

ON THE SYSTEM CL12 OF COMPUTABILITY LOGIC

GIORGI JAPARIDZE

Department of Computing Sciences, Villanova University, 800 Lancaster Avenue, Villanova, PA
19085, USA

URL: <http://www.csc.villanova.edu/~japaridz/>

e-mail address: giorgi.japaridze@villanova.edu

ABSTRACT. *Computability logic* (CoL) is a long-term project for redeveloping logic on the basis of a constructive game semantics, with games seen as abstract models of interactive computational problems. Among the fragments of CoL successfully axiomatized so far is **CL12** — a conservative extension of classical first-order logic, whose language augments that of classical logic with the so called *choice* (“constructive”) sorts of quantifiers and connectives. This system has already found fruitful applications as a logical basis for constructive and complexity-bound versions of Peano arithmetic, such as arithmetics for polynomial time computability, polynomial space computability, and beyond. The present paper introduces a third, indispensable complexity measure for interactive computations termed *amplitude complexity*, and establishes the adequacy (soundness/completeness) of **CL12** and the associated Logical Consequence mechanism with respect to (simultaneously) A amplitude, S space and T time computability under certain minimal conditions on the triples (A, S, T) of function classes. This result very substantially broadens the potential application areas of **CL12**, even when time and/or space complexity is the only concern. It would be sufficient to point out that, for instance, now **CL12** can be reliably used as a logical basis of systems for logarithmic space or exponential time computabilities — something that the earlier-known crude adequacy results for **CL12** were too weak to allow us to do. This paper is self-contained, and targets readers with no prior familiarity with the subject.

CONTENTS

1. Introduction	2
2. Constant games and “propositional” operations	8
3. Not-necessarily-constant games and “quantifier” operations	14
4. Interactive machines	20
5. Interactive complexity	23
6. Language of CL12 and its semantics	26

2012 ACM CCS: [Theory of computation]: Computational complexity and cryptography—Complexity classes; [Theory of computation]: Logic—Constructive mathematics.

2010 Mathematics Subject Classification: primary: 03F50; secondary: 03D75; 03D15; 03D20; 68Q10; 68T27; 68T30.

Key words and phrases: Computability logic; Interactive computation; Implicit computational complexity; Game semantics; Constructive logics; Efficiency logics.

7. Axiomatics of CL12	28
8. Soundness and completeness of CL12	32
8.1. \sqcup -Choose	34
8.2. \sqcap -Choose	36
8.3. \sqcup -Choose	36
8.4. \sqcap -Choose	37
8.5. Replicate	37
8.6. Wait	38
8.7. On the completeness of CL12	39
9. Logical Consequence	39
10. First preservation theorem	40
11. Second preservation theorem	46
11.1. Procedure UPDATE SKETCH	50
11.2. Procedure FETCH SYMBOL	52
11.3. Procedure MAKE HISTORY	52
11.4. The overall strategy and an example of its run	53
11.5. Complexity analysis	56
12. Final remarks	59
References	60
Appendix A. Proof of Lemma 10.1	62
Appendix B. Proof of Lemma 10.2	65
Appendix C. Proof of the completeness of CL12	66
Index	69

1. INTRODUCTION

Computability logic (CoL), introduced in [14, 18, 27], is a semantically conceived open-ended framework and long-term research project for redeveloping logic as a formal theory of *computability*. That is as opposed to the more traditional view of logic as a formal theory of *truth*. The main pursuit of the project is to provide ever more expressive and powerful formal tools for systematically telling what can be computed and how, just like classical logic is a systematic tool for finding what is true.

