y and f is recursive. If $y \notin K$ then by definition of K , TM_y does not halt on y . Then f cannot be recursive. If $y \in K$, then $f(y) = \text{TM}_y(y) + 1$. Then, $f \neq \text{TM}_y$. Therefore, f is not recursive.

Therefore, K is Σ_1^0 -complete and consequently, \overline{K} is Π_1^0 -complete. \square

Let P be the set of theorems of μMALL^∞ . By Corollary 6.3.2.1, \overline{K} is m -reducible to P . Therefore, P is a productive set with some computable productive functions f . Since μMALL° is Σ_1^0 -complete, the set of theorems of μMALL° is a W_y for some y ; $W_y \subset P$ by Theorem 6.4.3. Thus, $f(y) \in P \setminus W_y$; via inverses of encodings it is a sequent provable in μMALL^∞ but not in μMALL° . Theoretically, this example can be explicitly extracted from the reasoning presented above. In order to do so, one needs to track the m -reduction of K to P (via the reduction of Theorem 6.1.1). This yields a concrete algorithm for f , which can then be applied to y .

In this chapter, we showed that μMALL^* is Σ_1^0 -complete which straightforwardly implied that so are $\mu\text{MALL}^{\text{ind}}$ and μMALL° . On the other hand, we showed that μMALL^∞ is $(\Sigma_1^0 \cup \Pi_1^0)$ -hard and consequently proves strictly more theorems than μMALL° . As we conclude this part, we note that the study of μMALL systems is an interesting and complicated topic that can benefit techniques from automata theory related topics such as Petri nets and action lattices. Interestingly, wellfounded infinitely branching systems have appeared quite naturally in both chapters and seem to be a handle on grasping the non-wellfounded calculus. In the next part, we will move to a more intentional understanding of μMALL : the provability would not matter as much as the proofs themselves.

Part II


Infinets: the parallel syntax

Chapter 7

Proof-nets for systems of μ MLL

(A new hope)

The goal of this part is to develop a proof-net formalism (*i.e.* a parallel syntax) for the non-wellfounded calculus of μ MALL. Since proof-nets for additives and units are cumbersome, we will concentrate on the multiplicative fragment without units. We incrementally develop the theory of non-wellfounded proofs (*a.k.a.* *infinets*). In this first chapter, we first recall MLL proof-nets via an algebraic presentation due to Curien [Cur05], in Section 7.1. In non-wellfounded proof-nets, one needs to connect nodes by infinitely long paths. To formalise such concepts in infinitary graph theory, heavy topological machinery is necessary. We sacrifice the visual clarity of graphs to consider non-wellfounded proof-nets (or *infinets*) in the algebraic presentation. The development of *infinets* will be as follows.



In Section 7.2, we straightforwardly enhance this presentation to develop proof-nets for the finitary fragment of μ MLL $^\infty$ (*viz.* μ MLL *). We briefly revert back to the graphical presentation of proof-nets in Section 7.3 to discuss proof-nets for μ MLL $^{\text{ind}}$ and μ MLL $^{\text{circ}}$. In Section 7.4, we semi-formally discuss the several pitfalls of adapting proof-nets to the non-wellfounded setting and the various constructs appearing forthcoming chapter.




Figure 7.1: Illustration of partial syntax trees and proof-nets

7.1 MLL proof-nets: a closer look

In this section, we will discuss an alternate algebraic presentation of proof-nets due to Curien [Cur06]. In the context of MLL, it is a trivial reformulation of usual the graphical presentation.

Firstly, we will need to use the presentation of linear logic sequents as sets of formula occurrences. We begin by recalling that the syntax tree of an occurrence F induces a prefix closed language, $\mathcal{L}_F \subset \{l, r\}^*$ such that there is a natural bijection between the words in \mathcal{L} and the set of all simple paths starting from the root of the syntax tree.

We remind the readers that for a regular language \mathcal{L} , $\bar{\mathcal{L}}$ denotes its prefix-closure (cf. Chapter 2) and that negation is an involution over addresses such that $l^\perp = r$, $r^\perp = l$ and $i^\perp = i$ (cf. Section 3.1).

Definition 7.1.1. A **partial syntax tree**, F^U , is a subtree of the syntax tree of the formula occurrence, F , such that $U \subseteq \mathcal{L}_F$ and U represents a bar of the syntax tree of F i.e. any $u, u' \in U$ are pairwise disjoint¹ and for each $uav \in U$, there is a v' such that $ua^\perp v' \in U$. For $u \in \bar{U}$, we denote by (F, u) the unique suboccurrence of F with the address $\text{addr}(F).u$.

Example 7.1.1. Let $F = (a\wp a)\wp(a^\perp \otimes a^\perp)_\alpha$ be a formula occurrence. Then $\mathcal{L}_F = \alpha.\{\varepsilon, l, r, ll, lr, rl, rr\}$. The language $U = \alpha.\{lr, rl, rr\}$ is not a partial syntax tree but $U = \alpha.\{l, rl, rr\}$ is.

We illustrate a schematic partial syntax tree in Figure 7.1a. MLL proof-nets without cuts can be seen as a forest of partial syntax trees of the occurrences in the conclusion sequent and axiom links between their leaves (cf. Figure 7.1b). To incorporate cuts we need to add the partial syntax tree of the cut occurrences (along with the axioms links involving their leaves) and links between dual cut occurrences.

Definition 7.1.2. An MLL **proof-structure** is a 3-tuple $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ where:

- $\lambda \in \omega$.
- for all $i \in \lambda$, $F_i^{U_i}$ is a partial syntax tree; $\{F_i\}_{i \in \lambda}$ is called the set of **doors**.
- \mathfrak{K} is the set of cuts i.e. a (possibly empty) set of disjoint subsets of $\{F_i\}_{i \in \lambda}$ of the form $\{C, C^\perp\}$; and,
- Θ is the set of axiom links i.e. a partition of the set of **leaves**, $\mathcal{L} = \bigcup_{i \in \lambda} \{\alpha_i u_i \mid \text{addr}(F_i) = \alpha_i, u_i \in U_i\}$ such that each cell is pair of dual addresses i.e. of the form $\{\alpha_i u_i, \alpha_j u_j\}$ such that $\lceil (F_i, u_i) \rceil = \lceil (F_j, u_j) \rceil^\perp$.

Each cell of Θ represents an axiom, each element of \mathfrak{K} represents a cut, and $\{F_i\}_{i \in \lambda} \setminus \bigcup_{\theta \in \mathfrak{K}} \theta$ are the conclusions of the proof-structure. Observe that this presentation is more logic independent since we do not explicitly mention the operators involved in the logic. We will see this logic independence come in handy more clearly in the next chapter when we define μMLL^∞ proof-nets.

¹Defined in Chapter 2.

7.2 μMLL^* proof-nets

As a stepping stone to formulating proof-nets corresponding to μMLL^∞ , we first consider proof-nets in μMLL^* which is the proof system with the same inference rules as μMLL^∞ but with finite proofs. Recall that this logic is strictly weaker than μMLL^∞ (in particular $\vdash \nu x.x$ cannot be proved and the fixed points are interchangeable). We present them in the alternate syntax introduced in the previous section which is ultimately useful to lift μMLL^* proof-nets to μMLL^∞ . Consequently, μMLL^* proof-nets are a straightforward extension of MLL proof-nets discussed in the previous section (cf. Section 7.1).

Recall for μMALL , the syntax tree of a formula occurrence F is the (possibly infinite) unfolding tree of the Fischer-Ladner graph of F and a prefix closed language, $\mathcal{L}_F \subset \{l, r, i\}^\infty$ such that there is a natural bijection between the finite (respectively, infinite) words in \mathcal{L}_F and the finite (respectively, infinite) paths from the root in the syntax tree of F .


Example 7.2.1. Let $F = \mu x.x \wp (a^\perp \otimes a)_\alpha$. The syntax tree of F is the unfolding of $\mathfrak{G}(F)$ and induces the language $\alpha(i(li)^*r(l+r) + (il)^\omega)$. Furthermore, $F^{\{ili, irl, irr\}}$ is a partial syntax tree whereas $F^{\{ilil, irll, irrr\}}$ is not. If $u = ililir$ then $(F, u) = \alpha \wp a^\perp_{\alpha u}$.

Definition 7.2.1. A μMLL^* **proof-structure**, denoted $\mathcal{R}, \mathcal{S}, \dots$, is a 3-tuple $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ where:

- $\lambda \in \omega$;
- for all $i \in \lambda$, $F_i^{U_i}$ is a partial syntax tree with $U_i \subset \{l, r, i\}^*$; $\{F_i\}_{i \in \lambda}$ is called the set of **doors**.
- \mathfrak{K} is the set of cuts i.e. a (possibly empty) set of disjoint subsets of $\{F_i\}_{i \in \lambda}$ of the form $\{C, C^\perp\}$; and,
- Θ is the set of axiom links i.e. a partition of the set of leaves, $\mathcal{L} = \bigcup_{i \in \lambda} \{\alpha_i u_i \mid \text{addr}(F_i) = \alpha_i, u_i \in U_i\}$ such that each cell is pair of dual addresses i.e. of the form $\{\alpha_i u_i, \alpha_j u_j\}$ such that $\lceil (F_i, u_i) \rceil = \lceil (F_j, u_j) \rceil^\perp$.

An occurrence G is said to **occur** in a proof structure $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ if there exists $u \in U_i$ for some $i \in \lambda$ such that $G = (F_i, u)$.

Indeed, the only difference from Definition 7.1.2 is the shape of the partial syntax trees which are now allowed unary branching. This is the reason why we claimed that this presentation is more logic independent. Note that graphically this is basically adding more sorts of nodes to Definition 3.5.1 viz. nodes of the form:



We will now define desequentialisation, the translation of sequent proofs into proof-structures. Before that we need to extend the notion of address to proofs. The ultimate goal is to define infinit, hence we define the address of a μMLL^∞ pre-proof.

Definition 7.2.2. Given a pre-proof, π , $\text{addr}(\pi) \subseteq \{l, r, i\}^\infty$ is largest set of addresses such that if a finite address $\alpha \in \text{addr}(\pi)$ then for some φ , φ_α either occurs in an axiom or occurs infinitely often in an infinite branch in π with $\text{addr}(F) = \alpha$; and if an infinite address $\alpha \in \text{addr}(\pi)$ then there is an infinite branch β of π such that every finite prefix of α is an address of an occurrence appearing in β .

Definition 7.2.3. Let π be an MLL proof of the sequent $\vdash \Gamma$. The **desequentialisation** of π , denoted $\text{dsq}(\pi)$, is given by $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ such that the following holds.

- \mathfrak{K} is smallest set such that for any cut in π that introduces two occurrences, C and C^\perp , we have that $\{C, C^\perp\} \in \mathfrak{K}$;
- $\{F_i\}_{i \in \lambda} = \Gamma \cup \bigcup_{\kappa \in \mathfrak{K}} \kappa$ where $\lambda = |\Gamma| + |\mathfrak{K}|$;

- for every $i \in \lambda$, $U_i = \text{addr}(F_i)^{-1} \text{addr}(\pi)$; and,
- for every axiom in π of the form $\frac{\quad}{\varphi_\alpha, \varphi^{\perp\beta}} (\text{id})$, we have that $\{\alpha, \beta\} \in \Theta$.

Example 7.2.2. Consider the following proof π of the sequent $\vdash F$ where $F = \nu x.x \wp \mu x.x_\alpha$ (i.e. the same F from Example 7.2.1).

$$\frac{\frac{\frac{\quad}{\vdash \nu x.x_\alpha, \mu y.y_{\beta i}} (\text{id})}{\vdash \nu x.x_\alpha, \mu y.y_\beta} (\mu)}{\vdash \nu x.x_\alpha, \mu x.x_{\alpha r}} (\wp)}{\vdash \nu y.y_{\beta^\perp i}, \mu x.x_{\alpha r}} (\nu)}{\vdash \nu y.y_{\beta^\perp}, \mu x.x_{\alpha r}} (\text{cut})} (\wp)}{\vdash F} (\wp)$$

We choose β such that α and β are disjoint. We have that $\text{dsq}(\pi) = (\Gamma, \mathfrak{K}, \Theta)$ such that

$$\begin{aligned} \mathfrak{K} &= \{\{\mu y.y_\beta, \nu y.y_{\beta^\perp}\}\} & \Theta &= \{\{\alpha l, \beta i\}, \{\alpha r, \beta^\perp i\}\} \\ \Gamma &= \{F^{\{l,r\}}, \mu y.y_\beta^{\{i\}}, \nu y.y_{\beta^\perp}^{\{i\}}\} \end{aligned}$$

In the rest of section we rehash standard results of multiplicative proof-nets (some mentioned in Section 3.5) for μMLL^* .

7.2.1 Correctness criterion

We will lift the DR-correctness criterion to μMLL^* proof-nets. Fix a proof-structure $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ for the next set of definitions. Define

$$\mathbb{P} := \{(F_i, u) \mid \exists i \in \lambda. \exists u \in \overline{U_i} \text{ and } (F_i, u) = A \wp B\}$$

Essentially, \mathbb{P} is the set of \wp -occurrences. A **switching**, sw , of \mathcal{R} is a total function of the form $sw : \mathbb{P} \rightarrow \{l, r\}$. Fix a switching, sw , of \mathcal{R} .

Definition 7.2.4. Let $u = u_1 \dots u_n$ be a substring of a word w in $\overline{U_i}$ (u_i s being letters). Then, u is said to be **unbroken** if for all $j \in \{1, \dots, n-1\}$, either $(F_i, \nu u_1 \dots u_j) \in \mathbb{P}$ or $sw((F_i, \nu u_1 \dots u_j)) \neq u_{j+1}$ where νu is a prefix of w for some word v .

Definition 7.2.5. Define the **switching relation** $\text{SW} \subseteq \mathcal{L}^2$ such that $(x, y) \in \text{SW}$ iff either $x = y$ or one of the following holds:

- $w \neq \varepsilon$ and u and v are unbroken where $w = x \cap y$, $wu = x$, and $wv = y$; or,
- there exists $\{C, C^\perp\} \in \mathfrak{K}$ and $u, v \in \{l, r, i\}^*$ such that $\text{addr}(C) = \alpha$, $\text{addr}(C^\perp) = \alpha'$, $x = \alpha u$, $y = \alpha' v$ and u, v are unbroken.

Proposition 7.2.1. The switching relation SW is an equivalence relation.

Proof. By inspection of the definition, SW is reflexive and symmetric. Suppose $(x, y) \in \text{SW}$ and $(y, z) \in \text{SW}$. Assume that x, y , and z are distinct. Let $w = x \cap y$ and $w' = y \cap z$. There are three cases:

Case 1. $w \neq \varepsilon$ and $w' \neq \varepsilon$:

Let $x = wu$, $y = wv = w'v'$ and $z = w't$ with u, v, v' , and t unbroken. Assume $|w| \leq |w'|$. Then, by Levi's Lemma, there exists p such that $w'p = w$ and $v = pv'$. Therefore, $z = w'pt$. So, we have $x \cap z = w$. Note that pt is unbroken since v and t are unbroken. Therefore, $(x, z) \in \text{SW}$. Now assume $|w'| \geq |w|$. Then by Levi's Lemma, there exists p such that $w = w'p$ and $pv = v'$. Therefore, $x = w'pu$. So, we have $x \cap z = w'$. Note that pu is unbroken since v' and u are unbroken. Therefore, $(x, z) \in \text{SW}$.

Case 2. $w \neq \varepsilon$ and $w' = \varepsilon$:

Recall that $\text{addr}(C) \cap \text{addr}(C^\perp) = \varepsilon$. Let $x = wu$, $y = wv$, and $\{C, C^\perp\} \in \mathfrak{K}$ such that $\text{addr}(C) = \alpha$, $\text{addr}(C^\perp) = \alpha'$, $y = \alpha v'$, $z = \alpha' t$ and u, v, v' , and t are unbroken. Since α is the address of a door, $|\alpha| \leq |w|$. By Levi's Lemma, there exists p such that $w = \alpha p$ and $pv = v'$. Therefore, $x = \alpha pu$. Note that pu is unbroken since v' and u are unbroken. Therefore, $(x, z) \in \text{SW}$.

Note that the case when $w = \varepsilon$ and $w' \neq \varepsilon$ is symmetric to the previous case.




Figure 7.2: A tree with nodes n_i labelled $(v_i, \theta_i, \theta'_i)$ where $v_i \in [\text{SW}]$, $\theta_i, \theta'_i \in \Theta$, and $i \in \{0, 1\}^*$. Naturally, n_i can have children n_{i0} and n_{i1} . The child n_{i0} witnesses a disconnection of θ_i with some node while n_{i1} witnesses a disconnection of θ'_i with some node. The dashed lines represent potential connections.

Case 3. $w = \varepsilon$ and $w' = \varepsilon$:

Let $\{C, C^\perp\}, \{D, D^\perp\} \in \mathfrak{K}$ such that $\text{addr}(C) = \alpha$, $\text{addr}(C^\perp) = \alpha'$, $\text{addr}(D) = \beta$, $\text{addr}(D^\perp) = \beta'$, $x = \alpha u$, $y = \alpha' v = \beta v'$, $z = \beta' t$, and u, v, v' and t are unbroken. However, by construction $\alpha' \cap \beta = \varepsilon$ (since doors have disjoint addresses). Therefore, this case is not possible. Hence, **SW** is transitive. \square

If we see the elements of \mathcal{L} as the collection of leaves of the partial syntax trees of a proof net, cells of **SW** are the connected components of that proof net under the switching sw and without axiom links.

Definition 7.2.6. Given a switching sw of \mathcal{R} , the **orthogonal graph** of \mathcal{R} is defined as the undirected bipartite (multi)graph, $(\Theta, [\text{SW}], E)$, where Θ is the set of axioms of \mathcal{R} , $[\text{SW}]$ is the set of equivalence classes of **SW** and $(x, y) \in E$ iff $x \cap y \neq \emptyset$.

Let $(\Theta, [\text{SW}], E)$ be an orthogonal graph. Let $v \in \Theta$. Define:

$$\begin{aligned} \text{Reach}^0(v) &:= \{v' \mid (v, v') \in E \cap \Theta \times [\text{SW}]\}; \\ \text{Reach}^{2n+1}(v) &:= \text{Reach}^{2n}(v) \cup \{u \mid \exists v' \in \text{Reach}^{2n}(v); (v', u) \in E \cap [\text{SW}] \times \Theta\} \\ \text{Reach}^{2n+2}(v) &:= \text{Reach}^{2n+1}(v) \cup \{u \mid \exists v' \in \text{Reach}^{2n+1}(v); (u, v') \in E \cap \Theta \times [\text{SW}]\} \\ \text{Reach}^+(v) &:= \bigcup_{n \in \omega} \text{Reach}^n(v); \end{aligned}$$

Similarly, one can define $\text{Reach}^+(v)$ for $v \in [\text{SW}]$.

Definition 7.2.7. A proof-structure is said to be **DR-correct** if for every switching, for all vertices v of the orthogonal graph $(\Theta, [\text{SW}], E)$, $\text{Reach}^+(v) = (\Theta \cup [\text{SW}]) \setminus \{v\}$.

Lemma 7.2.1 (Soundness of correctness criterion). Let π be a μMLL^* proof. Then, $\text{dsq}(\pi)$ is DR-correct.

Proof. Let π be a μMLL^* proof. Let $G = (\Theta, [\text{SW}], E)$ be the orthogonal graph of $\text{dsq}(\pi)$ with respect to an arbitrary switching sw . We need to show that for all vertices $v \in G$ $\text{Reach}^+(v) = (\Theta \cup [\text{SW}]) \setminus \{v\}$. In other words, we need to show:

(Acyclicity of G) for all $v \in G$, $v \notin \text{Reach}^+(v)$;

(Connectedness of G) for all $v, v' \in G$ such that $v \neq v'$, $v' \in \text{Reach}^+(v)$.

Proof of acyclicity. Note that a bipartite graph does not have an odd cycle. Therefore, if $v \in \text{Reach}^+(v)$ then the smallest n such that $v \in \text{Reach}^n(v)$ is odd. Suppose there exists such an n . Observe that $n \neq 0$ since if $v \in \Theta$ then $\text{Reach}^0(v) \subseteq [\text{SW}]$ (and vice versa). Therefore there exists v_0, v_1, \dots, v_{n+1} such that $v = v_0 = v_{n+1}$ and for all $i \in [n]$, $E(v_i, v_{i+1})$. We further assume that it is a simple cycle *i.e.* v_0, v_1, \dots, v_n . This is enough since if there is a cycle, then there is a simple cycle.

Assume $v_{2k} \in \Theta$ for $k \in \{0, \dots, \frac{n+1}{2}\}$. The other case will be similar. Let $u_i \in v_i \cap v_{i+1}$ (it exists by construction). When i is even, either of the following is true:

1. $u_i \cap u_{i+1} = \varepsilon$ and there are partial syntax trees $F^U, F'^{U'}$ in \mathcal{R} such that $u_i \in U, u_{i+1} \in U'$, and $\{F, F'\} \in \mathfrak{R}$. By definition of desequentialisation, this corresponds to a cut inference in π .
2. $u_i \cap u_{i+1} \neq \varepsilon$ and there is a partial syntax tree F^U in \mathcal{R} such that $u_i, u_{i+1} \in U$, and $(F, u_i \cap u_{i+1})$ is a tensor formula. By definition of desequentialisation, this corresponds to a tensor rule in π .

We collect the set of $\lfloor \frac{n}{2} \rfloor$ such tensor or cut rules of π . Let r be the bottom-most rule. Note that this may not be unique. We take r to be any one of the minimal ones and let v_i be corresponding vertex in the orthogonal graph. Assume it is of the following form:

$$\frac{\begin{array}{c} \pi_0 \\ \nabla \\ \vdash \Gamma' \end{array} \quad \begin{array}{c} \pi_1 \\ \nabla \\ \vdash \Delta' \end{array}}{\vdash \Gamma, \Delta} (r)$$

Wlog, assume v_{i-1} and v_{i+1} corresponds to an axiom in π_0 and π_1 respectively. Now, $u_{i+1} \cap u_{i+2}$ corresponds to a forking rule (*i.e.* tensor or cut). This rule has to occur above or at the same height as r . If it is strictly above then it is in π_1 ; and, the only way it can be at the same height as r is when it is r , which is ruled out since we assume that the cycle is simple. This leads to contradiction as this traps the axioms and forking rules corresponding to v_{i+1}, v_{i+2}, \dots in π_1 but we know that at some point it needs to come back to π_0 since the axiom corresponding to v_{i-1} is in π_0 .

Proof of connectedness. Let $v, v' \in G$. Assume that $\theta' \in \text{Reach}^+(\theta)$ for all distinct $\theta, \theta' \in \Theta$. We will show that $v' \in \text{Reach}^+(v)$. If $v, v' \in \Theta$ we are done by assumption. If $v \in [\text{SW}]$ then there exists $\theta \in \Theta$ such that $E(v, \theta)$. If $v' \in \Theta$, then either $v' = \theta$ (in which case $v' \in \text{Reach}^0(v)$) or $v' \in \text{Reach}^+(\theta)$ (in which case $v' \in \text{Reach}^n(\theta)$ and $v' \in \text{Reach}^{n+1}(v)$ for some n). Therefore, wlog, we can assume that we are given two distinct disconnected vertices $\theta, \theta' \in \Theta$ *i.e.* $\theta \neq \theta'$. We will prove by contradiction and assume $\theta' \notin \text{Reach}^+(\theta)$. From the definition of desequentialisation, θ, θ' corresponds to two maximal branches of π . Let Δ be the greatest common prefix of these two branches. It must be a tensor rule (applied on the formula say F_ε) or a cut rule (introducing say F_ε and F_ε^\perp).

Assume that Δ is a tensor rule. There exists $v_\varepsilon \in [\text{SW}]$ in the switching partition of G and $w, w' \in v_\varepsilon$ such that $\text{addr}(F_\varepsilon)l$ is a prefix of w and $\text{addr}(F_\varepsilon)r$ is a prefix of w' . By construction of the orthogonal graph we have that there exist $\theta_\varepsilon, \theta'_\varepsilon \in \Theta$ such that $w \in \theta_\varepsilon$ and $w' \in \theta'_\varepsilon$. Clearly $E(v_\varepsilon, \theta_\varepsilon)$ and $E(v_\varepsilon, \theta'_\varepsilon)$. If Δ is a cut rule (on F_ε and F_ε^\perp) then $F_\varepsilon, F_\varepsilon^\perp$ are the doors of $\text{dsq}(\pi)$. There exist w, w' in the partial syntax trees of $F_\varepsilon, F_\varepsilon^\perp$ respectively, such that $w, w' \in v_\varepsilon$ for some $v_\varepsilon \in [\text{SW}]$. We choose $\theta_\varepsilon, \theta'_\varepsilon \in \Theta$ such that $w \in \theta_\varepsilon$ and $w' \in \theta'_\varepsilon$.

Now since there is a one-one correspondence between the maximal branches of π and Θ , $\theta, \theta', \theta_\varepsilon$ and θ'_ε can be arranged in the lexicographical ordering of the branches of the derivation tree. Assume wlog, $\theta <_{lex} \theta'$. Then by construction $\theta <_{lex} \theta_\varepsilon <_{lex} \theta'_\varepsilon <_{lex} \theta'$. Since we assume that θ, θ' are disconnected either $\theta_\varepsilon \notin \text{Reach}^+(\theta)$ or $\theta'_\varepsilon \notin \text{Reach}^+(\theta')$. Note that it could also be the case that $\theta_\varepsilon = \theta$ or $\theta'_\varepsilon = \theta'$ but both cannot hold. Assume $\theta_\varepsilon \notin \text{Reach}^+(\theta)$ which are not connected. We can repeat the same argument to get $v_0 \in [\text{SW}]$ and $\theta_0, \theta'_0 \in \Theta$ such that $E(v_0, \theta_0), E(v_0, \theta'_0), \theta_0 <_{lex} \theta'_0$, and either $\theta_0 \notin \text{Reach}^+(\theta)$ or $\theta'_0 \notin \text{Reach}^+(\theta_\varepsilon)$.

Since we assume that θ, θ' are disconnected, one essentially repeat this process *ad infinitum*. Therefore we have an infinite binary tree \mathcal{T} (cf. Figure 7.2) labelled by the $v_\sigma \in [\text{SW}]$. This is a contradiction since there cannot be infinitely many different occurrences in a proof. \square

7.2.2 Empires and Kingdoms

It is useful to study subgraphs of proof-nets which are proof-nets themselves. This leads one to the study of subnets, in particular, the largest and smallest subnets which contain a formula occurrence F as one of their doors called the empire and kingdom of F respectively. Girard's original proof [Gir87a] of sequentialisation goes through empires and we will recast that argument in the following subsection but obviously in the setting of μMLL^* . Our proofs and constructions on subnets are inspired by [BvdW95].

Definition 7.2.8. Let $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ and $\mathcal{R}' = (\{F'_i^{U'_i}\}_{i \in \lambda'}, \mathfrak{K}', \Theta')$ be proof-structures. \mathcal{R}' is a **substructure** of \mathcal{R} if there exists an injective map,

$$m : \{F'_i \mid i \in \lambda'\} \rightarrow \{(F_i, u) \mid u \in \overline{U_i}, i \in \lambda\}$$

such that the following hold:

- for all $i \in \lambda'$, $\lceil F'_i \rceil = \lceil m(F'_i) \rceil$;
- for all $i \in \lambda'$, $m(F'_i) = (F_j, u) \implies U'_i = u^{-1}U_j$;
- for all $\{C, C^\perp\} \in \mathfrak{K}$, if $C = m(D)$ then $C^\perp = m(D^\perp)$ and $\{D, D^\perp\} \in \mathfrak{K}'$;
- for all $\theta \in \Theta$, if $\text{addr}(m(F'))uu' \in \theta$ where $m(F') = (F_i, u)$ and $uu' \in U_i$ then every $w \in \theta$ is of the form $\text{addr}(m(G'))vv'$ where $m(G') = (F_j, v)$ and $v' \in U_j$ and $\theta' \in \Theta'$ where

$$\theta' = \{\text{addr}(F')u' \mid \text{addr}(m(F'))uu' \in \theta\}.$$

Furthermore, if the substructure \mathcal{R}' is DR-correct, it is called a **subnet**.

Wlog, one can assume that m is the identity map. This amounts to simply changing the address of the doors of \mathcal{R}' . In particular, assign the new address of F'_i to be $\text{addr}(m(F'_i))$. From now on, we will assume that m is an identity for any substructure or subnet unless otherwise mentioned. For the rest of this subsection fix a proof-net $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$.

Definition 7.2.9. Let $\mathcal{R}' = (\Gamma', \mathfrak{K}', \Theta')$ and $\mathcal{R}'' = (\Gamma'', \mathfrak{K}'', \Theta'')$ be two substructures of \mathcal{R} . The **join** of \mathcal{R}' and \mathcal{R}'' , denoted $\mathcal{R}' \cup \mathcal{R}''$, is defined as the substructure $(\Gamma' \cup \Gamma'', \mathfrak{K}' \cup \mathfrak{K}'', \Theta' \cup \Theta'')$. Dually, the **meet** of \mathcal{R}' and \mathcal{R}'' , denoted $\mathcal{R}' \cap \mathcal{R}''$, is defined as the substructure $(\Gamma' \cap \Gamma'', \mathfrak{K}' \cap \mathfrak{K}'', \Theta' \cap \Theta'')$. The meet is said to be **non-empty** if $\Gamma' \cap \Gamma'' \neq \emptyset$.

Definition 7.2.10. Let \mathcal{R}' and \mathcal{R}'' be two substructures of \mathcal{R} . We say that \mathcal{R}' is **included** in \mathcal{R}'' , denoted $\mathcal{R}' \subseteq \mathcal{R}''$, if G occurs in \mathcal{R}' implies G occurs in \mathcal{R}'' .

Note that \subseteq is a partial order on the set of all substructures of a proof-net \mathcal{R} . Note that for any two substructures $\mathcal{R}', \mathcal{R}''$ of \mathcal{R} , we have the following:

1. $\mathcal{R}' \cap \mathcal{R}'' \subseteq \mathcal{R}'$ and $\mathcal{R}' \cap \mathcal{R}'' \subseteq \mathcal{R}''$;
2. $\mathcal{R}' \subseteq \mathcal{R}' \cup \mathcal{R}''$ and $\mathcal{R}'' \subseteq \mathcal{R}' \cup \mathcal{R}''$.

Lemma 7.2.2. Let \mathcal{R}' and \mathcal{R}'' be two subnets of \mathcal{R} such that their meet is non-empty. Then, both their join and meet are subnets.

Proof. Fix a switching sw_j and sw_m of $\mathcal{R}' \cup \mathcal{R}''$ and $\mathcal{R}' \cap \mathcal{R}''$ respectively. Wlog, assume that sw_m is the restriction of sw_j to the par occurrences of $\mathcal{R}' \cap \mathcal{R}''$. Extend sw_j to a switching sw of \mathcal{R} ; moreover, let sw' and sw'' be the restrictions of sw_j to the par occurrences of \mathcal{R}' and \mathcal{R}'' , respectively. Let $G_{sw_j}^{\mathcal{R}' \cup \mathcal{R}''}$, $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$, $G_{sw}^{\mathcal{R}}$, $G_{sw'}^{\mathcal{R}'}$, and $G_{sw''}^{\mathcal{R}''}$ be the respective orthogonal graphs.

We first observe that $G_{sw_j}^{\mathcal{R}' \cup \mathcal{R}''}$ and $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$ are subgraphs of $G_{sw}^{\mathcal{R}}$ which is acyclic since \mathcal{R} is a proof-net. Therefore, $G_{sw_j}^{\mathcal{R}' \cup \mathcal{R}''}$ and $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$ are acyclic.

Now we will prove that $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$ is connected. Let v, v' be two of its vertices. Since $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$ is a subgraph of $G_{sw'}^{\mathcal{R}'}$ and $G_{sw''}^{\mathcal{R}''}$, and $\mathcal{R}', \mathcal{R}''$ are proof-nets there are paths ρ', ρ'' between v, v' in $G_{sw'}^{\mathcal{R}'}$ and $G_{sw''}^{\mathcal{R}''}$ respectively. But ρ', ρ'' are paths in $G_{sw}^{\mathcal{R}}$ which is acyclic. Therefore, $\rho' = \rho''$; and hence, a path in $G_{sw_m}^{\mathcal{R}' \cap \mathcal{R}''}$ as well which concludes our proof.

