

Game Semantics and the Geometry of Backtracking: a New Complexity Analysis of Interaction

Federico Aschieri *

Institut für Diskrete Mathematik und Geometrie
Technische Universität Wien
Wiedner Hauptstraße 8-10/104, 1040, Vienna, Austria

Abstract

We present abstract complexity results about Coquand and Hyland-Ong game semantics, that will lead to new bounds on the length of first-order cut-elimination, normalization, interaction between expansion trees and any other dialogical process game semantics can model and apply to. In particular, we provide a novel method to bound the length of interactions between visible strategies and to measure precisely the tower of exponentials defining the worst-case complexity. Our study improves the old estimates on average by several exponentials.

1 Introduction

Of all towers of the world, the one we would contemplate far more admiringly if it were lower rather than higher, is a tower in logic. It is the stack of exponentials defining the worst-case complexity of cut-elimination, normalization and dialogues between strategies interpreting first-order proofs [8, 31]:

$$\underbrace{2^{2^{\cdot^{2^k}}}}_{r+1}$$

The tower's height is measured by the rank r of the proof, which is in sequent calculus the highest among the logical complexities of the proof cut-formulas and in natural deduction the highest among the type complexities of the proof redexes; k is just the proof height. Given the use of complex cuts and redexes in logical derivations, it does not take too much to reach astronomical worst-case bounds.

Yet, thanks to the famous Curry-Howard correspondence [27], we know that logical proofs, both in natural deduction [7, 24, 30] and in sequent calculus [21, 16], can be seen as programs of some typed functional programming language. And in the New York of a programming language, many inhabitants are indeed to be found on the highest floors of the highest skyscrapers, but the vast majority of them live in far lower buildings. That's programming: one can write enormously complex programs, but also very simple and efficient ones. Having complexity bounds only for the very worst case of the very worst program is of no use when one looks for a more precise idea of how much time it takes a particular program to terminate its execution. Unless one is terribly unlucky, his actual programs will take much less time than the worst possible program to complete their computations. How much less? We need to measure precisely the tower of exponentials, because it is its height that is to blame for the unrealistic bounds it computes.

The reason why the worst-case analysis of cut-elimination fares so roughly on average is the complexity measure adopted: the rank of the proof is a superficial information, in a

* This work was funded by the Austrian Science Fund FWF Lise Meitner grant M 1930–N35 and previously supported by the FWF grant P25160–N25.

literal sense. Indeed, one looks only at the surface of a cut, the cut-formula, without caring about how the formula is used in the proof. It is like trying to determine the maximum speed of a car just by studying its outward appearance, while ignoring completely its internal engine! But what is the engine of a cut? What is the real responsible of the worst-case complexity bounds, the mechanism that determines the height of the tower of exponentials?

1.1 Coquand and Hyland-Ong Game Semantics

It turns out that the laws governing the complexity of cut-elimination are written in a game semantical language. If we want to answer our questions, we have no choice: we must translate them in terms of games and strategies. In the quest for a better complexity measure of cut-elimination and normalization, the single most important tool is then *game semantics*. Game semantics was in the background of logic at least since Herbrand's Theorem: this is clear by Miller's formulation of it using expansion trees [28], that is, game strategies, as remarked by Heijltjes [20]. But the idea of interpreting formulas as games and proofs as winning strategies, appeared time and again in the history of logic: in Gentzen's and Novikoff's proofs of consistency of Arithmetic (see [15]), in Gödel's and Kreisel's functional interpretations (see [25, 6, 29]), in Lorenzen's game semantics (see [18]).

The old, *static* game semantics was brought to a new life with the invention of the modern, *dynamical* game semantics, initiated by Coquand's [14, 15] discovery that cut-elimination in Arithmetic mirrors a debate between two strategies. As opposed to ordinary games, players may every now and then revise their moves and backtrack to a previous position. Coquand proved by elegant combinatorial means that the dialogues between strategies interpreting proofs always terminate, thus proving as well the termination of cut-elimination, as later shown by Herbelin [21]. The advantage of game semantics is that it ignores inessential aspects of proofs and turns the spotlight to their geometrical structure.

Hyland and Ong [22] took Coquand-style game semantics to a new level. The key move is realizing that one can drop altogether winning conditions: debates need not always terminate with the win or the loss of one of the two players. Game semantics can in such a way model pure dialogues, exchanges of answers and questions, as those occurring during the evaluation of a functional program. The linear head reduction strategy of typed lambda terms can thus be modeled as a debate between strategies [17]. Thanks to these and several other insights, Hyland and Ong discovered how to construct fully abstract models for typed functional languages (see also [1, 2]).

As plays correspond to cut-elimination [21] and normalization [17], one encounters in game semantics the same complexity estimates. The rank of a cut-formula becomes the depth of the game arena, and the very same tower of exponentials that bounds cut-elimination, bounds also the dialogues between strategies corresponding to first-order proofs. The result was proved by Clairambault [12, 13]; it constitutes an important demonstration that game semantics may serve as a viable complexity analysis tool. Nevertheless, the depth of game arenas is not a precise complexity measure for the same reasons we discussed above: it corresponds just to the rank of cut-formulas.

It may appear that game semantics is leading us nowhere. For the moment we hardly achieved anything but to translate doubts and questions from one language to another. What is the engine of a strategy? How to measure its complexity? We are missing a piece of the puzzle and to get it, we need to make two long jumps: one backward and the other forward.

1.2 Back to Coquand's Simple Backtracking

Debates in game semantics can be extremely involved; proving their termination, like escaping from an intricate labyrinth with no Ariadne's thread at hand. Coquand eventually overcame

these difficulties by inventing a notation with deep geometrical properties: pointer sequences. But in a preliminary effort [14] to establish his successive termination result [15], he settled for a weaker one: that debates between *simple* backtracking strategies, now called *1-backtracking strategies*, always terminate. According to Coquand, a 1-backtracking strategy “never changes its mind about a value it has previously considered wrong”. This kind of strategies are by design less involved than general ones and their debates are more predictable; that allowed Coquand to get away at least with some result from his first struggle.

1.3 Berardi-de’Liguoro Backtracking Level of a Strategy

The idea of classifying strategies according to some complexity measure had an extraordinary potential, yet it was completely abandoned. Only much later, did Berardi and de’Liguoro [9] generalize the concept of 1-backtracking to the more general *n-backtracking* for any $n \in \mathbb{N}$. It was a major conceptual advancement. Their definition is inspired by Gold’s learning theory [19, 23] and by the Berardi-Coquand-Hayashi [11] notion of *n-backtracking game* – an alternative to Coquand and Hyland-Ong games up to now only partially explored. However, it was not at all obvious how to formulate a notion of *n-backtracking strategy* in Coquand and Hyland-Ong games as well, which is what they did.

The definition is logically complex and has a geometrical and temporal flavour, but its meaning is simple to grasp. For example, a strategy is 2-backtracking if it is allowed to change its mind about previous mind changes of level 1. That reminds us of some juridical systems: only a judge of second level can change the decision of a judge of level one. Thus a judge of an appeal court can very well revise an ordinary judge’s decision, but two ordinary judges cannot question each other’s sentences. With Coquand-Hyland-Ong pointer notation, a mind change of level 1 is represented by a move pointing backward:



A mind change of level 2 can be depicted as a crossing:



Though they are *not* physically removed, we consider the moves between the ending points of a pointer to be *ideally* erased by the mind change. Thus, the move \circ_1 is a first mind change erasing the moves before itself and after \bullet_1 ; the move \circ_2 is a mind change about the first mind change and recovers one of the formerly retracted moves, \bullet_2 . There is a proviso: for the move \circ_2 to be considered a *real* mind change of level 2, at the moment \circ_2 is played the pointer from \circ_1 must be *active*: a pointer is active if it is crossed only by inactive pointers; a pointer is inactive if there is an active pointer crossing it. By considering mind changes about mind changes about mind changes... one arrives at the concept of *n-backtracking*. A *n-backtracking strategy* has backtracking level $b = n$. We shall have to describe precisely the notion and familiarize extensively with it in the following. For the moment it is enough to say that the backtracking level is a measure of how much involved and intricate the geometrical structure of a strategy is. The backtracking level is also easy to compute: if a strategy interprets a first-order proof or a simply typed lambda term, one can determine its backtracking level in plain quadratic time, and we conjecture linear.

1.4 Backtracking Level as Complexity Measure

We have found the complexity measure we were looking for: it is the backtracking level of a strategy. We shall prove that the height of the tower of exponentials bounding the length of

interactions between strategies is determined by the backtracking level of the strategies. The explanation of this complexity result is our paper's

Fundamental Thesis: *the real engine of a strategy is its backtracking mechanism; the power of this engine is measured by the strategy's backtracking level; the higher the backtracking level, the more involved the geometrical structure of the backtracking mechanism is and the longer it will take to conclude a debate with the strategy.*

As a consequence of our thesis, at first we conjectured a weaker result than the one presented here, which was: if we consider two strategies with backtracking level respectively n and m , $b = \max(m, n)$, and the size of the two strategies is bounded by k , then the complexity bound for the length of any interaction between them is

$$\underbrace{2^{2^{\cdot^{2^{k(\log k)^2}}}}}_{b+1}$$

This estimate already outperforms the old ones, because the backtracking level of a strategy is completely independent of the arena depth. A 2-backtracking strategy can live in an arena of depth 1000; then one has a new bound given by a tower of exponentials of height 3 versus an old one of height 1001. Moreover, as we shall see, the backtracking level of a strategy is bound by the depth of the arena, modulo a trivial transformation. That explains perfectly why the rank of a proof worked as a worst-case complexity measure; in a sense, one looked at the right place, without knowing the good reason.

Our fundamental thesis seemed to suggest that our first conjecture was already optimal. For example, one would expect that when a 5-backtracking strategy interacts with a 1-backtracking one, the debate should take in general no less time than prescribed by a tower of exponential of height 6.

We were struck when we found that the case it's not. The complexity bound in this situation is much tighter: a tower of height 2, that is just an exponential

$$2^{(k \log k)^2}$$

And 1-backtracking strategies are no fortunate exceptions. If a 5-backtracking strategy interacts with a 2-backtracking one, the complexity bound is just a tower of height 3, a double exponential:

$$2^{2^{(k \log k)^2}}$$

More in general, given two strategies of backtracking level respectively n and m and of size less than k , we shall prove that the length of their interaction is determined by the *minimum* among their backtracking levels, that is by $b = \min(m, n)$, and will be less than

$$\underbrace{2^{2^{\cdot^{2^{k(\log k)^2}}}}}_{b+1}$$

We can interpret this result as a *safety of interactions* result. No matter how unpredictable and complex the context in which a strategy is thrown in, the strategy will always have safe interactions, from the complexity point of view, whose rate of growth is determined by the strategy's own backtracking level.

1.5 n -Backtracking, n -Quantifiers Excluded-middle EM_n and Learning

One may acknowledge that indeed the concept of n -backtracking provides the means for a very refined analysis of the length of dialogues in logic, but still have the suspicions that the

analysis is of no practical consequence. Are there around enough 1-backtracking strategies, 2-backtracking strategies and so on, or are they odd entities appearing, like the 29th of February, once in every four years of logical and mathematical practice? Are almost all formulas of complexity n going to have only n -backtracking winning strategies? The answer is no. In Arithmetic the landscape is already partially studied, thanks to the correspondence between the excluded middle EM_n over formulas with n quantifiers and the concept of learning. In [4] it was shown, by means of Interactive realizability [5], that every arithmetical proof using only EM_1 as classical principle represents a winning recursive strategy for the related 1-backtracking game in the sense of [11]. In [11] was proved – in a less refined way – that any arithmetical formula derivable using EM_n as a classical principle, has a winning recursive strategy in n -backtracking games. As these results are very likely to hold also for first-order logic in the formulation of [7], it is clear that strategies of low backtracking level are the norm rather than the exception. Moreover, as a matter of fact many theorems in Arithmetic are proved just by using EM_1 and many theorems in Analysis with EM_2 .

The results of this work also definitely confirm a point that has been made several times [5, 3, 10]: proofs limited to EM_1 or EM_2 yield relatively efficient strategies, and in general the *learning theory* is the correct one to understand and improve the complexity of programs extracted from classical proofs. The surprise, here, is the connection with the complexity of intuitionistic proofs as well.

1.6 Content of the Paper

In this paper we establish the discussed bounds on the interaction lengths between visible bounded strategies in Coquand and Hyland-Ong game semantics. They will certainly lead to new bounds also for cut-elimination, normalization and dialogues between expansion trees in first-order logic; but given the great generality of game semantics, our results are interesting in their own right. For space reasons, we will not study in detail concrete applications and we delay the matter to forthcoming papers. As explained in [12, 13, 17], however, bounds in game semantics automatically translate into bounds for linear head reduction in typed lambda calculus. There is also Herbelin’s result [21], showing a step-by-step correspondence between plays in game semantics and the weak head reduction cut-elimination strategy in an infinitary sequent calculus for Arithmetic. The situation in first-order logic is a bit subtler, because atomic formulas are undecidable, thus winning conditions must be different; yet we do not see major obstacles to adapting Herbelin’s methods also to that setting.

We devote the introductory part of the paper to recalling the game semantical ideas of Coquand, Hyland-Ong and Berardi-de’Liguoro, with new examples and viewpoints that we feel are essential to understand the technical part of the paper. In particular, in Section §2, we describe the main concepts of Coquand and Hyland-Ong game semantics and study an example based on Miller’s expansion trees. In Section §3, we explain in detail the Berardi-de’Liguoro notion of backtracking level of a strategy. In Section §4, we outline our proof plan, which is carried out in detail in Sections §5, §6, §7, §8.

2 Coquand-Hyland-Ong Game Semantics

As Wadler [27] put it, “Curry-Howard is a double-barrelled name which ensures the existence of other double-barrelled names”. Since logical proofs are isomorphic to computer programs, many concepts happen to be discovered twice: once in logic, once in computer science. Girard-Reynolds system F is an example, Coquand-Hyland-Ong game semantics is another. Actually, most literature on game semantics fails to mention that Coquand and Hyland-Ong game semantics are by no means different concepts but just *one* semantics that turns out to

have several applications: cut-elimination in the case of Coquand, denotational semantics for typed lambda calculus in the case of Hyland-Ong. Since the latter has in turn as many variants as functional languages out there, to avoid the *one, none, one hundred thousand* Pirandellian fragmentation, we adopt the most abstract Hyland-Ong terminology (see [13]), calling it Coquand-Hyland-Ong game semantics.