Computational problems in CoL are understood in the most general, *interactive* sense. Interactive computational problems, in turn, are defined as *games* played by a *machine* (\top) against its *environment* (\perp), with computability meaning existence of a machine, called a *solution*, that always wins the game. As the name “machine” suggests, \top is a player that always follows determined, algorithmic strategies. On the other hand, there are no restrictions on the possible strategies of \perp , which represents a capricious user or the blind forces of nature. Classical propositions and predicates are seen as special sorts of games, called *elementary*. These are moveless games automatically won by \top (and hence lost by \perp) when true, and automatically lost by \top (and hence won by \perp) when false. The approach induces a rich and still-expanding collection of logical operators, standing for various basic operations on games. Those include all operators of classical logic, conservatively generalized from elementary games to all games. This makes CoL a conservative extension of classical

first-order logic. As such, the former is dramatically more expressive than the latter, containing a whole zoo of non-classical operators, some reminiscent of those of intuitionistic or linear logic, and some having no close relatives in the earlier literature. In a series of recent publications [5, 17, 24, 30, 35, 36, 43, 46], the expressive power of CoL has been further lifted to a qualitatively new level through generalizing formulas to so called *cirquents* — graph-style constructs allowing to explicitly account for *sharing* subcomponents between different components.

Since the formalism of CoL is inordinately expressive and, in fact, open-ended, attempts to axiomatize this semantically conceived logic can be reasonably expected to succeed only when focused on various limited yet interesting fragments rather than the whole logic. Recent years have seen rapid and sustained progress in this direction ([4],[15]-[17],[19]-[26],[28, 29, 35, 36, 40, 43]), and this trend is likely to continue in the near future.

Among the fragments of CoL successfully axiomatized so far is **CL12**, to which the present paper is exclusively devoted. This is a sequent calculus system. Every sequent in it looks like

$$E_1, \dots, E_n \circ - F,$$

where E_1, \dots, E_n ($n \geq 0$) and F are formulas. The language in which formulas are written is that of classical first-order logic with equality, function symbols and constants, augmented with the so called *choice* (“constructive”) operators. Namely, formulas are built from atoms in the standard way using the propositional connectives $\neg, \wedge, \vee, \sqcap, \sqcup$ and quantifiers $\forall, \exists, \sqcap, \sqcup$. Below we give a brief intuitive characterization of these logical operators of **CL12** as operations on games. It should be noted that, apparently for the exception of \forall and \exists , the basic ideas of those game operations, in one form or another, had surfaced well before CoL was officially introduced, in studies of game semantics by various authors, such as Lorenzen [39], Hintikka [10], Blass [3], Japaridze [11, 12, 13] and others. Connections with Girard’s [8] linear logic should also be easily noticeable. Since this is not a survey paper, we refer the reader to [14] for discussions of similarities and differences between the above-mentioned approaches and that of ours. It should also be noted that the overall logical vocabulary of CoL is much wider than the above, including operators for various additional natural operations on games, such as sequential operators (conjunction, disjunction, quantifiers and recurrences) studied in [25, 30], toggling operators (conjunction, disjunction, quantifiers and recurrences) studied in [29, 30, 42], parallel quantifiers and recurrences ([26, 27, 30, 32, 40, 45]), various flavors of branching recurrences ([26, 27, 29, 32, 33, 42, 40, 44]), and more. Those additional operators will not be discussed here as they are not relevant to what the present paper is focused on.

Negation \neg can be characterized as a role switch operation: the game $\neg A$ is the same from the point of view of a given player as what A is from the point of view of the other player. That is, the machine’s moves and wins become those of the environment, and vice versa. For instance, if *Chess* is the game of chess¹ as seen by the white player, then $\neg \text{Chess}$ is the same game as seen by the black player.

Next, $A \wedge B$ and $A \vee B$, called *parallel conjunction* and *parallel disjunction*, respectively, are games playing which means playing both A and B simultaneously. In $A \wedge B$, the machine is considered to be the winner if it wins in both components, while in $A \vee B$ winning in just one component is sufficient. In contrast, *choice conjunction* $A \sqcap B$ (resp. *choice disjunction* $A \sqcup B$) is a game where the environment (resp. machine) has to choose, at the

¹Modified so that ties are ruled out.

very beginning, one of the two components, after which the play continues according to the rules of the chosen component, with the failure to make an initial choice resulting in a loss for the corresponding player. To appreciate the difference, compare $\neg Chess \vee Chess$ with $\neg Chess \sqcup Chess$. The former is a two-board game, where the machine plays black on the left board and white on the right board. It is very easily won by the machine by just mimicking on either board the moves made by its adversary on the other board. On the other hand, $\neg Chess \sqcup Chess$ is not at all easy to win. Here the machine has to choose between playing black or white, after which the game continues as the chosen one-board game. Generally, the principle $\neg P \vee P$ is valid in CoL (in the sense of being “always winnable” by a machine) while $\neg P \sqcup P$ is not.