Now we will prove that $G_{sw_j}^{\mathcal{R}' \cup \mathcal{R}''}$ is connected. Let v, v' be two of its vertices. If v, v' both occur in $G_{sw'}^{\mathcal{R}'}$, then we are done by the connectedness of $G_{sw'}^{\mathcal{R}'}$. Similarly, if both occur in $G_{sw''}^{\mathcal{R}''}$, then we are done by the connectedness of $G_{sw''}^{\mathcal{R}''}$. The only case that we are left with is when v occurs only in $G_{sw'}^{\mathcal{R}'}$ and v' occurs only in $G_{sw''}^{\mathcal{R}''}$. Since $\mathcal{R}' \cap \mathcal{R}''$ is non-empty, $G_{sw'}^{\mathcal{R}'}$ and $G_{sw''}^{\mathcal{R}''}$ share at least one vertex, say s . By the connectedness of $G_{sw'}^{\mathcal{R}'}$, v is connected to s , and by the connectedness of $G_{sw''}^{\mathcal{R}''}$, s is connected to v' . Hence v, v' are connected. \square

Definition 7.2.11. Let G be a formula occurrence occurring in \mathcal{R} . The **empire** of G , denoted $\mathfrak{E}(G)$, is the largest subnet (in the \subseteq ordering) of \mathcal{R} with G as a door. The dual notion of empire is that of a **kingdom**, denoted $\mathfrak{K}(G)$, the smallest subnet of \mathcal{R} with G as a door.

Proposition 7.2.2. For all G occurring in \mathcal{R} , $\mathfrak{E}(G)$ and $\mathfrak{K}(G)$ exist.

Proof. Define

$$\mathfrak{S} = \{\mathcal{S} \mid \mathcal{S} \text{ is a subnet of } \mathcal{R} \text{ and has } G \text{ as a door}\}.$$

The crux of this proof is showing that \mathfrak{S} is non-empty. Assuming $\mathfrak{S} \neq \emptyset$, rest of the proof is straightforward: for any $\mathcal{S}, \mathcal{S}' \in \mathfrak{S}$, $\mathcal{S} \cap \mathcal{S}'$ is non-empty since they share a door. Using Lemma 7.2.2, $(\mathfrak{S}, \subseteq, \cup, \cap)$ is a lattice. So,

$$\mathfrak{E}(G) = \bigcup_{\mathcal{S} \in \mathfrak{S}} \mathcal{S} \quad \text{and} \quad \mathfrak{K}(G) = \bigcap_{\mathcal{S} \in \mathfrak{S}} \mathcal{S}$$

Now we will show that \mathfrak{S} is non-empty. If G is a door of \mathcal{R} , then $\mathcal{R} \in \mathfrak{S}$ and we are done. Otherwise, let $G = (F_i, va)$ where $a \in \{l, r, i\}$. Define the proof-structure

$$\mathcal{R}' := (\{F_j^{U_j}\}_{j \in \lambda \setminus \{i\}} \cup \{F_i^{U'_i}, (F_i, ua)^{U'}\}, \mathfrak{R}, \Theta)$$

where $U'_i = U_i \setminus v^{-1}U_i$ and $U' = u^{-1}U_i$. This is not necessarily a subnet of \mathcal{R} . We will extract a subnet from the orthogonal graph of \mathcal{R}' . Let sw' be a switching of \mathcal{R}' . This can be extended to sw , a switching of \mathcal{R} . We denote the corresponding orthogonal graphs of \mathcal{R}' and \mathcal{R} by $G_{sw'}^{\mathcal{R}'} = ([\Theta], [SW'], E')$ and $G_{sw}^{\mathcal{R}}$ respectively. Let \mathcal{G} be the subgraph induced by the set of $v \in [\Theta] \cup [SW']$ such that $v \cap U' \neq \emptyset$. We observe that since $G_{sw}^{\mathcal{R}}$ is connected and acyclic, then \mathcal{G} is also connected and acyclic. We take this connected component of $G_{sw'}^{\mathcal{R}'}$, and this induces substructure $\mathcal{R}'_{sw'}$. We repeat this process for all possible switchings of \mathcal{R}' . Let the set of all switchings of \mathcal{R}' be denoted by \mathfrak{s} . We claim that $\mathcal{R}'' = \bigcap_{sw' \in \mathfrak{s}} \mathcal{R}'_{sw'}$ is a subnet of \mathcal{R} with G as door.

This is relatively easy to see. Since every $\mathcal{R}'_{sw'}$ has G as a door, \mathcal{R}'' has G as one of its doors. Now, suppose for some switching sw , $G_{sw}^{\mathcal{R}''}$ is not connected. Then, sw can be extended to a switching sw' and sw'' of \mathcal{R}' and \mathcal{R} respectively. Then, the DR-correctness of \mathcal{R} induces a connected acyclic subgraph in the orthogonal graph of \mathcal{R}' as before. This contradicts $G_{sw}^{\mathcal{R}''}$ is disconnected. Arguing exactly similarly for acyclicity, we are done. \square

We define a relation \ll on the occurrences in \mathcal{R} by $F \ll G$ iff $\mathfrak{K}(F) \subseteq \mathfrak{K}(G)$. Let $\mathbb{F}_{\mathcal{R}}$ be the set of all occurrences in \mathcal{R} that do not occur in axiom *i.e.* the set of all (F_i, u) such that $u \notin \bar{U} \setminus U$.

Proposition 7.2.3. $\ll_{\mathcal{R}}$ is a partial order on \mathbb{F} .

Proof. We exhibit the three properties required to be a partial order.

Reflexivity By definition, X occurs in $\mathfrak{K}(X)$.

Anti-symmetry Suppose $X, Y \in \mathbb{F}$ are distinct occurrences such that $X \ll_{\mathcal{R}} Y$ (*i.e.* X occurs in $\mathfrak{D}(Y)$) and $Y \ll_{\mathcal{R}} X$ (*i.e.* Y occurs in $\mathfrak{D}(X)$). Using reflexivity, X, Y occur in both $\mathfrak{D}(X)$ and $\mathfrak{D}(Y)$. Therefore, $\mathfrak{D}(X) \cap \mathfrak{D}(Y)$ exists and is a subnet by Lemma 7.2.2. Note that X and Y are, in fact, the doors of $\mathfrak{D}(X) \cap \mathfrak{D}(Y)$. But $\mathfrak{D}(X)$ is supposed to be the smallest subnet with X as a door. Therefore, $\mathfrak{D}(X) = \mathfrak{D}(X) \cap \mathfrak{D}(Y)$. By a similar argument, $\mathfrak{D}(X) \cap \mathfrak{D}(Y) = \mathfrak{D}(Y)$. So, $\mathfrak{D}(X) = \mathfrak{D}(Y)$. Suppose $X \neq Y$. We will show that this induces a contradiction. Note that X is a door of $\mathfrak{D}(Y)$. Let U_X be its corresponding partial syntax tree in $\mathfrak{D}(Y)$. There are three cases based on the outermost connective of X .

Case 1: $X = A \wp B$. Replacing X^{U_X} by $\{A^{l^{-1}U_X}, B^{r^{-1}U_X}\}$ in $\mathfrak{D}(Y)$ gives us a subnet smaller than $\mathfrak{D}(Y)$ that contains Y as a door. Contradiction!


Case 2: $X = \mu x.F$. Replacing X^{U_X} by $\{F[\mu x.F/x]^{i^{-1}U_X}\}$ in $\mathfrak{D}(Y)$ gives us a subnet smaller than $\mathfrak{D}(Y)$ that contains Y as a door. Contradiction! Similarly for $X = \nu x.F$.

Case 3: $X = A \otimes B$. Y occurs in either $\mathfrak{W}(A)$ or $\mathfrak{W}(B)$. Wlog, assume it is $\mathfrak{W}(A)$. Consider $\mathfrak{D}(Y) \cap \mathfrak{W}(A)$. It is a subnet by Lemma 7.2.2. It contains Y as a door and it is strictly smaller than $\mathfrak{D}(A)$ since it does not contain X . Contradiction!

Transitivity Suppose $X, Y, Z \in \mathbb{F}$ are distinct occurrences such that $X \ll_{\mathcal{R}} Y$ (*i.e.* X occurs in $\mathfrak{D}(Y)$) and $Y \ll_{\mathcal{R}} Z$ (*i.e.* Y occurs in $\mathfrak{D}(Z)$). By reflexivity, Y occurs in $\mathfrak{D}(Y)$. Therefore, $\mathfrak{D}(Y) \cap \mathfrak{D}(Z)$ exists and is a subnet by Lemma 7.2.2. Note that Y is, in fact, a door of $\mathfrak{D}(Y) \cap \mathfrak{D}(Z)$. But $\mathfrak{D}(Y)$ is supposed to be the smallest subnet with Y as a door. Therefore, $\mathfrak{D}(Y) = \mathfrak{D}(Y) \cap \mathfrak{D}(Z)$. Therefore X occurs in $\mathfrak{D}(Y) \cap \mathfrak{D}(Z)$. Hence infer that X occurs in $\mathfrak{D}(Z)$. \square

7.2.3 Sequentialisation

The process of translating a proof-net into a proof is called *sequentialisation*. In this subsection we will show that the correctness criterion is indeed sufficient to ensure sequentialisation. Wlog, we can assume that we have cut-free proof-nets due to the following standard trick. So, in this subsection, we will write nets without their second component (which is always an empty set) *i.e.* write (Γ, Θ) as a shorthand for $(\Gamma, \emptyset, \Theta)$. We present the non-deterministic algorithm SEQUENTIALISE that produces a proof given a proof-net in Figure 7.3.



Lemma 7.2.3 (Correctness of sequentialisation algorithm). *If $G \otimes H$ is a door of a proof-net \mathcal{R} such that it is maximal in the $\ll_{\mathcal{R}}$ ordering then $\mathfrak{W}(G)$ and $\mathfrak{W}(H)$ do not interleave.*

Proof. Suppose $\mathfrak{W}(G) \cap \mathfrak{W}(H) \neq \emptyset$. Then, there exists C, D such that C occurs in \mathcal{R} , D is the immediate suboccurrence of C , D occurs in $\mathfrak{W}(G)$, and C does not occur in $\mathfrak{W}(G)$. (This is symmetric and one can also choose C, D such that C occurs in \mathcal{R} , D is the immediate suboccurrence of C , D occurs in $\mathfrak{W}(H)$, and C does not occur in $\mathfrak{W}(H)$.)

We now claim that $G \otimes H$ occurs in $\mathfrak{D}(C)$. Clearly, D occurs in $\mathfrak{W}(G) \cap \mathfrak{D}(C)$. Hence, by Lemma 7.2.2, $\mathfrak{W}(G) \cap \mathfrak{D}(C)$ and $\mathfrak{W}(G) \cup \mathfrak{D}(C)$ are subnets of \mathcal{R} . Suppose $G \otimes H$ does not occur in $\mathfrak{D}(C)$. Then, $\mathfrak{W}(G) \cup \mathfrak{D}(C)$ is a subnet of \mathcal{R} with G as door. But we assumed C does not occur in $\mathfrak{W}(G)$ whereas it occurs in $\mathfrak{W}(G) \cup \mathfrak{D}(C)$. So, $\mathfrak{W}(G) \cup \mathfrak{D}(C)$ is a strictly larger subnet than $\mathfrak{W}(G)$ that contains G as door. This contradicts the definition of $\mathfrak{W}(G)$.

There are two cases now. Either C is a door of \mathcal{R} in which case we have $G \otimes H \ll_{\mathcal{R}} C$. Otherwise C is a suboccurrence of a door Y of \mathcal{R} . Clearly, $\mathfrak{D}(C) \subset \mathfrak{D}(Y)$. Therefore, we have $G \otimes H \ll_{\mathcal{R}} Y$. This contradicts the fact that $G \otimes H$ is the maximal door in the $\ll_{\mathcal{R}}$ ordering. \square

Note that SEQUENTIALISE is a non-deterministic procedure. So can we say that sequentialisation and desequentialisation are inverses as functions? More concretely, is there a way to sequentialise $\text{dsq}(\pi)$ so as to get exactly π ? Conversely, given a proof-net \mathcal{R} , is $\text{dsq}(\pi) = \mathcal{R}$ where π is some sequentialisation of \mathcal{R} ?

```

function SEQUENTIALISE( $(\Gamma, \Theta)$ )
  if  $\Gamma = \{\varphi_\alpha^{\{\varepsilon\}}, \varphi_\beta^{\perp\{\varepsilon\}}\}$  and  $\Theta = \{\alpha, \beta\}$  then
     $\frac{}{\text{return } \vdash \varphi, \varphi^\perp}$  (id)
  else
    Choose  $F_i^{U_i} \in \Gamma$  such that  $U_i \neq \{\varepsilon\}$ .
    if  $F_i = G \wp H$  then
       $\frac{\text{SEQUENTIALISE}(\mathfrak{W}(G))}{\text{return } \vdash [\Delta], [G \wp H]}$  ( $\wp$ )
    else if  $F_i = G \otimes H$  and  $\exists j \neq i. F_i \ll F_j$  then
       $\frac{\text{SEQUENTIALISE}(\mathfrak{W}(G)) \quad \text{SEQUENTIALISE}(\mathfrak{W}(H))}{\text{return } \vdash [\Delta], [G \otimes H]}$  ( $\otimes$ )
    else if  $F_i = \eta x. G$  then  $\triangleright \eta \in \{\mu, \nu\}$ 
       $\frac{\text{SEQUENTIALISE}(\mathfrak{W}(G[\eta x. G/x]))}{\text{return } \vdash [\Gamma], [\eta x. G]}$  ( $\eta$ )
    end if
  end if
end function

```

Figure 7.3: The function SEQUENTIALISE

From Lemma 7.2.3, we have that

$$(\Gamma_G \setminus \{G\}) \cup (\Gamma \setminus \{H\}) \cup \{G \otimes H\} = \Gamma$$

where Γ_G, Γ_H , and Γ are the set of doors of $\mathfrak{W}(G)$, $\mathfrak{W}(H)$ and \mathcal{R} . Similarly for other connectives. Therefore, the sequentialisation procedure traverses through every node of every partial syntax. Therefore, if there is a proof π such that $\text{SEQUENTIALISE}(\mathcal{R}) = \pi$ then $\text{dsq}(\pi) = \mathcal{R}$. Therefore, we have the following.

Theorem 7.2.1. *A μMLL^* proof-structure is DR-correct iff it is a μMLL^* proof-net.*

The question now is if $\text{dsq}(\pi) = \text{dsq}(\pi')$ or equivalently, if there exists a proof-net \mathcal{R} such that $\text{SEQUENTIALISE}(\mathcal{R}) = \pi$ and $\text{SEQUENTIALISE}(\mathcal{R}) = \pi'$, can we say anything more about these proofs? Are these proofs computationally equivalent? We answer these questions in the next subsection.

7.2.4 Canonicity

Let \sim be a relation over all μMLL^* proofs that equates proofs with one permutation (Figure 7.4 and Figure 7.5). Let \sim^* be the reflexive transitive closure of \sim . Let \equiv_{PN} be the equivalence relation over μMLL^* proofs defined as $\pi \equiv_{\text{PN}} \pi'$ if $\text{dsq}(\pi) = \text{dsq}(\pi')$.

Theorem 7.2.2. *For all μMLL^* proofs $\pi, \pi', \pi \sim^* \pi'$ iff $\pi \equiv_{\text{PN}} \pi'$.*

Proof sketch. Let π and π' be two μMLL^* proofs such that $\pi \sim^* \pi'$. Clearly, nothing can be permuted below the root sequent. Therefore, $\text{dsq}(\pi)$ and $\text{dsq}(\pi')$ have the same set of doors. By inspecting each condition of \sim , one can infer that no inference rule (on a particular occurrence) is deleted. Moreover, it is also not duplicated². Therefore the partial syntax trees of $\text{dsq}(\pi)$ and $\text{dsq}(\pi')$ are also the same. For the same reason, the set of cuts of $\text{dsq}(\pi)$ and $\text{dsq}(\pi')$ are also the same. Imagine there is an axiom $\{A, B\}$ in π' that does not occur in π . Now, since same set of partial syntax trees appears on $\text{dsq}(\pi)$ and $\text{dsq}(\pi')$, A and B occur in some other axioms. Therefore, there is some tensor rule that took A and B to different subtrees. But this tensor cannot occur in π . But we proved that π and π' have the same set of inference rules. Contradiction! Therefore, $\text{dsq}(\pi)$ and $\text{dsq}(\pi')$ have the same set of axioms. Therefore, $\text{dsq}(\pi) = \text{dsq}(\pi')$.

Let π and π' be two distinct sequentialisation of a proof-net \mathcal{R} . Let π, π' differ at the inference rule r i.e. either the inference $\frac{\vdash A}{\vdash B}(r)$ occurs at that position for π but not for π' , or, r is a tensor rule on some $A \otimes B$ that occurs at that position for both π and π' but with different splittings of the

²This is not true in MALL which goes on to indicate why devising MALL proof-nets is a cumbersome business.

$$\begin{aligned}
& (\Gamma \cup \{\psi_\alpha^U, \varphi_\beta^{U'}, \varphi_{\beta^\perp}^{\{\varepsilon\}}\}, \mathfrak{R} \cup \{\{\varphi_\beta, \varphi_{\beta^\perp}\}\}, \Theta \cup \{\{\alpha u, \beta^\perp\}\}, \text{loc}) \rightarrow_{\mu\text{MLL}^*} \\
& (\Gamma \cup \{\psi_\alpha^{(U \setminus u(u^{-1}U)) \cup U'}\}, \mathfrak{R}, \Theta, \text{loc}[\varphi_{\alpha u} \mapsto \varphi_\beta]) \\
& (\Gamma, \mathfrak{R}, \Theta \cup \{\{\alpha, \beta\}\}, \text{loc}[\varphi_{\alpha'} \mapsto \varphi_\beta]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta \cup \{\{\alpha, \alpha'\}\}, \text{loc}) \\
& (\Gamma \cup \{\varphi_\alpha^U, \psi_\alpha^\perp \otimes \varphi_{\alpha^\perp}^{U'}\}, \mathfrak{R} \cup \{\{\varphi_\alpha, \psi_\alpha^\perp \otimes \varphi_{\alpha^\perp}\}\}, \Theta, \text{loc}) \rightarrow_{\mu\text{MLL}^*} \\
& (\Gamma \cup \{\varphi_{\alpha l}^{l^{-1}U}, \psi_{\alpha r}^{r^{-1}U}, \varphi_{\alpha^\perp l}^{l^{-1}U'}, \psi_{\alpha^\perp l}^{l^{-1}U'}\}, \mathfrak{R} \cup \{\{\varphi_{\alpha l}, \varphi_{\alpha^\perp l}\}, \{\psi_{\alpha r}, \psi_{\alpha^\perp l}\}\}, \Theta, \text{loc}) \\
& (\Gamma \cup \{\mu x. \varphi_\alpha^U, \nu x. \varphi_{\alpha^\perp}^{U'}\}, \mathfrak{R} \cup \{\{\mu x. \varphi_\alpha, \nu x. \varphi_{\alpha^\perp}\}\}, \Theta, \text{loc}) \rightarrow_{\mu\text{MLL}^*} \\
& (\Gamma \cup \{\varphi[\mu x. \varphi/x]_{\alpha i}^{i^{-1}U}, \varphi^\perp[\nu x. \varphi^\perp/x]_{\alpha^\perp i}^{i^{-1}U'}\}, \mathfrak{R} \cup \{\{\varphi[\mu x. \varphi/x]_{\alpha i}, \varphi^\perp[\nu x. \varphi^\perp/x]_{\alpha^\perp i}\}\}, \Theta, \text{loc})
\end{aligned}$$

Figure 7.6: Key cases of cut-elimination for μMLL^* proof-nets


$$\begin{aligned}
& (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi_\alpha \mapsto \varphi_\beta, \varphi_\beta \mapsto \varphi_\gamma]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi_\alpha \mapsto \varphi_\gamma]) \\
& (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi \otimes \psi_\alpha \mapsto \varphi \otimes \psi_\beta]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi_{\alpha l} \mapsto \varphi_{\beta l}, \psi_{\alpha r} \mapsto \psi_{\beta r}]) \quad [\forall \theta \in \Theta. \beta \notin \theta] \\
& (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi \wp \psi_\alpha \mapsto \varphi \wp \psi_\beta]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi_{\alpha l} \mapsto \varphi_{\beta l}, \psi_{\alpha r} \mapsto \psi_{\beta r}]) \quad [\forall \theta \in \Theta. \beta \notin \theta] \\
& (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\mu x. \varphi_\alpha \mapsto \mu x. \varphi_\beta]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi[\mu x. \varphi/x]_{\alpha i} \mapsto \varphi[\mu x. \varphi/x]_{\beta i}]) \quad [\forall \theta \in \Theta. \beta \notin \theta] \\
& (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\nu x. \varphi_\alpha \mapsto \nu x. \varphi_\beta]) \rightarrow_{\mu\text{MLL}^*} (\Gamma, \mathfrak{R}, \Theta, \text{loc}[\varphi[\nu x. \varphi/x]_{\alpha i} \mapsto \varphi[\nu x. \varphi/x]_{\beta i}]) \quad [\forall \theta \in \Theta. \beta \notin \theta]
\end{aligned}$$

Figure 7.7: Loc commutation cases of cut-elimination for μMLL^* proof-nets


Example 7.2.3. Consider the proof-net in example 7.2.2. We have,

$$\begin{aligned}
& (\{\nu x. x \wp \mu x. x_\alpha^{\{l, r\}}, \mu y. y_\beta^{\{i\}}, \nu y. y_{\beta^\perp}^{\{i\}}\}, \{\{\mu y. y_\beta, \nu y. y_{\beta^\perp}\}\}, \{\{\alpha l, \beta i\}, \{\alpha r, \beta^\perp i\}\}) \rightarrow_{\mu\text{MLL}^*} \\
& (\{\nu x. x \wp \mu x. x_\alpha^{\{l, r\}}, \mu y. y_{\beta i}^{\{\varepsilon\}}, \nu y. y_{\beta^\perp i}^{\{\varepsilon\}}\}, \{\{\mu y. y_{\beta i}, \nu y. y_{\beta^\perp i}\}\}, \{\{\alpha l, \beta i\}, \{\alpha r, \beta^\perp i\}\}) \rightarrow_{\mu\text{MLL}^*} \\
& (\{\nu x. x \wp \mu x. x_\alpha^{\{l, r\}}\}, \emptyset, \{\{\alpha l, \alpha r\}\})
\end{aligned}$$

Remark 7.2.1. The key rule involving locs graphically it looks as follows.



The other key rule which is new is obviously the one involving the fixed point operators. Graphically it looks like:



Definition 7.2.13. Let \mathcal{R} be μMLL^* proof-net (possibly with locs). Let $[\mathcal{R}] = \{\mathcal{R}' \mid \forall \mathcal{R}'' . \mathcal{R}'' \not\rightarrow_{\text{id/cut}} \pi' \rightarrow_{\text{id/cut}}^* \mathcal{R}\}$. \mathcal{R} is said to be DR-correct if every $\mathcal{R}' \in [\mathcal{R}]$ is DR-correct.

Lemma 7.2.4. Let $\mathcal{R} \rightarrow_{\mu\text{MLL}^*} \mathcal{R}'$ such that \mathcal{R} is a $\mu\text{MLL}_{\text{loc}}^*$ proof-net. Then, \mathcal{R}' is a $\mu\text{MLL}_{\text{loc}}^*$ proof-net.

Proof sketch. We do a case analysis on the rule applied. If it is a key rule involving locs and a loc commutation rules, then $[\mathcal{R}] = [\mathcal{R}']$ and we are done. Otherwise, suppose it is a key-rule involving a tensor and par, then note that the loc components of \mathcal{R} and \mathcal{R}' are the same. Since cut-reduction preserves DR-correctness in MLL proof-nets, we are done. Finally, the case when it is a key-rule involving a μ and ν is trivial, since the geometry of nets and the loc-components are unchanged by the reduction. Hence we are done. \square

Note that normal forms of $\mu\text{MLL}_{\text{loc}}^*$ proof-nets under the $\rightarrow_{\mu\text{MLL}^*}$ reduction are exactly cut-free μMLL^* proof-nets with an empty loc component.

Theorem 7.2.3. $\rightarrow_{\mu\text{MLL}^*}$ is confluent and terminating.

Proof. We define a termination measure. Let \mathfrak{K} be the set of cuts for some net \mathcal{R} . Define

$$\text{rk}(\mathfrak{K}) = \begin{cases} 0 & \text{if } \mathfrak{K} = \emptyset; \\ \frac{1}{\max_{\theta \in \mathfrak{K}} \{|\text{addr}(C)| \mid C \in \theta\}} & \text{otherwise.} \end{cases}$$

Then, $d(\mathcal{R}) = (|\mathfrak{K}|, \text{rk}(\mathfrak{K}))$ induces a lexicographic order on proof-nets. Also note, it strictly decreases for every key step of $\rightarrow_{\mu\text{MLL}^*}$ reduction. Similarly, the size of the addresses of relocated formulas increases with every loc commutation rule. But they cannot increase indefinitely since they are bounded by the size of the corresponding addresses of axioms. So, there cannot be infinitely many consecutive loc commutation rules in a reduction sequence. Therefore, $\rightarrow_{\mu\text{MLL}^*}$ terminates.

The only critical pairs are when there is an axiom $\{w, w'\}$ such that there exists two cuts $\{C, C^\perp\}$ and $\{D, D^\perp\}$ such that $\lceil C \rceil = \lceil D \rceil$, $\text{addr}(C) = w$, and $\text{addr}(D^\perp) = w'$. In these cases the reducts are isomorphic up to address renaming. Therefore, $\rightarrow_{\mu\text{MLL}^*}$ is confluent. \square

Remark 7.2.2. In our proof of confluence above, there is a certain notion of equality of proof-nets that is implicit. We say that two proof-nets $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta)$ and $(\{G_i^{U'_i}\}_{i \in \lambda}, \mathfrak{K}', \Theta')$ are equal if they are equal up to address renaming. For example, for all $i \in \lambda$, $\lceil F_i \rceil = \lceil G_i \rceil$ and $U_i = U'_i$, $\{F_i, F_j\} \in \mathfrak{K}$ iff $\{G_i, G_j\} \in \mathfrak{K}'$, and so on. However, we obtain confluence up to the usual equality of proof-nets when we reach the normal form.

7.3 $\mu\text{MLL}^{\text{ind}}$ and μMLL° proof-nets


In this section, we will discuss the parallel syntax of $\mu\text{MLL}^{\text{ind}}$ and μMLL° . As sequential proofs both wellfounded and circular derivations are finitely representable; consequently, their corresponding proof-nets are also finite objects. This section is independent of the developments of infinets, which are non-wellfounded objects that will be introduced in the next chapter. The algebraic presentation is helpful in the infinitary setting but since the development in this section is orthogonal to the proof-nets introduced in the next chapter on non-wellfounded proof-nets, we will use the graphical presentation of proof-nets in this section.

7.3.1 The statics of $\mu\text{MLL}^{\text{ind}}$ nets

The μ nodes are as before in Section 7.2. We need to devise the shape of the ν nodes. Recall the (ν) rule in $\mu\text{MALL}^{\text{ind}}$ sequent calculus:

$$\frac{\vdash \Gamma, S \vdash S^{\perp}, F[S/x]}{\vdash \Gamma, \mu x.F} (\nu)$$


As a start we provide the following definition of a ν node: if \mathcal{R} is a $\mu\text{MLL}^{\text{ind}}$ proof-net then so is the following.




There are a few issues with this formulation. Namely,

- The three outgoing edges are treated equivalently. However, what we have is rather $(S \otimes (S^{\perp} \wp F[S/x]))$. There is an easy fix. One refines the DR-correctness condition to account for the switching of $S^{\perp} \wp F[S/x]$.
- It is important that the tensor in $(S \otimes (S^{\perp} \wp F[S/x]))$ is splitting and split into two nets with Γ, S and $S^{\perp}, F[S/x]$ as doors respectively. We hardcode this information in the definition. This corresponds to a case considered in [Gir99, Cur06], that of extremal proof-structures, where \otimes nodes have exactly two splittings: everything either distributes to the left or to right.

We define ν nodes as follows: if $\mathcal{R}_1, \mathcal{R}_2$ are $\mu\text{MLL}^{\text{ind}}$ proof-nets then so is the following.



We justify our choice further using the second-order encoding of $\mu\text{MLL}^{\text{ind}}$. Recall that $\nu x.F$ is encoded as $\exists S.S \otimes (S^{\perp} \wp [F][S/x])$ where $[F]$ is the inductive encoding of $\mu\text{MLL}^{\text{ind}}$ formulas. The corresponding LL^2 proof-net is as follows.



Note that after three steps of sequentialisation we remove the box and after a fourth step we have exactly the same net as above up to the ν node (and $?\Delta$ doors). In fact, the $?\Delta$ are natural and occur in ν rules for $\mu\text{LL}^{\text{ind}}$:

$$\frac{\vdash \Gamma, S \quad \vdash S^\perp, F[S/x], ?\Delta}{\vdash \Gamma, \nu x.F, ?\Delta} (\nu)$$

To sum up, $\mu\text{MLL}^{\text{ind}}$ proof-nets are MLL with μ and ν as described above. DR-correctness is the same with one caveat that not only all \wp nodes are switched but also every second (or third) premisses of a ν node is switched. By correctness, the tensor hidden in the ν node is always splitting, so one can sequentialise any bottommost μ or ν node. This also corresponds to the fact that in $\mu\text{MALL}^{\text{ind}}$ proofs both μ and ν rules can be permuted down. Consequently, $\mu\text{MLL}^{\text{ind}}$ nets are canonical *i.e.* they exactly characterise the equivalence up to permutation of inference rules.

We will briefly discuss the dynamics of $\mu\text{MALL}^{\text{ind}}$ nets without delving into much detail. The reduction rule for the fixed point case is basically the desequentialisation of the corresponding key case in the sequent calculus. One can again justify this rule from the point of view of LL^2 proof-nets. In order to prove the cut-elimination, one needs to devise a measure on the set of $\mu\text{MLL}^{\text{ind}}$ proof-nets. A straightforward adaptation of the measure defined in the proof of Theorem 7.2.3 works.

7.3.2 μMLL° nets

The question of when are two proofs the same goes beyond the permutation of inference rules in the case of circular proofs. A non-wellfounded proof with finitely many subproofs admits several circular representations (*cf.* Figure 7.8). From an algorithmic perspective, there is a motivation to obtain the minimal representation of proof object. For instance, a regular tree can be seen as a deterministic automaton and one might choose the minimal deterministic automaton as its representation). However, such representations obfuscate proof-theoretic information.

In the following, we assume that two circular proofs that have the same unfoldings can be different. μMLL° proofs are finite objects. In fact, they are wellfounded trees with some extra information. We start off by treating backedges as generalised axioms. Consider the proof in Figure 7.8b and its desequentialisation Figure 7.9a. Note that to recover the backedges in the sequentialisation, one needs to ensure that in any sequentialisation there is a sequent of the form $\vdash F \wp F, G$ that can be the target of the backedge originating from the source *i.e.* the generalised axiom $\vdash F \wp F, G$. This is not necessarily true. The net can be sequentialised as follows:

$$\frac{\nabla}{\vdash F, G \otimes G} (\nu)$$

In this case, we never get the sequent $\vdash F \wp F, G$. The standard technique to restrict the parallelism of proof-nets is the introduction of boxes. Consider the proof-net in the Figure 7.9b. The box ensures that the μ node is sequentialised before the ν node.




Figure 7.8: Different representations of the same circular proof of $\vdash F, G$ where $[F] = [G^\perp] = \mu x. x \wp x_\alpha$.





Figure 7.9: Naive and faithful desequentialisations of the proof in Figure 7.8b.

Desequentialisation is a two-step process. Suppose we are given a μMLL° proof π . Construing sources of backedges as generalised axioms, π is a μMLL^* proof possibly with generalised axioms. We desequentialise it as in Definition 7.2.3 and call the derived net \mathcal{R} . Let S be the set of all sequents in π that are the target of a backedge in π . We will have a box \mathcal{B}_Γ for each $\Gamma \in S$. Note that Γ corresponds to the set of wires in \mathcal{R} . These are exactly the incoming wires of \mathcal{B}_Γ . The box is smallest substructure that contains every \bowtie node that is the source of a backedge that targets Γ . We now label these generalised axioms by the \mathcal{B}_Γ . Having as many boxes as target of backedges is enough sequentialisation information to ensure the faithfulness of backedges and labelling helps remove ambiguity during reconstructing backedges. Correctness therefore ensures if there are enough boxes along with DR-correctness.