Arenas are the structures where games take place. As in many real-life two-player games, there are a set of moves M and a binary relation of justification \vdash between moves. The elements of M may represent elementary actions of a player, like moving a single piece in chess or choosing a witness in logical dialogues; but they may also encode complex positions, like the entire chessboard or a full logical formula. Often a game can be formalized in both ways. The relation \vdash tells whether a move is a direct answer to another one or whether from a configuration of the game a player can pass to another. Every move is labeled by a function λ as either Player or Opponent move, and there are special initial moves that start the game.

- **Definition 1 (Arenas).** An **arena** \mathbb{A} is a structure (M, \vdash, λ, I) such that:
 - M, \vdash is a graph: M is a set of nodes, called **moves**, and \vdash is a binary relation over M , called **enabling**. If $m \vdash n$, then m is said to **justify** n and n is said to **answer** m .
 - $\lambda : M \rightarrow \{P, O\}$ is a function indicating whether any move m is a Player move or an Opponent move and is such that for every $m, n \in M$, $m \vdash n \implies \lambda(m) \neq \lambda(n)$, which is called the **alternation condition**.
 - I is a subset of M called the set of **initial moves** and is such that for every $m \in I$, $\lambda(m) = O$. We assume $*$ $\in I$.
 - A **simple play** over \mathbb{A} is any sequence of moves $m_1 m_2 m_3 \dots m_i$ such that $m_1 \in I$ and

$$m_1 \vdash m_2 \vdash m_3 \vdash \dots \vdash m_i$$

and represents a backtracking free play.

- \mathbb{A} is said to be of **finite depth** whenever there exists a $k \in \mathbb{N}$ such that no simple play over \mathbb{A} is of length greater than k ; in this case, k is said to be the **depth** of \mathbb{A} .

When we shall want to talk about a player without specifying who he is, we will call him player p, q, \dots and the other player will be \bar{p}, \bar{q}, \dots ; that is, if $p = P$, then $\bar{p} = O$ and if $p = O$, then $\bar{p} = P$. Often moves will be denoted with white \circ or black \bullet circles, that may be indexed; in a play, Player and Opponent moves will always be represented with different colors. Concatenations of sequences s and s' of moves will just be denoted as $s s'$. An **initial segment** of a sequence $s_0 s_1 \dots s_k$ is any prefix $s_0 s_1 \dots s_i$, with $0 \leq i \leq k$.


So far, we have described just a standard game: we have considered only *simple* plays, where players move in an alternating fashion and each player immediately answers the last move of the opponent. A simple play in chess would be: $*1.e4 e5 2.Nf3 Nf6 3.Nxe5$, which means moving in order a white pawn to e4, a black pawn to e5, a white knight on f3, a black knight on f6 and taking the pawn on e5 with the white knight. Graphically, with arrows representing the justification relation \vdash , one has:

$$* \longleftarrow e4 \longleftarrow e5 \longleftarrow Nf3 \longleftarrow Nf6 \longleftarrow Nxe5$$

When interpreting dialogues in logic and computer science, however, it is necessary to allow players to delay answers, take back moves, backtrack to a previous opponent move and answer differently. All of that can be represented by an act of answering a previous move with a pointer. For example, black, instead of answering the move $3.Nxe5$, may change his mind and answer differently the move $2.Nf3$ by placing his knight to c6:



Globally, there are now two plays on the board: *1.e4 e5 2.Nf3 Nf6 3.Nxe5 and *1.e4 e5 2.Nf3 Nc6. The play might go on like this



The plays on the board are now: *1.e4 e5 2.Nf3 Nf6 3.Nxe5 d6 and *1.e4 e5 2.Nf3 Nc6 3.e5. We can see in the example above the two different characters that backtracking moves may possess: the move d6 is a delayed first answer to Nxe5 and represents a mind change about which game line to play; the move Nc6 is a second answer to Nf3 and represents a mind change about how to answer a move that was already answered. As for our purposes it would be useless to take account of that difference, we shall regard each backtracking move as a mind change, without further distinctions.

Now we turn to a crucial concept, that of *view*: a backtracking strategy needs to know what part of the game history to consider when deciding a new move. In any game, what really matters from the perspective of a player is: starting with a move, waiting for an answer, making a move, waiting for an answer, making a move, waiting for an answer and so on... The kind of strategy that arises from the interpretation of proofs and programs, just needs this sequence of move-answer, move-answer, move-answer... in order to decide the next move. The view of white in the last game above is thus: * e4 e5 Nf3 Nf6 Nxe5 d6. In general, to determine the current view, a player looks at the last move of the opponent, then follows the pointer to the move that was answered, then consider the immediately preceding move, then follows the pointer to another move and so on... until an initial move is reached. In the picture



the displayed moves, where \bullet may or may not appear, represent the view of the player with the obligation to play after the last move \bullet_1 .

- **Definition 2 (Plays, Views).** Let $\mathbb{A} = (M, \vdash, \lambda, I)$ be an arena.
 - A finite sequence $s = s_0 s_1 \dots s_n$ of moves of M is said to be a **justified sequence** over \mathbb{A} if s_0 is initial and every move s_i of s , with $i > 0$, is equipped with a pointer to a previous move s_j justifying it: that is, $j < i$ and $s_j \vdash s_i$.
 - A justified sequence $s = s_0 s_1 \dots s_n$ is said to be a **play** over \mathbb{A} if satisfies the alternation condition: for every $0 \leq i < n$, $\lambda(s_i) \neq \lambda(s_{i+1})$.
 - The pairs of moves $e = (\bullet, \circ)$ such that $\bullet \vdash \circ$ are called **edges**; e is said to be the **edge from** \circ and is denoted as $\curvearrowright \circ$; if $\lambda(\circ) = P$, e is a **P-edge**, if $\lambda(\circ) = O$, e is an **O-edge**.
 - The **Player-view** $\ulcorner s \urcorner$ of a justified sequence s is defined recursively as follow:
 - $\ulcorner s \circ \urcorner = \ulcorner s \urcorner \circ$, if $\lambda(\circ) = P$.
 - $\ulcorner s \circ s' \bullet \urcorner = \ulcorner s \urcorner \circ \bullet$, if $\lambda(\bullet) = O$ and \bullet points to \circ .
 - $\ulcorner \bullet \urcorner = \bullet$, if $\bullet \in I$.
 Therefore, whenever $s = \bullet_0 \circ_1 s_1 \bullet_1 \circ_2 s_2 \bullet_2 \dots \circ_k s_k \bullet_k$ and \bullet_0 is an initial move and for each $i \geq 1$, \bullet_i is an Opponent move pointing to \circ_i , then $\ulcorner s \urcorner = \bullet_0 \circ_1 \bullet_1 \circ_2 \bullet_2 \dots \circ_k \bullet_k$.
 - The **Opponent-view** $\lrcorner s \lrcorner$ of a justified sequence s is defined recursively as follow:
 - $\lrcorner s \circ \lrcorner = \lrcorner s \lrcorner \circ$, if $\lambda(\circ) = O$.
 - $\lrcorner s \circ s' \bullet \lrcorner = \lrcorner s \lrcorner \circ \bullet$, if $\lambda(\bullet) = P$ and \bullet points to \circ .

– $\ulcorner s_j = \epsilon$, if s is the empty sequence ϵ .

Therefore, whenever $s = \bullet_0 s_0 \circ_0 \bullet_1 s_1 \circ_1 \dots \bullet_k s_k \circ_k$ and \bullet_0 is an initial move and for each $i \geq 0$, \circ_i is a Player move pointing to \bullet_i , then $\ulcorner s_j = \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_k \circ_k$.

Every play s over an arena \mathbb{A} comes with a set of edges, representing the justification pointers. We say that an edge e is in s whenever e is among the edges of s . As standard in the literature, we will not bother with how, formally, justification pointers are represented; it could be, for example, with a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $f(x) < x$ for every $x > 0$. However, we will just consider plays as sequences of moves with arrows pointing back to a previous move, as in the examples that we consider throughout the paper. In general, when taking subsequences of a play, we will maintain the arrows of the original play. We remark that, according to our definition, in any justified sequence every move different from the first points back to a previous move, like in [15, 22]; in particular, initial moves that occur in the middle of the play must obey the rule. There is no real technical need to enforce this restriction, but it is useful, for it allows to treat uniformly backtracking moves as those pointing before the immediately preceding move. Otherwise, an initial move in the middle of a play that does not point to anything should still be considered a backtracking to the start of the play, but this would not be recorded by an explicit pointer. In some literature, that behaviour is permitted; but one can always put in front of any play a pair of dummy moves, “Let’s begin” and “I’m ready”, the second of which justifies all Opponent moves, so every formerly initial move can point back to it.

The definition of strategy is standard: it is a set of plays, those considered to “follow” the strategy; alternatively, a strategy for a player can be interpreted as a function, mapping any play s ending with an adversary move to some suggested answers. A *bounded* strategy is the kind of strategy that is supposed to interpret a simply typed lambda term or a first-order proof: there is a bound on the length that the view of any play in the strategy can have. The idea is that these strategies always stop playing after a number of moves below a bound which is known in advance. In this paper we shall only be concerned with *visible* strategies: those only needing a view as the history to refer back to with their answers. Visible sequences represent interactions between visible strategies.

■ **Definition 3** (Visibility, Strategies, Interactions). Let $\mathbb{A} = (M, \vdash, \lambda, I)$ be an arena.

- Let $s = s_0 s_1 \dots s_n$ be a play over \mathbb{A} . s is said to be **visible** if for all $0 < i \leq n$, s_i points to a move in $\ulcorner s_0 s_1 \dots s_{i-1} \urcorner$ when $\lambda(s_i) = P$ and s_i points to a move in $\ulcorner s_0 s_1 \dots s_{i-1} \urcorner$ when $\lambda(s_i) = O$.
- A **Player strategy** over \mathbb{A} is any set σ of even-length plays over \mathbb{A} such that if $smn \in \sigma$, then $s \in \sigma$. σ is said to be **visible** if for every $sm \in \sigma$, m points to a move in $\ulcorner s \urcorner$.
- An **Opponent strategy** over \mathbb{A} is any set σ of odd-length plays over \mathbb{A} such that if $smn \in \sigma$, then $s \in \sigma$. σ is said to be **visible** if for every $sm \in \sigma$, with s non empty, m points to a move in $\ulcorner s \urcorner$.
- A visible Player strategy σ is said to be **bounded** by k if $k \in \mathbb{N}$ and for no $s \in \sigma$, $\ulcorner s \urcorner$ is of length greater than k . A visible Opponent strategy σ is said to be **bounded** by k if $k \in \mathbb{N}$ and for no $s \in \sigma$, $\ulcorner s \urcorner$ is of length greater than k .
- Let σ and τ be respectively a strategy for Player and a strategy for Opponent over \mathbb{A} . The set of **interactions** $\sigma \star \tau$ is defined as

$$\sigma \star \tau := \{sm \mid (\lambda(m) = P \implies sm \in \sigma \wedge s \in \tau) \wedge (\lambda(m) = O \implies sm \in \tau \wedge s \in \sigma)\}$$

It represents a match between σ and τ , i.e. a play which follows both the strategies.

The goal of this paper is to study the length of the interactions in $\sigma * \tau$, whenever σ and τ are bounded: by Coquand's termination Theorem [15], in that case, $\sigma * \tau$ contains only finite sequences. But first an example.


2.1 Example: Space, Time and Expansion Trees

Game semantics does more than just describe debates between strategies; it also sheds light on the roles that the very concepts of space and time play in proof representation. But what are space and time in logic?

Conventional wisdom suggests that proofs, on one hand, have a geometrical structure which is important to investigate: after all, a proof is a tree in a two dimensional *space*. On the other hand, conventional wisdom suggests that the temporal process of constructing a proof is as important as its geometrical appearance: proofs pertain also to *time*. Time is no issue: the act itself of drawing an inference takes place in time and to each logical step is possible to associate the moment in which it was made. But where exactly is the geometry of a proof to be sought? Here conventional wisdom is wrong.

Think about a cut-free proof of a prenex formula in a one-sided sequent calculus. The geometrical structure of the proof tree is trivial: it is a linear sequence of \exists and \forall inferences introducing respectively first-order terms and eigenvariables; after those introductions, a propositional tautology is proved, but at that very point the proof becomes no longer interesting! The geometrical structure of a proof does not coincide with its tree structure: it is concealed in expansion trees.

Miller's expansion trees [28] allow to formulate Herbrand's Theorem for what it is: a *game semantical interpretation of proofs*. For each cut-free Tait-style sequent calculus proof of a sequent Γ, A , one can trace how the formula A is used in the proof, keeping note of what first-order terms t the existential quantifiers of A are instantiated to and what eigenvariables a the universal quantifiers of A introduce. The result is an expansion tree. For example, the formula $\exists x \forall y \exists z P(x, y, z)$ may have the following expansion tree:



The representation of a proof as a two dimensional object, like an expansion tree, makes prominent space at the expense of time; the representation of a proof as a linear sequence, like in Hilbert's systems, gives priority to time at the cost of space. How to represent at once space and time?

We first transform a cut-formula of a proof under focus into an expansion tree; this operation extracts the geometrical information hidden in the proof. We then rearrange the expansion tree in linear form, according to the order in which its nodes were introduced in the proof, obtaining a pointer sequence: this operation unfolds the temporal information

coded in the proof. Pointer sequences then solve the problem of packing space and time information into a single object: the geometrical information is preserved by means of backwards pointers, while the temporal information is represented by the linear ordering of nodes! And finally something appears right before our eyes, something very clear, yet something we were completely blind to: the edges of the tree crossing each other. That is all we need to measure a proof's complexity. Let us study an example.