With the set of (canonical) *constants* throughout CoL being the set $\{0, 1, 10, 11, 100, \dots\}$ of binary numerals for natural numbers, the *choice universal quantification* $\sqcap x A(x)$ can now be defined as the infinite choice conjunction $A(0) \sqcap A(1) \sqcap A(10) \sqcap A(11) \sqcap A(100) \sqcap \dots$, and the *choice existential quantification* $\sqcup x A(x)$ defined as the infinite choice disjunction $A(0) \sqcup A(1) \sqcup A(10) \sqcup A(11) \sqcup A(100) \sqcup \dots$. So, for instance, where f is a unary function, $\sqcap x \sqcup y (y = f(x))$ is a game in which the first move is by the environment, consisting in choosing a particular constant m for x . Such a move, which intuitively can be seen as the question “What is the value of f at m ?” by the environment, brings the game down to $\sqcup y (y = f(m))$. Now, in this game/position, the machine is obligated to make a move (otherwise it loses). Such a move should be choosing a constant n for y , which further brings the game down to $n = f(m)$. The latter is an elementary game with no further moves, won by the machine if $n = f(m)$, i.e. if it correctly answered the question asked by the environment. From this explanation one can see that $\sqcap x \sqcup y (y = f(x))$, in fact, represents the problem of computing f , with the machine having an (algorithmic) winning strategy for $\sqcap x \sqcup y (y = f(x))$ iff f is computable in the standard sense. Similarly, where p is a unary predicate, $\sqcap x (p(x) \sqcup \neg p(x))$ is (represents) the problem of deciding p .

Next, in the *blind universal quantification* $\forall x A(x)$ (resp. *blind existential quantification* $\exists x A(x)$), no value for x is specified/chosen by either player. In order to win, the machine needs to play $A(x)$ “blindly” in a way that guarantees a win for every (resp. at least one) possible value of x from the universe of discourse. To compare the blind sorts of quantifiers with their choice counterparts, note that $\sqcap x (Even(x) \sqcup Odd(x))$ is a game easily won by the machine, while, on the other hand, $\forall x (Even(x) \sqcup Odd(x))$ is impossible to win: this is a one-move-deep game where only the machine has a move; such a move should consist in choosing one of the two \sqcup -disjuncts $Even(x)$ or $Odd(x)$; in order to win, the chosen disjunct should be true for every possible value of x , which is an unsatisfiable condition as long as the universe contains both even and odd numbers.

Finally, $A_1, \dots, A_n \circ - B$ can be characterized as the problem of *reducing* B to A_1, \dots, A_n . Several reduction operations emerge naturally within the framework of CoL (cf. [21, 26, 27, 28]), including \rightarrow defined by $A \rightarrow B = \neg A \vee B$. Among those, $\circ -$ stands out as the weakest, most general sort of algorithmic reduction. A play of $A_1, \dots, A_n \circ - B$ proceeds, in a parallel fashion, in all of its components. However, in the antecedental components A_1, \dots, A_n , the roles of the two players are switched. That is, from the machine’s perspective, they are $\neg A_1, \dots, \neg A_n$ rather than A_1, \dots, A_n . The machine is considered the winner if it wins in B as long as its adversary wins in each of the components of the antecedent. This game is thus similar to $\neg A_1 \vee \dots \vee \neg A_n \vee B$, i.e. $A_1 \wedge \dots \wedge A_n \rightarrow B$ (DeMorgan’s laws continue to hold in CoL). There is, however, a crucial difference between $A_1 \wedge \dots \wedge A_n \rightarrow B$ and $A_1, \dots, A_n \circ - B$. Namely, in the former, each of the antecedental games can be played only