Note that our use of boxes is similar to the ones in proof-nets with cycles considered in [Mon03]. In that work, a general fixed-point Y -combinator is added to polarised linear logic. In the sequent calculus, the combinator has the following rule that essentially types Y as $(P \rightarrow P) \rightarrow P$ where P is a positive formula and \rightarrow is classical implication.

$$\frac{\vdash \Gamma, P^\perp, P}{\vdash \Gamma, P} (Y)$$

The corresponding node in proof-nets introduces a so-called Y -box:



We end this section by discussing two issues of the μMLL° proof-nets that we introduced. Firstly, what we achieve is not entirely canonical. Not all circular proofs that are equivalent up to permutations of inference rules can be equated by a single proof-net. For example, the following proof-net has exactly two sequentialisations *viz.* one where every odd-numbered rule is a μ and every even-numbered rule is a ν ; another where every odd numbered rule is a ν and every even numbered rule is a μ .



Clearly, there are several other permutations that are regular (and the set of permutations which is regular is strictly smaller than the set of all permutations since for example, the proof where every prime step is a μ and every composite step is a ν is non-regular).

The second issue is that circular proofs are not closed under cut-elimination. So, one cannot explore the dynamics of μMLL° proof-nets. The solution to both these problems lies in devising more general structures *viz.* proof-nets for non-wellfounded proofs. We shall explore this in the next chapter.




Figure 7.10: Naive desequentialisations of some μMLL^∞ proofs. Back-edges are depicted using pointers (\star). Red and blue curves indicate *trips*.

7.4 Towards infinets

In this section, we discuss the potential pitfalls of extending the notion of proof-nets to the non-wellfounded setting. Let π_Δ be a proof of $\vdash \Delta, F$ where $F = \nu x.x \otimes (a \wp a^\perp)_\alpha$ for an arbitrary address α .

$$\begin{array}{c}
 \frac{}{\vdash a_{\alpha ir l}, a_{\alpha ir r}^\perp} \text{(id)} \\
 \frac{}{\vdash a_{\alpha ir l}, a_{\alpha ir r}^\perp} \text{(\wp)} \\
 \frac{\vdash \Delta, \psi_{\alpha il} \quad \vdash a \wp a_{\alpha ir}^\perp}{\vdash \Delta, \psi \otimes (a \wp a^\perp)_{\alpha i}} \text{(\otimes)} \\
 \frac{\vdash \Delta, \psi \otimes (a \wp a^\perp)_{\alpha i}}{\vdash \Delta, \psi_\alpha} \text{(\nu)}
 \end{array}$$


We first consider π_\emptyset . Now, if we naively translate it into a proof structure using the same recipe as Definition 7.2.3, we have $\text{dsq}(\pi_\emptyset) = (\Gamma, \emptyset, \Theta)$ where

$$\Gamma = \left\{ \psi_\alpha^{\{i(li)^* r(l+r) + (il)^\omega\}} \right\}; \quad \Theta = \{ \alpha i(li)^n r l, i(li)^n r r \}_{n \geq 0}.$$

Observe that $(il)^\omega$ is not in any partition. In fact, it represents a thread in an infinite branch and must be accounted for. Hence the partition should be equipped to account for the threads invariant by an infinite branch in a proof (in particular, in the example above there should be a singleton partition $\{(il)^\omega\}$). This is also the reason we will not use the graphical presentation for non-wellfounded proof-nets since we would potentially need to join two infinite paths by a node of which we are not aware of any rigorous graph-theoretic treatment. However, we will sometimes draw the “graph” of non-wellfounded proof-nets for ease of presentation by using ellipsis points to brush the technical difficulty under the carpet.

In Section 7.1, we formulated a proof-net as a set of formula occurrences, together with an order (the subformula ordering, representing which formula was deduced from which other formulas) and distinguished sets of formulas, representing the conclusions, axioms, and cuts. In particular, an axiom is just a set of two dual formulas. As discussed above, this set of data is not sufficient for de-sequentialising non-wellfounded proofs and as such, we need to introduce *infinite axioms*. Infinite axioms can be thought of as additional limit points at infinity to which the rays converge. We denote such infinite axioms by $\overline{\text{ax}}_\infty$ in our graphical presentation.

Now consider $\Delta = \{b_\beta\}$ where b is an atom and β is an address disjoint with α . The infinite branch of π_Δ has b_β occurring infinitely often. This information is germane to its desequentialisation *i.e.* a faithful translation of π_Δ would therefore look something like:



This can be justified in two ways. Firstly, without b_β being connected to the infinite axiom, the net could not be DR-correct. Secondly, without the information that b_β remains in the infinite branch forever, one can spuriously sequentialise to proofs where b_β goes to the right premiss of one of the tensor rules. Such proofs will not be permutatively equivalent with π_Δ and therefore, such a notion of proof-nets would quotient more than the commutation of inference rules.

7.4.1 The different types of infinite axioms

Therefore, at least as a first approximation, infinite axioms are the invariant of the infinite branches of pre-proofs, which we can picture graphically as a cell “above” a ray in a non-wellfounded proof-structure. This would correspond to a thread in the corresponding non-wellfounded proof.

Consider the proof π' of $\vdash F, \nu x. x_\alpha$ in Figure 7.10a where F is an arbitrary formula and α is an arbitrary address. Note that we do not explicitly mention the addresses in the proof (but they can be easily reconstructed by the reader). Imagine one desequentialised as usual and also takes into account the thread $\tau = \{\nu x. x_{\alpha i^n}\}_{n=0}^\infty$. Then, we have the proof structure in Figure 7.10b. It has two connected components, one with an infinite axiom “above” the ray of ν nodes (corresponding to τ) and second, an undirected ray, say ρ , of alternating axioms and cuts. This breaks DR-correctness.

Observe that every F introduced by a cut resides with τ in the infinite branch of π' . This information is lost in translation. Since we envisage infinite axioms as capturing the invariant of an infinite branch, ρ should be included in the infinite axiom in a correct desequentialisation. Paths like ρ alternating through axioms and cuts are called *visitable paths*. The dyadic notion in proofs corresponding to visitable paths are *trips* which we indicate by the blue curve in Figure 7.10a.

Cuts and tensors are geometrically quite similar; so, it is not very surprising that visitable paths can be formed using tensor nodes as well. In particular, a similar situation as above can be reproduced using tensors unfolded by some fixed point formula. Consider the proof π'' of $\vdash H, a_\beta$ in Figure 7.10c where $H = \nu x. x \wp (a \otimes a^\perp)_\alpha$, and α, β are arbitrary disjoint addresses. As before, we do not explicitly indicate the addresses of every suboccurrence occurring in the proof. Consider the proof structure in Figure 7.10d, the naive desequentialisation of π'' . The failure of DR-correctness is less obvious in this instance. Consider the switching where every \wp switches to the left (*i.e.* the right outgoing edge of every par is severed). Therefore, here as well, the visitable path of alternating axioms and tensors should be ported into the infinite axiom above the ray of alternating ν nodes and \wp nodes.




Figure 7.11: Why do we need virtual axioms?

It is not clear a priori that visitable paths are invariants of infinite branches which are supported by a thread. Consider the pre-proof π in Figure 7.11a. It proves a sequent $\vdash A, A^\perp$ by never operating on these formulas but delaying infinitely this treatment using cuts. From this perspective, it could be desequentialised naively as in Figure 7.11b. The pre-proof has two trips (following each formula as they are cut and introduced by axioms), just as the proof-structure has two visitable paths.

Nonetheless, we can argue as before that an infinite axiom should be atop the two visitable paths, representing that the two formulas A and A^\perp are infinitely pushed away together: viewed in this way, the pre-proof π represents an infinitely cut-expanded axiom.

However, we have no infinite branch in the proof-structure of Figure 7.11b to support an infinite axiom. We need to introduce a new kind of infinite axiom as in Figure 7.11c which is not “above” an infinite ray — we call it a *virtual axiom*. We will thus distinguish between infinite axioms that are supported by a straight thread (which we will call real axioms) and infinite axioms supported by visitable paths (virtual axioms). Just as real infinite axioms, virtual axioms can also contain formula occurrences with finite addresses (indeed, consider π with an arbitrary formula added in the conclusion sequent and pushed through all the cuts). In both cases, an infinite axiom is the invariant (under permutation of inference rules) of an infinite branch (of a pre-proof).

7.4.2 Higher-order trips

A final difficulty arises in the process of inventing infinitary proof structures. Consider the pre-proof in Figure 7.12a: it consists of an infinite sequence of the unfolding of a fixpoint such that, between two unfoldings, a cut introduces the infinite pre-proof π of conclusion $\vdash A, A^\perp$ studied in the last paragraph: as said there, it can be interpreted as an infinitely expanded axiom.

Let us imagine what the procedure to desequentialise the pre-proof in Figure 7.12a would look like, in particular, to compute the visitable paths and the infinite axioms. Typically one can imagine starting by tracing the sequents in such a way that they mimic the dynamics of a trip thereby recognising the infinitely many visitable paths from the infinitely many trips (each occurrence of π generating two maximal trips). The proof-net, at this point, is disconnected. In particular, the real axiom with the ν -thread is disconnected from all the cut occurrences, a geometry similar to the net in Figure 7.10b.

Recall π is akin to an axiom expansion, hence one can imagine bouncing on it going up via A and down via A^\perp as if it were a generalised axiom. Therefore, a second parse through the proof in Figure 7.12a reveals a *higher-order trip* (indicated in blue) which bounces through each copy of π by going up the blue trip and down the red trip of Figure 7.11a, and keeps going up. This trip corresponds to a *higher-order visitable path*. This path is grafted with the real axiom above the infinite ν -ray. We distinguish this trip/visitable path from the ones discussed above by using a stratification. Every trip/visitable path of a particular level ℓ , where $\ell \in \text{Ord}$.

We bookend this discussion by summarizing the terms introduced: they go by dyadic pairs, one in a pre-proof, and its corresponding notion in non-wellfounded proof-structures.




Figure 7.12: Exhibiting higher-order trips and visitable paths. Here π is the pre-proof in Figure 7.11a

Pre-proofs	Non-wellfounded Proof-structures
Axioms	Finite axiom links
Branches supported by threads	Real axioms
Branches not supported by threads	Virtual axioms
Trips	Visitable paths

Chapter 8

The genesis of infinets

(The kingdom strikes back)

In this chapter, we describe the first truly infinite class of μMLL proof-nets. We consider a fragment of μMLL^∞ *viz.* one which does not have *trips*. In Section 8.1, we formalise this fragment of μMLL^∞ . In Section 8.2, we define the appropriate notion of proof-nets for this fragment generalising the proof-nets in the previous chapter. The non-wellfoundedness provides several challenges, one of which is a more involved correctness condition. This correctness condition is introduced and shown to be complete with respect to sequentialisation in Section 8.3. In Section 8.4, we show that the objects we define are indeed canonical. Finally, in Section 8.5 and Section 8.6, we restrict and generalise this class of non-wellfounded proof-nets respectively. Section 8.5 considers a finitely presented fragment and proves some decidability results and connections with circular proofs. In Section 8.6, we introduce general non-wellfounded proof-nets that potentially contain objects corresponding to trips.

8.1 Trips and simple proofs

Girard's original correctness criterion for proof-nets was the long-trip criterion [Gir87a]. He envisioned each link of a proof-net as a router, having as ports the formulae that are premisses or conclusions of the link. Each link is associated with a set of routing rules that tell us from which port we come out when we enter from a given port. Axiom and cut have a fixed behaviour, while tensor and par have two possible behaviours, determined by a local switch. Starting from any node one travels along a path that visits a formula at-most once in each direction. These paths are called *trips*.

Trips can in fact be seen as some kind of operator acting on the proof. We define similar paths on proofs. Suspecting a connection with Girard's trips, we pre-emptively call our paths trips as well.

Definition 8.1.1. *Given a pre-proof π , a **pre-trip** starting from F is a sequence $\tau = \{(s_i, F_i, d_i)\}_{i \in \lambda}$, $\lambda \in \omega + 1$, where s_i is a sequent in π , $F_i \in s_i$ and $d_i \in \{\uparrow, \downarrow\}$ such that $F_0 = F$, $d_0 = \uparrow$ and for every $i \in \lambda$ exactly one of the following holds:*


- $d_i = d_{i+1} = \uparrow$, s_{i+1} is a premise of s_i and $F_{i+1} \sqsubseteq F_i$.
- $d_i = d_{i+1} = \downarrow$, s_i is a premise of s_{i+1} and $F_i \sqsubseteq F_{i+1}$.
- $d_i = \uparrow, d_{i+1} = \downarrow$, and $s_i = s_{i+1}$
- $d_i = \downarrow, d_{i+1} = \uparrow$, s_i and s_{i+1} are the premisses of a (cut) rule on F_i and F_{i+1} .
- $d_i = \downarrow, d_{i+1} = \uparrow$, s_i and s_{i+1} are the premisses of a \otimes rule with auxilliary occurrences F_i and F_{i+1} .

Furthermore τ satisfies that for every $i, j \in \lambda$, there does not exist a sequent, s , in π such that $F_i \wp F_j \in s$.

Definition 8.1.2. *Let τ be an infinite pre-trip of a pre-proof π . Let γ be an infinite branch of π . τ is said to be **associated** with γ if there exists an infinite subsequence $\{(s_i, F_i, d_i)\}_{i \in \omega}$ of τ such that $\{s_i\}_{i \in \omega}$ is a subsequence of γ .*

Proposition 8.1.1. *Let τ be an infinite pre-trip of a pre-proof π and γ_1, γ_2 be two infinite branches of π such that τ is associated with both of them. Then γ_1 and γ_2 coincide on infinitely many sequents. In other words, τ is associated with a unique branch modulo some finite prefix.*

Proof. Suppose not. Then it contains infinitely many sequents from two diverging infinite branches, γ_1 and γ_2 , of π . Let γ be the finite common prefix of γ_1 and γ_2 . By construction, there is a tensor or a cut rule in γ such that it has premises s_l and s_r respectively and it introduces F_l and F_r respectively such that s_l and s_r occur infinitely often in the pre-trip. By construction, during a downward travel via s_l (respectively s_r), in order to change directions, the trip must be through F_l (respectively F_r) rather than any other occurrence of s_l (respectively s_r). Similarly, during the immediately succeeding upward travel, the trip must be through (s_r, F_r, \uparrow) . So, there is a finite pre-trip starting from (s_r, F_r, \uparrow) and (s_r, F_r, \downarrow) . Schematically this finite trip looks like the following where $s_l^n = s_r$ and $F_l^n = F_r$.



We will show by induction on n that such a finite pre-trip cannot exist.

Base case. We have $F_l^1 = F_r$. Then, G_1 and G'_1 are suboccurrences of F_r . Let G be the suboccurrence which is the greatest common prefix of G_1 and G'_1 in the FL-graph of F_r . The outermost operator of G must be either a par or a tensor. In case it is a par, the finite pre-trip goes through two premises of G which is not allowed. If it is a tensor, then G_1 and G'_1 go to different sequents which is not possible.

Induction case. Then, G_1 and G'_n are suboccurrences of F_r . Let G be the suboccurrence which is the greatest common prefix of G_1 and G'_n in the FL-graph of F_r . The outermost operator of G' must be a tensor (parr is ruled out in the same way as in the base case). Let $G = H \otimes H'$. Since the finite pre-trip goes through G , $\{H, H'\} = \{F_l^m, F_r^m\}$ for some $m < n$. But then we can follow F_r^m down to F_r (since G is a suboccurrence of F_r). Hence we must have a shorter finite pre-trip of the above form. By induction hypothesis that does not exist.

Hence an pre-trip can be associated with at most one infinite branch. Further observe by the above argument every sequent is repeated at most finitely often in a pre-trip, so a trip visits higher and higher (or deeper depending on one's perspective) sequents. Hence there is exactly one infinite branch associated with every pre-trip. \square

Definition 8.1.3. *Let $\tau = \{(s_i, F_i, d_i)\}_{i \in \lambda}$, $\lambda \in \omega + 1$ be a pre-trip and ζ be some ordinal. A **trip** is a pre-trip τ such that for some ordinal ζ , τ is a trip of level ζ . A **trip of level ζ** is defined by transfinite induction on ζ as follows.*

- $\zeta = 0$ and for all $i \in \lambda$ such that $d_i = \uparrow$ and $d_{i+1} = \downarrow$, $s_i = s_{i+1} = \{F_i, F_{i+1}\}$ is the conclusion of an (id) rule.
- $\zeta \neq 0$ and for all $i \in \lambda$ such that $d_i = \uparrow$ and $d_{i+1} = \downarrow$, there exists trips $\tau' = \{(s'_i, F'_i, d'_i)\}_{i \in \lambda'}$, $\tau'' = \{(s''_i, F''_i, d''_i)\}_{i \in \lambda''}$ of level ζ' such that $\zeta' < \zeta$, $(s_i, F_i) = (s'_0, F'_0)$, $(s_{i+1}, F_{i+1}) = (s''_0, F''_0)$, and τ' and τ'' are associated with the same infinite branch of π .

Definition 8.1.4. *A trip with infinitely many terms of the form (s, F, \downarrow) for some s and F is said to be an **s-trip**. A **simple μMALL^∞ pre-proof** is a pre-proof that does not contain any s-trips.*

Remark 8.1.1. *Bouncing threads in [BDKS22] are special types of s-trips of level 0.*

Example 8.1.1. *Consider the proof π in Figure 7.10c. Let $\varphi = \nu x.x\wp(a \otimes a^\perp)$. Formally, the unique maximal s-trip in π is the following sequence.*

$$\begin{aligned}
& (\{\varphi_\alpha, a_\beta\}, a_\beta, \uparrow) \\
& (\{\varphi \wp (a \otimes a^\perp)_{\alpha i}, a_\beta\}, a_\beta, \uparrow) \\
& (\{\varphi_{\alpha i l}, (a \otimes a^\perp)_{\alpha i r}, a_\beta\}, a_\beta, \uparrow) \\
& (\{a_{\alpha i r r}^\perp, a_\beta\}, a_\beta, \uparrow) \\
& (\{a_{\alpha i r r}^\perp, a_\beta\}, a_{\alpha i r r}^\perp, \downarrow) \\
& (\{\varphi_{\alpha i r l}, a_{\alpha i r l}\}, a_{\alpha i r l}, \uparrow) \\
& \vdots
\end{aligned}$$

Example 8.1.2. Define π_1 to be the pre-proof in Figure 7.11. We inductively define π_n as follows:

$$\begin{array}{c}
\begin{array}{c} \pi_{n-1} \\ \nabla \\ \hline \vdash A, A^\perp \end{array} \quad \frac{\begin{array}{c} \pi_{n-1} \\ \nabla \\ \hline \vdash A, A^\perp \quad \vdash A, A^\perp \end{array}}{\vdash A, A^\perp} \text{(cut)} \\
\hline \vdash A, A^\perp \text{(cut)}
\end{array}$$

Now define π_ω as follows. Let $[C_i] = [A]$ for all $i \in \mathbb{N}$.

$$\begin{array}{c}
\begin{array}{c} \pi_1 \\ \nabla \\ \hline \vdash A, C_1^\perp \end{array} \quad \frac{\begin{array}{c} \pi_2 \\ \nabla \\ \hline \vdash C_1, C_2^\perp \end{array}}{\vdash A, C_1^\perp} \text{(cut)} \quad \frac{\begin{array}{c} \pi_3 \\ \nabla \\ \hline \vdash C_2, C_3^\perp \end{array} \quad \vdots \quad \vdash C_3}{\vdash C_2} \text{(cut)} \\
\hline \vdash A \text{(cut)}
\end{array}$$

We note that π_ω has the following s -trip of level ω .

$$\begin{aligned}
& (\{C_1\}, C_1, \uparrow) \\
& (\{C_1, C_2^\perp\}, C_1, \uparrow) \\
& (\{C_1, C_2^\perp\}, C_2^\perp, \downarrow) \\
& (\{C_2\}, C_2, \uparrow) \\
& (\{C_2, C_3^\perp\}, C_2, \uparrow) \\
& (\{C_2, C_3^\perp\}, C_3^\perp, \downarrow) \\
& (\{C_3\}, C_3, \uparrow) \\
& \vdots
\end{aligned}$$

The bulk of the theory that will be developed in this chapter and the subsequent one is for simple proofs. Note that on the one hand, Girard invented *Geometry of Interaction*, an operational semantics, by encoding the process of cut elimination at the level of trips [Gir89]. On the other hand, trips have been used to yield the coherent interpretation of a proof-net, without transiting through sequentialisation [Gir87a]. Thus, the exclusion of trips in our work possibly signifies the exclusion of some significant computational and denotational content.

8.2 Simple non-wellfounded proof structures

In this subsection, we define the geometric counterpart to simple proofs *viz.* simple proof-structures. The definitions have been designed in such a way that they look like an extension of the nets defined in Section 7.2. To our knowledge, these are the first non-wellfounded proof-nets in the literature. Notwithstanding, we note that our notion of infinite axioms is similar to the notion of *coaxioms* in [ADZ17a, ADZ17b] and that of *limit sequents* in [HK22].

Definition 8.2.1. A simple μMLL^∞ proof-structure is a 4-tuple $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta_f, \Theta_{inf})$ where:

- $\lambda \in \omega + 1$;
- for all $i \in \lambda$, $F_i^{U_i}$ is a partial syntax tree with $U_i \subset \{l, r, i\}^\omega$; $\{F_i\}_{i \in \lambda}$ is called the set of **doors**.
- \mathfrak{K} is the set of cuts i.e. a (possibly empty) set of disjoint subsets of $\{F_i\}_{i \in \lambda}$ of the form $\{C, C^\perp\}$;
- $\{F_i\}_{i \in \lambda} \setminus \bigcup_{\theta \in \mathfrak{K}} \theta$ is a finite set.
- Θ_f is the set of axiom links i.e. each element of Θ_f is pair of dual addresses;
- each element of Θ_{inf} are sets of words containing at least one infinite word (from $\{l, r, i\}^\omega$); and
- $\Theta_f \cup \Theta_{inf}$ is a partition of the set of leaves, $\mathcal{L} = \bigcup_{i \in I} \{\alpha_i u_i \mid \text{addr}(F_i) = \alpha_i, u_i \in U_i\}$

Definition 8.2.2. Let π be a simple μMLL^∞ pre-proof of the sequent $\vdash \Gamma$. The desquentialisation of π , denoted $\text{dsq}(\pi)$, is given by $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{K}, \Theta_f, \Theta_{inf})$ such that:

- for any cut in π that introduces two occurrences, C and C^\perp , we have that $\{C, C^\perp\} \in \mathfrak{K}$;
- $\{F_i\}_{i \in \lambda} = \Gamma \cup \bigcup_{\kappa \in \mathfrak{K}} \kappa$ where $\lambda = |\Gamma \cup \mathfrak{K}|$;
- for every $i \in \lambda$, $U_i = \text{addr}(F_i)^{-1} \text{addr}(\pi)$;
- for every axiom in π of the form $\frac{}{\vdash (F_i, u_i), (F_j, u_j)} \text{(id)}$, we have that $\{\text{addr}(F_i).u_i, \text{addr}(F_j).u_j\} \in \Theta_f$
- for every infinite branch γ in π , there is some $\theta \in \Theta_{inf}$ that is the largest subset of \mathcal{L} such that for every $\text{addr}(F).u \in \theta$:
 - either $u = u_1 u_2 \dots$ is an infinite word and $\{(F, u_1 \dots u_i)\}_{i \in \omega}$ is a thread of γ ; or,
 - u is a finite word and (F, u) occurs in infinitely many sequents along γ .

Note that simple proof-structures are not defined coinductively; therefore, desquentialisation cannot be defined as a coinductive process. In contrast, wellfounded proofs are inductive objects and desquentialisation of a wellfounded proof is a recursive process. In fact, as a dual of this, we will see that sequentialisation of a simple proof-structure into a non-wellfounded pre-proof (which are coinductive objects) is a corecursive process.

Proposition 8.2.1. Let π be a μMLL^∞ pre-proof with no virtual branches. Then, $\text{dsq}(\pi)$ is a simple μMLL^∞ proof-structure.

Proof. Since π has no virtual branches, every infinite branch is supported by a thread. Thus, by construction, for any $\theta \in \Theta_{inf}$ (if non-empty at all), θ satisfies the condition that it contains at least one infinite word.

We will now show that $\Theta_f \cup \Theta_{inf}$ is a partition of \mathcal{L} . Note that $\mathcal{L} = \text{addr}(\pi)$. Every finite word in $\text{addr}(\pi)$ is either an address of an occurrence in an axiom or an address of an occurrence that remains in an infinite branch forever. In both cases, it is in some $\theta \in \Theta_f \cup \Theta_{inf}$. Let $w \in \text{addr}(\pi)$ be an infinite word. Then, all the strict prefixes of w are addresses of occurrences appearing in π . It is easy to see that these occurrences form a thread. Then, by construction, there is some $\theta \in \Theta_{inf}$ such that $w \in \theta$. Now we show that for any $\theta, \theta' \in \Theta_f \cup \Theta_{inf}$, $\theta \cap \theta' = \emptyset$. Suppose not. Let $w \in \theta \cap \theta'$. If w is a finite word, this means the occurrence corresponding to w occurs two different branches of π which are not prefixes of each other. This is not possible in the multiplicative fragment. Then, w is infinite. By

construction, θ, θ' correspond to two branches β, β' in π . Let the sequent Γ be the greatest common prefix of β and β' . There is an occurrence $\varphi_{w'} \in \Gamma$ such that w' is a prefix of w . Furthermore, since $w \in \theta$ and $w \in \theta'$, there are two premisses immediately above Γ both containing the same occurrence. Contradiction! \square

Definition 8.2.3. *A simple infinit is a simple μMLL^∞ proof-structure that is the desequentialisation of some simple proof.*

We retain the progress condition from the sequent calculus and lift Definition 4.2.6 to simple μMLL^∞ proof-structures. Fix a simple μMLL^∞ proof-structure $\mathcal{R} = (\{F_i\}_{i \in \lambda}, \mathfrak{R}, \Theta_f, \Theta_{inf})$.

Definition 8.2.4. *Let $\theta \in \Theta_{inf}$ be an infinite axiom of \mathcal{R} and let $w = \alpha u_1 u_2 \dots$ be an infinite word in θ such that $\text{addr}(F) = \alpha$ for some door F of \mathcal{R} . Let $\sigma = \{[(F, u_1 u_2 \dots u_i)]\}_{i \in \omega}$. We say that w is **progressing** if the outermost connective of the smallest formula occurring infinitely often in σ is ν .*

Definition 8.2.5. *\mathcal{R} is said to be **progressing** if for all $\theta \in \Theta_{inf}$, there exists $w \in \theta$ such that w is progressing.*

8.3 Correctness criterion


A correctness criterion characterises the class of proof-nets within the class proof-structures. In particular, it is a sufficient condition for sequentialisation. What does it mean to be a necessary condition for sequentialisation? It means whatever algorithm one chooses to sequentialise, this is the bare minimum set of conditions required.

Since μMLL^* sits inside μMLL^∞ as sets of proofs, the correctness condition of μMLL^* proof structures is a necessary condition for sequentialising μMLL^∞ proof structures. We lift DR-correctness straightforwardly. In other words, Definitions 7.2.4, 7.2.6 and 7.2.7 are the same as in μMLL . Indeed, the larger class of μMLL^∞ infinets satisfies the acyclicity part of DR-correctness. The proof technique is exactly the same as the proof of Lemma 7.2.1. However, connectedness does not hold anymore for the whole class of μMLL^∞ proof-structures.

The problem is that, unconstrained, the class of pre-proofs is difficult to manage. Consider the following pre-proof.

$$\frac{\frac{\frac{\vdash \mu x.x, \nu x.x \quad \vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\mu)}{\vdash \mu x.x, \nu x.x} (\nu)}{\vdash \mu x.x, \nu x.x} (\text{cut})$$

The translation is critically disconnected:



One can wonder if this strange phenomenon is due to the fact that this pre-proof is non-progressing or due to infinitely many cuts. To nip that suspicion in its bud, we encode the infinitely many cuts using a greatest fixed point formula $\varphi = \nu y.(y \wp \nu x.x) \otimes (y \wp \mu x.x)$ as follows.

$$\pi = \frac{\frac{\frac{\frac{\vdash \mu x.x, \varphi, \nu x.x}{\vdash \mu x.x, \varphi \wp \nu x.x} (\wp)}{\vdash \mu x.x, (\varphi \wp \nu x.x) \otimes (\varphi \wp \mu x.x)} (\otimes)}{\vdash \mu x.x, \varphi, \nu x.x} (\nu)}{\vdash \mu x.x, \varphi, \nu x.x} (\mu), (\nu)$$

The reader can convince themselves that (i) π satisfies the progress condition and (ii) for *any* switching sw , the orthogonal $\mathcal{G}_{\mathcal{R}}^{sw}$ is not connected where $\mathcal{R} = \text{dsq}(\pi)$. Geometrically, this pre-proof is essentially encoding a sort of infinitary mix rule:


$$\frac{\frac{\frac{\frac{\vdash \mu x.x}{\vdash \mu x.x} (\mu)}{\vdash \mu x.x} (\text{mix}_2)}{\vdash \mu x.x} (\text{mix}_0)}{\vdash \mu x.x} (\text{mix}_2)}{\frac{\frac{\frac{\vdash \nu x.x}{\vdash \nu x.x} (\nu)}{\vdash \nu x.x} (\text{mix}_2)}{\vdash \nu x.x} (\text{mix}_0)}{\vdash \nu x.x} (\text{mix}_2)}{\vdash \mu x.x, \nu x.x} (\mu), (\nu)$$

However, we do have some understanding of desequentialisation of arbitrary simple pre-proofs.

Lemma 8.3.1. *Let π be a simple μMLL^∞ proof and let $\mathcal{R} = \text{dsq}(\pi)$. For all switchings sw , the orthogonal graph $\mathcal{G}_{\mathcal{R}}^{sw}$ is acyclic.*

We omit the proof as it is proved exactly the same way as Lemma 7.2.1. In terms of connectedness, this is as far as we can go.

This is an issue, since in order to define correctness one needs to make sure that the disconnectedness is indeed due to the encoding of an infinitary mix rule. Suppose a proof-structure has the following component:



With the absence of an explicit **mix**, it is not clear how to sequentialise this. In fact, it depends on the other components of the proof-structure if it is sequentialisable at all. The theory of proof-nets with **mix** is well developed [FR94, Bel97] but they do not scale to our situation. Consequently, we restrict our study to DR-correct μMLL^∞ proof-structures. Finally, we conjecture the following.

Open Question

Let π be a simple μMLL^∞ proof and let $\mathcal{R} = \text{dsq}(\pi)$. For all switchings sw , the orthogonal graph $G_{\mathcal{R}}^{sw}$ is either connected or has infinitely connected components.

For the rest of this thesis, we assume that we are working with DR-correct μMLL^∞ proof-structures. However, in this infinitary setting, DR-correctness is not sufficient to ensure sequentialisation. Let $F = \nu x.x\otimes x_\alpha$ and $G = a \otimes b_\beta$. Consider the μMLL^∞ proof-structure, $\mathcal{R} = (\{F^U, G^{\{l,r\}\}, \emptyset, \Theta)$ where,

$$U = (i(l+r))^\omega;$$

$$\Theta = \{\{\alpha(il)^\omega, \beta l\}, \alpha(U \setminus (il)^\omega) \cup \{\beta r\}\}.$$

Note that for any switching sw of \mathcal{R} , $[\text{SW}] = \{\{\alpha u\} \mid u \in U\} \cup \{\beta l, \beta r\}$. It is easy to check now that \mathcal{R} is DR-correct. But \mathcal{R} does not have a faithful sequentialisation *i.e.* for all proofs π that \mathcal{R} can be sequentialised into, we have that $\text{dsq}(\pi) \neq \mathcal{R}$. This is because any faithful sequentialisation must have a \otimes rule where G is the principal occurrence at some point. Now let us examine the left sequent of this rule. Either there are one (or more) suboccurrence(s) of F in the context of a in which case these suboccurrence(s) will produce infinitely many threads and $\{\alpha(il)^\omega, \beta l\}$ cannot be an axiom; or, there are no suboccurrences of F in the context of a in which case, it is not a provable sequent.

Definition 8.3.1. Let $\mathcal{R} = (\Gamma, \mathfrak{R}, \Theta_f, \Theta_{inf})$ be a DR-correct μMLL^∞ proof-structure and $\mathbb{F}_{\mathcal{R}}$ be the set of occurrences of \mathcal{R} such that their addresses do not occur in $\Theta_f \cup \Theta_{inf}$. \mathcal{R} is said to be **lock-free**¹ if for all occurrences $F \in \mathbb{F}_{\mathcal{R}}$, $\{F \ll G \mid G \in \mathbb{F}_{\mathcal{R}}\}$ is finite.