The arena for the Coquand-style backtracking game which corresponds to the formula $\exists x \forall y \exists z P(x, y, z)$ comprises as set of moves and justification relation:


$$M = \{*\} \cup \{x := t \mid t \text{ is a closed term}\} \cup \{y := t \mid t \text{ is a closed term}\} \cup \{z := t \mid t \text{ is a closed term}\}$$

$$\vdash = \{(*, x := t) \mid t \text{ closed}\} \cup \{(x := t_1, y := t_2) \mid t_1, t_2 \text{ closed}\} \cup \{(y := t_1, z := t_2) \mid t_1, t_2 \text{ closed}\}$$

As remarked by Heijltjes [20], the expansion tree above may be seen as a strategy: the term labels t_1, t_2, t_3, t_4 are the witnesses to be played by Player for the corresponding formula of the tree, while the eigenvariable labels a_1, a_2, a_3 are to be replaced by the answers of Opponent. In order for the tree to be a correct specification of a strategy, Player must know in which order to play the moves, because the terms t_i might contain eigenvariables a_j . The tree above is **correct** if there is a total ordering of its edge labels $\ell_1, \ell_2, \dots, \ell_n$ such that:

1. if $i < j$ and ℓ_i and ℓ_j are on the same branch, then ℓ_i comes before ℓ_j in that branch.
2. for every i, j , if ℓ_i is an eigenvariable and ℓ_j a term containing it, $i < j$.

An expansion tree extracted from a cut-free proof is correct by construction [20]. A correct ordering of the edge labels of the tree above could be: $t_1 a_1 t_2 a_2 t_3 a_3 t_4$, which would give the following strategy for Player: play first t_1 for the root formula, then wait for the answer a_1 , then backtrack and play t_2 (which can contain at most a_1 as variable) for the root formula, then wait for the answer a_2 , then backtrack and play t_3 (which can contain at most a_1, a_2 as variables) for the root formula, then wait for the answer a_3 and finally play t_4 (which can contain at most a_1, a_2, a_3 as variables). Formally, the strategy σ for Player is the set of even plays s such that $\ulcorner s \urcorner$ is an initial segment of a play of the following shape:



where $t'_2 = t_2[u_1/a_1]$, $t'_3 = t_3[u_1/a_1 \ u_2/a_2]$, $t'_4 = t_4[u_1/a_1 \ u_2/a_2 \ u_3/a_3]$ and u_1, u_2, u_3 are terms representing the values for a_1, a_2, a_3 selected by Opponent. The definition of σ might seem mysterious at first, but essentially there are no other choices: to determine his next move in a play s , Player must compute $\ulcorner s \urcorner$, which is the part of the history that concerns him, and then look at the corresponding initial segment of the linearized expansion tree. It is immediate to verify that σ is visible and bounded by 8.


A strategy for Opponent, likewise, is given by any expansion tree for the involutive

maximal length of these interactions. In this paper we do not investigate the exact connection between cut-elimination and strategy interaction, the situation in first-order logic being slightly different from the situation in Arithmetic [21].

2.2 A Simple Property of Views

The reader should make sure to be proficient with the concepts of view and visibility: several proofs will require computations on views and checking visibility of plays. To “warm up”, a simple property of views that we shall use very often is the following.

■ **Proposition 4** (Shape of Views). *Suppose $s = r \bullet r' \circ$ is a visible play over an arena \mathbb{A} and there is an edge (\bullet, \circ) in s . Then s is of the form:*



where, for $1 \leq i \leq k$, \bullet_{i+1} immediately precedes \circ_i and there is an edge (\circ_i, \bullet_i) in s .

Proof. Let $\lambda(\circ) = p$. We can write s as

$$s = a \circ_z \overset{\curvearrowright}{\bullet_z} \dots \circ_2 \overset{\curvearrowright}{\bullet_2} \circ_1 \overset{\curvearrowright}{\bullet_1} \circ$$

where, for $1 \leq i \leq z$, \bullet_{i+1} immediately precedes \circ_i and there is an edge (\circ_i, \bullet_i) in s and $a \in I$ or a is the empty sequence. Therefore the p -view of $r \bullet r'$ is

$$w = a \circ_z \bullet_z \dots \circ_2 \bullet_2 \circ_1 \bullet_1$$

Since s is visible, \circ must point to a move in w , so either $\bullet = a$ or $\bullet = \bullet_k$, with $1 \leq k < z$, which is what we wanted to show. ◀

Another well-known and simple-to-prove property of views ([22], pp. 330–332) that we recall is that for every visible play s and player p , the p -view of s is itself a visible play.

3 Berardi-de'Liguoro Backtracking Level

We now introduce the Berardi-de'Liguoro notion [9] of backtracking level of a strategy, which measures in a sophisticated and precise way the complexity of the backtracking moves made by the strategy. We shall have to extend the notion to arbitrary plays between strategies and investigate the resulting properties, thus continuing from where [9] stopped.

Our first goal is to describe the concept of “mind changing about a mind change” or equivalently “retraction of a retraction”, which will be formalized as the binary *inactivation relation* \triangleleft between edges of a play. The idea, here, is that in a play any edge that points to a move back in the past of the play, implicitly erases, *retracts*, all the moves between the two moves connected by the edge. A player act of pointing back to a past move, immediately implies that he is not satisfied with his former reactions to that move. It does not matter whether the reactions were ignoring the move altogether or answering it, the player has now changed his mind: either he has chosen to stop ignoring the move and answer it, or he decided to answer it differently. So, everything that was generated starting from his first former reaction no longer interests him, and he discards that part of the history. To continue the discussion we need the concept of crossing edges.

■ **Definition 5** (Crossing Edges). Let $s = s_0 s_1 \dots s_n$ be a play over an arena \mathbb{A} . We say that the edge (s_{m_1}, s_{n_1}) **crosses** the edge (s_{m_2}, s_{n_2}) , if


$$m_2 < m_1 < n_2 < n_1$$

which is represented in a picture as:



We write $e_2 \curvearrowright e_1$ whenever an edge e_1 crosses an edge e_2 , that is when e_2 is **crossed** by e_1 .

Now, how to interpret a situation in which a player edge is crossed by another edge of the same player? This means that the player no longer agrees with his previous act of discarding the history between the ending points of the first edge. He wants to access something that was erased, he retracts his retraction. It is tempting to say that if an edge is crossed by another edge, then the latter always represents a mind change about the retraction made by the former and thus that the latter *inactivates* the former. This definition would not technically work: we shall see this in the discussion before Proposition 9, which should be already enough of a reason to stop considering the idea. However, such a definition would be just wrong in itself and inconsistent with our discussion. Consider the following situation:




If we assume the second edge $\curvearrowright \circ_2$ inactivates the first edge $\curvearrowright \circ_1$, then the moves between \bullet_1 and \circ_1 may very well be accessible at the moment of the move \circ_3 ; in other words, since the second edge $\curvearrowright \circ_2$ has already retracted the first retraction $\curvearrowright \circ_1$, the third edge $\curvearrowright \circ_3$ retracts nothing at all, even if it crosses the first. The point is that the first edge is *active* at the moment the second edge is played, while it could be *inactive*, i.e. not “existing”, when the third edge is played. An edge is said to be inactive if there is an active edge crossing it: a real later retraction. An edge is said to be active if all edges crossing it are inactive: no real retractions. One edge is inactivated by another one, if the former is active at the very moment in which the latter is created. Since in order to determine whether an edge is active or inactive the future of the play is involved, the Berardi-de'Liguoro Definition 6 of activity and the inactivation relation is by backward induction.

■ **Definition 6** (Active and Inactive Edges of a Play, Inactivation). Let $s = s_0 s_1 \dots s_n$ be a play over an arena \mathbb{A} .

- We define by backward induction, for all $0 \leq i \leq n$, whether the edge from s_i is **active** or **inactive** in s .
 1. $\curvearrowright s_n$ is active.
 2. $\curvearrowright s_i$, with $i < n$, is inactive if there is a $j > i$ such that $\curvearrowright s_i \curvearrowright \curvearrowright s_j$ and $\curvearrowright s_j$ is active.
 3. $\curvearrowright s_i$, with $i < n$, is active if for all $j > i$, if $\curvearrowright s_i \curvearrowright \curvearrowright s_j$, then $\curvearrowright s_j$ is inactive.
- We say that $\curvearrowright s_i$ **inactivates** $\curvearrowright s_j$ **in** s or that $\curvearrowright s_j$ **is inactivated by** $\curvearrowright s_i$ **in** s , and write $s_j \triangleleft s_i$ in s , if $\curvearrowright s_j$ is active in $s_0 s_1 \dots s_{i-1}$ and $\curvearrowright s_j \curvearrowright \curvearrowright s_i$.

The combinatorial power of the definition lies in its logical complexity: the condition determining whether a edge is active or inactive in a play is a formula with possibly as many

alternating quantifiers as the length of the play. We remark that the concept of inactivation is always relative to a specified play and has a temporal connotation; we will always make sure to declare with respect to what play the inactivation relation is considered. As an example of active and inactive edges, let us consider the following visible play:



The edges $\varkappa \bullet_5$ are $\varkappa \circ_4$ active in the whole play, thus the edge $\varkappa \circ_3$ is inactive and the edge $\varkappa \circ_2$ is active in the whole play. Moreover, $\varkappa \circ_3$ is active in the initial segment of the play ending with the move \bullet_4 , therefore $\varkappa \circ_3 \triangleleft \varkappa \circ_4$. However, it is not the case that $\varkappa \circ_3 \triangleleft \varkappa \bullet_5$, for $\varkappa \circ_3$ is inactive in the initial segment of the play ending with the move \circ_4 .

The time has come to define the backtracking level of a move, play and strategy. The idea is to consider the chain of mind changes that originates from a move. A move might be a simple retraction, and thus have backtracking level 1, or it might be a retraction of a retraction, and thus have backtracking level 2, or it might be a retraction of a retraction of a retraction, and thus have backtracking level 3. In general, one measures the length of a maximal chain of such retractions of retractions. More formally, given a move a , if a does not retract anything because it answers the immediately previous move, then it is declared to have backtracking level 0. Otherwise, one defines the backtracking level of a as the length $k + 1$ of a maximal chain of the shape $e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_k \triangleleft \varkappa a$; in particular, if there is no e such that $e \triangleleft a$, the backtracking level of a is 1. The backtracking level of a play is the maximum backtracking level of its moves, while the backtracking level of a bounded strategy is the maximum among the backtracking levels of the views of the plays in it.

■ **Definition 7** (Backtracking level of a Move, Play and Strategy). Let $s = s_0 s_1 \dots s_n$ be a play over an arena \mathbb{A} .

■ We define by induction, for all $0 \leq i \leq n$, the **backtracking level in s of the move s_i** ($\text{Blevel}(s_i)$) as following.

If (s_{i-1}, s_i) is an edge of s or $i = 0$, then $\text{Blevel}(s_i) = 0$.

If (s_{i-1}, s_i) is not an edge of s and $i > 0$, then $\text{Blevel}(s_i) = k + 1$, where k is the length of the longest chain e_1, e_2, \dots, e_k of edges of s such that in s

$$e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_k \triangleleft \varkappa s_i$$

■ We define the **backtracking level of the play s** as:

$$\text{Blevel}(s) = \max\{\text{Blevel}(s_i) \mid 0 \leq i \leq n\}$$

We define the **backtracking level of the player p in the play s** as:

$$p\text{Blevel}(s) = \max\{\text{Blevel}(s_i) \mid 0 \leq i \leq n \wedge \lambda(s_i) = p\}$$

■ For any bounded Player strategy σ , we define its backtracking level as:


$$\text{Blevel}(\sigma) = \max\{\text{Blevel}(\ulcorner s \urcorner) \mid s \in \sigma\}$$

■ For any bounded Opponent strategy σ , we define its backtracking level as:


$$\text{Blevel}(\sigma) = \max\{\text{Blevel}(\lrcorner s \rceil) \mid s \in \sigma\}$$

■ If $\text{Blevel}(\sigma) = n$, we say that σ is a **n -backtracking strategy**.

Let us consider the following example:



The move o_1 is of backtracking level 0, o_2 of backtracking level 1, o_3 of backtracking level 2 and o_4 of backtracking level 3. But in the following play



o_4 is of backtracking level 1, since in the initial segment of the play terminating with \bullet_4 the edge $\varkappa o_3$ is active and thus $\varkappa o_2$ inactive.

4 Ideas of the Proof

We now have all the tools needed first to enter and then survive the agitated waters of complexity analysis. In order not to lose ourselves in the details before reaching the final destination, we trace a map with the principal objectives to achieve and their role in our proof. The lowest layer of our proof strategy is built upon some reflections about two crucial features of Coquand’s proof of debate termination [15], so it is wise to look back at it first.

4.1 Back to Coquand’s proof: Uncrossed Edges and the Trimming Operation

Coquand proved in [15] that in any arena of finite depth debates between well-founded strategies always terminate. The proof’s main idea is to consider the so called *orphan* moves, which are those moves that during the debate under focus are not directly answered by any later move. One can easily prove that in any visible play s , the edges from such moves can never be crossed. If we cut away from s all the moves that are in the interior of an edge from some orphan move, removing also the edge’s ending points, we obtain again a visible play s' . This play takes place by construction in an arena of depth strictly less than the original one, because all the moves of maximal depth are removed from the original play s . Therefore one can argue by induction that s' is of finite length and then easily calculate the length of the original play s . Interestingly, if we ignore all the other edges of the play, the edges from orphan moves can be seen from outside as 1-backtracking moves, because they never cross each other. Thus if we watch the play flowing from left to right like a sequence of photograms and, as soon as an edge from an orphan move appears, we erase all the moves in its interior, we see a 1-backtracking game [11], for example of the form:

$$o_1; \quad o_1 o_2; \quad o_1 o_2 o_3; \quad o_1 o_2 o_3 o_4; \quad o_1; \quad o_1 o_5; \quad o_1 o_5 o_6; \quad o_1 o_5 o_6 o_7; \quad o_1 o_5; \quad \dots$$

That is, each sequence of moves is either an extension or an initial segment of the previous sequence. These 1-backtracking plays are easy to bound, so the proof goes through.

4.2 The Max-Min Choice Technique

If one thinks carefully about the argument above, one can identify two crucial features:

1. a choice of edges with the property that they cannot be crossed by any other edge;

2. a trimming operation that cuts away some unwanted part of the original play – the moves in the intervals defined by the previously chosen edges – to produce a play in an arena of lower depth.