once, while, in the latter, they can be played and replayed any number of times at the machine’s discretion. Furthermore, at any time, the machine is allowed to split/fork any already reached position of any A_i and thus create several threads continuing from that position (rather than restart A_i from the very beginning). This way, from the machine’s perspective, A_1, \dots, A_n are computational *resources* that can be used and reused, in the strongest algorithmic sense possible, in the process of playing/solving B . In other words, for the machine, solving $A_1, \dots, A_n \circ - B$ means solving B while the environment providing (interactive) oracles for A_1, \dots, A_n . It is therefore no surprise that $\circ -$ turns out to be a conservative generalization of the well-known concept of Turing reduction from traditional, input-output sorts of problems to all interactive problems. Namely, when A_1, \dots, A_n, B are “traditional” kinds of problems such as computing a function or deciding a predicate, \top has a winning strategy (algorithmic solution) for $A_1, \dots, A_n \circ - B$ if and only if B is Turing reducible to A_1, \dots, A_n .

As promised earlier, the semantical meanings of $\neg, \wedge, \vee, \forall, \exists$ (and hence \rightarrow as well) are exactly classical when these operations are applied to elementary (moveless) games. For instance, when A and B are elementary, then so is $A \wedge B$, which is automatically won by the machine, i.e. true, iff so are both A and B . Furthermore, when all games A_1, \dots, A_n, B are elementary, $A_1, \dots, A_n \circ - B$ can be seen to be equivalent to $A_1 \wedge \dots \wedge A_n \rightarrow B$. It is this fact that eventually makes **CL12** a conservative extension of classical first-order logic.

The system **CL12** was proven in [34] to be sound and complete in the sense that a sequent is **CL12**-provable if and only if it has a uniform (“purely logical”) solution, i.e. algorithmic strategy of \top that wins the game/sequent under any interpretation of its nonlogical components such as predicate and function letters. Furthermore, such a strategy can be effectively extracted from a proof of the sequent. Both soundness and completeness, in fact, were shown in [34] to hold in a significantly stronger sense. Namely, for the completeness part, it was shown that if a sequent S is not provable in **CL12**, then there is simply no strategy — whether algorithmic or non-algorithmic — that wins S under every interpretation. As for the soundness part, it was shown that the strategies extracted from **CL12**-proofs run in polynomial time and polynomial space. What the time and space complexity concepts exactly mean in the context of interactive problems represented by the formulas and sequents of **CL12** will be seen later in Section 5. For now, it would be sufficient to note that those are natural conservative generalizations of the usual complexity-theoretic concepts.

While the above adequacy theorem of [34] establishes the completeness of **CL12** in an extreme — strongest possible — sense, the soundness part, as it turns out, can be significantly sharpened. Among the results of the present paper (Theorem 8.2) is showing that the strategies extracted from **CL12**-proofs, in fact, run in linear time and constant space, essentially meaning that such strategies are as efficient as they could possibly be. Also, a conceptual novelty of the present contribution is introducing (Definition 5.2) a third kind of a complexity measure for interactive computations, termed *amplitude complexity*. The latter is concerned with the sizes of \top ’s moves relative the sizes of the moves made by its adversary. As it happens, in terms of amplitude complexity, **CL12** is again as efficient as it could possibly be. Namely, strategies extracted from **CL12**-proofs run in identity (non-size-increasing) amplitude. Amplitude complexity proves itself to be an indispensable measure when it comes to interactive computation, interesting not only in its own right but also as a means for analyzing time and space complexities at a much finer level than previously possible.

CL12 induces a rule of inference that we call *Logical Consequence*. The latter allows us to jump to conclusion F from premises E_1, \dots, E_n whenever **CL12** proves the sequent $E_1, \dots, E_n \multimap F$. The adequacy of **CL12** extends to the adequacy of this rule. Namely, a formula F is a logical consequence of formulas E_1, \dots, E_n (i.e., the former follows from the latter by Logical Consequence) if and only if a solution for F can be extracted from solutions for E_1, \dots, E_n in a purely logical, i.e. interpretation-independent, way. Furthermore, as shown in [34], whenever Ω is a class of functions containing all polynomial functions and closed under composition, Logical Consequence preserves both Ω time and Ω space computabilities. So, for instance, if F is a logical consequence of E_1, \dots, E_n and, under a given interpretation of nonlogical symbols, each E_i has a polynomial time solution, then so does F , and such a solution for F can be extracted from the solutions for E_1, \dots, E_n . The same holds for space instead of time, as well as for any bigger classes of functions closed under composition, such as the classes of elementary or primitive recursive functions.