The reader is encouraged to convince themselves that the notions of subnets, empires, and kingdoms lift straightforwardly from Chapter 7 to DR-correct μMLL^∞ proof-structures. We check if the μMLL^∞ proof-structure described above is lock-free. Note that for all formulas $F' \in \text{FL}(F) \cap \mathbb{F}_{\mathcal{R}}$, we have that $G \ll F'$. There are infinitely many such F' (one for each word of $(i(l+r))^*$) hence the proof-structure is not lock-free.

Theorem 8.3.1. Let π be a simple μMLL^∞ proof and let $\mathcal{R} = \text{dsq}(\pi)$. Then, \mathcal{R} is lock-free.

Proof. Let F be an occurrence in π such that neither F occurs in an axiom in π nor is it ultimately inactive in an infinite branch of π . Consequently, $F \in \mathbb{F}_{\mathcal{R}}$. Let G be an occurrence such that $F \ll G$. In other words, $F \in \mathfrak{S}(G)$. Let Δ be the sequent in π (possibly the conclusion of π) where G is introduced. Let π' be the subproof rooted at Γ . Then, $\text{dsq}(\pi')$ is a subnet of \mathcal{R} . Furthermore, it contains G as a door; so, $\mathfrak{S}(G) \subseteq \text{dsq}(\pi')$. Therefore, $F \in \text{dsq}(\pi')$. Consequently, F occurs in π' . If there are infinitely many such G , then they are introduced higher and higher in π . Therefore, F occurs in higher and higher subproofs of π . This contradicts the fact that F is not ultimately inactive in an infinite branch of π . \square

There is an alternate characterisation of lock-freeness for nets. Fix a DR-correct μMLL^∞ proof-structure $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R}, \Theta_f, \Theta_{inf})$ and let \mathbb{F} be the set of its occurrences.

Definition 8.3.2. For any $u_i \in U_i, u_j \in U_j$, we say that (u_i, u_j) is a **coherent pair** if there exists $\theta \in \Theta_f \cup \Theta_{inf}$ such that $\{\alpha_i u_i, \alpha_j u_j\} \subseteq \theta$ where $\text{addr}(F_i) = \alpha_i$ and $\text{addr}(F_j) = \alpha_j$.

¹We borrow the terminology from concurrent programs.

Definition 8.3.3. A **switching path** is an undirected path in a partial syntax tree such that it does not go consecutively through the two premises of a \wp formula occurrence. Switching paths that do not go consecutively through the two premises of a \otimes formula occurrence are called **straight switching paths**. Two straight switching paths ρ and ρ' are said to be **coherent** if there exists $u, u' \theta \in \Theta_f \cup \Theta_{inf}$ such that $\{u\rho, u'\rho'\} \subseteq \theta$.

Definition 8.3.4. A **switching sequence** is a finite sequence $\sigma = \{(\gamma_i, \delta_i)\}_{i \in [n]}$ of pairs of switching paths such that:

- for all $i \in [n]$, γ_i and δ_i are coherent, and
- for all $i \in [n-1]$, there exists a switching path ρ connecting δ_i and γ_{i+1} .

Two occurrences, $F = (F_i, u)$ and $G = (F_j, u')$, are said to be **connected** by the switching sequence, σ , if $\text{src}(\gamma_1) = u$ and $\text{tgt}(\gamma_n) = u'$.

For ease of presentation, we restrict ourselves to the case when \mathfrak{K} is finite. Define

$$\begin{aligned} \mathbb{T} &= \{F \in \mathbb{F}_{\mathcal{R}} \mid F = G \otimes H\} \\ \mathbb{P} &= \{F \in \mathbb{F}_{\mathcal{R}} \mid F = G \wp H\} \end{aligned}$$

Definition 8.3.5. Let $A \wp B \in \mathbb{P}$ and $C \otimes D \in \mathbb{T}$. $A \wp B$ is said to be **t-connected** to C if there exists switching sequences σ, σ' such that they do not go through $A \wp B$ and either one of the following holds:

- A is connected to C by σ and B is connected to D by σ' ; or,
- A is connected to D by σ and B is connected to C by σ' .

If we wanted to incorporate cuts, we would need to update \mathbb{T} as $\{F \in \mathbb{F}_{\mathcal{R}} \mid F = G \otimes H\} \cup \mathfrak{K}$ and modify the definition of t-connectedness accordingly.

Definition 8.3.6. The **dependency graph** of \mathcal{R} , denoted $\text{DepGrph}(\mathcal{R})$, is the directed graph (V, E) such that:

- $V = \mathbb{T} \cup \mathbb{P}$;
- for every $F, G \in V$, $E(F, G)$ if $F \in \mathbb{T}$, $G \in \mathbb{P}$, and G is t-connected to F ; and,
- for every $F, G \in V$, $E(F, G)$ if F is a suboccurrence of G .

Theorem 8.3.2. \mathcal{R} is lock-free iff $\text{DepGrph}(\mathcal{R})$ has a finite degree and does not contain a ray.

Proof. Let $E(F, G)$ in $\text{DepGrph}(\mathcal{R})$. We claim that $F \in \wp(G)$. If F is a suboccurrence of G then we are done since kingdoms are upward closed. Otherwise, let $F = A \otimes B \in \mathbb{T}$ and $G = C \wp D \in \mathbb{P}$ such that G is t-connected to F . Wlog, assume A and C are connected and B and D are connected. So, $C \in \wp(A)$ (and $D \in \wp(B)$). Therefore, $C \in \wp(A) \cap \wp(G)$. By Lemma 7.2.2, $\wp(A) \cup \wp(G)$ is a subnet of \mathcal{R} . Suppose $F \notin \wp(G)$ and $G \in \wp(A)$. Then $\wp(A) \cup \wp(G)$ is a subnet with door A , which is larger than $\wp(A)$, since it contains G , contradicting the definition of empires. Therefore, either $F \in \wp(G)$ or $G \in \wp(A)$. Reasoning similarly as above using the fact $D \in \wp(B)$, we have $G \in \wp(B)$. So, $G \in \wp(A) \cap \wp(B)$ which implies $D \in \wp(A)$ and $C \in \wp(B)$. Then, $C \in \wp(A)$ and $D \in \wp(B)$; noting, $F \notin \wp(A)$, this contradicts the DR-correctness of $\wp(A)$ (and similarly also that of $\wp(B)$). Therefore, $F \in \wp(G)$.

Finally, the fact that \ll is a partial order gives us that $\text{DepGrph}(\mathcal{R})$ is a directed acyclic graph. Conclude by König's lemma that lock-freeness is equivalent to $\text{DepGrph}(\mathcal{R})$ having a finite degree and no rays. \square

Definition 8.3.7. A μMLL^∞ proof-structure is said to be an **infinet** if it is DR-correct and lock-free.

We will now show that this notion of correctness is sufficient to ensure sequentialisation. We first note that the standard technique of treating cuts as tensors can now potentially lead to sequents with infinitely many occurrences (due to infinitely many cuts). We generalise our setting to “quasi” proof-structures with potentially infinitely many conclusions.

Note that in this infinitary setting, we need to strengthen the sequentialisation algorithm in Figure 7.3 by adding a notion of *fairness* since we may never explore one thread by forever prioritising the


```

function FAIRSEQUENTIALISE( $(\Gamma, \Theta), \tau$ )
  if  $\Gamma = \{\varphi_\alpha^{\{\varepsilon\}}, \varphi_\beta^\perp^{\{\varepsilon\}}\}$  and  $\Theta = \{\alpha, \beta\}$  then
    return  $\frac{}{\vdash \varphi, \varphi^\perp}$  (id)
  else
    Choose  $F_i^{U_i} \in \Gamma$  such that  $U_i \neq \{\varepsilon\}$  and  $\tau(F_i) = \min(\tau)$ .
    if  $F_i = G \wp H$  then
      return  $\frac{\text{FAIRSEQUENTIALISE}(\mathfrak{W}(G), \tau_G)}{\vdash [\Delta], [G \wp H]}$  ( $\wp$ )
    else if  $F_i = G \otimes H$  and  $\exists j \neq i. F_i \ll F_j$  then
      return  $\frac{\text{FAIRSEQUENTIALISE}(\mathfrak{W}(G), \tau_G) \text{ FAIRSEQUENTIALISE}(\mathfrak{W}(H), \tau_H)}{\vdash [\Delta], [G \otimes H]}$  ( $\otimes$ )
    else if  $F_i = \eta x.G$  then  $\triangleright \eta \in \{\mu, \nu\}$ 
      return  $\frac{\text{FAIRSEQUENTIALISE}(\mathfrak{W}(G[\eta x.G/x]), \tau_{G[\eta x.G/x]})}{\vdash [\Gamma], [\eta x.G]}$  ( $\eta$ )
    end if
  end if
end function

```

Figure 8.1: The function FAIRSEQUENTIALISE

sequentialisation of another thread. For example, consider the proof-structure $(\{\mu x.x_\alpha^{i\omega}, \nu x.x_\beta^{i\omega}\}, \{\alpha i^\omega, \beta i^\omega\})$. At each step of sequentialisation, suboccurrences of both $\mu x.x_\alpha$ and $\nu x.x_\beta$ can be chosen to be sequentialised. If one does not choose suboccurrences of $\nu x.x_\beta$ after some finite number of steps, we will miss a thread in the resulting pre-proof (and in this the only progressing thread). Fairness is ensured by timestamping the occurrences of the inifinet with elements of $\mathbb{N} \cup \{\infty\}$, which dictates that at any particular step the occurrences with the least timestamp is to be sequentialised.

Given a proof-structure \mathcal{R} , we define a **timestamping function** as any function of the form $\tau : \mathbb{F}_{\mathcal{R}} \rightarrow \mathbb{N} \cup \{\infty\}$. We extend the natural number ordering by asserting that for all $n \in \mathbb{N}$, $n < \infty$. We define the **minimal finite image** of τ , denoted $\min(\tau)$, as the minimum of $\text{Im}(\tau)$. For any occurrence F in \mathcal{R} we define the non-deterministic function τ_F as follows such that t is an arbitrary natural number greater than $\tau(F')$ where F is the immediate suboccurrence of F' in \mathcal{R} .

$$\tau_F(G) = \begin{cases} t & \text{if } G \text{ is maximal in the } \ll \text{ ordering in } \mathfrak{W}(F); \\ \tau(G) & \text{otherwise.} \end{cases}$$

We describe the fair sequentialisation algorithm in Figure 8.1. It describes a corecursive function FAIRSEQUENTIALISE that takes as input a μMLL^∞ proof-structure (Γ, Θ) and a timestamping function τ . We initialise τ by assigning arbitrary natural numbers to doors that are maximal in \ll ordering or not tensor formula occurrences and ∞ to the other occurrences in $\mathbb{F}_{\mathcal{R}}$.

Remark 8.3.1. *There is a catch here. The above initial condition is enough to justify fairness when there are only finitely many cuts. However, in the presence of infinitely many cuts, it is possible that infinitely many of them are maximal in the \ll ordering after translation into tensors. We have to carefully initialise the timestamping function such that infinitely many natural numbers are free to be used as timestamps at later stages of the sequentialisation. For example, one can consider τ which injectively timestamps every maximal door in the \ll ordering and every non-tensor door by powers of two and every other occurrence by ∞ . Note that any cofinite sequence in place of powers of two would also work.*

Lemma 8.3.2. *The timestamping assigns a finite natural number to every occurrence of the inifinet that one starts with after some finite iterations of the sequentialisation process.*

Proof. We will prove by contradiction. Suppose there are occurrences which are never assigned a finite natural number by the time stamping algorithm. Let $F_0 = (F_i, u)$ be the minimal such occurrence i.e. for all such other occurrences (F_j, u') , we have $|u| \leq |u'|$.

Then, by construction, after finite iterations of the sequentialisation process, it becomes a door (otherwise we would have found a node with even a lesser distance). Since it is not assigned a finite number, it is a tensor occurrence that is not maximal in the kingdom ordering. Consider $S = \{F \mid F \ll F_0\}$. Since t is not maximal S is non-empty. By lock-freeness S is finite.

If every $F \in S$ can be assigned a finite natural number after finitely many steps of sequentialisation, then F_0 can be assigned a finite natural number after finitely many steps of sequentialisation, and we have a contradiction. Therefore there exists $F_1 \in S$ such that it is never assigned a finite natural number. We repeat the same argument for F_1 as we did for F_0 . By continuing like this *ad infinitum*, we obtain an infinite sequence F_0, F_1, F_2, \dots ; but observe that for all $i \geq 0$, $F_i \ll F_{i+1}$. Recall that \ll is transitive. Therefore, S is infinite contradicting the lock-freeness of \mathcal{R} . \square

Theorem 8.3.3. FAIRSEQUENTIALISE is *productive and correct*.

Proof. The productivity of FAIRSEQUENTIALISE follows from the fact by lock-freeness, one can always find a door which is maximal in the kingdom ordering. The correctness follows from Lemma 8.3.2 and the straightforward extension of Lemma 7.2.3. \square

Corollary 8.3.3.1. π is a proof iff $\text{dsq}(\pi)$ is progressing.

8.4 Canonicity

Recall the \sim relation over proofs from Section 7.2.4 that quotients proofs up to *finite* permutation of inference rules. We did not explicitly mention the finiteness for μMLL^* proof-nets since the objects themselves were finite. However, in simple infinets, it is indeed possible to permute a rule infinitely up. Consider the following two proofs, π and π' , of $\vdash \mu x.x, \nu x.x$.

$$\left(\frac{\frac{\vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\mu)}{\vdash \mu x.x, \nu x.x} (\nu) \right) \quad \left(\frac{\frac{\vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\nu)}{\vdash \mu x.x, \nu x.x} (\mu) \right)$$

Note that they have the same infinite $(\{\mu x.x_{\alpha}^{i^{\omega}}, \nu x.x_{\beta}^{i^{\omega}}\}, \emptyset, \{\{\alpha i^{\omega}, \beta i^{\omega}\}\})$ but $\pi \not\sim^* \pi'$. We need to consider a larger equivalence class. Suppose we allow infinite permutations.

Let $\pi \sim^{\omega} \pi'$ if there exists a infinite sequence of proofs $\{\pi_i\}_{i=0}^{\infty}$ such that $\pi_0 = \pi$, $\pi_i \sim \pi_{i+1}$, and $\lim_{i \rightarrow \infty} d(\pi_i, \pi') = 0$. There are two issues here. \sim^{ω} is not transitive (so not an equivalence) and potentially quotients more than proof-net equality.

Consider the proofs π above. One can permute down the second (ν) , third (ν) , and so on. In other words, we have the sequence $\{\pi_i\}_{i=0}^{\infty}$ such that $\pi_0 = \pi$ and

$$\begin{array}{c} \pi \\ \triangle \\ \frac{\vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\nu^i) \\ \pi_i = \end{array}$$

Therefore, $\pi \sim^{\omega} \pi''$ where π'' is:

$$\frac{\vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\nu)$$

Similarly we get that $\pi \sim^{\omega} \pi'''$ where π''' is:

$$\frac{\vdash \mu x.x, \nu x.x}{\vdash \mu x.x, \nu x.x} (\mu)$$

Now observe that π''' is not progressing, the desquentialisations of π'' and π''' are different, and indeed one cannot obtain π''' from π'' by infinitary permutations (or vice versa). Consequently, \sim^{ω} is too large.

As a first step in restricting \sim^{ω} , we annotate \sim with a formula occurrence. In particular, we write $\pi \sim_F \pi'$ if F is the active formula of the inference rule r that has been permuted up (*i.e.* the height of r in π' is greater than its height in π).

Definition 8.4.1. *Given two pre-proofs π and π' , we define $\pi \sim_F^{\omega} \pi'$ if there exists a infinite sequence of proofs $\{\pi_i\}_{i=0}^{\infty}$ such that the following holds.*

- $\pi_0 = \pi$,
- $\pi_i \sim_{F_i} \pi_{i+1}$ for some F_i ,
- $\lim_{i \rightarrow \infty} d(\pi_i, \pi') = 0$, and
- $\text{Inf}(\{F_i\}_{i=0}^{\infty}) = \emptyset$.

Finally, define \sim^{∞} as $\sim^* \cup \sim_F^{\omega}$

The last condition essentially means that infinite permutations are allowed only if every occurrence is permuted upwards at most finitely many times.

Theorem 8.4.1. *For all simple μMLL^{∞} proofs π, π' , $\pi \sim^{\infty} \pi'$ iff $\pi \equiv_{\text{PN}} \pi'$.*

Proof sketch. The proof is similar to that of Theorem 7.2.2. There are some minor observations to be made regarding infinite steps of permutation. Firstly, since the infinite sequence of permutations is *fair i.e.* every instance of a rule is permuted only finitely many steps, we have that if $\pi \sim^{\infty} \pi'$ then the same set of instances of inference rules occurs in both π and π' . The fairness of the sequentialisation preserves this property. \square

8.5 Regular infinets

In non-wellfounded sequent calculi, one obtains the finitely presentable by considering a fragment of pre-proofs that only have finitely many distinct pre-proofs *i.e.* regular pre-proofs. This notion, however, is crucially related to the sequent presentation and a necessary condition is that the non-wellfounded proof must have finitely many distinct sequents occurring in it. There is no proper proof-net counterpart to this. Moreover, regularity is not preserved under the permutation equivalence defined in the previous section. For example, in the following, inferences can be permuted in a circular proof π_0 to obtain a non-wellfounded proof π_1 such that $\pi_0 \sim_f^\omega \pi_1$.

$$\pi_0 = \frac{\frac{\frac{}{\vdash \mu x.x, \nu x.x} (\mu)}{\vdash \mu x.x, \nu x.x} (\nu)}{\vdash \mu x.x, \nu x.x} (\mu)}{\vdash \mu x.x, \nu x.x} (\nu) \quad \pi_1 = \frac{\frac{\frac{}{\vdash \nu x.x, \mu x.x} (\mu^i)}{\vdash \nu x.x, \mu x.x} (\nu)}{\vdash \nu x.x, \mu x.x} (\mu^i)}{\vdash \nu x.x, \mu x.x} (\nu)$$

π_{i+1}

Consequently, we will define a fragment of μMLL^∞ proof-structures, called *regular μMLL^∞ proof-structure*, such that the desequentialisation of μMLL° proofs are regular μMLL^∞ proof-structures but they do not necessarily sequentialise to circular proofs. These objects are more robust than μMALL° proof-nets designed in Section 7.3 since two different presentations of the same regular proofs desequentialise to the same regular μMLL^∞ proof-structure.

Definition 8.5.1. A regular μMLL^∞ proof-structure is a 3-tuple $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R}, \Theta)$ where:

- $\lambda \in \omega$,
- for all $i \in \lambda$, $F_i^{U_i}$ is a partial syntax tree and U_i is a regular language;
- \mathfrak{R} is a finite (possibly empty) set of disjoint subsets of $\{F_i\}_{i \in \lambda}$ of the form $\{C, C^\perp\}$; and,
- $\Theta \subset (\{l, r, i\}^\infty)^2$ is a regular language such that it is an equivalence over $\{l, r, i\}^\infty$ and the set of equivalence classes is a partition of $\mathcal{L} = \bigcup_{i \in I} \{\alpha_i u_i \mid \text{addr}(F_i) = \alpha_i, u_i \in U_i\}$ such that for all $w \in \mathcal{L}$ the following hold.
 - Either the equivalence class containing w is of the form $\{\alpha_i u, \alpha_j u'\}$ such that $[(F_i, u)] = [(F_j, u')]^\perp$ where $\text{addr}(F_i) = \alpha_i, u \in U_i, \text{addr}(F_j) = \alpha_j$, and $u' \in U_j$.
 - Or, the equivalence class containing w contains at least one infinite word.

Definition 8.5.2. Let $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R}, \Theta)$ be a regular μMLL^∞ proof-structure. The **unfolding** of \mathcal{R} is defined as the μMLL^∞ proof-structure $U(\mathcal{R}) = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R}, \Theta_f, \Theta_{inf})$ where $\Theta_f \cup \Theta_{inf}$ is the partition induced by Θ such that if an equivalence class θ contains an infinite word then $\theta \in \Theta_{inf}$; otherwise, $\theta \in \Theta_f$.

Theorem 8.5.1. Let π be a simple μMLL° pre-proof such that there are no cycles containing a cut. Then, $\text{dsq}(\pi) = U(\mathcal{R})$ for some regular μMLL^∞ proof-structure \mathcal{R} .

Proof. Since there are no cycles containing a cut, the unfolding of π has finitely many cuts. Therefore there are finitely many doors (say, n) of $\text{dsq}(\pi) = (\{F_i^{U_i}\}_{i \in [n]}, \mathfrak{R}, \Theta_f, \Theta_{inf})$.

Now, we claim that for each F^U in $\text{dsq}(\pi)$, U is a regular language. We will provide a finite state automata \mathcal{A}_F and a Büchi automata \mathcal{B}_F such that $\mathcal{L}(\mathcal{A}_F) \cup \mathcal{L}(\mathcal{B}_F) = U$. Note that there are only finitely formula occurrences in π . So, for each door F , there is a bar B of U such that every address on B_F corresponds to either an axiom or the source of a backedge in π . Let U_F be the finite prefix of U such that the leaves of U_F are exactly B . Define the following.

- $Q = \bigcup_{F \in \Gamma} U_F$
- $\Sigma = \{l, r, i\}$
- $\Delta = \bigcup_{F \in \Gamma} \{(w, a, wa) \mid w \in U_F, wa \in W_F, a \in \Sigma\} \cup \Delta_\varepsilon$ where

$$\Delta_\varepsilon = \bigcup_{F, G \in \Gamma} \{(w, \varepsilon, w') \mid w \in U_F, w' \in U_G, \text{rename}(\alpha w) = \alpha' w', \text{addr}(F) = \alpha, \text{addr}(G) = \alpha'\}$$

- $Q_F = \text{Root of } U_F$
- $\mathcal{F} = \bigcup_{F \in \Gamma} B_F$

Define $\mathcal{A}_F = (Q, \Sigma, \Delta, Q_F, \mathcal{F})$ construed as a finite state automaton and $\mathcal{B}_F = (Q, \Sigma, \Delta, Q_F, \mathcal{F})$ construed as a Büchi automaton. It is easy to check that $\mathcal{L}(\mathcal{A}_F) \cup \mathcal{L}(\mathcal{B}_F) = U$. We will now show that $\Theta = \{(w, w') \mid \{w, w'\} \subseteq \Theta_f \cup \Theta_{in_f}\}$ is a regular language. Let S be the set of all sequents occurring in π . Define the following.

- $Q = S \times \bigcup_{F \in \Gamma} U_F \times S \times \bigcup_{F \in \Gamma} U_F$ restricted to all (s, α, s', α') such that there exists F, G such that $F \in s, G \in s', \text{addr}(F) = \alpha$, and $\text{addr}(G) = \alpha'$.
- $\Sigma = \{l, r, i\}$
- $((s, \alpha, s', \alpha'), (a, b), (t, \beta, t', \beta')) \in \Delta$ if the following holds:
 - either $\alpha' = \alpha a$ or there is a back-edge in π from s to t such that φ_α is renamed $\varphi_{\alpha'}$ for some formula φ .
 - either $\beta' = \beta b$ or there is a back-edge in π from s' to t' such that φ_β is renamed $\varphi_{\beta'}$ for some formula φ .
- $Q_i = \{(s_0, \alpha, s_0, \beta) \mid s_0 \text{ is the conclusion of } \pi\}$
- $\mathcal{F} = \{(s, \alpha, s, \beta) \mid s \text{ is the conclusion of an (id) rule}\}$.
- $\mathcal{F}^\omega = \{(s, \alpha, s, \beta) \mid (s, \alpha, s, \beta) \in Q\}$.

Let $\mathcal{A} = (Q, \Sigma, \Delta, Q_i, \mathcal{F})$ be a finite state automaton and $\mathcal{B} = (Q, \Sigma, \Delta, Q_i, \mathcal{F}^\omega)$ be a Büchi automaton. It is routine to check that $\mathcal{L}(\mathcal{A}) \cup \mathcal{L}(\mathcal{B}) = \Theta$. Therefore, $\mathcal{R} = (\{F_i^{U_i}\}_{i \in [n]}, \mathfrak{R}, \Theta)$ is a regular μMLL^∞ proof-structure such that $U(\mathcal{R}) = \text{dsq}(\pi)$. \square

However, the sequentialisation of any regular infinit is not necessarily a regular derivation. Consider the following regular infinit where $\varphi = \nu x.x\wp x$.

$$(\{\varphi_\alpha^{i(l+r)^\omega}\}, \emptyset, \alpha(i(l+r)^\omega) \times \alpha(i(l+r)^\omega))$$

Every sequentialisation of its unfolding is non-regular. For example, the following proof.

$$\begin{array}{c} \vdots \\ \frac{\vdash \varphi, \varphi, \varphi}{\vdash \varphi \wp \varphi, \varphi} (\wp) \\ \frac{\vdash \varphi \wp \varphi, \varphi}{\vdash \varphi, \varphi} (\wp) \\ \frac{\vdash \varphi, \varphi}{\vdash \varphi \wp \varphi} (\wp) \\ \frac{\vdash \varphi \wp \varphi}{\vdash \varphi} (\nu) \end{array}$$

A regular μMLL^∞ proof-structure \mathcal{R} is DR-correct if $U(\mathcal{R})$ is DR-correct. It is progressing if $U(\mathcal{R})$ is progressing.

Theorem 8.5.2. *Checking the progress condition on regular μMLL^∞ proof-structures is decidable.*

Proof. Let $\mathcal{R} = (\Gamma, \mathfrak{R}, \Theta)$ be a regular μMLL^∞ proof-structure. Let $F^U \in \Gamma$. Since U is regular language, it can be construed as a regular tree. Therefore, we can assume that the automata \mathcal{A} over finite and infinite words (the accepting condition for finite words being that their runs end in the final state and for infinite words being the Büchi condition) has a finite prefix of \bar{U} as the set of states such that it has transitions of the form $w^a w a$ or $w^\varepsilon w'$ such that $\llbracket (F, w) \rrbracket = \llbracket (F, w') \rrbracket$.

We will define a parity automata \mathcal{P}_F which is nothing but \mathcal{A} with a parity accepting condition. We assign colours to Q_F by the map $\chi : Q_F \rightarrow \mathbb{N}$ such that:

- If (F, w) is a μ -occurrence (respectively, ν -occurrence), then $\chi(w)$ is odd (respectively, even).
- If $(F, w) \leq (F, w')$ then $\chi(w) \leq \chi(w')$.

Let $\mathcal{L}_p = \bigcup_{F^U \in \Gamma} \mathcal{L}(\mathcal{P}_F)$. Then, \mathcal{R} is progressing iff $\Theta \subseteq (\mathcal{L}_p \times \mathcal{L}) \cup (\mathcal{L}_p \times \mathcal{L})$. Since this is decidable, we are done. \square

Open Question

Checking DR-correctness of regular μMLL^∞ proof-structures is decidable.

8.6 Generalising to non-simple proofs

In the next series of definitions, we define general non-wellfounded proof-structures. We disassemble some components of simple μMLL^∞ proof-structures so that it is easier to talk about them. We first define a *chain of syntax trees* which is essentially a simple μMLL^∞ proof-structure without the axiom components. Then we define pre-visitable paths which are the object corresponding to pre-trips in proof-structures. This allows us to define non-wellfounded proof-structures. Thereafter, we define visitable paths to be pre-visitable paths which can be assigned a level consistently, just like we define trips from pre-trips. Subsequently, we define a sanity check on proof-structures and their desequentialisation. We conclude the section with some details on the correctness condition on these structures.

Definition 8.6.1. A chain of syntax trees is a tuple of the form $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R})$ such that:

- $\lambda \in \omega + 1$
- $F_i^{U_i}$ is a partial syntax tree for every $i \in \lambda$.
- \mathfrak{R} contains sets of formulas of the form $\{F_i, F_j\}$ such that $F_i = F_j^\perp$.
- $\{F_i\}_{i \in \lambda} \setminus \cup_{\theta \in \mathfrak{R}} \{F \mid F \in \theta\}$ is finite.

Definition 8.6.2. Let $\mathcal{R} = (\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R})$ be a chain of syntax trees. A **pre-visitable path** is a sequence of triples of the form $\{(t_i, g_i, d_i)\}_{i < \omega}$ such that for all $i < \omega$, $g_i \neq d_i$ and exactly one of the following holds:

- t_i is a tensor formula occurrence (F_j, u) for some $j \in \lambda$ such that $u \in \overline{U_j}$, $g_i \in U_j$ and $d_{i+1} \in U_j$ are incomparable, and u is their greatest common prefix.
- t_i is an element $\{F_j, F_{j'}\}$ of \mathfrak{R} such that $g_i \in U_j$ and $d_{i+1} \in U_{j'}$.

If $\rho = (t_i)_{i < \omega}$ and $\rho' = (t'_i)_{i < \omega}$ are two pre-visitable paths such that there exists n such that for all $i < \omega$, $t_i = t'_{i+n}$, then we write $\rho \sqsubseteq \rho'$. We denote by $\mathfrak{P}^{\mathcal{R}}$ the set of maximal pre-visitable paths of \mathcal{R} in the \sqsubseteq ordering.

Definition 8.6.3. A μMLL^∞ **proof-structure** is a 5-tuple $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R}, V, \Theta_f, \Theta_r, \Theta_v)$ such that:

- $(\{F_i^{U_i}\}_{i \in \lambda}, \mathfrak{R})$ is a chain of syntax trees;
- $V \subseteq \mathfrak{P}^{\mathcal{R}}$;
- each element of Θ_f is pair of dual addresses i.e. of the form $\{\alpha_i u_i, \alpha_j u_j\}$ such that $\text{addr}(F_i) = \alpha_i$, $\text{addr}(F_j) = \alpha_j$, and $\lceil (F_i, u_i) \rceil = \lceil (F_j, u_j) \rceil^\perp$;
- each element of Θ_r contains at least one infinite word;
- each element of Θ_v contains no infinite word and at least one element of V ;
- $\Theta_f \cup \Theta_r \cup \Theta_v$ is a partition of $\mathcal{L} \cup \mathfrak{P}^{\mathcal{R}}$.

Definition 8.6.4. Given a μMLL^∞ proof-structure, a **visitable path of level ζ** for some ordinal ζ is a pre-visitable path $\rho = \{(t_i, g_i, d_i)\}_{i < \omega} \in V$ such that:

- $\zeta = 0$ and for all $i < \omega$, $\{d_i, g_{i+1}\} \in \Theta_f \cup \Theta_r$.
- $\zeta \neq 0$ and for all $i < \omega$, there are visitable paths $\rho' = \{(t'_i, g'_i, d'_i)\}_{i < \omega}$ and $\rho'' = \{(t''_i, g''_i, d''_i)\}_{i < \omega}$ of level ζ' and ζ'' respectively such that $(t_i, g_i, d_i) = (t'_0, g'_0, d'_0)$, $(t_{i+1}, g_{i+1}, d_{i+1}) = (t''_0, g''_0, d''_0)$, and $\rho', \rho'' \in \theta$ for some $\theta \in \Theta_r \cup \Theta_v$ and $\zeta', \zeta'' < \zeta$.

A **visitable path** is a visitable path of level ζ for some ordinal ζ .

Definition 8.6.5. A μMLL^∞ proof-structure is said to be **well-formed** if V is exactly the set of maximal visitable paths.

Definition 8.6.6. Let π be a μMLL^∞ pre-proof of the sequent $\vdash \Gamma$. The **desequentialization** of π , denoted $\text{dsq}(\pi)$, is given by $(\{F_i^{U_i}\}_{i \in I}, \mathfrak{R}, V, \Theta_f, \Theta_r, \Theta_v)$ such that

- for any cut in π that introduces two occurrences, C and C^\perp , $\{C, C^\perp\} \in \mathfrak{K}$.
- $\{F_i\}_{i \in I} = \Gamma \cup \bigcup_{\kappa \in \mathfrak{K}} \kappa$.
- for every $i \in I$, $U_i = \text{addr}(F_i)^{-1} \text{addr}(\pi)$.
- for every axiom linking (F_i, u_i) to (F_j, u_j) , $\{\text{addr}(F_i).u_i, \text{addr}(F_j).u_j\} \in \Theta_f$.
- for every maximal s -trips in π , we collect the points of alternation of directions that gives us a sequence of cuts or tensors and axioms. This gives us a set of maximal visitable paths V .
- for every real infinite branch γ in π , $\theta \in \Theta_r$ is the largest subset of $\mathcal{L} \cup V$ such that:
 - for every $\text{addr}(F)u \in \theta$, either $u = u_1 u_2 \dots$ is an infinite word and $\{(F, u_1 \dots u_i)\}_{i \in \omega}$ is a straight thread of γ ; or u is a finite word and (F, u) occurs in infinitely many sequents along γ .
 - for every $v \in \theta \cap V$ there exists a trip ρ associated with γ such that v is obtained from ρ .
- For every virtual infinite branch, γ , in π , $\theta \in \Theta_v$ is the largest subset of $\mathcal{L} \cup V$ such that:
 - for every $\text{addr}(F)u \in \theta$, u is a finite word and (F, u) occurs in infinitely many sequents along γ .
 - for every $v \in \theta \cap V$ there exists a trip ρ associated with γ such that v is obtained from ρ .