The problem is that the induction that one winds up to employ is on the depth of the arena. Thus Coquand’s argument yields again the old, well-known tower of exponentials, as high as the depth of the arena. If we want to obtain finer results, we need a more sophisticated induction. The idea is to choose differently the edges to remove and thus make use of a different trimming operation:

1. we choose as edges to remove those from the moves of *maximal* backtracking level among those made by the player with *minimal* backtracking level in the play;
2. the trimming operation cuts away some unwanted part of the original play – those moves contained in the interior of the previously chosen edges – to produce a play in which the previously chosen player has backtracking level strictly less than in the original play.

These concepts inspire the following definition.

- **Definition 8** (Real Backtracking Level of a Play, Principal Player). Let s be a visible play.
 - Let n be the backtracking level of Player in s and m the backtracking level of Opponent in s . We define the **real backtracking level of s** as $\min(n, m)$. If $\min(n, m) = n$, then Player is called the **principal player of s** ; otherwise, Opponent is the **principal player of s** .

According to this terminology, the proof strategy described above identifies the moves by the principal player of a play that have backtracking level equal to the real backtracking level of the play. We speak of *real* backtracking level, because thanks to the deep symmetry between the two players of the game, one can choose to look at a play *with the eyes* of the principal player. If one looks instead at the very same play from the perspective of the other player, one sees an *apparent* backtracking level higher than the real one and thus the play looks, wrongly, more complex than really is.

With this approach, we can make a refined induction on the backtracking level of the principal player of the play, thus obtaining a tower of exponential possibly much lower than the old, well-known one. However, we shall need a lot of technical preparatory study to make our proof strategy succeed. The fact that one can focus just on the principal player of a play is surprising. Actually, it may seem a stroke of luck that such an argument can work, because it faces many obstructions that we are going to describe in a moment. As a matter of fact, we discovered our results by means of a more intuitive geometrical proof; as it was much more difficult to formalize, we discarded it and turned to the present approach.

4.3 Obstacles

The **first obstacle** to address is the following. Given a visible play, consider the edges from the moves of *maximal* backtracking level among those made by the principal player p , that is the one with *minimal* backtracking level in the play: can these edges be crossed? It is not trivial that they do not – let alone that they can assume the role that orphan moves have in Coquand’s proof. Since the other player \bar{p} may have backtracking level greater than that of p , the edge from some \bar{p} -move could in principle cross and inactivate the edge from a p -move of maximal backtracking level *among those of p* . Therefore, we need to prove that this cannot happen: it will be a task of Section §5.

The **second obstacle** to address is: if two strategies interact together, what is the backtracking level of the resulting play? Could it be, for example, that if a Player strategy has backtracking level 2 and an Opponent strategy has backtracking level 5, then in the

play resulting from their interaction Player has backtracking level 5 or even $5 + 2$, due to the interference of the Opponent moves? In Section §5, we rule out this possibility and show that in the resulting play the backtracking level of Player is still 2 and that of Opponent still 5. This is not obvious, because chains of crossing edges of length $5 + 2$ could pop up at any moment in the play.


The **third obstacle** to address is whether the Trimming operation really transforms a visible play, first of all, into a visible play, secondly, into a play in which the principal player of the initial play has backtracking level strictly smaller. In Coquand’s proof the Trimming operation is easily shown to transform visible plays into visible plays; in our proof we shall have to investigate in depth the shape of views to check that it is the case. The study will be carried out in Section §6.

There is a last tricky **fourth obstacle** to address. What if a strategy has backtracking level greater than the depth of the arena? This is a curious situation, but the answer is essentially: such a strategy is *badly defined*. Section §8 will settle the matter: it will be shown that inserting some dummy moves in a strategy, that form a sort of switch-on/switch-off graph-theoretic gadget, one can get a strategy of backtracking level less than or equal to the depth of the arena. The same can be done with an arbitrary play. This provides conclusive evidence that in principle, a strategy should never have backtracking level greater than the depth of the correspondent arena: if it does, it has been constructed in a careless manner.

5 Maximal Edges and The Max-Min Theorem of Backtracking Levels Interaction

In this Section we establish several key properties of the inactivation relation between edges, removing the first two obstacles in our way to the main results. Namely, we shall see that the edges from the moves of maximal backtracking level, among those of one player, cannot be crossed, let alone inactivated. Moreover, we will show that any interaction between two strategies of backtracking level respectively n and m is a play of backtracking level less than or equal to the maximum among n and m and of real backtracking level less than or equal to the minimum among n and m . We shall call this property the Max-Min Theorem of Backtracking Levels Interaction.

We start off by proving that only the very same player that made a still active backtracking can retract it with another backtracking: in other terms, he alone has the power to recover the moves that he himself “erased”. Unless he decides that those moves have to be reconsidered, the other player will never be able to see them again. More precisely, if e is a p -edge inactivated by the edge e' in a visible play s , then e' is a p -edge as well. Moreover, the inactivation relation in s reflects exactly the inactivation relation in the views of s : if two p -edges are in the inactivation relation in s and the second is in the p -view of s , then they are in the inactivation relation also in the p -view of s , and vice versa. These properties are due to the special nature of the inactivation relation and are absolutely essential for the rest of the paper, so we check them carefully. Already the first would not hold with other naive definitions of \triangleleft , namely if \triangleleft had been taken to be equal to the crossing relation \curvearrowright . For example, whenever in some visible play $\curvearrowright \circ_1 \curvearrowleft \curvearrowright \circ_2$, the move \bullet immediately following the move \circ_2 is of opposite polarity and the corresponding player can see in the interior of the first edge and thus point to a move in it, enabling the crossing relation $\curvearrowright \circ_1 \curvearrowleft \curvearrowright \bullet$:



From now on we fix an arena $\mathbb{A} = (M, \vdash, \lambda, I)$ and we shall consider only plays over \mathbb{A} .

■ **Proposition 9** (Preservation of Inactivations). *Suppose s is a visible play, e is a p -edge of s and $e' = (\bullet, \circ)$ is another edge of s .*

1. *If $e \triangleleft e'$ in s , then e' is a p -edge.*
2. *Assume e' is in the p -view w of s . Then $e \triangleleft e'$ in s if and only if $e \triangleleft e'$ in w .*

Proof. Moves by p and \bar{p} will be represented by different colors: for the moment we do not assume that white is for p , nor that black is for p . We need first to establish some facts. By hypothesis $s = s' \bullet s'' \circ s'''$; by Proposition 4, $s'' = \circ_k \dots \bullet_k \circ_{k-1} \dots \bullet_{k-1} \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1$ and

$$s = s' \bullet \overset{e'}{\curvearrowright} \circ_k \overset{\curvearrowright}{\dots} \bullet_k \overset{\curvearrowright}{\circ_{k-1}} \dots \bullet_{k-1} \dots \circ_2 \overset{\curvearrowright}{\dots} \bullet_2 \overset{\curvearrowright}{\circ_1} \dots \bullet_1 \circ s'''$$

We argue, by induction on i , that each edge (\circ_i, \bullet_i) , with $1 \leq i \leq k$ is active in $s' \bullet s''$. For $i = 1$, that holds, since (\circ_1, \bullet_1) is crossed by no edge in $s' \bullet s''$, being it the last edge. So, suppose $i > 1$. Then

$$s = s' \bullet \overset{e'}{\curvearrowright} \circ_k \overset{\curvearrowright}{\dots} \bullet_k \dots \circ_i \overset{\curvearrowright}{\dots} \bullet_i \dots \circ_2 \overset{\curvearrowright}{\dots} \bullet_2 \overset{\curvearrowright}{\circ_1} \dots \bullet_1 \circ s'''$$

By induction hypothesis, for any $j < i$, (\circ_j, \bullet_j) is active in $s' \bullet s''$; therefore every edge from every move a lying between the moves \circ_j and \bullet_j and such that $\curvearrowright a \curvearrowleft (\circ_j, \bullet_j)$, is by Definition 6 inactive in $s' \bullet s''$. Then, the only active edges that could cross (\circ_i, \bullet_i) are those of the form $\curvearrowleft \circ_j$, for some $j < i$; but that cannot happen, because by visibility of s , \circ_j must point to an element in the $\lambda(\circ_j)$ -view v of the play

$$s' \bullet \circ_k \dots \bullet_k \dots \circ_i \dots \bullet_i \dots \circ_{j+1} \dots \bullet_{j+1}$$

and since $v = v' \bullet \circ_k \bullet_k \dots \circ_i \bullet_i \dots \circ_{j+1} \bullet_{j+1}$, it is not the case that $(\circ_i, \bullet_i) \curvearrowleft \curvearrowleft \circ_j$ in s .

1. Now, assume $e \triangleleft e' = (\bullet, \circ)$ in s . By Definition 6 of the inactivation relation, e is active in $s' \bullet s''$ and $e \curvearrowleft e'$. For every move a lying between \circ_i and \bullet_i for some $1 \leq i \leq k$ and such that $\curvearrowleft a \curvearrowleft (\bullet, \circ)$, we deduce that $\curvearrowleft a$ is inactive in $s' \bullet s''$, since $\curvearrowleft a \curvearrowleft (\circ_i, \bullet_i)$ and (\circ_i, \bullet_i) is active in $s' \bullet s''$. Therefore, it must be the case that $e = \curvearrowleft \circ_l$, for some $1 \leq l \leq k$, which establishes that white is the color for p moves and that e' is a p -edge.
2. Let w be the p -view of s . If e' is not a p -edge, it is impossible both that $e \triangleleft e'$ in s , by 1., and that $e \triangleleft e'$ in w , because every \bar{p} move in w points to the immediately preceding move; thus the thesis is trivial. We can hence assume that e' is a p -edge. Since e' is in w , for some sequences r, r' we have

$$w = r \bullet \circ_k \bullet_k \circ_{k-1} \bullet_{k-1} \dots \circ_2 \bullet_2 \circ_1 \bullet_1 \circ r'$$

and for $1 \leq i \leq k$, the edge (\circ_i, \bullet_i) is in w . Let

$$w' = r \bullet \circ_k \bullet_k \circ_{k-1} \bullet_{k-1} \dots \circ_2 \bullet_2 \circ_1 \bullet_1$$

We argue, by induction on i , that each edge $\curvearrowleft \circ_i$, with $1 \leq i \leq k$, is active in w' if it is active in $s' \bullet s''$ and is inactive in w' if it is inactive in $s' \bullet s''$.

For $i = 1$, that holds, since $\varkappa \circ_1$ is not crossed by any edge in w' and is thus active in w' ; moreover, $\varkappa \circ_1$ is active in $s' \bullet s''$, because (\circ_1, \bullet_1) is active in $s' \bullet s''$ and therefore whenever $\varkappa \circ_1 \vDash d$ with d in $s' \bullet s''$, surely $d \neq (\circ_1, \bullet_1)$, so one has $d \vDash (\circ_1, \bullet_1)$ and so d is inactive in $s' \bullet s''$.

So, suppose $i > 1$.

- If $\varkappa \circ_i$ is active in $s' \bullet s''$, all the edges $\varkappa \circ_j$, such that $j < i$ and $\varkappa \circ_i \vDash \varkappa \circ_j$, are inactive in $s' \bullet s''$; by induction hypothesis, they are inactive in w' as well, and since there is no other edge crossing $\varkappa \circ_i$ in w' , we have that $\varkappa \circ_i$ is active in w' .
- If $\varkappa \circ_i$ is inactive in $s' \bullet s''$, then there is an active edge d in $s' \bullet s''$ such that $\varkappa \circ_i \vDash d$; but d must be equal to $\varkappa \circ_j$, for some $j < i$, since all the other edges crossing $\varkappa \circ_i$ are inactive, because they are crossed by some active edge (\circ_l, \bullet_l) ; by induction hypothesis, $\varkappa \circ_j$ is active in w' , moreover $\varkappa \circ_i \vDash \varkappa \circ_j$ in w' and so $\varkappa \circ_i$ is inactive in w' .

Finally, let us prove the thesis. Assume $e \triangleleft e'$ in s . By the proof of 1., $e = \varkappa \circ_l$, with $1 \leq l \leq k$, and e is active in $s' \bullet s''$ and thus in w' . Therefore $e \triangleleft e'$ in w . Assume now $e \triangleleft e'$ in w . Then $e = \varkappa \circ_l$, with $1 \leq l \leq k$, and e is active in w and thus in $s' \bullet s''$. Therefore $e \triangleleft e'$ in s .

◀

Proposition 10 below provides a simple way to compute the backtracking level of a play: it is enough to determine the backtracking level of the views of its initial segments. This is possible because the inactivation chains of a visible play are exactly those already contained in the views of the play's initial segments.

■ **Proposition 10** (Preservation of Inactivation Chains). *Let $s = s_0 s_1 \dots s_k$ be a visible play. Let e_1, e_2, \dots, e_n be edges of s , assume $e_n = \varkappa s_i$ is a p -edge, for some player p , and w is the p -view of $s_0 s_1 \dots s_i$. Then*

$$e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_n \text{ in } s \iff e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_n \text{ in } w$$

Proof. \implies) Assume $e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_n$ in s . By definition of view and visibility of s , e_n belongs to the p -view w of $s_0 s_1 \dots s_i$. By repeated application of Proposition 9, $e_{n-1} \triangleleft e_n$ in w , $e_{n-2} \triangleleft e_{n-1}$ in w , ..., $e_1 \triangleleft e_2$ in w , which is the thesis.

\impliedby) As in the other direction.

◀

We prove now the key property of maximal backtracking moves and of the edges from them, which states that if a player p of backtracking level n in some visible play has made a move of backtracking level n , then the edge from that move cannot be crossed by any other edge, let alone inactivated. These are the edges that will play the role that the edges from unanswered moves played in Coquand's termination proof [15].

■ **Proposition 11** (Maximal Backtracking Edges). *Suppose $s = s_0 s_1 \dots s_k$ is a visible play and p a player of backtracking level n in s . Assume $\lambda(s_i) = p$ and $\text{Blevel}(s_i) = n$. Then there is no j such that $\varkappa s_i \vDash \varkappa s_j$ in s .*

Proof. Let us consider, for the sake of contradiction, the smallest j such that $\varkappa s_i \vDash \varkappa s_j$ in s . By choice of j , $\varkappa s_i$ is not crossed by any edge in $s_0 s_1 \dots s_{j-1}$, therefore $\varkappa s_i$ is active in $s_0 s_1 \dots s_{j-1}$ and so $\varkappa s_i \triangleleft \varkappa s_j$ in s . By Proposition 9, $\varkappa s_j$ is a p -edge. Since $\text{Blevel}(s_i) = n$, there is in s an inactivation chain

$$e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_{n-1} \triangleleft \varkappa s_i \triangleleft \varkappa s_j$$

But this means that s_j and thus p are of backtracking level $n + 1$ in s , which contradicts the hypothesis.

◀

The next Theorem is the first major result of this paper: if two strategies respectively of backtracking level n and m play one against the other, they generate a visible play of backtracking level equal to the maximum among n and m . However, the strategy with backtracking level equal to the minimum among n and m will maintain its backtracking level also in the generated play. Therefore, the real backtracking level of the generated play is equal to the minimum among n and m .

■ **Theorem 12** (Max-Min Property of Backtracking Levels Interaction). *Let σ be a Player n -backtracking strategy, τ an Opponent m -backtracking strategy and assume $s \in \sigma \star \tau$. Then s is a visible play of backtracking level less than or equal to $\max(n, m)$ and of real backtracking level less than or equal to $\min(n, m)$.*

Proof. In order to prove that s is visible, suppose that $s = s' a$. Assume $\lambda(a) = P$. Since by Definition 3 of interaction, $s' a \in \sigma$ and $s' \in \tau$, and σ and τ are closed by even and odd prefixes, every prefix of s belongs to σ or τ . Since σ and τ are visible, we are done. If $\lambda(a) = O$, exchange σ and τ in the previous argument.

In order to prove the thesis, it is enough to show that Player and Opponent have backtracking level in s less than or equal to respectively n and m . Suppose that $c = e_1 \triangleleft e_2 \triangleleft \dots \triangleleft e_i$ is an inactivation chain in s , where e_i is a p -edge, for some player p . By Proposition 10, c is an inactivation chain in the p -view w of the initial segment s' of s terminating with the p -move corresponding to e_i . Since $s' \in \sigma$ when $p = P$, and $s' \in \tau$ when $p = O$, by Definition 7, w has backtracking level less than or equal to n if $p = P$ and less or equal than m if $p = O$. We conclude that $i \leq n$ if $p = P$, and $i \leq m$ if $p = O$. ◀

6 The Trimming Operation

In this Section we study the Trimming operation sketched in Section §4, we show that it transforms visible plays into visible plays and we determine the resulting backtracking level, thus removing the third obstacle in our way to the main results.

We will show in Proposition 18 that any visible play of real backtracking level n can be interpreted as a sequence of visible plays of real backtracking level *less* than n . These plays cannot grow infinitely and are compressed and stretched like springs by the player of backtracking level n , who, having a finite amount of energy, at some point must stop. In order to identify such sequence of plays, we want to describe an operation of identifying in any visible play s of real backtracking level n a subsequence which represents a visible play s' of real backtracking level strictly less than n . The idea is to erase whatever has been retracted by moves of backtracking level n : the sequence s' is obtained from s by removing all the moves that are in the interior of an edge from any move of backtracking level n made by the principal player of s . More generally, we define an operation removing every move which is in the interior of an edge from a move of some fixed backtracking level made by a fixed player.

■ **Definition 13** (Trimming of a Play). Suppose $s = s_0 s_1 \dots s_k$ is any visible play, $n \in \mathbb{N}$, $n > 0$ and p a player. Let $\mu = (n, p)$. We define s^μ as the subsequence of s obtained from s by removing all the moves s_l such that for some i, j it holds: a) $j < l < i$ and $\curvearrowright s_i = (s_j, s_i)$; b) $\text{Blevel}(s_i) = n$ and $\lambda(s_i) = p$.

It will be useful to exploit the following simple characterization of the Trimming operation.

■ **Lemma 14.** *Suppose $s = s_0 s_1 \dots s_k$ is any visible play, p a player of backtracking level n in s and $\mu = (n, p)$.*

1. If $\text{Blevel}(s_i) \neq n$ or $\lambda(s_i) \neq p$, then

$$(s_0 s_1 \dots s_i)^\mu = (s_0 s_1 \dots s_{i-1})^\mu s_i$$

2. If $\text{Blevel}(s_i) = n$ and $\lambda(s_i) = p$, assuming $\prec s_i = (s_j, s_i)$, then

$$(s_0 s_1 \dots s_i)^\mu = (s_0 s_1 \dots s_j)^\mu s_i$$

3. Assume

$$s = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 r'$$

and that: r is the empty sequence or $r \in I$; $r' = s_i \dots s_k$ and either \bullet_1 is in s^μ or s_i in s^μ ; for $1 \leq i \leq z$, \bullet_{i+1} immediately precedes \circ_i and there is an edge (\circ_i, \bullet_i) in s . Then

$$s^\mu = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} r''$$

where r'' is a subsequence of r' and $1 \leq i_1 < i_2 \dots < i_m \leq z$.

Proof.

1. By hypothesis $\text{Blevel}(s_i) \neq n$ or $\lambda(s_i) \neq p$. Therefore the moves s_l , with $l < i$, that are in $(s_0 s_1 \dots s_i)^\mu$, are exactly those in $(s_0 s_1 \dots s_{i-1})^\mu$; since s_i is in $(s_0 s_1 \dots s_i)^\mu$, we obtain the thesis.
2. For no l, l' , with $l' < j < l < i$, it is the case that $\text{Blevel}(s_l) = n$ and a p -edge $(s_{l'}, s_l)$ is in s , otherwise $\prec s_l \prec \prec s_i$, which is impossible by Proposition 11. Therefore the moves s_l , with $l \leq j$, that are in $(s_0 s_1 \dots s_i)^\mu$, are exactly those in $(s_0 s_1 \dots s_j)^\mu$. Moreover all the moves s_l , with $j < l < i$ are not in $(s_0 s_1 \dots s_i)^\mu$, which is the thesis.
3. By hypothesis

$$s = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 r'$$

where, for $1 \leq j \leq z$, \bullet_{j+1} immediately precedes \circ_j and there is an edge (\circ_j, \bullet_j) in s . We must show

$$s^\mu = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} r''$$

By hypothesis $r' = s_i \dots s_k$ and either \bullet_1 is in s^μ or $\circ_0 := s_i$ is in s^μ ; therefore, every p move s_l of s , such that $\text{Blevel}(s_l) = n$ and $i < l$, points in s to a move $s_{l'}$, with $i - 1 \leq l'$. What we have to show now is that for every l , with $1 \leq l \leq z$: i) if \bullet_l is in s^μ , then \circ_l is in s^μ as well, and ii) if \bullet_l is not in s^μ , then \circ_l and all the moves between \circ_l and \bullet_l are not in s^μ either. s is represented as:

$$r \circ_z \overset{\curvearrowright}{\dots} \bullet_z \dots \circ_l \overset{\curvearrowleft}{\dots} \bullet_l \dots \circ_2 \overset{\curvearrowright}{\dots} \bullet_2 \circ_1 \overset{\curvearrowleft}{\dots} \bullet_1 r'$$

To show i), assume \bullet_l is in s^μ and, for the sake of contradiction, that \circ_l is not in it. Then there is a p move a in s such that $\text{Blevel}(a) = n$, a points in s to a move on the left of \circ_l , a is on the right of \circ_l but it is not the case that a is on the right of \bullet_l . This means that a lies between \circ_l and \bullet_l . But then $\prec a \prec (\circ_l, \bullet_l)$, impossible by Proposition 11.

To show ii), assume \bullet_l is not in s^μ and, for the sake of contradiction, that either \circ_l or some move between \circ_l and \bullet_l is in it. Then there is a p move a such that $\text{Blevel}(a) = n$, and $(\circ_l, \bullet_l) \prec \prec a$ or a points to \circ_l . As before, a cannot lie between \circ_h and \bullet_h , for any $h < l$, otherwise $\prec a \prec (\circ_h, \bullet_h)$, impossible by Proposition 11. We have already showed that a cannot lie on the right of the first move of r' . Therefore $a = \circ_h$, for some $0 \leq h < l$. But by visibility of s , this is impossible: \circ_h should point to one of the moves $\bullet_1, \bullet_2, \dots, \bullet_z, r$.

◀

The Trimming operation takes a visible play s and returns a sequence of moves s^μ . It is not evident that s^μ is a visible play, for we have not yet studied in detail the shape of the views in s^μ , nor it is immediate to assess what is the backtracking level of s^μ . This study will be carried out in Proposition 15 below, where we establish that views in s^μ are nothing but subsequences of views in s . Moreover, if $\mu = (n, p)$ and p has backtracking level n in s , then p has backtracking level less than or equal to $n - 1$ in s^μ . Thus s^μ is exactly the subsequence of s we were looking for: a visible play where p plays only moves with backtracking level strictly less than n .

■ **Proposition 15** (Properties of Trimming). *Suppose $s = s_0 s_1 \dots s_k$ is any visible play, p a player of backtracking level $n > 0$ in s and $\mu = (n, p)$. Then:*

1. *For any player q , the q -view of any initial segment IM of s^μ is a subsequence of the q -view of the initial segment IS of s that ends with the last move of IM .*
2. *s^μ is a visible play.*
3. *p is in s^μ of backtracking level less than or equal to $n - 1$.*

Proof.

1. By induction on s . Let IM be any initial segment of s^μ and let $s_i = \circ$ its last move. Let us fix also a player q . We have two cases:
 - $\lambda(\circ) = q$. We can write s as

$$s = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 \circ \dots s_k$$

where, for $1 \leq i \leq z$, \bullet_{i+1} immediately precedes \circ_i and there is an edge (\circ_i, \bullet_i) in s and $r \in I$ or r is the empty sequence. By Lemma 14, point 3,

$$s^\mu = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} \circ \dots s_k$$

for some sequence of numbers $1 \leq i_1 < i_2 \dots < i_m \leq z$. Therefore

$$IM = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} \circ$$

$$IS = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 \circ$$

Now, assuming that r' is the q -view of r , the q -view of IM is

$$r' \circ_{i_m} \bullet_{i_m} \dots \circ_{i_2} \bullet_{i_2} \circ_{i_1} \bullet_{i_1} \circ$$

while the q -view of IS is

$$r' \circ_z \bullet_z \dots \circ_2 \bullet_2 \circ_1 \bullet_1 \circ$$

which is what we wanted to show.

- $\lambda(\circ) \neq q$. Letting $\circ_1 = \circ$, we can write s as

$$s = r \bullet_z \dots \circ_z \dots \bullet_2 \dots \circ_2 \bullet_1 \dots \circ_1 \dots s_k$$

where, for $1 \leq i \leq z$, \circ_{i+1} immediately precedes \bullet_i and there is an edge (\bullet_i, \circ_i) in s and $r \in I$ or r is the empty sequence. As \circ_1 is in s^μ , by Lemma 14, point 3, we have

$$s^\mu = r \bullet_{i_m} \dots \circ_{i_m} \dots \bullet_{i_1} \dots \circ_{i_1} \bullet_1 \dots \circ_1 \dots s_k$$

for some sequence of numbers $1 < i_1 < i_2 \dots < i_m \leq z$. Therefore

$$IM = r \bullet_{i_m} \dots \circ_{i_m} \dots \bullet_{i_1} \dots \circ_{i_1} \bullet_1 \dots \circ_1$$

$$IS = r \bullet_z \dots \circ_z \dots \bullet_2 \dots \circ_2 \bullet_1 \dots \circ_1$$

Now, assuming that r' is the q -view of r , the q -view of IM is

$$r' \bullet_{i_m} \circ_{i_m} \dots \bullet_{i_1} \circ_{i_1} \bullet_1 \circ_1$$

while the q -view of IS is

$$r' \bullet_z \circ_z \dots \bullet_2 \circ_2 \bullet_1 \circ_1$$

which is what we wanted to show.

2. First of all, s^μ is a justified sequence: if in s there is an edge (s_j, s_i) and s_i is in s^μ , then s_j must be in s^μ as well, otherwise there would be an index l such that $j < l < i$, $\lambda(s_l) = p$, $\text{Blevel}(s_l) = n$ in s and $\varkappa s_l \vDash (s_j, s_i)$; by Proposition 11, that is impossible. To show that s^μ is visible, assume that \circ is the last move of an initial segment IM of s^μ , with $\lambda(\circ) = q$, and that IS is the initial segment of s ending with the last move of IM . We can write s as

$$s = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 \circ \dots s_k$$

where, for $1 \leq i \leq z$, \bullet_{i+1} immediately precedes \circ_i and there is an edge (\circ_i, \bullet_i) in s and $r \in I$ or r is the empty sequence. By Lemma 14, point 3.,

$$s^\mu = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} \circ \dots s_k$$

for some sequence of numbers $1 \leq i_1 < i_2 \dots < i_m \leq z$. Therefore

$$IM = r \circ_{i_m} \dots \bullet_{i_m} \dots \circ_{i_2} \dots \bullet_{i_2} \circ_{i_1} \dots \bullet_{i_1} \circ$$

$$IS = r \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 \circ$$

Now, assuming that r' is the q -view of r , by point 1., the q -view of IM is


$$r' \circ_{i_m} \bullet_{i_m} \dots \circ_{i_2} \bullet_{i_2} \circ_{i_1} \bullet_{i_1} \circ$$

while the q -view of IS is


$$r' \circ_z \bullet_z \dots \circ_2 \bullet_2 \circ_1 \bullet_1 \circ$$

We have to show that \circ points in IM to r or some \bullet_{i_n} , with $1 \leq n \leq m$. Since s is visible, \circ points in IS to r or some \bullet_l , with $1 \leq l \leq z$. But since IM is justified, \bullet_l is in IM , and thus $l = i_n$, with $1 \leq n \leq z$.

3. We first show that for every edge $e = (s_j, s_i)$ of s^μ , e is active in an initial segment r' of s^μ if and only if e is active in the initial segment r of s ending with the last move of r' . We proceed by induction on $k - i$. If $i = k$, then $e = (s_j, s_k)$, $r' = s^\mu$ and $r = s$, so e must be active both in r' and r , for it is the last edge of both plays. If $i < k$, we prove separately the two directions of the implication.
 - For one direction, suppose that e is active in r' : we want to show that e is active also in r , and thus that if $e' = (s_{j'}, s_{i'})$ is an edge of r such that $e \vDash e'$, then e' is inactive in r . Now, if e' is in r' , then e' is inactive in r' , and by induction hypothesis e' is inactive also in r . So suppose e' is not in r' . Then, there is a p -edge $e'' = (s_{j''}, s_{i''})$ in s such that $\text{Blevel}(s_{i''}) = n$ in s and $j'' < i' < i''$. A picture may be useful to summarize how s looks like:



We are first going to determine where $s_{i''}$ points to, that is, where $s_{j''}$ lies. Since s_i is in r' , it holds $i \leq j''$ and we already know that $j'' < i'$; therefore $e' \prec e''$. Hence s is of the shape



Finally, let s_l be the last move of r . By hypothesis, s_l is also in r' . Moreover, since $i' \leq l$, it is not the case that $l < i''$: otherwise, s_l would not be in r' . We conclude that e'' is in r and, since by Proposition 11 e'' is active in r , by definition e' is inactive in r .