We will lift the notion of orthogonal graphs from Section 8.3 to our setting. The orthogonal graph is bipartite graph $(\Theta, [\text{SW}], E)$ where $\Theta = \Theta_f \cup \Theta_r \cup \Theta_v$ and $[\text{SW}], E$ is defined as usual. Consequently, Θ is a partition of $\mathcal{L} \cup V$ whereas $[\text{SW}]$ is a partition of \mathcal{L} . We generalise the notion of a path in the orthogonal graph as follows.

Definition 8.6.7. Let $(\Theta, [\text{SW}], E)$ be an orthogonal graph. A **path** is defined as a sequence $\{u_i\}_{i \in \lambda}$ ($\lambda < \omega + 1$) such that for all i :

- if $u_i \in \Theta$ and $i + 1 \leq \lambda$ then $u_{i+1} \in [\text{SW}]$;
- if $u_i \in [\text{SW}]$ and $i + 1 \leq \lambda$ then $u_{i+1} \in \Theta$; and,
- either one of the following holds:
 - $(u_i, u_{i+1}) \in E$
 - there exists $\rho \in u_i \cap V$ such that there is $(t, g, d) \in \rho$ with $\{vg, vd\} \subseteq u_{i+1}$ for some v .
 - there exists $\rho \in u_{i+1} \cap V$ such that there is $(t, g, d) \in \rho$ with $\{g, d\} \subseteq u_i$

Proposition 8.6.1. Let $\mathcal{R} = (\Gamma, \mathfrak{K}, V, \Theta_f, \Theta_r, \Theta_v)$ be well-formed μMLL^∞ proof-structure. Then, $V = \emptyset$ iff for all switchings sw , $G_{\mathcal{R}}^{sw}$ does not contain an infinite path.

Proof. We will prove the contrapositive in both directions. Let $\rho = \{(t_i, g_i, d_i)\}_{i < \omega}$ be a visitable path. We will show that there exists a switching sw such that $G_{\mathcal{R}}^{sw}$ contains an infinite path. Choose a switching such that for all i , $(g_i, d_i) \in \text{SW}$. Now, consider the path $u_0 u_1 u_2 \dots$ in $G_{\mathcal{R}}^{sw}$ such that for all $i = 2k$, $\{g_k, d_k\} \subseteq u_i$ and for all $i = 2k + 1$, either $\{d_k, g_{k+1}\} \in u_i$ or there exists visitable paths $\rho, \rho' \in u_i$ containing (t_k, g_k, d_k) and $(t_{k+1}, g_{k+1}, d_{k+1})$ respectively. It is routine to check that $u_0 u_1 \dots$ is indeed a path in $G_{\mathcal{R}}^{sw}$.

For the opposite direction, suppose there is a switching sw such that $G_{\mathcal{R}}^{sw}$ contains an infinite path $u_0 u_1 \dots$ assuming wlog that $u_0 \in [\text{SW}]$. One can now construct a pre-visitable path $\rho = \{(t_k, g_k, d_k)\}_{k < \omega}$. For all $k > 0$, let define (t_k, g_k, d_k) as follows:

- if $(u_{2k-1}, u_{2k}) \in E$ and (u_{2k}, u_{2k+1}) then t is either a tensor of the form (F, u) where $u = u_{2k-1} \cap u_{2k} \cap u_{2k+1}$ or t is a cut $\{C, C^\perp\}$ such that $u_{2k-1} \cap u_{2k} \in U$ and $u_{2k} \cap u_{2k+1} \in U'$ for $C^U, C^{\perp U'} \in \Gamma$.
- if u_{2k-1} or u_{2k+1} has visitable paths which have terms (t, g, d) such that $\{vg, vd\} \subseteq u_{2k}$ for some v then $(t_k, g_k, d_k) = (t, g, d)$.

It is routine to check that ρ is in fact a visitable path. \square




Figure 8.2: A DR-correct μMLL^∞ proof-structure with infinitely ascending chain of kingdoms.

Let ρ be a visitable path in \mathcal{R} . Fix a switching sw . If there is a ray r in the orthogonal graph $G_{sw\rho}$ that corresponds to ρ , we denote this path by $\text{pth}_{sw}(\rho)$. By Proposition 8.6.1, we know that there is at least one such switching. Furthermore, define $\overline{\text{pth}_{sw}(\rho)}$ to be $\text{pth}_{sw}(\rho)$ construed as a set of vertices. Let $v \in \Theta$. Define:

$$\begin{aligned} \text{PreReach}^0(v) &:= \{v' \mid (v, v') \in E \cap \Theta \times [\text{SW}]\}; \\ \text{PreReach}^{\lambda \cdot 2 + 1}(v) &:= \text{PreReach}^{\lambda \cdot 2}(v) \cup \{u \mid \exists v' \in \text{PreReach}^{\lambda \cdot 2}(v); (v', u) \in E \cap [\text{SW}] \times \Theta\} \\ \text{PreReach}^{\lambda \cdot 2 + 2}(v) &:= \text{PreReach}^{\lambda \cdot 2 + 1}(v) \cup \{u \mid \exists v' \in \text{PreReach}^{\lambda \cdot 2 + 1}(v); (u, v') \in E \cap \Theta \times [\text{SW}]\} \\ \text{PreReach}^\lambda(v) &:= \bigcup_{\lambda' < \lambda} \text{PreReach}^{\lambda'}(v); \end{aligned}$$

Similarly, one can define $\text{PreReach}^\lambda(v)$ for $v \in [\text{SW}]$. Now, let $\lambda + 1$ be a successor ordinal. We will define $\text{Reach}^{\lambda+1}(v)$ from $\text{PreReach}^{\lambda+1}(v)$. Let $S = (\text{PreReach}^{\lambda+1}(v) \setminus \text{PreReach}^\lambda(v)) \cap \Theta$. Noting that the set of visitable paths of \mathcal{R} is denoted by V , we define $\text{Reach}^{\lambda+1}(v)$ as follows.

$$\text{Reach}^{\lambda+1}(v) = \begin{cases} \text{PreReach}^{\lambda+1}(v) & \text{if } S = \emptyset; \\ \text{PreReach}^{\lambda+1}(v) \cup \bigcup_{\theta \in S} \overline{\text{pth}_{sw}(\rho)} \mid \rho \in \theta \cap V & \text{otherwise.} \end{cases}$$

Let λ be a limit ordinal. We will define $\text{Reach}^\lambda(v)$ from $\text{PreReach}^\lambda(v)$. Wlog, assume that $v \in \text{SW}$. Let $\{v_i\}_{i \in \omega}$ such that $v_0 = v$, for all $i \in \omega$, $v_i \in \text{PreReach}^\lambda(v)$, $v_0 v_1 \cdots = \text{pth}_{sw}(\rho)$ for some visitable path ρ . Let $\theta \in \Theta$ such that $\rho \in \theta$, then $\theta \in \text{Reach}^\lambda(v)$.

Clearly, $\text{Reach}(\bullet)$ is a monotonic operation over the set of vertices. Therefore, for all vertices v , there exists an ordinal λ such that $\text{Reach}^\lambda(v) = \text{Reach}^{\lambda+1}(v)$. Let λ_0 be the smallest ordinal such that for all v , $\text{Reach}^{\lambda_0+1}(v) = \text{Reach}^{\lambda_0}(v)$.

Definition 8.6.8. A well-formed μMLL^∞ proof-structure is said to be **DR-correct** if for any switching sw , in the orthogonal graph G_{sw} the following holds:

- $\forall v, v', v' \in \text{Reach}^{\lambda_0}(v)$
- $\forall v, v \notin \text{Reach}^{\lambda_0}(v)$

Note that productivity of the sequentialisation procedure is guaranteed if it is guaranteed that at each step of the sequentialisation one either finds a \wp or a fixed point formula or a splitting \otimes . This is ensured because by lock-freeness there can be no infinite ascending chain in \ll . Furthermore, if one just wants to ensure productivity and *not* fairness of sequentialisation, lock-freeness can be relaxed to the ascending chain condition. Interestingly, in this case, lock-freeness becomes moot for simple proof-structures.

Proposition 8.6.2. *A simple DR-correct proof-structure $\mathcal{R} = (\Gamma, \mathfrak{R}, \Theta_f, \Theta_{inf})$ has no infinite ascending chain in the \ll ordering.*

Proof. Suppose there is an infinite ascending chain in the \ll ordering. Then, by Theorem 8.3.2, $\text{DepGrph}(\mathcal{R}) = (V, E)$ has a ray $\beta = v_0 v_1 v_2 \dots$; furthermore, β has a infinite subword $v_{i_0} v_{i_1} v_{i_2} \dots$ for some infinite index set $\{i_j\}_{j \in \omega}$ such that v_{i_j} is tensor formula occurrence for all $j \in \omega$. Fix a switching sw and let $\mathcal{G}_{\mathcal{R}}^{sw}$ be the corresponding orthogonal graph. Let $v_{i_0} = (F, u)$ for some $F^U \in \Gamma$ and $u \in \bar{U}$. Then, u belongs to some \bar{v}^0 for $v^0 \in [\text{SW}]$. Similarly, $v_{i_1} = (G, u')$ and $u' \in \bar{v}^1$ for $v^1 \in [\text{SW}]$. By construction, v^0 and v^1 are distinct. Since \mathcal{R} is DR-correct, $\mathcal{G}_{\mathcal{R}}^{sw}$ is connected. Therefore, v^0 is connected to v^1 through at least one axiom θ_0 . Similarly, v^1 is connected to v^2 through at least one axiom θ_1 . We have that $\theta_0 \neq \theta_1$ because this contradicts the acyclicity of $\mathcal{G}_{\mathcal{R}}^{sw}$. Continuing like this, we have a ray in $\mathcal{G}_{\mathcal{R}}^{sw}$. By Proposition 8.6.1, we have that \mathcal{R} has a visitable path. Contradiction! \square

Finally, we give an example of a μMLL^∞ proof-structure (*cf.* Figure 8.2) with visitable paths that is DR-correct but has an infinite ascending chain in the \ll ordering. The kingdoms of the nodes p_1 and t_2 , $\hat{\wp}(p_1)$ and $\hat{\wp}(t_2)$, are shaded in cyan and magenta areas respectively. Observe that $t_1 \ll p_1 \ll t_2 \ll p_2 \ll \dots$

In conclusion. In this chapter, we generalised the finite proof-nets of μMLL^* to non-wellfounded objects. The correctness condition of the finite situation does not necessarily hold any more for de-sequentialisations of (pre-)proofs; on the other hand, we had to devise new conditions to guarantee sequentialisation. We developed this theory for proofs with trips and then concluded with pointer to generalisations to general proofs.

Chapter 9

The dynamics of infinets

(Revenge of the sequent)

In this chapter, we study the dynamics of infinets. We studied the dynamics of μMLL^* proof-nets in Section 7.2. Since inference rules are the same for μMLL^* and μMLL^∞ , the reduction rules for infinets are a superset of the reduction rules μMLL^* proof-nets. Fortunately, we developed them in the algebraic presentation, which will make it easier to lift them to the infinitary setting. We digress to remind the reader that we also developed the dynamics of $\mu\text{MLL}^{\text{ind}}$ proof-nets in Section 7.3.

Coming back to this chapter, in Section 9.1, we treat infinets as a metric rewriting system. We first guess the normal form (*big step*) and then show that a infinite reduction sequence of small steps converges to the big-step in the limit. To guess the limit, one has to sacrifice some structure *viz.* η -expand all axioms rendering the calculus without atoms. In Section 9.1, we treat simple infinets in full generality. However, our proof is not completely independent of the sequent calculus. In fact, to obtain limits of infinite reduction sequences, we go via a cut-elimination result in sequent calculus that we prove in Section 9.2.

9.1 Cut-elimination in infinets without axioms

In this section, we provide a cut elimination result for simple infinets with no finite axioms and no atoms. Note that this restriction is also crucially used to prove cut-elimination in [BDS16].

Definition 9.1.1. A simple μMLL^∞ proof-structure $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta_f, \Theta_{inf})$ is said to be **axiom-free** if $\Theta_f = \emptyset$ and for all $\theta \in \Theta_{inf}$, $\theta \subset \{l, r, i\}^\omega$.

From here on, we will write axiom-free infinets with just three components $(\Gamma, \mathfrak{K}, \Theta_{inf})$, the set of partial syntax trees, the set of cuts, and the set infinite axioms respectively.

Proposition 9.1.1. Let $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta_{inf})$ be an axiom-free infinet. For all $F^U \in \Gamma$, for all partial syntax trees U' of F , we have that $\overline{U'} \subseteq \overline{U}$.

Proof. Let $u \in \overline{U'}$. Let u' be the largest prefix of u such that $u' \in \overline{U}$. If $u' = u$, then we are done. Otherwise, u' is a strict (and hence, finite) prefix of u . Since \mathcal{R} is axiom-free, $U \subseteq \{l, r, i\}^\omega$ *i.e.* it does not contain finite words. Therefore, there exists $u'' \sqsupseteq u$ such that $u'' \in U$. Since u' is largest prefix of u in \overline{U} , we have that $u'' \cap u = u'$. Consequently $u = u'av$ and $u'' = u'bv'$ for distinct $a, b \in \{l, r\}$ and $v, v' \in \{l, r, i\}^\omega$. But since U is a partial syntax tree, we have $u'b \in \overline{U}$ implies that $u'a \in \overline{U}$. This contradicts the maximality of u' . \square

As a consequence, for all $\{C, C^\perp\} \in \mathfrak{K}$, we have that $u \in U$ iff $u^\perp \in U'$ where U, U' are the partial syntax trees of C and C^\perp respectively. For two such pairs of orthogonal addresses (u, u^\perp) , define their corresponding **orthogonal axioms** to be (θ, θ') such that $\theta, \theta' \in \Theta_{inf}$, $\alpha u \in \theta$, and $\alpha^\perp u^\perp \in \theta'$ where $\alpha = \text{addr}(C)$.

Definition 9.1.2. Given an axiom-free infinet \mathcal{R} the **cut-connection graph** $G_{\mathcal{R}}$ is defined as the undirected graph (Θ, E) where E is the set of all pairs of orthogonal axioms.

We will prove cut-elimination on axiom-free simple infinets. Imagine we are given a (possibly infinite) cut-reduction sequence. Our game plan is as follows. We will guess the normal form and then

show that the sequence converges to this normal form for some notion of convergence. Essentially, we define a big-step semantics [Plo81, Kah87] and show that the small-step reduction converges to the final result obtained from the big-step.


Definition 9.1.3. Let $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta_{inf})$ be an axiom-free μMLL^∞ proof-structure. The **normal form** of \mathcal{R} , denoted $\langle\!\langle\mathcal{R}\rangle\!\rangle$, is defined as the axiom-free μMLL^∞ proof-structure $(\Gamma', \emptyset, \Theta'_{inf})$ where,

- $\Gamma' = \{F^U \in \Theta \mid \{F, F^\perp\} \notin \mathfrak{K}\}$
- For each connected component \mathcal{C} of the cut-connected graph $G_{\mathcal{R}}$, $\theta_{\mathcal{C}} \in \Theta'_{inf}$ where $\theta_{\mathcal{C}}$ is defined as follows.

$$\theta_{\mathcal{C}} = \left(\bigcup_{\theta \in \mathcal{C}} \theta \right) \cap \left(\bigcup_{F^U \in \Gamma'} \text{addr}(F).U \right)$$

Readers familiar with game semantics will notice the similarities between our definition of the normal form and the composition of strategies [AJ92]. The union over cut-connected axioms is the *parallel composition*, the deletion of sub-occurrences of cut occurrences is the *hiding*. This is not surprising as indeed cuts are computationally compositions; however, any deeper connections need further investigation.

Example 9.1.1. Let $F = \nu x.x_\alpha$, $G = (\mu x.x \otimes \nu x.x)_\beta$, and $H = \mu x.x_\gamma$. Let $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta)$ such that $\Gamma = \{F^{i^\omega}, G^{(l+r)^{i^\omega}}, G^{\perp(l+r)^{i^\omega}}, H^{i^\omega}, H^{\perp i^\omega}\}$, $\mathfrak{K} = \{\{G, G^\perp\}, \{H, H^\perp\}\}$, and $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$ where $\theta_1 = \{\alpha^{i^\omega}, \beta l^{i^\omega}\}$, $\theta_2 = \{\beta r^{i^\omega}, \gamma^{i^\omega}\}$, $\theta_3 = \{\gamma^{\perp i^\omega}\}$, and $\theta_4 = \{\beta^\perp l^{i^\omega}, \beta^\perp r^{i^\omega}\}$. We have $G_{\mathcal{R}}$ is as follows.



Finally, $\langle\!\langle\mathcal{R}\rangle\!\rangle = \{\{F^{i^\omega}\}, \emptyset, \{\{\alpha^{i^\omega}\}\}\}$.

Lemma 9.1.1. Let \mathcal{R} be a progressing axiom-free inffinet. Then, $\langle\!\langle\mathcal{R}\rangle\!\rangle$ is also a progressing axiom-free inffinet.

Proof. We will first show that $\langle\!\langle\mathcal{R}\rangle\!\rangle$ is indeed a μMLL^∞ proof-structure. Fix a connected component \mathcal{C} of $G_{\mathcal{R}}$. We need to show, $\theta_{\mathcal{C}} \in \Theta'_{inf}$ is non-empty. Suppose not. Then, for all $\theta \in \mathcal{C}$, θ does not contain any word from a non-cut partial syntax tree *i.e.* $\theta \cap (\bigcup_{F^U \in \Gamma'} \text{addr}(F).U) = \emptyset$. Now recall that \mathcal{R} is progressing, therefore for all $\theta \in \Theta_{inf}$, there exists $w \in \theta$ such that w is progressing. Choose arbitrary $\theta_0 \in \mathcal{C}$. Following the discussion above, we infer there exists a progressing $w \in \theta_0$ such that w comes from a cut *i.e.* $w = \alpha u$ for some $\{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}$. Subsequently, there exists $\theta_1 \in \mathcal{C}$ such that $w^\perp \in \theta_1$. Now, since w is progressing, w^\perp is *not*. Therefore, there exists $w' = w^\perp \in \theta_1$ that is progressing. If we cannot continue like this *ad infinitum* then we have contradiction otherwise we have a sequence of axioms $\theta_0, \theta_1, \dots$ such that (θ_i, θ_{i+1}) is an edge in $G_{\mathcal{R}}$. Let (w_i, w_i^\perp) be the pair of orthogonal words that witnesses the orthogonality of (θ_i, θ_{i+1}) . Now, consider a switching sw such that for all i , w_i and w_i^\perp are in the same vertex in the switching component. There are two cases now.

Case 1. Every θ_i is distinct. Then, $\mathcal{G}_{\mathcal{R}}^{sw}$ contains an infinite path $\theta_0 v_0 \theta_1 v_1 \theta_2 \dots$ where v_i is the vertex in the switching component containing w_i and w_{i+1} . This contradicts the fact that \mathcal{R} is simple.

Case 2. There exists $i < j$ such that $\theta_i = \theta_j$. Then, $\mathcal{G}_{\mathcal{R}}^{sw}$ contains a cycle $\theta_i v_i \theta_{i+1} v_{i+1} \dots \theta_j$ where v_k is the vertex in the switching component containing w_k and w_{k+1} for $k \in \{i, i+1, \dots, j-1\}$. This contradicts the correctness of \mathcal{R} .

Now, we will show that $\langle\!\langle\mathcal{R}\rangle\!\rangle$ is correct. Fix a switching sw of $\langle\!\langle\mathcal{R}\rangle\!\rangle$. Note that sw can be extended to a switching sw' of \mathcal{R} such that they coincide on switching suboccurrences of non-cut doors. Let $\mathcal{G}_{\langle\!\langle\mathcal{R}\rangle\!\rangle}^{sw}$ and $\mathcal{G}_{\mathcal{R}}^{sw'}$ be the corresponding orthogonal graphs. Let $v_1 \theta_{\mathcal{C}_1} v_1 \theta_{\mathcal{C}_1} v_3 \dots v_n$ be a path in $\mathcal{G}_{\langle\!\langle\mathcal{R}\rangle\!\rangle}^{sw}$ such that for all $i \in [n-1]$, \mathcal{C}_i is a connected component of $G_{\mathcal{R}}$. For all $i \in [n]$, there exists $\theta_i, \theta'_i \in \mathcal{C}_i$ such that

$v_i \cap \theta_i = v_i \cap \theta_{C_i}$ and $v_{i+1} \cap \theta'_i = v_{i+1} \cap \theta_{C_i}$. Since $\mathcal{G}_{\mathcal{R}}^{sw'}$ is connected, there is a simple path between θ_i, θ'_i for all $i \in [n-1]$. This induces a path between v and v in $\mathcal{G}_{\mathcal{R}}^{sw'}$. One can extract a simple path from there, which contradicts the acyclicity of $\mathcal{G}_{\mathcal{R}}^{sw'}$.

Let $v_1 \theta_1 v_2 \theta_2 \dots v_n$ be a path in $\mathcal{G}_{\mathcal{R}}^{sw'}$. For all $i \in [n-1]$, $\theta_i \in C_i$ for some connected component C_i of $\mathcal{G}_{\mathcal{R}}$. Therefore, there is a path $v_1 \theta_{C_1} v_1 \theta_{C_1} v_3 \dots v_n$ be a path in $\mathcal{G}_{(\mathcal{R})}^{sw}$. Similarly as above, one can extract a simple path from there. Therefore, connectedness of $\mathcal{G}_{(\mathcal{R})}^{sw}$ follows from connectedness of $\mathcal{G}_{\mathcal{R}}^{sw'}$.

Finally, we are left to show that (\mathcal{R}) is lock-free. We will first show that $\text{DepGrph}((\mathcal{R})) = (V, E)$ is a subgraph of $\text{DepGrph}(\mathcal{R}) = (V', E')$. Clearly, $V \subseteq V'$ and if $E(F, G)$ such that F is a suboccurrence of G , then $E'(F, G)$. Assume $E(A \otimes B, C \wp D)$ such that $C \wp D$ is t -connected to $A \otimes B$. We need to show that $E'(A \otimes B, C \wp D)$. Wlog, assume A is connected to C and B to D .

We will induct on the length of switching sequences. Suppose A is connected to C by (γ, δ) , a switching sequence of length 1. If the axiom connecting γ and δ is present in \mathcal{R} , we are done. Otherwise, it is an axiom of the form θ_C for some connected component of $\mathcal{G}_{\mathcal{R}}$. There exists θ, θ' in C containing γ and δ respectively (technically, they contain words whose suffixes are γ and δ respectively). Since θ and θ' are connected in $\mathcal{G}_{\mathcal{R}}$ and since by Proposition 8.6.1 there cannot be an infinite path θ in any orthogonal graph of \mathcal{R} , we infer that there is a switching sequence connecting γ and δ in \mathcal{R} (consequently connecting A and C). The reasoning for B and D is symmetric. Therefore, $E'(F, G)$.

For the induction case, assume the length of the switching sequence is $n+1$. Then there exists a formula C' such that A is connected to C' by a switching sequence of length n , and C' is connected to C by a switching sequence of length 1. Arguing as before, C' is connected to C in \mathcal{R} and by hypothesis, A is connected to C' in \mathcal{R} . Therefore, A is connected to C in \mathcal{R} . Following a symmetric argument for B and D , we have $E'(F, G)$.

Now, by Theorem 8.3.2, if (\mathcal{R}) is not lock-free, then $\text{DepGrph}((\mathcal{R}))$ either has a ray or a vertex with infinite degree. Since $\text{DepGrph}((\mathcal{R}))$ is a subgraph of $\text{DepGrph}(\mathcal{R})$, this means \mathcal{R} is not lock-free. Hence done. \square

Definition 9.1.4. *The cut elimination relation $\rightarrow_{\mu\text{MLL}^\infty}$ is the binary relation over μMLL^∞ proof structures generated by the following rules.*

- $(\Gamma \cup \{\varphi \wp \psi_\alpha^U, \psi^\perp \otimes \varphi_{\alpha^\perp}^{U'}\}, \mathfrak{R} \cup \{\{\varphi \wp \psi_\alpha, \psi^\perp \otimes \varphi_{\alpha^\perp}\}\}, \Theta_{inf}) \rightarrow_{\{\varphi \wp \psi_\alpha, \psi^\perp \otimes \varphi_{\alpha^\perp}\}}$
- $(\Gamma \cup \{\varphi_{\alpha l}^{l^{-1}U}, \psi_{\alpha r}^{r^{-1}U}, \varphi_{\alpha^\perp r}^{r^{-1}U'}, \psi_{\alpha^\perp l}^{l^{-1}U'}\}, \mathfrak{R} \cup \{\{\varphi_{\alpha l}, \varphi_{\alpha^\perp r}\}, \{\psi_{\alpha r}, \psi_{\alpha^\perp l}\}\}, \Theta_{inf})$
- $(\Gamma \cup \{\mu x. \varphi_\alpha^U, \nu x. \varphi_{\alpha^\perp}^{U'}\}, \mathfrak{R} \cup \{\{\mu x. \varphi_\alpha, \nu x. \varphi_{\alpha^\perp}\}\}, \Theta_{inf}) \rightarrow_{\mu\text{MLL}^\infty}$
- $(\Gamma \cup \{\varphi[\mu x. \varphi/x]_{\alpha i}^{i^{-1}U}, \varphi^\perp[\nu x. \varphi^\perp/x]_{\alpha^\perp i}^{i^{-1}U'}\}, \mathfrak{R} \cup \{\{\varphi[\mu x. \varphi/x]_{\alpha i}, \varphi^\perp[\nu x. \varphi^\perp/x]_{\alpha^\perp i}\}\}, \Theta_{inf})$

By simple inspection of the rules, we have following proposition.

Proposition 9.1.2. *Let $\mathcal{R}, \mathcal{R}'$ be μMLL^∞ proof-structures such that $\mathcal{R} \rightarrow_{\mu\text{MLL}^\infty} \mathcal{R}'$. If \mathcal{R} is axiom-free, then so is \mathcal{R}' . Furthermore, if \mathcal{R} is correct, so is \mathcal{R}' .*

Note that the elimination relation defines a rewriting system on axiom-free μMLL^∞ proof-structures. Cut reduction sequences in such infinitary settings could potentially be infinite.

Definition 9.1.5. *A sequence of infinets, $\{\mathcal{R}_i\}_{i \in \omega}$, is called a **reduction sequence** if for every $i \in \omega$, $\mathcal{R}_i \rightarrow_{\mu\text{MLL}^\infty} \mathcal{R}_{i+1}$ by the cut reduction rules in definition 9.1.4. A reduction sequence is said to be **fair** if for every $i \in \omega$ and $\{C, C^\perp\}$ such that $\mathcal{R}_i \rightarrow_{\{C, C^\perp\}} \mathcal{R}'$, there is some $j \geq i$, such that $\{C, C^\perp\}$ cannot be reduced in \mathcal{R}_j , i.e. there is no infinet \mathcal{R}'' such that $\mathcal{R}_j \rightarrow_{\{C, C^\perp\}} \mathcal{R}''$.*

If a transfinite cut reduction sequence is convergent then it is strongly convergent. By the compression lemma, this implies that there is a cut reduction sequence bounded by ω that converges to the same limit. Fairness is a sufficient condition for convergence within ω . Note that, if there are finitely many cuts then fairness can be relaxed to the following condition known as **weak fairness**.

Definition 9.1.6. *A reduction sequence $\{\mathcal{R}_i\}_{i \in \omega}$, is said to be **weakly fair** if for every cut $\{C, C^\perp\}$ in \mathcal{R}_i , there is a $j > i$ such that $\mathcal{R}_j \rightarrow_{\{C', C'^\perp\}} \mathcal{R}_{j+1}$ for some suboccurrence C' of C .*

It is easy to see that fairness implies weak fairness. We will show that the cut elimination relation defines a metric rewriting system and we will show that a fair cut reduction sequence from an infinit \mathcal{R} strongly converges to (\mathcal{R}) . In order to do that we need to construct a metric on axiom-free infinits.

In order to define the distance function d between two infinits $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta_{inf})$ and $\mathcal{R}' = (\Gamma', \mathfrak{K}', \Theta'_{inf})$, we individually define distance functions (denoted by d_1 , d_2 , and d_3 respectively) on each of its components. The product metric is defined as the L_1 norm. To define the distance functions d_1 , d_2 , and d_3 we need to define the symmetric difference of two sets. The **symmetric difference** of two sets A and B , denoted by $A\Delta B$, is defined as $A\Delta B = (A \setminus B) \cup (B \setminus A)$. We state a few basic properties of symmetric difference.

- $A\Delta B = B\Delta A$
- $A\Delta B = \emptyset$ iff $A = B$.
- $A\Delta\emptyset = A = \emptyset\Delta A$.

Finally, we are ready to define the distance between sets of partial syntax trees, cuts, and axioms respectively.

$$d_1(\Gamma, \Gamma') = \begin{cases} 0 & \text{if } \Gamma = \Gamma'; \\ (\min(\{|\alpha| \mid \varphi_\alpha^U \in \Gamma\Delta\Gamma'\}))^{-1} & \text{otherwise.} \end{cases}$$

$$d_2(\mathfrak{K}, \mathfrak{K}') = \begin{cases} 0 & \text{if } \mathfrak{K} = \mathfrak{K}'; \\ (\min(\{|\alpha| \mid \{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}\Delta\mathfrak{K}'\}))^{-1} & \text{otherwise.} \end{cases}$$

$$d_3(\Theta_{inf}, \Theta'_{inf}) = \begin{cases} 0 & \text{if } \Theta_{inf} = \Theta'_{inf}; \\ (\min(\{|\alpha| \mid \alpha u \in (\bigcup_{\theta \in \Theta_{inf}} \theta) \Delta (\bigcup_{\theta' \in \Theta'_{inf}} \theta')\}))^{-1} & \text{otherwise.} \end{cases}$$

Remark 9.1.1. *There is a bit of hand-waving in the definition of d_3 . It has not been mentioned how one is supposed to determine α from any word w in $(\bigcup_{\theta \in \Theta_{inf}} \theta) \Delta (\bigcup_{\theta' \in \Theta'_{inf}} \theta')$. Technically, we extend the alphabet with the letter $\#$ and ensure that words in this set are of the form $\alpha\#u$ i.e. every (possibly infinite) word has exactly one finite prefix ending at $\#$. Equality on such words are defined as $\alpha\#u = \alpha'\#u'$ if $\alpha u = \alpha' u'$.*

Lemma 9.1.2. *The set of all sets of partial syntax trees (respectively, of cuts, and axioms) equipped with d_1 (respectively d_2 and d_3) is a metric space.*

Proof. Fix arbitrary sets of partial syntax trees Γ, Γ' , and Γ'' . First note that $d_1(\Gamma, \Gamma') \geq 0$. Symmetry follows from the symmetry of Δ .

Identity of indiscernibles. By definition, $d(\Gamma, \Gamma) = 0$. Also note $d_1(\Gamma, \Gamma') = 0$ is only possible if $\Gamma = \Gamma'$ (in all other case it is a real number in $(0, 1]$).