- For the other direction, suppose that e is active in r : we want to show that e is active also in r' , and thus that if e' is an edge of r' such that $e \prec e'$, e' is inactive in r' . Indeed, since e' is in r' , it is as well in r ; hence, e' is inactive in r . By induction hypothesis, e' is inactive in r' as well.

Suppose now that $e \triangleleft e'$ in s^μ ; we want to show that $e \triangleleft e'$ in s as well. Indeed, let us consider the shortest initial segment r' of s^μ that contains e and e' ; let r be the initial segment of s ending with the last move of r' . By Proposition 9, $e = (\bullet, \circ)$ and $e' = (\bullet', \circ')$, with $\lambda(\circ) = \lambda(\circ')$. Then for some move \bullet'' , we have

$$r = \dots \bullet \dots \bullet' \dots \circ \dots \bullet'' \circ'$$

and we claim that r' must be of the shape

$$r' = \dots \bullet \dots \bullet' \dots \circ \dots \bullet'' \circ'$$

possibly with fewer moves than r : we have just to check that \bullet'' must be in r' . Indeed, we have $\lambda(\circ') \neq p$ or $\text{Blevel}(\circ') \neq n$ in s , otherwise \circ would not be in r' ; moreover, no move in s different from \circ' can point to the left of \bullet'' without pointing to the left of \circ' , which implies that if \bullet'' was not in r' , also \circ' would not be in r' , a contradiction. Since $e \triangleleft e'$ in s^μ , e is active in the sequence obtained from r' by removing the last move: $\dots \bullet \dots \bullet' \dots \circ \dots \bullet''$; therefore, by what we have just proved, e is active in the sequence obtained from r by removing the last move: $\dots \bullet \dots \bullet' \dots \circ \dots \bullet''$. We conclude that $e \triangleleft e'$ in r and thus in s . We are finally able to prove the thesis: suppose e is a p -edge from a move of backtracking level $n-1$ in s^μ ; we must show that for no edge e' , $e \triangleleft e'$ in s^μ . Indeed, suppose $e \triangleleft e'$ in s^μ . Then, there is in s^μ an inactivation chain $c = e_1 \triangleleft \dots \triangleleft e_{n-2} \triangleleft e \triangleleft e'$. We have proved that c must be an inactivation chain also in s . Therefore, by Proposition 9, e' is a p -edge from a move of backtracking level n in s ; since $e \prec e'$, e should not be in s^μ , contradicting our hypothesis. ◀

7 Upper Bounds on the Length of Interactions

In this Section, we establish our main results about the length of visible plays and interactions between bounded strategies.

We first introduce another numerical measure $\text{comp}(s, i)$ of how complex is a move s_i in a play s , which is very different from the backtracking level of s_i : it represents how advanced

the move is in the current view. Given this association between moves of a play and numbers, we can map each visible play s into a sequence of numbers, which represents its complexity. The complexity $\text{comp}_\mu(s)$ of s^μ is instead defined as the restriction of the complexity of s to s^μ : in other terms, $\text{comp}_\mu(s)$ is the sequence of the complexities that the moves of s^μ have in s . The lexicographic ordering $<$ of finite sequences of numbers is then used to order the complexities of plays.

■ **Definition 16** (Complexity of Moves and Plays, Lexicographic Ordering $<$). Let $s = s_0 s_1 \dots s_k$ be a visible play.

- Assume $0 \leq i \leq k$. We define the complexity of s_i in s , denoted with $\text{comp}(s, i)$, as the length of $\uparrow s_0 s_1 \dots s_i$ if $\lambda(s_i) = P$, as the length of $\downarrow s_0 s_1 \dots s_i$ if $\lambda(s_i) = O$.
- Let $n \in \mathbb{N}$, p be a player and $\mu = (n, p)$. Suppose $s^\mu = s_{i_0} s_{i_1} \dots s_{i_m}$. We define

$$\text{comp}_\mu(s) = \text{comp}(s, i_0) \text{comp}(s, i_1) \dots \text{comp}(s, i_m)$$

- Let $s = a_0 a_1 \dots a_n$ and $r = b_0 b_1 \dots b_m$ be two sequences of natural numbers. We define that $s < r$ if and only if either

$$n < m \wedge \forall i. 0 \leq i \leq n \implies a_i = b_i$$

or

$$\exists k. 0 \leq k \leq n \wedge (\forall i. 0 \leq i < k \implies a_i = b_i) \wedge a_k < b_k$$

The next Proposition 17 is simple, but is really crucial: it says that if a player answers an adversary move, this answer is more complex than the player's first reaction to that move and than any of his other previous non-retracted answers to that move. Intuitively, changing one's mind and backtracking costs, and this cost prevents the backtracking during a play between bounded strategies to be unlimited. This phenomenon was also noticed and exploited by Coquand ([15], pp. 332).

■ **Proposition 17.** Let $s = s_0 s_1 \dots s_k$ be a visible play and (s_j, s_i) is a p -edge of s , with $i \neq j + 1$. Then:

1. $\text{comp}(s, j + 1) < \text{comp}(s, i)$.
2. Suppose that $j < l < i$, that (s_j, s_l) is an edge of s and for no edge e of s , $(s_j, s_l) \curvearrowright e$. Then $\text{comp}(s, l) < \text{comp}(s, i)$.

Proof. By Proposition 4

$$s_0 \dots s_i = s_0 \dots s_j \circ_z \dots \bullet_z \dots \circ_2 \dots \bullet_2 \circ_1 \dots \bullet_1 s_i$$

where, for $1 \leq n \leq z$, \bullet_{n+1} immediately precedes \circ_n and there is an edge (\circ_n, \bullet_n) in s . Since (s_j, s_l) is crossed by no edge of s , it must be the case that $s_l = \circ_h$, with $1 \leq h \leq z$. Thus to prove 1. and 2. simultaneously, it is enough to assume $s_n = \circ_h$, with $1 \leq h \leq z$, and show $\text{comp}(s, n) < \text{comp}(s, i)$. Indeed, let v be the p -view of $s_0 \dots s_j$; s_i and \circ_h are p moves, so the p -view of $s_0 \dots s_i$ is

$$v \circ_z \bullet_z \dots \circ_{h+1} \bullet_{h+1} \circ_h \bullet_h \dots \circ_2 \bullet_2 \circ_1 \bullet_1 s_i$$

and the p -view of $s_0 \dots \circ_h$ is

$$v \circ_z \bullet_z \dots \circ_{h+1} \bullet_{h+1} \circ_h$$

As $s_n = \circ_h$, we get $\text{comp}(s, n) < \text{comp}(s, i)$. ◀

We now prove that any visible play $s_0 s_1 \dots s_k$ in which some player p has backtracking level n can be seen as a sequence of visible plays in which p has backtracking level strictly less than n . The plays of the sequence are obtained just by applying the Trimming operation, with $\mu = (n, p)$, from left to right to all initial segments of s . The plays are each less complex than the next one, according to the relation $<$. More precisely, define $r_i = s_0 s_1 \dots s_i$. If we imagine to watch the sequence $(r_0)^\mu, (r_1)^\mu, \dots, (r_k)^\mu$ as a succession of photograms, we would see that for some time the $(r_i)^\mu$ are extended by one element, then they are reduced by removing some final segment, then they are extended again, and so forth. Something like a spring, which is stretched for some time, then suddenly is compressed only to be stretched again and so on: a sort of 1-backtracking game [11], as we observed in Section §4. Moreover, this series of extensions and compressions come with a cost: $\text{comp}_\mu(r_0) < \text{comp}_\mu(r_1) < \dots < \text{comp}_\mu(r_k)$.

■ **Proposition 18** (Spring Property). *Let $s = s_0 s_1 \dots s_k$ be a visible play, p a player of backtracking level $n > 0$ in s and $\mu = (n, p)$. Then, whenever $1 \leq i < k$, it holds*

$$\text{comp}_\mu(s_0 s_1 \dots s_i) < \text{comp}_\mu(s_0 s_1 \dots s_{i+1})$$

Proof. There are two possibilities: either $\text{Blevel}(s_{i+1}) = n$ and $\lambda(s_{i+1}) = p$ or not.

i) Suppose $\text{Blevel}(s_{i+1}) = n$ and $\lambda(s_{i+1}) = p$. Assume $\prec s_{i+1} = (s_j, s_{i+1})$. Then $j \neq i$, since $n > 0$, and by Lemma 14 (point 2. and 1.)

$$(s_0 s_1 \dots s_{i+1})^\mu = (s_0 s_1 \dots s_{j-1})^\mu s_j s_{i+1}$$

Moreover, $\lambda(s_i) = \lambda(s_j) \neq p$ and for every l such that $j < l \leq i$, $\lambda(s_l) = p$ and $\text{Blevel}(s_l) = n$ in s , it is not the case that $\prec s_l = (s_h, s_l)$ with $h < j$, otherwise $\prec s_l \prec \prec s_{i+1}$ and by Proposition 11 it is impossible. Therefore, since $i \neq j + 1$, there is an l such that $j < l < i$

$$(s_0 s_1 \dots s_i)^\mu = (s_0 s_1 \dots s_{j-1})^\mu s_j s_l \dots s_i$$

Now, either $l = j + 1$ or $\lambda(s_l) = p$, $\text{Blevel}(s_l) = n$ in s and (s_j, s_l) is in s (because if we assume $l \neq j + 1$, then s_{j+1} is not in $(s_0 s_1 \dots s_i)^\mu$, thus it must hold that for some h , with $j + 1 < h \leq l$, $\lambda(s_h) = p$, $\text{Blevel}(s_h) = n$ and (s_j, s_h) is in s ; moreover, $h = l$, because s_h must be in $(s_0 s_1 \dots s_i)^\mu$, since $\prec s_h$ is crossed by no edge by Proposition 11); in either case, by Propositions 11 and 17, $\text{comp}(s, l) < \text{comp}(s, i + 1)$, thus we get

$$\begin{aligned} \text{comp}_\mu(s_0 s_1 \dots s_i) &= \text{comp}_\mu(s_0 s_1 \dots s_{j-1}) \text{comp}(s, j) \text{comp}(s, l) \dots \text{comp}(s, i) \\ &< \text{comp}_\mu(s_0 s_1 \dots s_{j-1}) \text{comp}(s, j) \text{comp}(s, i + 1) \\ &= \text{comp}_\mu(s_0 s_1 \dots s_{i+1}) \end{aligned}$$

ii) Suppose $\text{Blevel}(s_{i+1}) \neq n$ or $\lambda(s_{i+1}) \neq p$. Then, by Lemma 14 (point 1.)

$$(s_0 s_1 \dots s_{i+1})^\mu = (s_0 s_1 \dots s_i)^\mu s_{i+1}$$

Therefore

$$\text{comp}_\mu(s_0 s_1 \dots s_i) < \text{comp}_\mu(s_0 s_1 \dots s_i) \text{comp}(s, i + 1) = \text{comp}_\mu(s_0 s_1 \dots s_{i+1})$$

◀

Finally, we are able to prove our first main Theorem, which tells that the length of any visible play of real backtracking level b is bounded by a tower of exponentials of height $b + 1$. More precisely, define the **hyperexponential function** a_n^m , with $a, n, m \in \mathbb{N}$, by: $a_0^m = m$ and $a_{n+1}^m = a_n^m$; then we prove the following statement.

■ **Theorem 19** (Abstract Min-Backtracking-Theorem). *Suppose $s = s_0 s_1 \dots s_l$ is a visible play of real backtracking level b , p is the principal player of s , k is the length of the longest p -view or \bar{p} -view of an initial segment of s . Then*

$$l \leq \underbrace{k^{k^{k^k}}}_{b+1} \leq \underbrace{2^{2^{\dots 2^{k(\log k) \cdot 2}}}}_{b+1} = 2_b^{k(\log k)2}$$

Proof. We prove the first inequality, by induction on b .

- Case $b = 0$. For every p move s_j , with $j > 0$, $\text{Blevel}(s_j) = 0$, so one has by definition $\prec s_j = (s_{j-1}, s_j)$. Therefore, for every i , the \bar{p} -view of $s_0 s_1 \dots s_i$ is $s_0 s_1 \dots s_i$. Since the length of \bar{p} -views is bounded by k , we conclude $l = k$, which is what we wanted to show.
- Case $b > 0$. Let $\mu = (b, p)$. For every i , with $0 \leq i \leq l$, let $r_i = s_0 s_1 \dots s_i$. By Proposition 18

$$\text{comp}_\mu(r_0) < \text{comp}_\mu(r_1) < \text{comp}_\mu(r_2) < \dots < \text{comp}_\mu(r_l)$$

By Proposition 15, for every i , p has in $(r_i)^\mu$ backtracking level less than or equal to $b-1$ and no view of no initial segment of $(r_i)^\mu$ is longer than k , thus by induction hypothesis $(r_i)^\mu$ is no longer than

$$n = \underbrace{k^{k^{k^k}}}_b$$

Moreover, for every $0 \leq m \leq l$, $\text{comp}(s, m) \leq k$. Therefore the set of sequences

$$\{\text{comp}_\mu(r_i) \mid 0 \leq i \leq l\}$$

has cardinality less than or equal to k^n . By Definition 16 of $<$, we conclude that $l \leq k^n$, which is the first inequality we wanted to show.

Now, let us prove by induction on $i \in \mathbb{N}$ that $k_i^k \leq 2_i^{(k+i) \log k}$. If $i = 0$, then $k_i^k = k$ and $2_i^{(k+i) \log k} = k \log k$, and we are done. Moreover, by induction hypothesis and standard inequalities:

$$\begin{aligned} k_{i+1}^k &= k^{k_i^k} = 2^{k_i^k \log k} \leq \\ &\leq 2^{2_i^{(k+i) \log k} \log k} \leq 2^{2_i^{(k+i) \log k} 2_i^{\log k}} \leq 2^{2_i^{(k+i) \log k + \log k}} = 2^{2_i^{(k+i+1) \log k}} = 2_{i+1}^{(k+i+1) \log k} \end{aligned}$$

Since by Proposition 10, $b \leq k$, we conclude

$$l \leq k_b^k \leq 2_b^{(k+b) \log k} \leq 2_b^{(k+k) \log k}$$

◀

We now readily prove our second main Theorem: bounded strategies have only finite interactions, whose length is bounded by a tower of exponentials of height equal to the minimum among the backtracking levels of the strategies plus one.

■ **Theorem 20** (Min-Backtracking-Theorem). *Let σ be a Player n -backtracking strategy bounded by k , τ an Opponent m -backtracking strategy bounded by k and $b = \min(n, m)$. Assume $s \in \sigma \star \tau$ and $s = s_0 s_1 \dots s_l$. Then*

$$l \leq \underbrace{k^{k^{k^k}}}_{b+1} \leq \underbrace{2^{2^{\dots 2^{k(\log k) \cdot 2}}}}_{b+1} = 2_b^{k(\log k)2}$$


Proof. By Theorem 12, s is visible and of real backtracking level b . By Definition 3 of $\sigma \star \tau$, every initial segment s' of s is in σ or τ and since both strategies are bounded by k , $\lceil s' \rceil \leq k$ and $\lfloor s' \rfloor \leq k$. By Theorem 19,

$$l \leq \underbrace{k^{k^{\cdot^{\cdot^{\cdot^k}}}}}_{b+1} \leq \underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2^{k(\log k) \cdot 2}}}}}}}_{b+1}$$

◀

8 Backtracking Level and Arena Depth

In this Section, we study the relationship between backtracking level and arena depth, removing the fourth and last obstacle in our way to our main result. One would expect that all plays and strategies relative to an arena of depth d have backtracking level at most equal to d . But the following example shows that a badly specified strategy may have backtracking level 6 in an arena of depth 5:



For the sake of legibility we have not displayed the edges from the moves $\bullet_3, \bullet_4, \bullet_5$, which are all supposed to point to $*$. We have $\lceil \bullet_2 \triangleleft \lceil \bullet_6 \triangleleft \lceil \bullet_7 \triangleleft \lceil \bullet_8 \triangleleft \lceil \bullet_9 \triangleleft \lceil \bullet_{10}$, therefore \bullet_{10} is of backtracking level 6. Of course, there is nothing special about the number 6: this construction can be generalized to produce strategies with as great a backtracking level as one wishes. The pattern is: the first n black moves point to $*$, the following n black moves point to the move that comes exactly $2(n-1)$ moves before.

This example notwithstanding, we are going to show that, by inserting some dummy moves into any play relative to an arena of depth $d \geq 2$, we can lower the backtracking level of the play to $d-2$.


How? For energy saving reasons, turning off the light when leaving a room is in order. Likewise, to avoid that the backtracking level of a strategy or of a play grows arbitrarily, it is enough to inactivate backtracking edges immediately after they have been played. This idea can be implemented by inserting in any play some administrative moves, whose only purpose is to lower the backtracking level of a selected player, in our case O . We insert a pair of moves $\circledast \circledast$ immediately after every P move. In this way, before any non-starting O move \bullet there will be the move \circledast and immediately after the P move \circ that follows \bullet there will be the move \circledast . The sequence $\circledast \bullet \circ \circledast$ forms what we call an **inactivation gadget**: \circledast “turns on” the edge from \bullet , while \circledast “turns it off”. We would want the move \circledast to simply point back to the nearest occurrence of \circledast . That works when the play is an O -view, but in general, to preserve visibility, \circledast must refer back to an occurrence of \circledast which is placed much before.

■ **Definition 21** (Inactivation Gadget). We add to our fixed arena \mathbb{A} special moves \circledast and \circledast which, for simplicity, are justified by any other move of opposite polarity and $\lambda(\circledast) = P$ and $\lambda(\circledast) = O$. We denote with \mathbb{A}^+ the enlarged arena. Given a play $s = \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_k \circ_k \bullet_{k+1}$, we define

$$s^\nabla = \bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_k \circ_k \circledast \circledast \bullet_{k+1}$$

which is always assumed to be a play such that: the moves \bullet_i and \circ_i for $0 \leq i \leq k+1$ point to the same moves they point to in s ; the first occurrence of \circledast points to \circ_0 ; for every edge

(\bullet_i, \circ_j) in s , with $j > 0$, the occurrence of the move \circledast immediately following \circ_j , points to the occurrence of the move \circledast immediately preceding \bullet_i ; each occurrence of the move \circledast points to the immediately preceding move \circledast . In a picture, when $j > 0$



We shall denote with $\circledast^{(i)}$ and $\circledast^{(j)}$ respectively the i -th occurrence of \circledast and the j -th occurrence of \circledast in s^∇ .

We now characterize the views of s^∇ and prove that s^∇ is a visible play over \mathbb{A}^+ whenever s is a visible play over \mathbb{A} . In particular we show that the operation of O -view commutes with ∇ , which means: $\llcorner s^\nabla \lrcorner = \llcorner s \lrcorner^\nabla$. Moreover, the operation of P -view eliminates all the gadgets: $\lrcorner s^\nabla \lrcorner = \lrcorner s \lrcorner$.

■ **Proposition 22** (Views and Visibility of s^∇). *Let $s = \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_k \circ_k \bullet_{k+1}$ be a visible play over \mathbb{A} .*

1. *Assume $0 \leq j \leq k$ and there are indexes $i_1, j_1, \dots, i_m, j_m$ such that*

$$\llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \lrcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{i_m} \circ_{j_m}$$

Then

$$\llcorner (\bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_j \circ_j) \lrcorner = \bullet_{i_1} \circ_{j_1} \circledast \circledast \bullet_{i_2} \circ_{j_2} \circledast \circledast \dots \bullet_{i_m} \circ_{j_m}$$

$$\llcorner (\bullet_0 \circ_0 \circledast \circledast \dots \bullet_j \circ_j \circledast \circledast \bullet_{j+1}) \lrcorner = \bullet_{i_1} \circ_{j_1} \circledast \circledast \dots \bullet_{i_m} \circ_{j_m} \circledast \circledast \bullet_{j+1}$$

$$\lrcorner (\bullet_0 \circ_0 \circledast \circledast \dots \bullet_j \circ_j \circledast \circledast \bullet_{j+1}) \lrcorner = \lrcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \bullet_{j+1} \lrcorner$$

$$\lrcorner (\bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_j \circ_j \circledast \circledast) \lrcorner = \circledast^{(i_1)} \circledast^{(j_1)} \dots \circledast^{(i_m)} \circledast^{(j_m)}$$

2. s^∇ is a visible play over \mathbb{A}^+ .

Proof.

1. By induction on j . Let

$$r = \bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_j \circ_j$$

and let $\varkappa \circ_j = (\bullet_i, \circ_j)$. Then

$$r = \bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_{i-1} \circ_{i-1} \circledast \circledast \bullet_i \dots \circ_j$$

By definition of O -view

$$\llcorner r \lrcorner = \llcorner (\bullet_0 \circ_0 \circledast \circledast \bullet_1 \circ_1 \circledast \circledast \dots \bullet_{i-1} \circ_{i-1}) \lrcorner \circledast \circledast \bullet_i \circ_j$$

Letting

$$\llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_{i-1} \circ_{i-1} \lrcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{j_n} \circ_{j_n}$$

we have by definition of O -view and induction hypothesis

$$\llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \lrcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{j_n} \circ_{j_n} \bullet_i \circ_j$$

$$\llcorner r \lrcorner = \bullet_{i_1} \circ_{j_1} \circledast \circledast \dots \bullet_{i_n} \circ_{j_n} \circledast \circledast \bullet_i \circ_j$$

which establishes the first equation of the thesis. The second one is established by

$$\begin{aligned} & \llcorner (\bullet_0 \circ_0 \otimes \otimes \dots \bullet_j \circ_j \otimes \otimes \bullet_{j+1}) \llcorner \\ &= \llcorner (\bullet_0 \circ_0 \otimes \otimes \dots \bullet_j \circ_j) \llcorner \otimes \otimes \bullet_{j+1} \\ &= \bullet_{i_1} \circ_{j_1} \otimes \otimes \dots \bullet_{i_m} \circ_{j_m} \otimes \otimes \bullet_{j+1} \end{aligned}$$

We now prove the third equation. Assume $\varkappa \bullet_{j+1} = (\circ_i, \bullet_{j+1})$. Then by induction hypothesis

$$\begin{aligned} & \lrcorner (\bullet_0 \circ_0 \otimes \otimes \dots \bullet_j \circ_j \otimes \otimes \bullet_{j+1}) \lrcorner \\ &= \lrcorner (\bullet_0 \circ_0 \otimes \otimes \dots \otimes \otimes \bullet_i) \lrcorner \circ_i \bullet_{j+1} \\ &= \lrcorner (\bullet_0 \circ_0 \dots \bullet_i) \lrcorner \circ_i \bullet_{j+1} \\ &= \lrcorner (\bullet_0 \circ_0 \dots \bullet_j \circ_j \bullet_{j+1}) \lrcorner \end{aligned}$$

Finally, let us prove the fourth equation. Let

$$r = \bullet_0 \circ_0 \otimes \otimes \bullet_1 \circ_1 \otimes \otimes \dots \bullet_j \circ_j \otimes$$

and let $\varkappa \circ_j = (\bullet_i, \circ_j)$. Then

$$r = \bullet_0 \circ_0 \otimes \otimes \dots \bullet_{i-1} \circ_{i-1} \otimes \otimes^{(i)} \bullet_i \dots \circ_j \otimes^{(j)}$$

and there is an edge $(\otimes^{(i)}, \otimes^{(j)})$ in r . By definition of P -view

$$\lrcorner r \lrcorner = \lrcorner (\bullet_0 \circ_0 \otimes \otimes \dots \bullet_{i-1} \circ_{i-1} \otimes) \lrcorner \otimes^{(i)} \otimes^{(j)}$$

Letting

$$\llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_{i-1} \circ_{i-1} \llcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{j_n} \circ_{j_n}$$

we have by definition of O -view and induction hypothesis

$$\begin{aligned} & \llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \llcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{j_n} \circ_{j_n} \bullet_i \circ_j \\ & \lrcorner r \lrcorner = \otimes^{(i_1)} \otimes^{(j_1)} \dots \otimes^{(i_n)} \otimes^{(j_n)} \otimes^{(i)} \otimes^{(j)} \end{aligned}$$

which is the thesis.

2. By Definition 21

$$s^\nabla = \bullet_0 \circ_0 \otimes \otimes \bullet_1 \circ_1 \otimes \otimes \dots \bullet_k \circ_k \otimes \otimes \bullet_{k+1}$$

Let ra be any initial segment of s^∇ , with r non empty. We must show that a points to a move in $\lrcorner r \lrcorner$ if $\lambda(a) = P$, and that a points to a move in $\llcorner r \llcorner$ if $\lambda(a) = O$. The thesis is trivial when a points to the immediately previous move, thus when $a = \otimes$ or $a = \otimes^{(0)}$. We have other three cases, according to the remaining possible shapes of a .

- If $a = \otimes^{(j)}$, with $j > 0$, then $\lambda(a) = O$ and by Definition 21

$$r = \dots \otimes^{(i)} \bullet_i \dots \circ_j$$

and the edges (\bullet_i, \circ_j) and $(\otimes^{(i)}, \otimes^{(j)})$ are in s^∇ . Since $\llcorner r \llcorner = \dots \otimes^{(i)} \bullet_i \circ_j$, we are done.

- If $a = \circ_0$, then $\lambda(a) = P$, thus $\lrcorner r \lrcorner = \bullet_0$ and we are done. If $a = \circ_{j+1}$, then $\lambda(a) = P$ and by the third equation of point 1., we have

$$\lrcorner r \lrcorner = \lrcorner (\bullet_0 \circ_0 \otimes \otimes \dots \bullet_j \circ_j \otimes \otimes \bullet_{j+1}) \lrcorner = \lrcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \bullet_{j+1} \lrcorner$$

Since s is visible, \circ_{j+1} points in s – and thus in s^∇ – to some \bullet_i in $\lrcorner r \lrcorner$, with $0 \leq i \leq j+1$, and we are done.

- If $a = \bullet_{j+1}$, then $\lambda(a) = O$ and as a consequence of the first equation of point 1., we have

$$\llcorner r_{\lrcorner} = \llcorner (\bullet_0 \circ_0 \otimes \otimes \dots \otimes \bullet_j \circ_j \otimes \otimes)_{\lrcorner} = \bullet_{i_1} \circ_{j_1} \otimes \otimes \dots \otimes \bullet_{j_m} \circ_{j_m} \otimes \otimes$$

with

$$\llcorner \bullet_0 \circ_0 \dots \otimes \bullet_j \circ_j \lrcorner = \bullet_{i_1} \circ_{j_1} \dots \otimes \bullet_{j_m} \circ_{j_m}$$

Since s is visible, \bullet_{j+1} points in s – and thus in s^{∇} – to some \circ_{j_n} , with $1 \leq n \leq m$, and we are done. \blacktriangleleft

Thanks to the results of Section §5, we know that in order to compute the backtracking level of a visible play, it is sufficient to compute the backtracking level of its views. Thus we start by proving that for every O -view s , the transformation mapping s into s^{∇} effectively lowers the backtracking level of O to $d - 2$, assuming the arena is of depth $d \geq 2$.