Triangle inequality. We will in fact, prove a stronger statement. We will show that d_1 is an *ultrametric* i.e. $\max(d_1(\Gamma, \Gamma'), d_1(\Gamma', \Gamma'')) \geq d_1(\Gamma, \Gamma'')$. If $\Gamma = \Gamma'$, then $\max(d(\Gamma, \Gamma'), d(\Gamma', \Gamma'')) = d(\Gamma', \Gamma'') = d(\Gamma, \Gamma'')$. Hence done. Similarly, if $\Gamma' = \Gamma''$. If $\Gamma = \Gamma''$, then the inequality is trivially true. Therefore, assume Γ, Γ' and Γ'' are distinct. So, we need to show that

$$\max\left(\frac{1}{f(\Gamma, \Gamma')}, \frac{1}{f(\Gamma', \Gamma'')}\right) \geq \frac{1}{f(\Gamma, \Gamma'')}$$

Wlog assume $\max\left(\frac{1}{f(\Gamma, \Gamma')}, \frac{1}{f(\Gamma', \Gamma'')}\right) = \frac{1}{f(\Gamma, \Gamma')} = \frac{1}{|\alpha|}$. Therefore,

$$f(\Gamma, \Gamma') \leq f(\Gamma', \Gamma'') \tag{9.1}$$

From the definition of d , we have that $|\alpha| \leq |\beta|$ for all $\varphi_\beta^U \in \Gamma\Delta\Gamma'$. From Equation (9.1), $|\alpha| \leq |\beta|$ for all $\varphi_\beta^U \in \Gamma\Delta\Gamma'$. Combining the two, we have $|\alpha| \leq |\beta|$ for all $\varphi_\beta^U \in (\Gamma \cup \Gamma' \cup \Gamma'') \setminus (\Gamma \cap \Gamma' \cap \Gamma'')$. Therefore, $|\alpha| \leq f(\Gamma, \Gamma'')$. Hence $\frac{1}{|\alpha|} \geq \frac{1}{f(\Gamma, \Gamma'')}$. The proofs for d_2 and d_3 are similar. \square

Finally, define $d(\mathcal{R}, \mathcal{R}') = d(\Gamma, \Gamma') + d(\mathfrak{K}, \mathfrak{K}') + d(\Theta_{inf}, \Theta'_{inf})$.

Theorem 9.1.1. *The set of all axiom-free μMLL^∞ proof-structures equipped with d is a metric space.*

Proof. Since L_p norm defines a product metric for any p , this follows immediately from Lemma 9.1.2. \square

Theorem 9.1.2. *Let $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \omega}$ be a fair reduction sequence such that \mathcal{R}_0 is valid. Then, \mathcal{S} strongly converges to (\mathcal{R}_0) .*

Proof. Weak fairness (which is implied by fairness) already ensures that cuts of larger and larger addresses are reduced. This ensures that a weak converging sequence is indeed strongly converging. In order to prove weak convergence, we will prove the following two claims:

1. $d(\mathcal{R}_{i+1}, (\mathcal{R}_0)) \leq d(\mathcal{R}_i, (\mathcal{R}_0))$;
2. for all $\varepsilon > 0$, there exists N such that $d(\mathcal{R}_N, (\mathcal{R}_0)) < \varepsilon$.

Note that by Lemma 9.1.1, (\mathcal{R}_0) is indeed a valid axiom-free infinnet. In particular, this ensures that the distance function can be applied meaningfully above. Combining these two claims, we have that for all $\varepsilon > 0$, there exists N such that for all $i \geq N$, $d(\mathcal{R}_i, (\mathcal{R}_0)) < \varepsilon$ proving that \mathcal{S} weakly converges to (\mathcal{R}_0) . Let $\mathcal{R}_i = (\Gamma_i, \mathfrak{K}_i, \Theta_i)$ and $(\mathcal{R}_0) = (\Gamma, \emptyset, \Theta)$. Since no cuts are removed in finitely many steps, for all i , $\Gamma_i \neq \Gamma$, $\mathfrak{K}_i \neq \emptyset$, and $\Theta_i \neq \Theta$. Moreover, the only difference between \mathcal{R}_i and (\mathcal{R}_0) is in the cut occurrences, hence $d_1(\Gamma_i, \Gamma) = d_2(\mathfrak{K}_i, \emptyset) = d_3(\Theta_i, \Theta)$.

Proof of claim 1. It suffices to show that $d_2(\mathfrak{K}_{i+1}, \emptyset) \leq d_2(\mathfrak{K}_i, \emptyset)$. In other words, we need to show that $\min(\{|\alpha| \mid \{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}_i\}) \leq \min(\{|\alpha| \mid \{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}_{i+1}\})$ since $A\Delta\emptyset = A$. Noting that the cut reduction rules in Definition 9.1.4 increases the size of the addresses of cut formulas, this inequality holds.


Proof of claim 2. Let $n = \lfloor \frac{1}{\varepsilon} \rfloor - \min(\{|\alpha| \mid \{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}_0\})$. If n is negative then choose $N = 0$. Otherwise, we will compute N by induction on n . The base case is $n = 0$. Note that there are only finitely many distinct addresses of length $|\alpha|$, therefore there are finitely many cuts whose addresses are of size $|\alpha|$. Let $N \in \omega$ be the least index such that \mathcal{R}_N has no cuts of $|\alpha|$. We have $d(\mathcal{R}_N, (\mathcal{R}_0)) < \varepsilon$. By fairness, N exists. For the induction case assume $n = m + 1$. By a similar reasoning as above, we have that there exists N such that $\lfloor \frac{1}{\varepsilon} \rfloor - \min(\{|\alpha| \mid \{\varphi_\alpha, \varphi_{\alpha^\perp}\} \in \mathfrak{K}_N\}) = m$. By IH, there exists, N' such that $d(\mathcal{R}_{N'}, (\mathcal{R}_0)) < \varepsilon$. Hence done. \square

Our reduction sequences are peculiar. We neither allow finite reduction sequences nor transfinite ones. Therefore, in particular, standard notions of confluence cannot be imported. However, we note that, even without the fairness condition, we can show that reduction sequences are convergent (one needs to generalise the notion of norm in order to do this). Consequently, we have the following notion of **confluence**.

Corollary 9.1.2.1. *Let \mathcal{S} and \mathcal{S}' be two reduction sequences starting from a valid axiom-free infinnet \mathcal{R} such that \mathcal{R} and \mathcal{S}' converges to \mathcal{R}_1 and \mathcal{R}_2 respectively. Then, all fair reduction sequences starting from \mathcal{R}_1 and \mathcal{R}_2 converge to (\mathcal{R}) .*

Infinitary η -expansion

The notion of axiom expansion in proofs can be lifted to proof-nets in MLL by way of the following graph rewrite rules.



In the rest of this section, we will show that the η -expansion of MLL nets can be lifted to some special types of infinets. This will serve two purposes: firstly, it will exhibit the utility of the metric we define on infinets outside of the proof of Theorem 9.1.2, secondly, it will help us show that the axiom-free restriction is not *ad-hoc* as it may appear.

Definition 9.1.7. An *infinet* $(\Gamma, \mathfrak{K}, \Theta_f, \Theta_{inf})$ is said to be **non-stationary** if for all $\theta \in \Theta_{inf}$, $\theta \subset \{l, r, i\}^\omega$. An *infinet* which is not non-stationary is called a **halting infinet**.

The reason behind the nomenclature is that in a non-stationary infinet, no occurrence is stationary in an infinite axiom. We will define infinitary η -expansion for non-stationary infinets. Since we work with the algebraic presentation of proof-nets, we will first provide the η -expansion rules in this presentation.

Definition 9.1.8. The **axiom expansion relation** \longrightarrow_η is a binary relation over infinets generated by the following rules.

1. $(\Gamma \cup \{A^U, B^V\}, \mathfrak{K}, \Theta \cup \{\{\alpha u, \beta v\}\}) \longrightarrow_\eta (\Gamma \cup \{A^{U'}, B^{V'}\}, \mathfrak{K}, \Theta \cup \{\{\alpha ul, \beta vr\}, \{\alpha ur, \beta vl\}\})$
 where $u \in U, v \in V, (A, u) = F \otimes G, (B, u') = G^\perp \wp H^\perp, U' = (U \setminus \{u\}) \cup \{ul, ur\}$ and $V' = (V \setminus \{v\}) \cup \{vl, vr\}$.
2. $(\Gamma \cup \{A^U, B^V\}, \mathfrak{K}, \Theta \cup \{\{\alpha u, \beta v\}\}) \longrightarrow_\eta (\Gamma \cup \{A^{U'}, B^{V'}\}, \mathfrak{K}, \Theta \cup \{\{\alpha ui, \beta vi\}\})$
 where $u \in U, v \in V, (A, u) = \mu x.F, (B, u') = \nu x.F^\perp, U' = (U \setminus \{u\}) \cup \{ui\}$ and $V' = (V \setminus \{v\}) \cup \{vi\}$.

If an infinet is halting, η -expansion can be crucially non-confluent. Consider an infinite axiom which has two stationary formulas $F \otimes G$ and $G^\perp \wp H^\perp$. Then, there are several choices for the next step of η -expansion. If we are not careful, we can break the DR-correctness (it is reasonable to preserve correctness by η -expansion) as in Figures 9.1a and 9.1b. If we are more careful, we do have infinets, but they are crucially different *i.e.* they cannot be confluent under η -expansion as in Figures 9.1c and 9.1d.

An **infinitary η -expansion sequence** $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \omega}$ is a sequence of infinets such that for all $i \in \omega, \mathcal{R}_i \longrightarrow_\eta \mathcal{R}_{i+1}$. A **fair η -expansion sequence** is the one for every finite axiom, an η -expansion rule is applied to it after a finite number of steps (analogously, in a fair cut reduction, for every cut, a cut-reduction rule is applied on it after finite time). As before, we will guess the limit of an infinitary η -expansion sequence.

Definition 9.1.9. Let $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta_f, \Theta_{inf})$ be a non-stationary infinet. Let $\{\alpha u, \beta u'\} \in \Theta_f$ such that $\lceil (F, u) \rceil = \varphi$ and $\lceil (G, u') \rceil = \varphi^\perp$ for some $F^U, G^{U'} \in \Gamma$. Let V, V^\perp be the syntax tree of φ and φ^\perp respectively. Let $\Theta_{\{\alpha u, \beta u'\}} = \{\{\alpha uv, \beta u'v^\perp\} \mid v \in U\}, U_{\{\alpha u, \beta u'\}} = \alpha uV$ and $U'_{\{\alpha u, \beta u'\}} = \beta u'V^\perp$. Define $\eta^\infty(\mathcal{R})$ as the infinet $(\Gamma', \mathfrak{K}, \Theta'_f, \Theta'_{inf})$ where $\Gamma' = \{F^{\eta^\infty(U)} \mid F^U \in \Gamma\}, \Theta'_f \cup \Theta'_{inf} = \Theta_{inf} \cup \bigcup_{\theta \in \Theta_f} \Theta_\theta$, and $\eta^\infty(U) = (U \setminus (\bigcup_{\theta \in \Theta_f} \theta)) \cup \bigcup_{\theta \in \Theta_f} U_\theta$.

Lemma 9.1.3. For all infinets $\mathcal{R}, \eta^\infty(\mathcal{R})$ is also an infinet.

Proof Sketch. For any $F^U, F^{\eta^\infty(U)}$ is indeed a partial syntax tree such that $\overline{U} \subseteq \overline{\eta^\infty(U)}$. Note that $\Theta'_f \cup \Theta'_{inf}$ is indeed a partition of the leaves of Γ . So basically we need to show that $\eta^\infty(\mathcal{R})$ is DR-correct and lock-free.

The basic idea is that every finite axiom θ in \mathcal{R} has been replaced by a proof-structure (say, \mathcal{R}_θ) that has exactly two doors: the occurrences in θ . Moreover, \mathcal{R}_θ is not arbitrary, it is maximal. Now that we have shown \mathcal{R}_θ is an infinet for all θ , the result follows from the fact that \mathcal{R} is an infinet. \square

Now, we need to show fair η -expansion sequences starting from \mathcal{R} strongly converge to $\eta^\infty(\mathcal{R})$. We will slightly modify the metric from before:

$$d'_1(\Gamma, \Gamma') = \begin{cases} 0 & \text{if } \Gamma = \Gamma'; \\ \left(\min(\{|\alpha u| \mid u \in U \Delta U', \varphi_\alpha^U \in \Gamma, \varphi_\alpha^{U'} \in \Gamma'\}) \right)^{-1} & \text{otherwise.} \end{cases}$$

$$d'_3(\Theta_{inf}, \Theta'_{inf}) = \begin{cases} 0 & \text{if } \Theta_{inf} = \Theta'_{inf}; \\ \left(\min \left(\left\{ |\alpha u| \mid \alpha u \in \left(\bigcup_{\theta \in \Theta_{inf}} \theta \right) \Delta \left(\bigcup_{\theta' \in \Theta'_{inf}} \theta' \right) \right\} \right) \right)^{-1} & \text{otherwise.} \end{cases}$$

Note that these distances are well-defined if the underlying set of doors of \mathcal{R} and \mathcal{R}' are identical. Verifying that they are indeed a metric is similar to the proof of Lemma 9.1.2 and we shall not recast it here. By abuse of notation, the metric on infinets, $\mathcal{R}, \mathcal{R}'$ defined as $d'_1(\Gamma, \Gamma') + d_2(\mathfrak{K}, \mathfrak{K}') + d'_3(\Theta, \Theta')$ is still referred to as $d(\mathcal{R}, \mathcal{R}')$.




Figure 9.1: Different choices for η -expanding an infinite axiom atop $F \otimes G, G^\perp \otimes H^\perp$ and an infinite thread of $\nu x.x$ unfolding.

Theorem 9.1.3. *Let $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \omega}$ be a fair η -expansion sequence. Then, \mathcal{S} strongly converges to $(\eta^\infty(\mathcal{R}_0))$.*

Proof. As in Theorem 9.1.2, fairness ensures that finite axioms on larger and larger addresses are reduced. This ensures that a weakly converging sequence is indeed strongly converging. In order to prove weak convergence, we will prove the following two claims:

1. $d(\mathcal{R}_{i+1}, \eta^\infty(\mathcal{R}_0)) \leq d(\mathcal{R}_i, \eta^\infty(\mathcal{R}_0))$;
2. for all $\varepsilon > 0$, there exists N such that $d(\mathcal{R}_N, \eta^\infty(\mathcal{R}_0)) < \varepsilon$.

By Lemma 9.1.3, $\eta^\infty(\mathcal{R}_0)$ is an infinet which ensures that the distance function can be applied meaningfully above. As in Theorem 9.1.2, combining these two claims, we have that \mathcal{S} weakly converges to $\eta^\infty(\mathcal{R}_0)$. Let $\mathcal{R}_i = (\Gamma_i, \mathfrak{K}_i, \Theta_i)$ and $\eta^\infty(\mathcal{R}_0) = (\Gamma, \emptyset, \Theta)$. We observe that η -expansion does not touch cuts, hence for all i, j , $\mathfrak{K}_i = \mathfrak{K}_j = \mathfrak{K}$. So, for all i , $d_2(\mathfrak{K}_i, \mathfrak{K}) = 0$. Since, \mathcal{R}_i and $\eta^\infty(\mathcal{R}_0)$ only differ in axioms, we have $d'_1(\Gamma_i, \Gamma) = d'_3(\Theta_i, \Theta)$. Furthermore, since no axioms are fully expanded in finitely many steps, $\Gamma_i \neq \Gamma$ and $\Theta_i \neq \Theta$.

Proof of claim 1. This follows from noting the η -expansion rules increase the size of the finite words in the partial syntax trees of the same occurrence.

Proof of claim 2. Let $n = \lfloor \frac{1}{\varepsilon} \rfloor - \min \left(\left\{ |\alpha u| \mid \alpha u \in \left(\bigcup_{\theta \in \Theta_{inf}} \theta \right) \Delta \left(\bigcup_{\theta' \in \Theta'_{inf}} \theta' \right) \right\} \right)$. If n is negative then choose $N = 0$. Otherwise, we will compute N by induction on n . The base case is $n = 0$. Noting that there are only finitely many distinct addresses of length $|\alpha u|$, therefore there are finitely many finite axioms whose addresses are of size $|\alpha u|$. Let $N \in \omega$ be the least index such that \mathcal{R}_N has no finite axioms of $|\alpha u|$. We have $d(\mathcal{R}_N, \eta^\infty(\mathcal{R}_0)) < \varepsilon$. By fairness, N exists. The induction case goes similarly. \square

We have thus established a notion of infinitary η -expansion on non-stationary infinets. Infinitary η -expansions have been studied in the context of Böhm trees [Nak75]. But how does this connect to axiom-free infinets? Imagine in LK, we fix a particular valuation of atoms: every positive atom is substituted by \top and every negative atom is substituted by \perp . Every proof of a formula is also a proof

under the substitution map. One can carry forward this intuition to linear logic and μMALL^∞ . Now recall from Chapter 5 that $\llbracket \mu x.x \rrbracket = \mathbf{0}$ and $\llbracket \nu x.x \rrbracket = \top$. Therefore, if every positive atom is substituted by $\nu x.x$ and every negative atom is substituted by $\mu x.x$, then the geometry of the proofs (and the proof-nets) does not change. Therefore, there is a trivial geometry-preserving map from infinets with atoms to infinets without axioms. Finally, note that the infinitary η -expansion of atom-free infinets is axiom-free.

9.2 μMALL^∞ cut-elimination revisited

While the multicut brings uniformity in the treatment of cut-elimination in sequent calculus, it is not well-suited for our purpose of developing a canonical and parallel treatment of cuts in non-wellfounded proof systems. It is indeed better suited to use the usual cut-rule to draw a comparison between cut-reductions in sequent systems and in proof-nets, as we will do in the next section of this chapter. To serve this purpose, we develop here an alternative approach to cut-elimination for non-wellfounded proof which avoids the use of the multicut but relies on the standard cut instead and we will prove a new cut-elimination result in this case. Note that our proof is not independent of Theorem 4.5.2.

In order to formally work with multicuts, we need to switch to occurrences from formulas. We denote the non-wellfounded system with multicuts over occurrences by μMALL_m^∞ . Formally the multicut rule comes with a function ι which shows how the occurrences of the conclusion are distributed over the premises (modulo renaming), and a relation $\perp\!\!\!\perp$ specifying which occurrences are cut-connected. We recall the formal definition from [BDS16, Dou17, BDKS22].

Definition 9.2.1. *Given sequents $\Gamma, \Gamma_1, \dots, \Gamma_n$ where $n > 0$ and such that Γ_i, Γ_j are disjoint for all $i \neq j$, a **multicut** of conclusion $\vdash \Gamma$ and premisses $(\Gamma_i)_{i \in [1:n]}$ is given by an injection $\iota : \Gamma \mapsto \cup_{i=1}^n \Gamma_i$ and a binary relation $\perp\!\!\!\perp \subseteq (\cup_{i=1}^n \Gamma_i)^2$ such that:*

- For all $F \in \Gamma$, $\lceil \iota(F) \rceil = \lceil F \rceil$.
- For all $F, G \in \cup_{i=1}^n \Gamma_i$, $F \perp\!\!\!\perp G$ implies $\lceil F \rceil = \lceil G \rceil$.
- If $F \in \Gamma_i$ and $G \in \Gamma_j$ such that $F \perp\!\!\!\perp G$ then $i \neq j$.
- $\text{dom}(\perp\!\!\!\perp) = (\cup_{i=1}^n \Gamma_i) \setminus \text{im}(\iota)$.
- Given two sequents Γ_i and Γ_j , we say that they are $\perp\!\!\!\perp$ -connected on a pair of formula occurrences (F, G) when $F \in \Gamma_i$ and $G \in \Gamma_j$ such that $F \perp\!\!\!\perp G$. We say that they are $\perp\!\!\!\perp$ -connected, and we write $\Gamma_i \perp\!\!\!\perp \Gamma_j$, when they are $\perp\!\!\!\perp$ -connected on some (F, G) . The relation $\perp\!\!\!\perp$ on sequents must satisfy two conditions:
 - two sequents must be $\perp\!\!\!\perp$ -connected on at most one pair of occurrences F, G ;
 - the graph of the relation $\perp\!\!\!\perp$ must be connected and acyclic.

We write this multicut rule as:

$$\frac{\vdash \Gamma_1 \quad \dots \quad \vdash \Gamma_n}{\vdash \Gamma} (\text{mcut}(\iota, \perp\!\!\!\perp))$$

If clear from the context, we omit to specify ι and $\perp\!\!\!\perp$ in the rule name.

In the rest of this section, we will detail the cut-elimination procedure for μMALL^∞ with the standard cut rule. In other words, we will avoid the use of the multicut rule. We shall simply retain however a degenerate case of the multi-cut *viz.* the *unary* case, used to lazily perform the cut-axiom reduction and relocate addresses. Indeed, as we work with explicit occurrences, the cut/id case is as follows:

$$\frac{\frac{}{\vdash F, G^\perp} (\text{id}) \quad \frac{\pi}{\vdash G, \Gamma}}{\vdash F, \Gamma} (\text{cut})$$

with $\lceil F \rceil = \lceil G \rceil$, which cannot simply be reduced to $\frac{\pi}{\vdash F, \Gamma}$ as the occurrences do not match (in fact, the addresses of F and G are disjoint). Instead of substituting occurrences in π (which is a non-wellfounded object), we treat this substitution lazily, in the form of an explicit substitution (cf. [ACCL91]) adding the following unary inference rule: $\frac{\vdash \Gamma'}{\vdash \Gamma} (\text{Loc}(\iota))$ where ι is a one-to-one map from Γ to Γ' such that for all $F \in \Gamma$, $\lceil \iota(F) \rceil = \lceil F \rceil$. In the rest of this section, when writing μMALL^∞ , we mean μMALL^∞ extended with the $\text{Loc}(\iota)$ rule.

Definition 9.2.2. *The cut elimination relation \longrightarrow_c is the binary relation over proofs generated by extending the key rules of MALL (cf. Figure 3.2, with the exception of the cut-axiom rule) with fixed point reductions and extending the commutation rules of MALL (cf. Figure 3.3) with commutation rules for fixed points and Loc (cf. Figure 9.2) and the following cut-axiom rule:*

$$\frac{\frac{}{\vdash A, B^\perp} \text{ (id)} \quad \frac{\pi}{\vdash B, \Gamma}}{\vdash A, \Gamma} \text{ (cut)} \quad \rightarrow \quad \frac{\pi}{\vdash B, \Gamma} \text{ (Loc}(\iota))$$

where $\iota(A) = B$, $\iota(H) = H$ for $H \in \Gamma$.

We use the following notations:

- \longrightarrow_m for the multicut reduction rules of [BDS16];
- $\longrightarrow_{\text{merge}}$ for the mcut merge reduction (defined in Section 4.5).
- $\longrightarrow_{\text{comm}}$ for the cut commutation reduction.

Definition 9.2.3. Let π be a μMALL_m^∞ proof. Define the **intermediary translation** of π , denoted $[\pi]$, as the proof where each unary mcut has been replaced by an appropriate Loc and each binary mcut has been replaced by a cut and (possibly) a Loc rule. Formally, in π every occurrence of

- $\frac{\vdash \Gamma'}{\vdash \Gamma} \text{ (mcut}(\iota, \emptyset))$ is replaced by $\frac{\vdash \Gamma'}{\vdash \Gamma} \text{ (Loc}(\iota))$
- $\frac{\vdash \Gamma_1 \quad \vdash \Gamma_2}{\vdash \Gamma} \text{ (mcut}(\text{id}, \perp))$ is replaced by $\frac{\vdash \Gamma_1 \quad \vdash \Gamma_2}{\vdash \Gamma} \text{ (cut)}$
- $\frac{\vdash \Gamma_1 \quad \vdash \Gamma_2}{\vdash \Gamma} \text{ (mcut}(\iota, \perp))$ is replaced by $\frac{\vdash \Gamma_1 \quad \vdash \Gamma_2}{\vdash \Gamma} \text{ (cut)}$ where $\Gamma' = \iota(\Gamma)$

Finally, define $\text{CSeq}(\pi) = \{[\pi'] \mid \exists \pi'. \forall \pi''. \pi'' \not\rightarrow_{\text{merge}} \pi' \rightarrow_{\text{merge}}^* \pi\}$.

For the rest of this section, fix a μMALL_m^∞ proof π_m . In general, $[\pi_m]$ is a proof in the hybrid system with both multicuts and Loc rules. However, elements of $\text{CSeq}(\pi_m)$ are all proofs in μMALL^∞ (i.e. without any multicuts).

Lemma 9.2.1. If $\pi \in \text{CSeq}(\pi_m)$, then $\pi \in \mu\text{MALL}^\infty$. Furthermore, every sequent in π_m also occurs in π .

Proof. By definition, $\pi = [\pi']$ such that for all proofs $\pi'' \in \mu\text{MALL}_m^\infty$, $\pi'' \not\rightarrow_{\text{merge}} \pi' \rightarrow_{\text{merge}}^* \pi_m$. In other words, there does not exist any π'' such that $\pi'' \rightarrow_{\text{merge}} \pi'$. Therefore, π' only has unary or binary (mcut) inferences. Since $[\bullet]$ transforms unary or binary (mcut) inferences into (cut) and (Loc) inferences, $\pi \in \mu\text{MALL}^\infty$. The subsequent claim is trivial. Note that for any two proofs π_0, π_1 such that $\pi_0 \rightarrow_{\text{merge}} \pi_1$, then every sequent in π_1 occurs in π_0 (the converse doesn't necessarily hold). Since $[\bullet]$ does not change sequents (rather it simply rearranges sequents with possibly different inferences), we conclude. \square

Lemma 9.2.2. If π_m contains an instance of the mcut rule with premisses $\vdash \Gamma_1$ and $\vdash \Gamma_2$ such that $\Gamma_1 \perp \Gamma_2$, then there exists some $\pi \in \text{CSeq}(\pi_m)$ containing an instance of the (cut) rule with Γ_1 and Γ_2 as premisses.

$$\begin{array}{c}
\frac{\pi}{\frac{\frac{\frac{\pi}{\vdash \Gamma', F'[\eta x.F'/x]}{(\eta)}}{\vdash \Gamma', \eta x.F'}{(\text{Loc}(\iota))}}{\vdash \Gamma, \eta x.F'}} \quad \longrightarrow_r \\
\text{with } \iota \text{ such that } \iota(\eta x.F) = \eta x.F' \quad \text{and} \quad \iota(\Gamma) = \Gamma' \text{ and } \kappa = \begin{cases} G \mapsto \iota(G) & \text{if } G \in \Gamma \\ F[\eta x.F'/x] \mapsto F'[\eta x.F'/x] \end{cases}
\end{array}$$

$$\frac{\frac{\pi_1}{\vdash \Gamma', F'} \quad \frac{\pi_2}{\vdash \Delta', G'}}{\frac{\frac{\frac{\frac{\pi_1}{\vdash \Gamma', F'}{(\text{Loc}(\kappa))} \quad \frac{\frac{\pi_2}{\vdash \Delta', G'}}{(\text{Loc}(\lambda))}}{\vdash \Gamma, \Delta, F \otimes G}}{(\otimes)}}{\vdash \Gamma, \Delta, F \otimes G}} \quad \longrightarrow_r \\
\text{with } \kappa = \begin{cases} H \mapsto \iota(H) & \text{if } H \in \Gamma \\ F \mapsto F' \end{cases} \quad \text{and} \quad \lambda = \begin{cases} H \mapsto \iota(H) & \text{if } H \in \Delta \\ G \mapsto G' \end{cases}$$

$$\frac{\frac{\pi}{\frac{\frac{\frac{\pi}{\vdash \Gamma', F', G'}}{(\wp)}}{\vdash \Gamma', F' \wp G'}}{(\text{Loc}(\iota))}} \quad \longrightarrow_r \\
\text{with } \iota = \begin{cases} H \mapsto \iota(H) & \text{if } H \in \Gamma \\ F \mapsto F' \\ G \mapsto G' \end{cases}$$

$$\frac{\frac{\frac{\frac{\frac{\frac{\pi_1}{\vdash \Gamma', C}}{(\text{Loc}(\kappa))} \quad \frac{\frac{\frac{\pi_2}{\vdash \Delta', C^\perp}}{(\text{Loc}(\lambda))}}{\vdash \Gamma, C}}{(\text{cut})}}{\vdash \Gamma, \Delta}}{(\text{Loc}(\iota))}}{\vdash \Gamma, \Delta}} \quad \longrightarrow_r \\
\text{with } \kappa = \begin{cases} H \mapsto \iota(H) & \text{if } H \in \Gamma \\ C \mapsto C \end{cases} \quad \text{and} \quad \lambda = \begin{cases} H \mapsto \iota(H) & \text{if } H \in \Delta \\ C^\perp \mapsto C^\perp \end{cases}$$

$$\frac{\frac{\frac{\frac{\frac{\pi}{\vdash \Gamma''}}{(\text{Loc}(\kappa))}}{\vdash \Gamma'}{(\text{Loc}(\iota))}}{\vdash \Gamma}}{(\text{id})}}{\vdash F', G'} \quad \longrightarrow_r \quad \frac{\frac{\frac{\frac{\pi}{\vdash \Gamma''}}{(\text{Loc}(\kappa \circ \iota))}}{\vdash \Gamma}}{(\text{id})}}{\vdash F', G'}$$

Figure 9.2: Commutation of logical rules with relocations where $\eta \in \{\mu, \nu\}$. (The case for $(\&)$ and (\oplus) , (\perp) is similar to that of (\otimes) and (\wp) respectively. The case for (\perp) and $(\mathbf{1})$ is similar to that of (id) .) Here r is the formula which is principal after the rule application.

Proof. Consider some $\pi_0 \in \text{CSeq}(\pi_m)$ such that the number of (cut) inferences (say n_c) on a path from Γ_1 to Γ_2 is minimal. If $n_c = 1$, we are done. Otherwise, we claim that there exists π_1 such that $\pi_0 \xrightarrow{\text{c}_{\text{comm}}} \pi_1$, $\pi_1 \in \text{CSeq}(\pi_m)$, and the distance between Γ_1 and Γ_2 less than n_c . This immediately contradicts the minimality of n_c . Therefore the only thing left to show is that such a π_1 exists.

Since $n_c > 1$, we observe that in π_0 , Γ_1 and Γ_2 are introduced by different cuts (since they are cut-connected, they are necessarily introduced by cuts). Wlog, assume that Γ_2 is introduced by a cut inference c that is higher than the cut inference c' introducing Γ_1 . Since there is a multicut with premisses $\vdash \Gamma_1$ and $\vdash \Gamma_2$, c' can be permuted down until it is exactly above a cut rule c'' (possibly c). Consider π_1 where c' permutes below c'' . Note that π_1 satisfies all requisite conditions. \square

Lemma 9.2.3. *Suppose π_m has finitely many mcuts. Let $\pi, \pi' \in \text{CSeq}(\pi_m)$. Then, $\pi \xrightarrow{\text{c}_{\text{comm}}^*} \pi'$.*

Proof. Since there are finitely many mcuts, wlog, there is exactly one mcut given by $(\iota, \perp\!\!\!\perp)$. Let $\{\vdash \Gamma_i\}_{i \in [n]}$ be its premisses. Then, by Lemma 9.2.1, $\vdash \Gamma$ occurs in π and π' . Since there can be no merge reductions above them, the subtree rooted under $\vdash \Gamma_i$ is identical in π, π' , and π_m for all $i \in [n]$. We will induct on the size of the of graph of $\perp\!\!\!\perp$. Let $\Gamma_1, \Gamma_2, \Gamma'_1$ and Γ'_2 be such that the lowest cut in π

(respectively, π') has premisses $\vdash \Gamma_1$ and $\vdash \Gamma_2$ (respectively $\vdash \Gamma'_1$ and $\vdash \Gamma'_2$). If $\{\Gamma_1, \Gamma_2\} = \{\Gamma'_1, \Gamma'_2\}$, we consider the subproofs rooted at Γ_1 and Γ_2 and apply the induction hypothesis. Otherwise, wlog, assume $\Gamma_1 \neq \Gamma'_1$. Then, $\vdash \Gamma_1$ occurs above $\vdash \Gamma'_1$ in π' . We permute the cut with premiss $\vdash \Gamma_1$ below the lowermost cut (by way of a finite sequence of $\longrightarrow_{\text{comm}}$ rules), so that we are in the same situation as before. \square

Lemma 9.2.4. *Let π_m, π'_m such that $\pi_m \longrightarrow_m \pi'_m$. Then, there exists $\pi, \pi' \in \mu\text{MALL}^\infty$ such that the following holds.*

1. Either $\pi' = \pi$ or $\pi \longrightarrow_c \pi'$;
2. $\pi' \in \text{CSeq}(\pi'_m)$;
3. $d(\pi_m, \pi'_m) \geq d(\pi, \pi')$.

Proof. If $\pi_m \longrightarrow_{\text{merge}} \pi'_m$ then, choose any $\pi \in \text{CSeq}(\pi_m)$ and $\pi' = \pi$. Otherwise, there is a formula occurrence F such that F is principal after the rule application that takes π_m to π'_m . Then, there exists Γ, Γ' that is cut-connected on (F, F^\perp) . By Lemma 9.2.2, there exists $\pi \in \text{CSeq}(\pi_m)$ which has a cut c on Γ and Γ' . By simple case analysis on each multicut reduction rule (other than merge , observe that one can apply the corresponding rule \longrightarrow_c on c). Let π' be the result of this rule application. Checking $\pi' \in \text{CSeq}(\pi'_m)$ and $d(\pi_m, \pi'_m) \geq d(\pi, \pi')$ is immediate. \square

Lemma 9.2.5. *Let $\{\pi_m^i\}_{i \in \omega}$ be a strongly convergent sequence of μMALL_m^∞ proofs such that $\pi_m^0 \in \mu\text{MALL}^\infty$, $\pi_m^i \longrightarrow_m \pi_m^{i+1}$, and π_m as limit. Then, there exists a strongly convergent sequence $\{\pi^i\}_{i \in \omega}$ of μMALL^∞ proofs such that*

1. for all $i \in \omega$, $\pi^i \longrightarrow_c \pi^{i+1}$,
2. $\pi \in \text{CSeq}(\pi_m)$ where π is its limit, and
3. there exists a subsequence $\{\pi^{\sigma(i)}\}_{i \in \omega}$ such that $\pi^{\sigma(0)} = \pi^0 = \pi_m^0$ and for all $i \in \omega$, $\pi^{\sigma(i)} \in \text{CSeq}(\pi^i)$.

Proof. We will construct $\{\pi^{\sigma(i)}\}_{i \in \omega}$ by induction on i . One can duly obtain $\{\pi^i\}_{i \in \omega}$ from our construction. If $i = 0$, then let $\pi^0 = \pi^{\sigma(0)} = \pi_m^0$. Suppose $i > 0$. Since $\pi_m^{i-1} \longrightarrow_m \pi_m^i$, by Lemma 9.2.4, there exists π, π' such that $\pi \longrightarrow_c \pi'$, $\pi \in \text{CSeq}(\pi_m^{i-1})$, $\pi' \in \text{CSeq}(\pi_m^i)$, $d(\pi_m^{i-1}, \pi_m^i) \geq d(\pi, \pi')$ and either $\pi = \pi'$ or $\pi \longrightarrow_c \pi'$. By induction hypothesis, we have $\pi^{\sigma(i-1)} \in \text{CSeq}(\pi_m^{i-1})$. Now, by Lemma 9.2.3, $\pi^{\sigma(i-1)} \longrightarrow_{\text{comm}} \pi$. Choose $\pi^{\sigma(i)} = \pi'$. Note that, the $d(\pi_m^{i-1}, \pi_m^i) \geq d(\pi, \pi')$ condition implies that $\{\pi^i\}_{i \in \omega}$ inherits the strong convergence of $\{\pi_m^i\}_{i \in \omega}$.

The only thing left to show is that the limits are the same. Note that $\pi^{\sigma(i)} \in \text{CSeq}(\pi_m^i)$ has the same cuts as π_m^i up to several merge rules. Therefore, if a multicut is removed in the limit π_m , then it goes higher and higher in $\{\pi_m^i\}_{i \in \omega}$. Consequently, the cuts corresponding to this multicut are also removed in π , the limit of $\{\pi^i\}_{i \in \omega}$. Similarly, if a multicut is not removed in π_m , the cuts corresponding to it are not removed in π . Therefore, $\pi \in \text{CSeq}(\pi_m)$. \square

Theorem 9.2.1. *If π^0 is a μMALL^∞ proof, then, there is a sequence of μMALL^∞ proofs $\{\pi^i\}_{i \in \omega}$ with $\pi^i \longrightarrow_c \pi^{i+1}$ strongly converging to a cut-free μMALL^∞ proof π' .*

Proof. Let $\{\pi_m^i\}_{i \in \omega}$ be a μMALL_m^∞ fair reduction sequence such that $\pi_m^0 = \pi^0$. By Theorem 4.5.2, it is a strongly convergent sequence with a cut-free limit π' . By Lemma 9.2.5, there exists a strongly convergent sequence $\{\pi^i\}_{i \in \omega}$ with limit $\pi'' \in \text{CSeq}(\pi')$. Since, π' is cut-free, $\pi'' = \pi'$. Hence done. \square

Corollary 9.2.1.1. *Let $\mathcal{S}, \mathcal{S}'$ be fair reduction sequences from a proof π_0 . Then, they strongly converge to the same limit.*

This same investigation can be done for cut elimination with respect to the bouncing thread cut-elimination (cf. Theorem 4.5.3). The details are outside the scope of this thesis. We end this section by showing that Figure 4.1b indeed has a productive cut-elimination using the \longrightarrow_c reduction rules.




Figure 9.5: Discovering the subnet erasure rule

$$\pi = \frac{\frac{\frac{\pi'}{\triangle} \quad \frac{\vdots}{\vdash \nu X.X, A}(\nu)}{\vdash \Gamma, A^\perp \quad \vdash \nu X.X, A}(\text{cut})}{\vdash \Gamma, \nu X.X} \xrightarrow{\mu\text{MALL}^\infty} \frac{\frac{\frac{\pi'}{\triangle} \quad \frac{\vdots}{\vdash \nu X.X, A}(\nu)}{\vdash \Gamma, A^\perp \quad \vdash \nu X.X, A}(\text{cut})}{\vdash \Gamma, \nu X.X} \xrightarrow{\mu\text{MALL}^\infty} \frac{\frac{\frac{\pi'}{\triangle} \quad \frac{\vdots}{\vdash \Gamma, A^\perp}(\text{cut})}{\vdash \Gamma, \nu X.X}(\nu)}{\vdash \Gamma, \nu X.X} \xrightarrow{\mu\text{MALL}^\infty} \dots$$

Figure 9.6: A productive infinite reduction

9.3 Cut-elimination in infinets with axioms



In Section 9.1 cuts were eliminated only at the limit. In this section, we have finite axioms which interact with cuts by annihilating one another. In Section 7.2, we saw that, in order to have such cut/ax reduction rules in proof-nets with explicit handling of addresses, we need relocation cells. Consequently, infinets in this section have the *loc* component.

9.3.1 The subnet-erasure rule

The nets in Section 9.1 not only have no finite axioms, but they are also axiom-free. In particular, this means that every word in an infinite axiom is infinite. This is not necessarily the case in general. Consider the infinet $\mathcal{R} = (\{F^{i^\omega}, A^{\{\varepsilon\}}, A^\perp^{\{\varepsilon\}}, G^{i^\omega}\}, \{\{A, A^\perp\}\}, \{\{\alpha i^\omega, \beta\}, \{\gamma i^\omega, \beta^\perp\}\})$ where $F = \nu x.x_\alpha$, $G = \nu x.x_\gamma$, and $A = \varphi_\beta$ for any arbitrary formula φ . Unless we devise new reduction rules, we have no cut reduction rule that can be applied on this net. Therefore, we need to define new reduction rules for infinite axioms.

Note that a straightforward adaptation of the rule for finite axioms makes no sense. Imagine we reduce the cut construing the infinite axiom on the right as a finite axiom. It would result in reducing \mathcal{R} to the object in Figure 9.5c which is not an infinet. The type of the infinite path on the right changes to A^\perp at the limit *viz.* the relocation cell changes not just the addresses but the formula itself. Now, imagine we reduce the cut and the infinite axiom on the left. Similarly, it would reduce to an untyped object as above.

To justify a better rule, consider the template of an infinet where such a rule is potentially applicable in Figure 9.5a. Note that \mathcal{S} is an arbitrary infinet with doors Γ and φ^\perp . To get an intuition for the rule, we go back to the sequent calculus. Let π be a sequentialisation of the infinet above such that $\text{dsq}(\pi') = \mathcal{S}$. The infinite axiom is represented in π by the infinite branch and the only way to make it interact with π' (corresponding to \mathcal{S} interacting with the infinite axiom in the net) using the rules in Definition 9.2.2 is by commuting the cut with the ν -rule. Iterating such permutations yields infinite reduction sequence such that every proof in this sequence desequentialises to \mathcal{R} . The sequence converges to the cut-free proof in Figure 9.5b where π' has been deleted and Γ is supported by the infinite branch. Desequentialised, this yields the proof-structure in Figure 9.5c. So, an infinitary axiom and a cut interact by removing the whole subinfinet “above” the cut. Consequently, we have the following rule where \mathcal{S} is a subnet with φ^\perp as door.

Figure 9.7: Instances of \mathcal{S} from Example 9.3.1

Although this operation will be represented by a single rule it does not correspond to one step of cut-elimination in the sequent calculus but to an infinite sequence of permutations.

Proposition 9.3.1. *The subnet-erasure rule preserves correctness.*

Proof sketch. We use the notation in the graph rewriting rule above and denote the infinet on the left to be \mathcal{R} and proof-structure on the right to \mathcal{R}' . We first note that it is impossible for the rule to break connectedness. Potentially there could be a cycle involving the infinite axiom. The wires in and out of the infinite axiom in the cycle cannot be of Δ or we could reproduce the same cycle in \mathcal{R} . Therefore, they are from Γ . Now since \mathcal{R} is DR-correct, these wires are not connected in \mathcal{S} which breaks the connectedness of \mathcal{S} . Finally observe that the set of tensors, cuts, and pars in \mathcal{R} is a subset of the set of tensors, cuts, and pars in \mathcal{R}' and same dependencies. Therefore, if $\text{DepGrph}(\mathcal{R}')$ has a ray then so does $\text{DepGrph}(\mathcal{R})$. Hence done. \square

Example 9.3.1. *We provide two concrete examples of \mathcal{S} . In the first one, we observe that \mathcal{S} must be a subnet. In the second example, we observe that any subnet works.*

1. Consider the net in Figure 9.7a. Note that the smallest subnet with $a^\perp \wp b^\perp$ as door i.e. $\mathfrak{C}(a^\perp \wp b^\perp)$ has the door $a \otimes b$. If we consider the smaller substructure with doors a and b , then a, b both get grafted into the infinite axiom, thereby breaking DR-correctness.
2. Let $\varphi = a \otimes b$ and \mathcal{S} be the net in Figure 9.7b. Note that there are two distinct subnets with $a^\perp \otimes b^\perp$ as door viz. $\mathfrak{C}(a^\perp \otimes b^\perp)$ and $\mathfrak{P}(a^\perp \otimes b^\perp)$. We can choose to delete either one of them and obtain a correct net. Both of them correspond to two different reduction sequences in sequent proofs (the reduction sequence corresponding to $\mathfrak{P}(a^\perp \otimes b^\perp)$ commutes down the par-rule at some point and the reduction sequence corresponding to $\mathfrak{C}(a^\perp \otimes b^\perp)$ doesn't). However, this is non-confluent as we obtain different cut-free normal forms.

Remark 9.3.1. *In Example 9.3.1, when we considered Figure 9.7b, the reduction sequence corresponding to the erasure of $\mathfrak{P}(a^\perp \otimes b^\perp)$ infinitely commutes down the par-rule. Therefore, it is not a fair reduction sequence. Consequently, we consider a smaller set of reduction rules: instead of deleting any subnet, we specifically delete kingdoms. This not only helps us restore confluence but it also restores fairness.*

Definition 9.3.1. *The cut elimination relation $\rightarrow_{\mu\text{MLL}^\perp}$ is the binary relation over proof structures generated by the rules in Definition 7.2.12 and the following rule:*


$$(\Gamma \cup \{F^\perp{}^U, F^{\{\varepsilon\}}\}, \mathfrak{R} \cup \{\{F, F^\perp\}\}, \Theta \cup \{\theta \cup \{\alpha\}\}) \rightarrow_{\{F, F^\perp\}} (\Gamma', \mathfrak{R}', \Theta')$$

where $\text{addr}(F) = \alpha$. Let $\mathfrak{G}(F^\perp) = (\Gamma'', \mathfrak{K}'', \Theta'')$. Then, $\mathfrak{K}' = \mathfrak{K} \setminus \mathfrak{K}''$. For all F^U , if there exists $u \in U$ such that (F, u) is a door in Γ'' then $F^{U'} \in \Gamma'$ where U' is pruned at u . Finally, $\Theta' = (\Theta \setminus \Theta'') \cup \{\theta''\}$ where $\theta'' = \theta \cup \{\text{addr}(G) \mid G \in \Gamma''\}$.

Proposition 9.3.2. *Cut-reduction on infinets is confluent.*

Proof sketch. Let \mathcal{R} be an infinet that reduces to \mathcal{R}_1 and \mathcal{R}_2 by the rule r_1 and r_2 respectively. If $r_1 = r_2$, then $\mathcal{R}_1 = \mathcal{R}_2$ and we are done. Assume that $r_1 \neq r_2$. If r_1, r_2 are not kingdom erasure rules, then there cannot be any critical pairs. Therefore, one can apply r_2 on \mathcal{R}_1 and r_1 on \mathcal{R}_2 to obtain an infinet \mathcal{R}' .

Now without loss of generality, assume that r_1 is a kingdom erasure that deletes the $\mathfrak{G}(F)$ in \mathcal{R} . If the cut on which r_2 acts not in $\mathfrak{G}(F)$ then this is not a critical pair and one can apply r_2 on \mathcal{R}_1 and r_1 on \mathcal{R}_2 as before. Otherwise, note that r_2 cannot be applied on \mathcal{R}_1 but r_1 can be applied on \mathcal{R}_2 to give \mathcal{R}_1 as follows.



In any rewriting system, the diamond property implies confluence. Hence done. \square

Finally, reduction sequences and fair reduction sequences are defined in the same way as Definition 9.1.5.

9.3.2 Limits of reduction sequences

It is difficult to make a big-step guess of the normal form in this case. There is a good reason for that. Axiom-free nets are computationally less expressive than general μMLL^∞ nets. Construing cuts as computation via the CH correspondence, guessing the normal form essentially amounts to guessing the value of some function at some input without going through the function's small-step semantics. The more complicated a function is, the harder it is to guess the normal form.

Here is what we can salvage from the technique in Section 9.1: let the distance between two infinets $\mathcal{R} = (\Gamma, \mathfrak{K}, \Theta)$ and $\mathcal{R}' = (\Gamma', \mathfrak{K}', \Theta')$ be defined as $d(\mathcal{R}, \mathcal{R}') = d_1(\Gamma, \Gamma') + d_2(\mathfrak{K}, \mathfrak{K}')$ where d_2 is the distance between sets of cuts as defined in Section 9.1. Note that this distance d is reflexive and transitive but does not satisfy the identity of indiscernibles. So, one can prove that *w.r.t.* d , any fair reduction sequence $\{\mathcal{R}_i\}_{i \in \omega}$ converges to a net with the same non-cut doors as any term \mathcal{R}_i and no cuts. However, there can be several such nets and this is as far as we can go.

Instead, we define a metric on the set of infinets by appealing the heights of the cuts in a sequentialisation: basically, we use a sequentialisation to give a tree-like ordering to a proof-structure, and hence, a notion of distance compatible with the reduction. This method works for progressing infinets, as we have Theorem 9.2.1 for the sequent calculus.

Thus, infinitary cut-elimination is carried out in *correct* and *progressing* μMLL^∞ proof-structures: correctness allows us to use the tree topology of the sequentialisations; while the progress condition ensures productivity. However, we do not have a straightforward one-one correspondence between reduction sequences in proofs and proof-nets because proof-nets quotient several commutation steps. In fact, the kingdom-erasure rule corresponds to an infinite reduction sequence.

Definition 9.3.2. *We define the family of relations $\{\Rightarrow_h \mid h \in \mathbb{N}\}$ on μMLL^∞ proofs such that $\pi_0 \Rightarrow_h \pi'$ if the prefixes of π_0 and π' of height h are identical and one of the following holds.*

- either π' is the limit of an infinite sequence $(\pi_i)_{i \geq 0}$ such that for all $i \geq 0$, π_{i+1} is obtained from π_i by a permutation;
- or, there exists a finite sequence $(\pi_i)_{i \leq n}$ such that for all $i \leq n-1$, π_{i+1} is obtained from π_i by a permutation of an inference rule, and π' can be obtained from π_n by an external cut-reduction.

Proposition 9.3.3. *Let $\mathcal{R}, \mathcal{R}'$ be progressing infinets such that $\mathcal{R} \rightarrow_\kappa \mathcal{R}'$. Then, there exists π, π' such that $\text{dsq}(\pi) = \mathcal{R}$, $\text{dsq}(\pi') = \mathcal{R}'$, and $\pi \Rightarrow_h \pi'$. Diagrammatically, we have the following.*

$$\begin{array}{ccc}
\mathcal{R} & \xrightarrow{\quad \kappa \quad} & \mathcal{R}' \\
\text{dsq}(\bullet) \uparrow & & \text{dsq}(\bullet) \uparrow \\
\pi & \xrightarrow{\quad h \quad} & \pi'
\end{array}$$

Proof. Fix a sequentialisation π of \mathcal{R} . We will now construct π' . There are two cases.

Case 1. κ labels a kingdom erasure reduction. As discussed earlier, this can be simulated by an infinite number of permutations above the cut rule in π (while not modifying anything below).

Case 2. κ labels a reduction that is not a kingdom erasure. Then, it corresponds to one external step of cut-elimination in π after a finite (possibly none) number of commutation steps. \square

Definition 9.3.3. Let $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \lambda}$ be a reduction sequence. A **sequentialisation** of \mathcal{S} is defined as a sequence of proofs $\{\pi_i\}_{i \in \lambda}$ such that for all $i \in \lambda$ such that $i + 1 \in \lambda$, $\text{dsq}(\pi_i) = \mathcal{R}_i$, $\text{dsq}(\pi_{i+1}) = \mathcal{R}_{i+1}$, $\pi_i \Rightarrow_{h_i} \pi_{i+1}$ and for all π' such that $\pi_i \Rightarrow_{h'} \pi'$, we have $h' \leq h$. Diagrammatically we have the following.

$$\begin{array}{ccccccc}
\mathcal{R}_0 & \longrightarrow & \mathcal{R}_1 & \longrightarrow & \mathcal{R}_2 & \longrightarrow & \cdots \longrightarrow \mathcal{R}_k \longrightarrow \cdots \\
\uparrow & & \uparrow & & \uparrow & & \uparrow \\
\pi_0 & \xrightarrow{h_0} & \pi_1 & \xrightarrow{h_1} & \pi_2 & \xrightarrow{h_2} & \cdots \xrightarrow{h_{k-1}} \pi_k \xrightarrow{h_k} \cdots
\end{array}$$

Proposition 9.3.4. Let $\{\pi_i\}_{i \in \omega}$ be a sequentialisation of a reduction sequence such that for all $i \in \omega$, $\pi_i \Rightarrow_{h_i} \pi_{i+1}$. Then, $\{\pi_i\}_{i \in \omega}$ has a limit.

Proof. Fix an arbitrary i . Let $\mathcal{R}_i = \text{dsq}(\pi_i)$, $\mathcal{R}_{i+1} = \text{dsq}(\pi_{i+1})$, and $\mathcal{R}_i \rightarrow_{\kappa} \mathcal{R}_{i+1}$ for some κ . By construction, κ is a cut that occurs in π_i but not in π_{i+1} . Therefore, π_i and π_{i+1} differs at the height κ occurs. Since for all π' such that $\pi_i \Rightarrow_{h'} \pi'$, we have $h' \leq h$, we have that κ occurs in π_i at height h . Now, there can be at most finitely many cuts at a particular height, therefore for all $m \in \mathbb{N}$, there exists $n > 0$ such that for all $i \geq n$, $h_i > m$. Therefore, $\{h_i\}_{i \in \omega}$ diverges. Consequently, the sequence $\{\pi_i\}_{i \in \omega}$ is Cauchy with respect to the standard distance on infinite trees. Furthermore, we recall that the class of infinite trees with respect to this distance is a complete metric space. Therefore, every Cauchy sequence has a limit. \square

Lemma 9.3.1. Let $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \omega}$ be a fair reduction sequence such that \mathcal{R}_0 is progressing. Then, every sequentialisation of \mathcal{S} converges to a cut-free proof.

Proof. Let $\{\pi_i\}_{i \in \omega}$ be a sequentialisation of \mathcal{S} . By Proposition 9.3.4, it strongly converges to a pre-proof π . Since every \Rightarrow_h rule is a (finite or infinite) sequence of cut-reduction rules, we have a transfinite reduction sequence $\mathcal{S}' = \{\pi'_i\}_{i < \alpha}$ (for some $\alpha \in \text{Ord}$) such that:

1. $\{\pi_i\}_{i \in \omega}$ is its subsequence;
2. $\pi'_\beta \rightarrow_c \pi'_{\beta+1}$ for all $\beta < \alpha$;
3. $\pi'_\lambda = \sup\{\pi'_\beta \mid \beta < \lambda\}$ for limit ordinals $\lambda < \alpha$;

Clearly, for all $i < \alpha$ there exists $j < \omega$ such that $\text{dsq}(\pi'_i) = \mathcal{R}_j$ and \mathcal{S}' also converges to π . Therefore, by the compression lemma there is a subsequence $\mathcal{S}'' = \{\pi''_i\}_{i < \omega}$ of length at most ω which converges to π . We claim that \mathcal{S}' is fair. Suppose not. Then, there exists $i < \omega$ such that there is a cut C in π_i such that for all $k > i$, no cut reduction rule is applied on C in π_k . Let $\text{dsq}(\pi_i) = \mathcal{R}_j$. Then, no cut reduction rule is applied on C in \mathcal{R}_k for all $k > j$. Therefore, \mathcal{S}' is fair. Since \mathcal{R}_0 is progressing, π'_0 is a proof. Therefore, by Theorem 9.2.1, π is cut-free and satisfies the progress condition. \square

Lemma 9.3.2. Let $\mathcal{S} = \{\mathcal{R}_i\}_{i \in \omega}$ be a fair reduction sequence such that \mathcal{R}_0 is progressing. Let \mathcal{S}_1 and \mathcal{S}_2 be two sequentialisations of \mathcal{S} such that they converge to π_1 and π_2 respectively. Then, $\text{dsq}(\pi_1) = \text{dsq}(\pi_2)$.

Proof. This is trivial. By Lemma 9.3.1, π_1 and π_2 are cut-free and by Corollary 9.2.1.1 they are equivalent up to permutations. By Theorem 8.4.1, $\text{dsq}(\pi_1) = \text{dsq}(\pi_2)$. \square




Figure 9.8: Cut reduction on the desequentialisation of the pre-proofs in Figures 4.1a and 4.1b

It may not be obvious why fairness is necessary for Lemma 9.3.2. Suppose we consider a reduction sequence starting from $\mathcal{R}_0 = (\Gamma, \mathfrak{R}, \Theta)$. In the absence of fairness, two sequentialisations can reduce two different strict subsets \mathfrak{R}' and \mathfrak{R}'' of \mathfrak{R} . So in the limit of one will contain the cuts from $\mathfrak{R} \setminus \mathfrak{R}'$ while the other limit will contain cuts from $\mathfrak{R} \setminus \mathfrak{R}''$. Therefore the cannot desequentialise to the same infinnet.

Definition 9.3.4. Let $S = \{\mathcal{R}_i\}_{i \in \lambda}$ be a reduction sequence such that \mathcal{R}_0 is progressing. The **limit** of S is defined as the desequentialisation of the limit of a sequentialisation of S .

By Lemma 9.3.2, Definition 9.3.4 is well-defined. The following theorems follows.

Theorem 9.3.1. Let $S = \{\mathcal{R}_i\}_{i \in \lambda}$ be a reduction sequence such that \mathcal{R}_0 is progressing. Then, its limit is a cut-free infinnet.

Consider the desequentialisation of the pre-proofs in Figures 4.1a and 4.1b and the steps of cut reduction on it in Figure 9.8 using the rules in Definition 9.3.1. Observe it closely resembles the reduction sequence in Figure 9.3. However, providing the same reasoning as above for bouncing threads is not possible bouncing thread validity is not stable under permutation of inference rules. Therefore, although we have a reduction sequence over infinnets, proving that it has a limit is not clear.

In conclusion. In this part, we developed a theory of non-wellfounded proof-nets and in this chapter, the theory developed so far culminated to realise (at least a part of) its raison d'être. We proved cut-elimination for simple infinnets with the help of the cut-elimination result in the sequent calculus and proved cut-elimination for axiom-free infinnets completely independently. The main stumbling block in this chapter has been to find limits of infinite reduction sequence of proof-nets.

Chapter 10

Conclusion

We pitched this thesis on the following premise.

Linear logic with fixed points (μ MALL) has an intricate provability relation and the study of its computational content can be done much more systematically in the framework of proof-nets.

In Part I, we developed the provability semantics of several wellfounded calculi for μ MALL and showed that decision problems for all established systems for μ MALL are all undecidable but some are more undecidable than others. In particular we showed that the circular system is Σ_1^0 -complete while the non-wellfounded system is $(\Sigma_1^0 \cup \Pi_1^0)$ -hard which helped us to separate these systems based on the set of theorems they prove. In Part II, we developed proof-nets for μ MALL systems and it turned out to be especially challenging for the non-wellfounded situation. However, the study revealed interesting objects on non-wellfounded proofs *viz.* trips that promise connections with GoI. Furthermore, we proved cut-elimination on non-wellfounded proof-nets that required interesting non-local cut-reduction rules involving classic gadgets from proof-net theory such as kingdoms. We believe that we have justified our pitch and we conclude with several directions that have opened up during our investigation.

Perspectives on Part I

Provability relation of the non-wellfounded calculus. We sketched a proof idea for obtaining truth semantics for the non-wellfounded calculus and conjectured that provability of a formula in this calculus is in the analytic hierarchy. For both ideas, it is crucial to consider intermediary systems that are non-wellfounded and infinitely branching. Thus, to understand the exact provability relation of the non-wellfounded calculus, it is imperative to understand its relation with infinitely branching systems.

Understand the relation of the non-wellfounded system with infinitely branching systems.

Since the calculus is highly undecidable and non-regularisable, it is also interesting to obtain fragments such that they are decidable or regularisable or both. We showed that the finite fragment is already undecidable and the $\&$ -free finite fragment is equivalent to MELL. On the other hand, the additive fragment is easily decided. It is interesting to consider the decidability of the multiplicative fragment of the non-wellfounded calculus. We showed that the additive fragment is regularisable and also certain non-wellfounded proofs of multiplicative formulas (like that of $\nu x.x\wp x$) are regularisable. Therefore, it is interesting to obtain a natural fragment that would contain both additive formulas and $\nu x.x\wp x$ at the least.

Obtain natural non-trivial fragments of the non-wellfounded system that are decidable or regularisable or both.

Provability of the circular calculus. We showed that the non-wellfounded calculus and circular calculus prove different sets of theorems. Hence they cannot have the same provability semantics. In fact, we believe the technique of going through intermediary infinitely branching systems will not work since infinitely branching systems are inherently non-uniform.

However, the truth semantics of the circular system is really interesting. In particular, the completeness proof for the circular system will be technically interesting: circular proofs do not admit cuts whereas the completeness of phase semantics usually gives cut admissibility. Moreover, the result will have deep implications. Note that we have already obtained the truth semantics of the wellfounded system and proved that it has the same complexity as the circular system. Therefore, if these systems do not prove the same set of theorems, then that cannot be exhibited by a complexity argument. However, if we interpret formulas provable in the circular calculus in the same mathematical model as that for the wellfounded system, we can check if they prove the same set of theorems.

Obtain the truth semantics for the circular system in order to answer the Brotherston-Simpson conjecture for μ MALL *viz.* if the wellfounded and circular system prove the same set of theorems.

Refining the completeness of wellfounded systems. Completeness of phase semantics usually gives cut admissibility and that was also the case in the proof systems we considered. It was especially interesting in the infinitely branching systems since cut admissibility for such systems is usually proved using cut-ranks. Techniques involving cut ranks used to obtain cut-admissibility also provide upper bounds on the size of the cut-free proof. It would be interesting to see if our completeness proof can be refined to obtain such bounds.

The wellfounded system with Park's (co)induction rules has the focusing property but assigning polarities to fixed point operators is not a priori clear. In fact, it holds for both possible assignments (the proofs being quite different). In the circular and non-wellfounded case, one can syntactically argue that μ has to be positive (consequently ν should be negative). Categorical semantics of the polarised wellfounded system also informs us that μ should indeed be positive. Can phase semantics also shed light on the polarities of fixed points?

Refine the completeness of wellfounded systems to obtain upper bounds on sizes of cut-free proofs and polarities of fixed point formulas.

Perspectives on Part II

Strengthen the results for infinets. We discussed the correctness and sequentialisation of connected infinets. However, as we saw that infinets can be inherently disconnected and it is important to lift our results for disconnected infinets. Secondly, a lot of our results were obtained for simple infinets which lack crucial computational content and it is imperative to lift our results for infinets with visitable paths.

Develop the statics of infinets for more generalised structures that are potentially disconnected and might have visitable paths.

On the dynamics side, our most general result relies on the cut-elimination in the sequent calculus. To really reap the benefits of the canonicity of proof-nets, it is necessary to prove this independent of the sequent calculus result. Furthermore, the bouncing thread progress condition cannot be immediately lifted to proof-nets as it is inherently non-canonical *i.e.* it is not stable under permutation of inference rules. It is also related to the previous problem since bouncing threads on infinets would be special types of visitable paths. The ultimate goal would thus be to define a corresponding bouncing thread progress condition on infinets and show that it is a sufficient condition for ensuring the productivity of cut-elimination.

Lift bouncing thread progress condition to infinets and prove cut-elimination for bouncing thread progressing infinets.

Non-wellfounded and circular natural deduction. The translation between sequent calculus and natural deduction is usually immediate in wellfounded systems. However, the progress condition does not scale to natural deduction. In fact, preliminary investigations suggest that the natural

progress condition for natural deduction is the bouncing thread condition. Furthermore, natural deduction proofs are also canonical objects like proof-nets. Therefore, the study of non-wellfounded and circular natural deduction will possibly shed light on the bouncing thread progress condition for infinites.

Develop infinitary proof theory in the natural deduction setting.



Bibliography

- [Abr90] V. Michele Abrusci. A comparison between lambek syntactic calculus and intuitionistic linear propositional logic. *Math. Log. Q.*, 36(1):11–15, 1990.
- [ACCL91] Martín Abadi, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. Explicit substitutions. *Journal of Functional Programming*, 1(4):375–416, 1991.
- [Acz77] Peter Aczel. An introduction to inductive definitions. In Jon Barwise, editor, *Handbook of Mathematical Logic*, volume 90 of *Studies in Logic and the Foundations of Mathematics*, pages 739–782. Elsevier, 1977.
- [AD93] Gerard Allwein and J. Michael Dunn. Kripke models for linear logic. *Journal of Symbolic Logic*, 58(2):514–545, 1993.
- [ADZ17a] Davide Ancona, Francesco Dagnino, and Elena Zucca. Generalizing inference systems by coaxioms. In Hongseok Yang, editor, *Programming Languages and Systems*, pages 29–55, Berlin, Heidelberg, 2017. Springer Berlin Heidelberg.
- [ADZ17b] Davide Ancona, Francesco Dagnino, and Elena Zucca. Reasoning on divergent computations with coaxioms. *Proc. ACM Program. Lang.*, 1(OOPSLA), oct 2017.
- [AH89] R. Alur and T.A. Henzinger. A really temporal logic. In *30th Annual Symposium on Foundations of Computer Science*, pages 164–169, 1989.
- [AJ92] Samson Abramsky and Radha Jagadeesan. Games and full completeness for multiplicative linear logic. In Rudrapatna Shyamasundar, editor, *Foundations of Software Technology and Theoretical Computer Science*, pages 291–301, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.
- [AL13] Bahareh Afshari and Graham E. Leigh. On closure ordinals for the modal mu-calculus. In Simona Ronchi Della Rocca, editor, *Computer Science Logic 2013 (CSL 2013)*, volume 23 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30–44, Dagstuhl, Germany, 2013. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [All89] Lloyd Allison. Circular programs and self-referential structures. *Software: Practice and Experience*, 19(2):99–109, 1989.
- [And92] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *J. Log. Comput.*, 2(3):297–347, 1992.
- [AR02] Andrea Asperti and Luca Roversi. Intuitionistic light affine logic. *ACM Trans. Comput. Logic*, 3(1):137–175, jan 2002.
- [Asp87] Andrea Asperti. A logic for concurrency. Technical report, Technical report, Dipartimento di Informatica, Università di Pisa, 1987.
- [Asp98] A. Asperti. Light affine logic. In *Symposium on Logic in Computer Science*, page 300, 1998.
- [AU79] Alfred V. Aho and Jeffrey D. Ullman. Universality of data retrieval languages. In *Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '79, page 110–119, New York, NY, USA, 1979. Association for Computing Machinery.