■ **Lemma 23.** *Let $s = \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_k \circ_k \bullet_{k+1}$ be a visible play over the arena \mathbb{A} and $d \geq 2$ the depth of \mathbb{A} . Assume that for all i , the edge (\bullet_i, \circ_i) is in s . Then the backtracking level of O in s^{∇} is less than or equal to $d - 2$.*


Proof. By Definition 21

$$s^{\nabla} = \bullet_0 \circ_0 \otimes \otimes \bullet_1 \circ_1 \otimes \otimes \dots \otimes \bullet_k \circ_k \otimes \otimes \bullet_{k+1}$$


We first prove that if $0 < i < k + 1$, then $\varkappa \bullet_j$ is inactive for every $j \geq i$ in the following initial segment of s^{∇} :

$$r = \bullet_0 \circ_0 \otimes \dots \otimes \bullet_i \circ_i \otimes \dots \otimes \bullet_j \circ_j \otimes \otimes$$


We argue by induction on $j - i$. By Definition 21, for $i \leq l \leq j$, the sequence $\otimes \bullet_l \circ_l \otimes$, with justification pointers, must have the following shape:



In particular the sequence $\otimes \bullet_j \circ_j \otimes$ has the following shape:



Since the above j -th occurrence of $\varkappa \otimes$ is active in r , we infer that $\varkappa \bullet_j$ is inactive in r . Moreover, for every l such that $i < l < j$, by induction hypothesis $\varkappa \bullet_l$ is inactive in r . Now consider the sequence $\otimes \bullet_i \circ_i \otimes$:



Since every edge e of r such that $\varkappa \otimes \varkappa e$ must be of the form $e = \varkappa \bullet_l$, with $i < l \leq j$ and e is inactive in r , we conclude that $\varkappa \otimes$ is active in r , and finally that $\varkappa \bullet_i$ is inactive in r . From what we have just proved we conclude that for no i, j it is the case that $\varkappa \bullet_i \triangleleft \varkappa \bullet_j$ in s^{∇} . Hence, every maximal inactivation chain c in s^{∇} that determines the backtracking level of O is contained in a chain having one of the two following shapes:

$$\varkappa \bullet_{i_1} \triangleleft \varkappa \otimes^{(i_1)} \triangleleft \varkappa \bullet_{i_2} \triangleleft \varkappa \otimes^{(i_2)} \triangleleft \dots \triangleleft \varkappa \bullet_{i_m} \triangleleft \varkappa \otimes^{(i_m)}$$

$$\blacktriangleleft \circlearrowleft^{(i_1)} \triangleleft \blacktriangleleft \bullet_{i_2} \triangleleft \blacktriangleleft \circlearrowleft^{(i_2)} \triangleleft \dots \triangleleft \blacktriangleleft \bullet_{i_m} \triangleleft \blacktriangleleft \circlearrowleft^{(i_m)}$$

In the first case, there is a play

$$\bullet_{i_1} \leftarrow \circ_{i_1} \leftarrow \bullet_{i_2} \leftarrow \circ_{i_2} \leftarrow \dots \leftarrow \bullet_{i_m} \leftarrow \circ_{i_m}$$

Since \bullet_{i_1} points backward, it cannot be the first move and thus points to a Player move, which in turn must point to an Opponent move; therefore, \bullet_{i_1} must be the last move of a simple play over \mathbb{A} of at least 3 moves. Since the arena depth is d , we conclude $2m \leq d - 2$, which is the thesis. In the second case, there is a play

$$\circ_{i_1} \leftarrow \bullet_{i_2} \leftarrow \circ_{i_2} \leftarrow \dots \leftarrow \bullet_{i_m} \leftarrow \circ_{i_m}$$

Now, $i_1 = 0$ is impossible, for $\circlearrowleft^{(0)}$ points to the immediately preceding move \circ_0 and it cannot be inactivated by any edge. Therefore, \circ_{i_1} points to the immediately preceding Opponent move, which points to a Player move pointing to yet another Opponent move; this means that \circ_{i_1} must be the last move of a simple play over \mathbb{A} of at least 4 moves. We conclude $2m - 1 \leq d - 2$, which is the thesis. \blacktriangleleft

We are now in a position to prove that for any visible play s over \mathbb{A} , the backtracking level of O in s^∇ is less than or equal to the depth d of the arena minus 2.

■ **Theorem 24** (Backtracking Level Normalization). *Let $s = \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_k \circ_k \bullet_{k+1}$ be a visible play over the arena \mathbb{A} and $d \geq 2$ the depth of \mathbb{A} . Then the backtracking level of O in s^∇ is less than or equal to $d - 2$.*

Proof. By Definition 21

$$s^\nabla = \bullet_0 \circ_0 \circlearrowleft \circlearrowright \bullet_1 \circ_1 \circlearrowleft \circlearrowright \dots \bullet_k \circ_k \circlearrowleft \circlearrowright \bullet_{k+1}$$

By Proposition 22, the O -views of the initial segments of s^∇ that we have to consider are of the form

$$\begin{aligned} & \llcorner (\bullet_0 \circ_0 \circlearrowleft \circlearrowright \dots \bullet_j \circ_j \circlearrowleft \circlearrowright \bullet_{j+1}) \lrcorner \\ & = \bullet_{i_1} \circ_{j_1} \circlearrowleft \circlearrowright \dots \bullet_{i_m} \circ_{j_m} \circlearrowleft \circlearrowright \bullet_{j+1} \end{aligned}$$

and

$$\begin{aligned} & \llcorner (\bullet_0 \circ_0 \circlearrowleft \circlearrowright \dots \bullet_j \circ_j \circlearrowleft) \lrcorner \\ & = \bullet_{i_1} \circ_{j_1} \circlearrowleft \circlearrowright \dots \bullet_{i_m} \circ_{j_m} \circlearrowleft \end{aligned}$$

with

$$\llcorner \bullet_0 \circ_0 \bullet_1 \circ_1 \dots \bullet_j \circ_j \lrcorner = \bullet_{i_1} \circ_{j_1} \dots \bullet_{i_m} \circ_{j_m}$$

and thus, for $1 \leq n \leq m$, the edges $(\bullet_{j_n}, \circ_{j_n})$ are in s and so in s^∇ . By Proposition 10, every inactivation chain in s^∇ ending with an O -edge is already contained in the O -view of the initial segment of s^∇ ending with that edge. Moreover, by Lemma 23, every O -view of any initial segment of s^∇ terminating with \bullet_{j+1} or \circlearrowleft is of backtracking level less than or equal to $d - 2$. We conclude that the backtracking level of O in s^∇ is less than or equal to $d - 2$. \blacktriangleleft

We conclude by proving a strengthened version of Theorem 20, which guarantees that the tower of exponentials defining the complexity of strategy interactions is never higher than the arena depth minus 2.

■ **Theorem 25** (Refined-Min-Backtracking-Theorem). *Let σ be a Player n -backtracking strategy bounded by k , τ an Opponent m -backtracking strategy bounded by k , $d \geq 2$ the depth of the arena \mathbb{A} and $b = \min(n, m, d - 2)$. Assume $s \in \sigma \star \tau$ and $s = s_0 s_1 \dots s_l$. Then*

$$l \leq \underbrace{(2k)^{(2k)^{\cdot^{(2k)^{2k}}}}}_{b+1} \leq \underbrace{2^{2^{2^{2k(\log 2k)-2}}}}_{b+1} = 2_b^{2k(\log 2k)-2}$$

Proof. By Theorem 12, s is visible and of real backtracking level b . If $n < d - 2$ or $m < d - 2$, the thesis follows by Theorem 20. Otherwise, by Theorem 24, O is of backtracking level no greater than $d - 2$ in s^∇ . By Definition 3 of $\sigma \star \tau$, every initial segment s' of s is in σ or τ and since both strategies are bounded by k , $\lceil s' \rceil \leq k$ and $\lfloor s' \rfloor \leq k$. From the characterization of views of Proposition 22, it easily follows that for every initial segment s' of s^∇ , it holds $\lceil s' \rceil \leq 2k$ and $\lfloor s' \rfloor \leq 2k$. By Theorem 19 applied to s^∇ , and since s is shorter than s^∇ , we obtain the thesis. ◀

Concluding remarks. Theorem 25 can also be derived in another way. When the minimum among the backtracking levels of the strategies is smaller than the depth of the arena minus 2, one uses Theorem 19. When it is not the case, one can use Coquand's Trimming technique; with little adjustments, our main calculations would work just fine and one would obtain the correct upper bounds. The techniques of this section, however, are sharper: they open the possibility of optimizing strategies. For example, if just a proper initial segment of a strategy has excessive backtracking level, using our techniques one may very well take the overall backtracking level even below the arena depth, thus obtaining better bounds.

Our bounds, moreover, are sharp. This easily follows from Clairambault's lower bounds [13] or any others. Indeed, whenever a strategy debate takes as many steps as computed by a tower of exponentials having height $d - 2$, we know that the minimum among the backtracking levels of the players must be at least $d - 2$, so it can be taken to be exactly $d - 2$.

Acknowledgments. This work was inspired by interesting conversations with Stefan Hetzl and Daniel Weller on their expansion trees with cut [26], which provided my motivation for studying in depth game semantical debates between strategies. I would also like to thank the first for his academic support during my first year at TU Wien.

References

- 1 S. Abramsky, R. Jagaadesan, P. Malacaria, *Full Abstraction for PCF*, Information and Computation, vol. 163, n. 2, pp. 409–470, 2000.
- 2 S. Abramsky, G. McCusker, *Game Semantics*, in H. Schwichtenberg and U. Berger eds, Logic and Computation: Proceedings of the 1997 Marktoberdorf Summer School, Springer-Verlag, 1998.
- 3 F. Aschieri, *Constructive Forcing, CPS Translations and Witness Extraction in Interactive Realizability*, Mathematical Structures in Computer Science, 2015.
- 4 F. Aschieri, *Learning Based Realizability for HA + EM1 and 1-Backtracking Games: Soundness and Completeness*, Annals of Pure and Applied Logic, vol. 164, n. 6, pp. 591–671, 2013.
- 5 F. Aschieri, S. Berardi, *Interactive Learning-Based Realizability for Heyting Arithmetic with EM1*, Logical Methods in Computer Science, vol. 6, n. 3, 2010.
- 6 F. Aschieri, M. Zorzi, *A “Game Semantical” Intuitionistic Realizability Validating Markov's Principle*, Proceedings of TYPES 2013, LIPIcs, vol. 26, pp. 24–44, 2014.

- 7 F. Aschieri, M. Zorzi, *On Natural Deduction in Classical First-Order Logic: Curry-Howard Correspondence, Strong Normalization and Herbrand's Theorem*, Theoretical Computer Science, to appear.
- 8 A. Beckmann, *Exact Bounds for Lengths of Reductions in Typed Lambda-Calculus*, Journal of Symbolic Logic, vol. 66, n. 3, pp. 1277–1285, 2001.
- 9 S. Berardi, U. de'Liguoro, *Toward the Interpretation of non-Constructive Reasoning as non-Monotonic Learning*, Information and Computation, vol. 207, n. 1, pp. 63–81, 2009.
- 10 S. Berardi, U. de'Liguoro, *Knowledge Spaces and the Completeness of Learning Strategies*, Logical Methods in Computer Science, vol. 10, n. 1, 2014.
- 11 S. Berardi, T. Coquand, S. Hayashi, *Games with 1-Backtracking*, Annals of Pure and Applied Logic, vol. 161, n. 10, pp. 1256–1269, 2010.
- 12 P. Clairambault, *Estimation of the Length of Interactions in Arena Game Semantics*, Proceedings of FOSSACS 2011, LNCS, vol. 6604, 2011.
- 13 P. Clairambault, *Bounding Linear Head Reduction and Visible Interaction through Skeletons*, Logical Methods in Computer Science, vol. 11, n. 2, 2015.
- 14 T. Coquand, *A Semantics of Evidence for Classical Arithmetic*, in: G. Huet, C. Jones eds., Proceedings of the Second Workshop on Logical Frameworks 1991, pp. 87–99, 1991.
- 15 T. Coquand, *A Semantics of Evidence for Classical Arithmetic*, Journal of Symbolic Logic, vol. 60, n. 1, pp. 325–337, 1995.
- 16 P.-L. Curien, H. Herbelin, *The Duality of Computation*, Proceedings of ICFP '00, ACM, pp. 233–243, 2000.
- 17 V. Danos, L. Regnier, H. Herbelin, *Game Semantics & Abstract Machines*, Proceedings of LICS 1996, pp. 394–495, 1996.
- 18 W. Felscher, *Dialogues as a Foundation for Intuitionistic Logic*, in D. Gabbay, F. Guenther eds., Handbook of Philosophical Logic, vol. 5, 2002.
- 19 E. M. Gold, *Language Identification in the Limit*, Information and Control, vol. 10, n. 5, pp. 447–474, 1967.
- 20 W. Heijltjes, *Classical Proof Forestry*, Annals of Pure and Applied Logic, vol. 161, n. 11, pp. 1346–1366, 2010.
- 21 H. Herbelin, *Séquents qu'on calcule: de l'interprétation du calcul des séquents comme calcul de lambda-termes et comme calcul de stratégies gagnantes*, Ph.D thesis, Université Paris 7, 1995.
- 22 M. Hyland, L. Ong, *On Full Abstraction for PCF: I, II and III*, Information and Computation, vol. 163, n. 2, pp. 285–408, 2000.
- 23 S. Jain, A. Sharma, *Mind Change Complexity of Learning Logic Programs*, Theoretical Computer Science, vol. 284, n. 1, pp. 143–160, 2002.
- 24 J.-L. Krivine, *Classical Realizability*, in Interactive models of computation and program behavior. Panoramas et synthèses, pp. 197–229, Société Mathématique de France, 2009.
- 25 G. Kreisel, *On Weak Completeness of Intuitionistic Predicate Logic*, Journal of Symbolic Logic, vol. 27, n. 2, pp. 139–158, 1962.
- 26 S. Hetzl, D. Weller, *Expansion Trees with Cut*, preprint, arXiv:1308.0429, 2013.
- 27 P. Wadler, *Propositions as Types*, Communications of the ACM, 2015.
- 28 D. Miller, *A Compact Representation of Proofs*, Studia Logica, vol. 46, n. 4, pp. 347–370, 1987.
- 29 P. Oliva, T. Powell, *A Game-Theoretic Computational Interpretation of Proofs in Classical Analysis*, in Gentzen's Centenary: the Quest for Consistency, Springer, 2015.
- 30 M. Parigot, *Lambda-My-Calculus: An Algorithmic Interpretation of Classical Natural Deduction*. LPAR 1992, LNCS, vol. 624, pp. 190–201, 1992.
- 31 A. Troesltra, H. Schwichtenberg, *Basic Proof Theory*, Tracts in Theoretical Computer Science, vol. 43, CUP, 2000.