- [Avi01] Jeremy Avigad. Algebraic proofs of cut elimination. *The Journal of Logic and Algebraic Programming*, 49(1):15–30, 2001.
- [Avr88] Arnon Avron. The semantics and proof theory of linear logic. *Theoretical Computer Science*, 57(2):161–184, 1988.
- [BACJ98] Frank A. Bäuerle, David Albrecht, John N. Crossley, and John S. Jeavons. Curry-howard terms for linear logic. *Studia Logica*, 61(2):223–235, Sep 1998.
- [Bae08] David Baelde. *A linear approach to the proof-theory of least and greatest fixed points*. PhD thesis, Ecole Polytechnique, December 2008.
- [Bae12] David Baelde. Least and greatest fixed points in linear logic. *ACM Transactions on Computational Logic*, 13(1), January 2012.
- [BDKS22] David Baelde, Amina Doumane, Denis Kuperberg, and Alexis Saurin. Bouncing threads for circular and non-wellfounded proofs: Towards compositionality with circular proofs. In Christel Baier and Dana Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, pages 63:1–63:13. ACM, 2022.
- [BDS15] David Baelde, Amina Doumane, and Alexis Saurin. Least and Greatest Fixed Points in Ludics. In Stephan Kreutzer, editor, *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*, volume 41 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 549–566, Dagstuhl, Germany, 2015. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [BDS16] David Baelde, Amina Doumane, and Alexis Saurin. Infinitary proof theory: the multiplicative additive case. In *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 - September 1, 2016, Marseille, France*, volume 62 of *LIPIcs*, pages 42:1–42:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [Bel97] Gianluigi Bellin. Subnets of proof-nets in multiplicative linear logic with mix. *Mathematical Structures in Computer Science*, 7(6):663–669, 1997.
- [Ber34] George Berkeley. *The Analyst: A Discourse addressed to an Inſidel Mathematician*. J. Tonson (London), 1734.
- [BG90] C. Brown and D. Gurr. A categorical linear framework for petri nets. In *Proceedings. Fifth Annual IEEE Symposium on Logic in Computer Science*, pages 208–218, 1990.
- [BG95] C. Brown and D. Gurr. A categorical linear framework for petri nets. *Information and Computation*, 122(2):268–285, 1995.
- [Bir84] R. S. Bird. Using circular programs to eliminate multiple traversals of data. *Acta Informatica*, 21(3):239–250, Oct 1984.
- [BM96] Jon Barwise and Lawrence Moss. *Vicious Circles: On the mathematics of non-wellfounded phenomena*. CSLI, Stanford California, 1996.
- [BM07] David Baelde and Dale Miller. Least and greatest fixed points in linear logic. In Nachum Dershowitz and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning, 14th International Conference, LPAR 2007, Yerevan, Armenia, October 15-19, 2007, Proceedings*, volume 4790 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2007.
- [Bra08] H. Brandes. *Über die axiomatische Einfachheit mit besonderer Berücksichtigung der auf Addition beruhenden Zerlegungsbeweise des Pythagoräischen Lehrsatzes*. Friedrich Vieweg und Sohn, 1908.
- [Bro89] Carolyn Brown. Relating Petri nets to formulae of linear logic. Technical report, University of Pennsylvania Laboratory for Foundations of Computer Science, 1989.

- [Bro05] James Brotherston. Cyclic proofs for first-order logic with inductive definitions. In Bernhard Beckert, editor, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 78–92, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [Bro06] James Brotherston. *Sequent Calculus Proof Systems for Inductive Definitions*. PhD thesis, University of Edinburgh, November 2006.
- [Bru20] Paul Brunet. A complete axiomatisation of a fragment of language algebra. In Maribel Fernández and Anca Muscholl, editors, *28th EACSL Annual Conference on Computer Science Logic, CSL 2020, January 13-16, 2020, Barcelona, Spain*, volume 152 of *LIPICs*, pages 11:1–11:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [BS07] James Brotherston and Alex Simpson. Complete sequent calculi for induction and infinite descent. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 51–62. IEEE, 2007.
- [BS10] James Brotherston and Alex Simpson. Sequent calculi for induction and infinite descent. *Journal of Logic and Computation*, 21(6):1177–1216, 10 2010.
- [BS11] James Brotherston and Alex Simpson. Sequent calculi for induction and infinite descent. *Journal of Logic and Computation*, 21(6):1177–1216, 2011.
- [BS12] Kai Brünnler and Thomas Studer. Syntactic cut-elimination for a fragment of the modal μ -calculus. *Annals of Pure and Applied Logic*, 163(12):1838–1853, 2012.
- [BT17a] Stefano Berardi and Makoto Tatsuta. Classical system of martin-löf’s inductive definitions is not equivalent to cyclic proof system. In *FoSSaCS*, volume 10203 of *Lecture Notes in Computer Science*, pages 301–317, 2017.
- [BT17b] Stefano Berardi and Makoto Tatsuta. Equivalence of inductive definitions and cyclic proofs under arithmetic. In *LICS*, pages 1–12. IEEE Computer Society, 2017.
- [BT19] Stefano Berardi and Makoto Tatsuta. Classical system of martin-löf’s inductive definitions is not equivalent to cyclic proofs. *Log. Methods Comput. Sci.*, 15(3), 2019.
- [Bus06] Wojciech Buszkowski. On Action Logic: Equational Theories of Action Algebras. *Journal of Logic and Computation*, 17(1):199–217, 10 2006.
- [BvdW95] G. Bellin and J. van de Wiele. Subnets of proof-nets in MLL^\perp . In *Proceedings of the Workshop on Advances in Linear Logic*, page 249–270, USA, 1995. Cambridge University Press.
- [Car37] Rudolf Carnap. *Logical Syntax of Language*. Kegan Paul, Trench and Truber, 1937.
- [CC79] Patrick Cousot and Radhia Cousot. Constructive versions of Tarski’s fixed point theorems. *Pacific Journal of Mathematics*, 82(1):43 – 57, 1979.
- [CG14] Kaustuv Chaudhuri and Nicolas Guenot. Equality and fixpoints in the calculus of structures. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, CSL-LICS ’14, New York, NY, USA, 2014. Association for Computing Machinery.
- [CGP99] Edmund M Clarke, Orna Grumberg, and Doron A. Peled. Model checking for the μ -calculus. In *Model checking*, chapter 7, pages 97–108. MIT Press, London, Cambridge, 1999.
- [Chu36a] Alonzo Church. A note on the entscheidungsproblem. *Journal of Symbolic Logic*, 1(1):40–41, 1936.
- [Chu36b] Alonzo Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, 1936.

- [Cla09a] Pierre Clairambault. Least and greatest fixpoints in game semantics. In Ralph Matthes and Tarmo Uustalu, editors, *6th Workshop on Fixed Points in Computer Science, FICS 2009, Coimbra, Portugal, September 12-13, 2009*, pages 39–45. Institute of Cybernetics, 2009.
- [Cla09b] Pierre Clairambault. Least and greatest fixpoints in game semantics. In Luca de Alfaro, editor, *Foundations of Software Science and Computational Structures, 12th International Conference, FOSSACS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5504 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2009.
- [Cla10] Pierre Clairambault. *Logique et Interaction : une Étude Sémantique de la Totalité. (Logic and Interaction : a Semantic Study of Totality)*. PhD thesis, Paris Diderot University, France, 2010.
- [CO21] Wojciech Czerwinski and Lukasz Orlikowski. Reachability in Vector Addition Systems is Ackermann-complete. April 2021.
- [Col82] Alain Colmerauer. Prolog and infinite trees. *Logic programming*, pages 231 – 251, 1982.
- [Coq] Coq reference manual. <https://coq.inria.fr/distrib/current/refman/>. Accessed: 2022-08-28.
- [Coq94] Thierry Coquand. Infinite objects in type theory. In Henk Barendregt and Tobias Nipkow, editors, *Types for Proofs and Programs*, pages 62–78, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.
- [Cou82] Bruno Courcelle. *Fundamental Properties of Infinite Trees*, pages 417–471. Springer Netherlands, Dordrecht, 1982.
- [CP10] Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *International Conference on Concurrency Theory*, pages 222–236. Springer, 2010.
- [CP22] Luca Ciccone and Luca Padovani. An infinitary proof theory of linear logic ensuring fair termination in the linear π -calculus, 2022. Accepted for publication at CONCUR 2022.
- [CR36] Alonzo Church and J. B. Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39(3):472–482, 1936.
- [Cur52] Haskell B. Curry. The permutability of rules in the classical inferential calculus. *Journal of Symbolic Logic*, 17(4):245–248, 1952.
- [Cur05] Pierre-Louis Curien. Introduction to linear logic and ludics, part I. *Advances in Mathematics (China)*, 34(5):513 – 544, 2005.
- [Cur06] Pierre-Louis Curien. Introduction to linear logic and ludics, part II. *Advances in Mathematics (China)*, 35(1):1 – 44, 2006.
- [Cza10] Marek Czarnecki. How fast can the fixpoints in modal μ -calculus be reached? In Luigi Santocanale, editor, *7th Workshop on Fixed Points in Computer Science, FICS 2010, Brno, Czech Republic, August 21-22, 2010*, pages 35–39. Laboratoire d’Informatique Fondamentale de Marseille, 2010.
- [DBHS16] Amina Doumane, David Baelde, Lucca Hirschi, and Alexis Saurin. Towards completeness via proof search in the linear time μ -calculus: The case of büchi inclusions. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS ’16*, page 377–386, New York, NY, USA, 2016. Association for Computing Machinery.
- [dC18] Daniel de Carvalho. Taylor expansion in linear logic is invertible. *Log. Methods Comput. Sci.*, 14(4), 2018.

- [DDP18] Anupam Das, Amina Doumane, and Damien Pous. Left-handed completeness for Kleene algebra, via cyclic proofs. In *LPAR*, volume 57 of *EPIc Series in Computing*, pages 271–289. EasyChair, 2018.
- [DDS22] Anupam Das, Abhishek De, and Alexis Saurin. Decision problems for linear logic with least and greatest fixed points. In Amy P. Felty, editor, *7th International Conference on Formal Structures for Computation and Deduction, FSCD 2022, August 2-5, 2022, Haifa, Israel*, volume 228 of *LIPICs*, pages 20:1–20:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [dF94] Pierre Simon de Fermat. *Oeuvres de Fermats, T.2*. Gauthier-Villars et Fils, Paris, 1894.
- [DG02] Anuj Dawar and Yuri Gurevich. Fixed Point Logics. *Bulletin of Symbolic Logic*, 8(1):65 – 88, 2002.
- [dGGS04] P. de Groote, B. Guillaume, and S. Salvati. Vector addition tree automata. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.*, pages 64–73, 2004.
- [DHL06] Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ -calculus. In *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS'06*, page 273–284, Berlin, Heidelberg, 2006. Springer-Verlag.
- [Die08] Reinhard Diestel. *Graph Theory*. Springer, New York, NY, October 2008.
- [DJS22] Abhishek De, Farzad Jafarrahmani, and Alexis Saurin. Phase semantics of linear logic with least and greatest fixed point. Accepted for publication at FSTTCS 2022, Dec 2022.
- [DM04] Ugo Dal Lago and Simone Martini. Phase semantics and decidability of elementary affine logic. *Theoretical Computer Science*, 318(3):409–433, 2004.
- [Dou17] Amina Doumane. *On the infinitary proof theory of logics with fixed points. (Théorie de la démonstration infinitaire pour les logiques à points fixes)*. PhD thesis, Paris Diderot University, France, 2017.
- [DP17] Anupam Das and Damien Pous. A cut-free cyclic proof system for Kleene algebra. In Renate A. Schmidt and Cláudia Nalon, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 261–277, Cham, 2017. Springer International Publishing.
- [DP18] Amina Doumane and Damien Pous. Completeness for identity-free kleene lattices. In Sven Schewe and Lijun Zhang, editors, *29th International Conference on Concurrency Theory, CONCUR 2018, September 4-7, 2018, Beijing, China*, volume 118 of *LIPICs*, pages 18:1–18:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [DP22] Farzaneh Derakhshan and Frank Pfenning. Circular proofs as session-typed processes: A local validity condition. *Log. Methods Comput. Sci.*, 18(2), 2022.
- [DPPDD09] A.K. Doxiadēs, C.H. Papadimitriou, A. Papadatos, and A. Di Donna. *Logicomix: An Epic Search for Truth*. Bloomsbury, 2009.
- [DPS21] Abhishek De, Luc Pellissier, and Alexis Saurin. Canonical proof-objects for coinductive programming: infinets with infinitely many cuts. In Niccolò Veltri, Nick Benton, and Silvia Ghilezan, editors, *PPDP 2021: 23rd International Symposium on Principles and Practice of Declarative Programming, Tallinn, Estonia, September 6-8, 2021*, pages 7:1–7:15. ACM, 2021.
- [DR89] Vincent Danos and Laurent Regnier. The structure of multiplicatives. *Archive for Mathematical Logic*, 28(3):181–203, Oct 1989.

- [DS19] Abhishek De and Alexis Saurin. Infinets: The parallel syntax for non-wellfounded proof-theory. In Serenella Cerrito and Andrei Popescu, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 28th International Conference, TABLEAUX 2019, London, UK, September 3-5, 2019, Proceedings*, volume 11714 of *Lecture Notes in Computer Science*, pages 297–316. Springer, 2019.
- [EJ21] Thomas Ehrhard and Farzad Jafarrahmani. Categorical models of linear logic with fixed points of formulas. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2021.
- [EJS21] Thomas Ehrhard, Farzad Jafarrahmani, and Alexis Saurin. On relation between totality semantic and syntactic validity. In *5th International Workshop on Trends in Linear Logic and Applications (TLLA 2021)*, Rome (virtual), Italy, June 2021.
- [Euc56] Euclid. *The Elements of Euclid*. Dover Publications Inc., New York, 2nd edition, 1956.
- [FL79] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.
- [For03] Thomas Forster. *Logic, Induction and Sets*. London Mathematical Society Student Texts. Cambridge University Press, 2003.
- [For14] Jérôme Fortier. *Puissance expressive des preuves circulaires. (Expressive Power of Circular Proofs)*. PhD thesis, Aix-Marseille University, Aix-en-Provence, France, 2014.
- [FR94] Arnaud Fleury and Christian Retoré. The mix rule. *Mathematical Structures in Computer Science*, 4(2):273–285, 1994.
- [Fro95] Jacob Frost. A case study of co-induction in Isabelle. Technical report, University of Cambridge, Computer Laboratory, February 1995.
- [FS13] Jérôme Fortier and Luigi Santocanale. Cuts for circular proofs: semantics and cut-elimination. In *CSL*, 2013.
- [Gen36] Gerhard Gentzen. Die widerspruchsfreiheit der reinen zahlentheorie. *Mathematische Annalen*, 112(1):493–565, Dec 1936.
- [Geu92] Herman Geuvers. Inductive and coinductive types with iteration and recursion. In *Proceedings of the 1992 Workshop on Types for Proofs and Programs, Bastad*, pages 193–217, 1992.
- [GG89] Carl A. Gunter and Vijay Gehlot. Nets as tensor theories. Technical Report MS-CIS-89-68, University of Pennsylvania Department of Computer and Information Science, October 1989.
- [GGS16] Silvio Ghilardi, Maria João Gouveia, and Luigi Santocanale. Fixed-point elimination in the intuitionistic propositional calculus. In Bart Jacobs and Christof Löding, editors, *Foundations of Software Science and Computation Structures*, pages 126–141, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [GGS19] Silvio Ghilardi, Maria João Gouveia, and Luigi Santocanale. Fixed-point elimination in the intuitionistic propositional calculus. *ACM Trans. Comput. Logic*, 21(1), September 2019.
- [Gir71] J. Y. Girard. Une extension de l’interpretation de Godel a l’analyse, et son application a l’elimination des coupures dans l’analyse et la theorie des types. *Proceedings of the second Scandinavian logic symposium*, 63:63–92, 1971.
- [Gir87a] Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.
- [Gir87b] Jean-Yves Girard. Linear logic and parallelism. In Marisa Venturini Zilli, editor, *Mathematical Models for the Semantics of Parallelism*, pages 166–182, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg.

- [Gir89] Jean-Yves Girard. Towards a geometry of interaction. *Contemporary Mathematics*, 92(69-108):6, 1989.
- [Gir92] Jean Yves Girard. A fixpoint theorem in linear logic. <https://www.seas.upenn.edu/~sweirich/types/archive/1992/msg00030.html>, 1992. An email posting to the mailing list linear@cs.stanford.edu.
- [Gir95] Jean-Yves Girard. Linear logic: Its syntax and semantics. In Jean-Yves Girard, Yves Lafont, and Laurent Regnier, editors, *Advances in Linear Logic*, pages 222–1. Cambridge University Press, 1995.
- [Gir96] Jean-Yves Girard. Proof-nets: The parallel syntax for proof-theory. In *Logic and Algebra*, pages 97–124. Marcel Dekker, 1996.
- [Gir98] Jean-Yves Girard. Light linear logic. *Information and Computation*, 143(2):175–204, 1998.
- [Gir99] Jean-Yves Girard. On the meaning of logical rules i: syntax versus semantics. In *Computational logic*, pages 215–272. Springer, 1999.
- [Gir11] J.Y. Girard. *The Blind Spot: Lectures on Logic*. European Mathematical Society, 2011.
- [Göd31] Kurt Gödel. Über formal unentscheidbare sätze der principia mathematica und verwandter systeme i. *Monatshefte für Mathematik und Physik*, 38(1):173–198, Dec 1931.
- [GS17] Maria João Gouveia and Luigi Santocanale. Aleph1 and the Modal μ -Calculus. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 38:1–38:16, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [GTW03] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games*. Lecture notes in computer science. Springer, New York, NY, 2002 edition, February 2003.
- [Hag87] Tatsuya Hagino. *A categorical programming language*. PhD thesis, University of Edinburgh, UK, 1987.
- [Har16] Robert Harper. *Practical foundations for programming languages*. Cambridge University Press, Cambridge, England, 2 edition, April 2016.
- [Her95] Hugo Herbelin. A λ -calculus structure isomorphic to gentzen-style sequent calculus structure. In Leszek Pacholski and Jerzy Tiuryn, editors, *Computer Science Logic*, pages 61–75, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.
- [HH16] Willem Heijltjes and Robin Houston. Proof equivalence in MLL is pspace-complete. *Log. Methods Comput. Sci.*, 12(1), 2016.
- [HK22] Emile Hazard and Denis Kuperberg. Cyclic proofs for transfinite expressions. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPIcs*, pages 23:1–23:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [HM19] Quentin Heath and Dale Miller. A proof theory for model checking. *J. Autom. Reason.*, 63(4):857–885, 2019.
- [HP79] John Hopcroft and Jean-Jacques Pansiot. On the reachability problem for 5-dimensional vector addition systems. *Theoretical Computer Science*, 8(2):135–159, 1979.
- [JDNM11] Paulin Jacobé De Naurois and Virgile Mogbil. Correctness of Linear Logic Proof Structures is NL-Complete. *Theoretical Computer Science*, 412(20):pp. 1941–1957, April 2011. 23 pages.

- [JKS08] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Canonical completeness of infinitary μ . *The Journal of Logic and Algebraic Programming*, 76(2):270–292, 2008. Logic and Information: From Logic to Constructive Reasoning.
- [Kah87] G. Kahn. Natural semantics. In Franz J. Brandenburg, Guy Vidal-Naquet, and Martin Wirsing, editors, *4th Annual Symposium on Theoretical Aspects of Computer Sciences on STACS 87*, pages 22–39, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg.
- [Kil73] Gary A. Kildall. A unified approach to global program optimization. In *Proceedings of the 1st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '73, page 194–206, New York, NY, USA, 1973. Association for Computing Machinery.
- [Kle51] Stephen Cole Kleene. Permutability of inferences in gentzen's calculi lk and lj. In *Bulletin Of The American Mathematical Society*, volume 57, pages 485–485, 1951.
- [Kop01] Alexei Kopylov. Decidability of linear affine logic. *Inf. Comput.*, 164:173–198, 01 2001.
- [Koz88] Dexter Kozen. A finite model theorem for the propositional μ -calculus. *Stud Logica*, 47(3):233–241, 1988.
- [Koz91] D. Kozen. A completeness theorem for kleene algebras and the algebra of regular events. In *[1991] Proceedings Sixth Annual IEEE Symposium on Logic in Computer Science*, pages 214–225, 1991.
- [Koz97] Dexter Kozen. *Automata and computability*. Undergraduate texts in computer science. Springer, New York, 1997.
- [Koz02] Dexter Kozen. On the complexity of reasoning in kleene algebra. *Inf. Comput.*, 179(2):152–162, dec 2002.
- [KPP21] Denis Kuperberg, Laureline Pinault, and Damien Pous. Cyclic proofs, system t, and the power of contraction. *Proc. ACM Program. Lang.*, 5(POPL), jan 2021.
- [Kro91] Daniel Krob. Complete systems of b-rational identities. *Theoretical Computer Science*, 89(2):207–343, 1991.
- [KS17] Dexter Kozen and Alexandra Silva. Practical coinduction. *Mathematical Structures in Computer Science*, 27(7):1132–1152, 2017.
- [KSSdV05] Richard Kennaway, Paula Severi, Ronan Sleep, and Fer-Jan de Vries. *Infinitary Rewriting: From Syntax to Semantics*, pages 148–172. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- [KT27] Bronisław Knaster and Alfred Tarski. Un théorème sur les fonctions d'ensembles. *Annales de la Société Polonaise de Mathématique*, 6:133–134, 1927.
- [Kuz18] Stepan Kuznetsov. *-continuity vs. induction: Divide and conquer. In Guram Bezhanishvili, Giovanna D'Agostino, George Metcalfe, and Thomas Studer, editors, *Advances in Modal Logic 12, proceedings of the 12th conference on "Advances in Modal Logic," held in Bern, Switzerland, August 27-31, 2018*, pages 493–510. College Publications, 2018.
- [Kuz20] Stepan Kuznetsov. Complexity of commutative infinitary action logic. In Manuel A. Martins and Igor Sedlár, editors, *Dynamic Logic. New Trends and Applications*, pages 155–169, Cham, 2020. Springer International Publishing.
- [Laf97] Yves Lafont. The finite model property for various fragments of linear logic. *J. Symb. Log.*, 62(4):1202–1208, 1997.
- [Laf04] Yves Lafont. Soft linear logic and polynomial time. *Theoretical Computer Science*, 318(1):163–180, 2004. Implicit Computational Complexity.
- [Lam58] Joachim Lambek. The mathematics of sentence structure. *The American Mathematical Monthly*, 65(3):154–170, 1958.

- [Ler21] Jérôme Leroux. The Reachability Problem for Petri Nets is Not Primitive Recursive. May 2021.
- [Lib04] Leonid Libkin. *Elements of Finite Model Theory*. Springer, August 2004.
- [Lin92] Patrick Lincoln. *Computational Aspects of Linear Logic*. PhD thesis, Stanford University, 1992.
- [Lin95] Patrick Lincoln. Deciding provability of linear logic formulas. In *Proceedings of the Workshop on Advances in Linear Logic*, page 109–122, USA, 1995. Cambridge University Press.
- [LMAS92] Patrick Lincoln, John Mitchell, Scedrov Andre, and Natarajan Shankar. Decision problems for propositional linear logic. *Annals of Pure and Applied Logic*, 56(1):239–311, 1992.
- [LPM94] François Leclerc and Christine Paulin-Mohring. Programming with streams in coq: A case study: The sieve of eratosthenes. In *Proceedings of the International Workshop on Types for Proofs and Programs*, TYPES '93, page 191–212, Berlin, Heidelberg, 1994. Springer-Verlag.
- [LS14] Ranko Lazić and Sylvain Schmitz. Non-elementary complexities for branching vass, mell, and extensions. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, CSL-LICS '14, New York, NY, USA, 2014. Association for Computing Machinery.
- [LS15] Ranko Lazić and Sylvain Schmitz. Nonelementary complexities for branching vass, mell, and extensions. *ACM Trans. Comput. Logic*, 16(3), June 2015.
- [Mae91] Shoji Maehara. Lattice-valued representation of the cut-elimination theorem. *Tsukuba journal of mathematics*, 15:509–521, 1991.
- [MD97] Lawrence S. Moss and Norman Danner. On the Foundations of Corecursion. *Logic Journal of the IGPL*, 5(2):231–257, 03 1997.
- [Mel09] Paul-André Melliès. *Categorical semantics of linear logic*, pages 1 – 196. Number 27 in Panoramas et Synthèses. Société Mathématique de France, 2009.
- [Mil93] Dale Miller. The π -calculus as a theory in linear logic: Preliminary results. In E. Lamma and P. Mello, editors, *Extensions of Logic Programming*, pages 242–264, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.
- [Min67] Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., USA, 1967.
- [Min78] Grigori Mints. Finite investigations of transfinite derivations. *Journal of Soviet Mathematics*, 10:548–596, 1978.
- [ML71] Per Martin-Löf. Hauptsatz for the intuitionistic theory of iterated inductive definitions. In J.E. Fenstad, editor, *Proceedings of the Second Scandinavian Logic Symposium*, volume 63 of *Studies in Logic and the Foundations of Mathematics*, pages 179–216. Elsevier, 1971.
- [ML75] Per Martin-Löf. An intuitionistic theory of types: Predicative part. In H.E. Rose and J.C. Shepherdson, editors, *Logic Colloquium '73*, volume 80 of *Studies in Logic and the Foundations of Mathematics*, pages 73 – 118. Elsevier, 1975.
- [Mon03] Raphaël Montelatici. Polarized proof nets with cycles and fixpoints semantics. In Martin Hofmann, editor, *Typed Lambda Calculi and Applications*, pages 256–270, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [MT91] Robin Milner and Mads Tofte. Co-induction in relational semantics. *Theoretical Computer Science*, 87(1):209–220, 1991.

- [Nak75] Reiji Nakajima. Infinite normal forms for the λ -calculus. In C. Böhm, editor, *λ -Calculus and Computer Science Theory*, pages 62–82, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg.
- [Nas50] J. F. Nash. Equilibrium points in N -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(48-49), 1950.
- [Niw97] Damian Niwiński. Fixed point characterization of infinite behavior of finite-state systems. *Theor. Comput. Sci.*, 189(1–2):1–69, December 1997.
- [NST18] Rémi Nollet, Alexis Saurin, and Christine Tasson. Local validity for circular proofs in linear logic with fixed points. In Dan R. Ghica and Achim Jung, editors, *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4–7, 2018, Birmingham, UK*, volume 119 of *LIPICs*, pages 35:1–35:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018.
- [NST19] Rémi Nollet, Alexis Saurin, and Christine Tasson. Pspace-completeness of a thread criterion for circular proofs in linear logic with least and greatest fixed points. In Serenella Cerrito and Andrei Popescu, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 28th International Conference, TABLEUX 2019, London, UK, September 3–5, 2019, Proceedings*, volume 11714 of *Lecture Notes in Computer Science*, pages 317–334. Springer, 2019.
- [NW96] Damian Niwinski and Igor Walukiewicz. Games for the μ -calculus. *Theor. Comput. Sci.*, 163(1&2):99–116, 1996.
- [Oka96] Mitsuhiro Okada. Phase semantics for higher order completeness, cut-elimination and normalization proofs. *Electron. Notes Theor. Comput. Sci.*, 3:154, 1996.
- [Oka99] Mitsuhiro Okada. Phase semantic cut-elimination and normalization proofs of first- and higher-order linear logic. *Theoretical Computer Science*, 227(1):333–396, 1999.
- [Pal07] Ewa Palka. An infinitary sequent system for the equational theory of $*$ -continuous action lattices. *Fundam. Inf.*, 78(2):295–309, April 2007.
- [Pam04] Victor Pambuccian. Early examples of resource-consciousness. *Studia Logica: An International Journal for Symbolic Logic*, 77(1):81–86, 2004.
- [Par69] David Park. Fixpoint induction and proofs of program properties. *Machine intelligence*, 5, 1969.
- [Pau97] Lawrence C. Paulson. Mechanizing Coinduction and Corecursion in Higher-order Logic. *Journal of Logic and Computation*, 7(2):175–204, 04 1997.
- [Pet62] Carl Adam Petri. *Kommunikation mit Automaten*. PhD thesis, University of Bonn, 1962.
- [Plo81] Gordon D Plotkin. *A structural approach to operational semantics*. Aarhus university, 1981.
- [Poi05] Henri Poincaré. *Science and Hypothesis*. London W. Scott, 1905.
- [Pra65] Dag Prawitz. *Natural Deduction: A Proof-Theoretical Study*. Dover Publications, 1965.
- [Rab69] Michael O Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
- [Rab72] Michael Oser Rabin. *Automata on Infinite Objects and Church’s Problem*. American Mathematical Society, USA, 1972.
- [RT92] Jan J. M. M. Rutten and Daniele Turi. On the foundation of final semantics: Non-standard sets, metric spaces, partial orders. In *Proceedings of the REX Workshop on Semantics: Foundations and Applications*, page 477–530, Berlin, Heidelberg, 1992. Springer-Verlag.

- [San02] Luigi Santocanale. A calculus of circular proofs and its categorical semantics. In Mogens Nielsen and Uffe Engberg, editors, *Foundations of Software Science and Computation Structures*, pages 357–371, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [San09] Davide Sangiorgi. On the origins of bisimulation and coinduction. *ACM Trans. Program. Lang. Syst.*, 31(4), may 2009.
- [Sau08] Alexis Saurin. *Une étude logique du contrôle (appliquée à la programmation fonctionnelle et logique)*. PhD thesis, Ecole Polytechnique, 2008.
- [Sch60] Kurt Schütte. Syntactical and semantical properties of simple type theory. *The Journal of Symbolic Logic*, 25(4):305–326, 1960.
- [Sch68] Kurt Schütte. *Vollständige Systeme modaler und intuitionistischer Logik*. Springer Berlin Heidelberg, 1968.
- [Sch91] Harold Schellinx. Some Syntactical Observations on Linear Logic. *Journal of Logic and Computation*, 1(4):537–559, 09 1991.
- [Sco70] Dana Scott. Outline of a mathematical theory of computation. Technical Report PRG02, OUCL, November 1970.
- [SD03] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: Circular and tree-shaped proofs in the μ calculus. In Andrew D. Gordon, editor, *Foundations of Software Science and Computation Structures*, pages 425–440, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [Sha14] D. S. Shamkanov. Circular proofs for the gödel-löb provability logic. *Mathematical Notes*, 96(3):575–585, Sep 2014.
- [Sim17] Alex Simpson. Cyclic arithmetic is equivalent to peano arithmetic. In *FoSSaCS*, volume 10203 of *Lecture Notes in Computer Science*, pages 283–300, 2017.
- [Soa14] Robert I Soare. *Recursively enumerable sets and degrees*. Perspectives in Mathematical Logic. Springer, Berlin, Germany, January 2014.
- [SS02] Ulrich Schöpp and Alex K. Simpson. Verifying temporal properties using explicit approximants: Completeness for context-free processes. In *FoSSaCS*, volume 2303 of *Lecture Notes in Computer Science*, pages 372–386. Springer, 2002.
- [Str19] Lutz Straßburger. On the decision problem for mell. *Theoretical Computer Science*, 768:91–98, 2019.
- [Tai75] William W. Tait. A realizability interpretation of the theory of species. In Rohit Parikh, editor, *Logic Colloquium*, pages 240–251, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg.
- [Tar33] Alfred Tarski. The concept of truth in the languages of deductive systems. 1933.
- [Tar55] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.
- [Thi03] Rüdiger Thiele. Hilbert’s twenty-fourth problem. *The American Mathematical Monthly*, 110(1):1–24, 2003.
- [Tro95] A.S. Troelstra. Natural deduction for intuitionistic linear logic. *Annals of Pure and Applied Logic*, 73(1):79–108, 1995. A Tribute to Dirk van Dalen.
- [Tur37] A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, 1937.
- [VH67] Jean Van Heijenoort. *Logic as Calculus and Logic as Language*, pages 440–446. Springer Netherlands, Dordrecht, 1967.
- [Wal94] Igor Walukiewicz. *A complete deductive system for the μ -calculus*. PhD thesis, Warsaw University, 1994.

- [Weh21] Dominik Wehr. An abstract framework for the analysis of cyclic derivations. Master's thesis, Universiteit van Amsterdam, 2021.
- [WRK78] L. Wittgenstein, R. Rhees, and A. Kenny. *Philosophical Grammar*. Philosophical Grammar. University of California Press, 1978.