The adequacy result established in [34] opened a whole new area of applications for **CL12**: this logic can be used as an appealing alternative to classical or intuitionistic logics as a logical basis for complexity-oriented applied formal theories, such as Peano arithmetic. The papers [28, 31, 37, 38] constructed the series **CLA1-CLA10** of CoL-based — more specifically, **CL12**-based — versions of arithmetic, generically named “*clarithmetics*”. Below we briefly discuss four of those: **CLA4**, **CLA5**, **CLA6** and **CLA7**.

All of the above clarithmetical theories have the same language — the language of classical first-order Peano arithmetic augmented with the choice operators $\sqcap, \sqcup, \sqcap, \sqcup$. They are **CL12**-based in the sense that the sole logical rule of inference² of all those systems is Logical Consequence. The set of nonlogical axioms of all those systems consists of the ordinary Peano axioms (including the induction axiom scheme restricted to the ordinary, $\sqcap, \sqcup, \sqcap, \sqcup$ -free formulas) plus the single extra-Peano axiom $\sqcap x \sqcup y (y = x + 1)$, expressing the computability of the successor function. The only exception is **CLA4**, which has $\sqcap x \sqcup y (y = 2x)$ as an additional extra-Peano axiom (this axiom is derivable/redundant in all other systems). Finally, the only nonlogical rule of inference of all systems is (constructive) induction. It is exactly the induction rule where the systems differ from each other.

The induction rule of **CLA4** is

$$\frac{F(0) \quad \sqcap x (F(x) \rightarrow F(2x)) \quad \sqcap x (F(x) \rightarrow F(2x+1))}{\sqcap x F(x)},$$

where $F(x)$ is any *polynomially bounded* formula, meaning a formula where every subformula $\sqcap y E(y)$ looks like $\sqcap y (|y| \leq t \rightarrow G(y))$ and every subformula $\sqcup y E(y)$ looks like $\sqcup y (|y| \leq t \wedge G(y))$, where t is any $0, ', +, \times$ -combination ($'$ stands for the successor function, i.e., a' is interpreted as $a + 1$) of $|z_1|, \dots, |z_n|$, where z_1, \dots, z_n are any variables different from y and not bound by \forall or \exists within $F(x)$,³ and $|u|$ is a (pseudo)term for the function “the size of the binary representation of x ” (an integer approximation of the base 2 logarithm of u). This is, in fact, nothing but an adaptation of Buss’s [6] PIND axiom to the environment in which **CLA4** operates. Note that, however, the boundedness requirement applies only to the choice quantifiers \sqcap, \sqcup and not the blind quantifiers \forall, \exists . This clear, natural and (as expected) very beneficial separation of constructive and non-constructive operators was metaphorically characterized in [31] as “giving Caesar what belongs to Caesar and God

²This includes logical axioms as special cases of logical rules of inference.

³The condition “not bound by \forall or \exists within $F(x)$ ” is missing on page 1330 of [31], which is a technical error.

what belongs to God”. As shown in [31], **CLA4** is sound and complete with respect to polynomial time computability. Sound in the sense that there is an effective procedure that takes any **CLA4**-proof of any formula F and constructs a polynomial time solution for the problem represented by F under the standard arithmetical meanings of its nonlogical symbols $0, ', +, \times$. And complete in the sense that any arithmetical problem (meaning the problem represented by some formula of the language of **CLA4**) with a polynomial time solution is represented by some theorem of **CLA4**.

The induction rule of **CLA5** is

$$\frac{F(0) \quad \Box x(F(x) \rightarrow F(x+1))}{\Box x F(x)}, \quad (1.1)$$

where $F(x)$, again, is any polynomially bounded formula. As proven in [38], **CLA5** is sound and complete in the same sense as **CLA4**, but with respect to polynomial space (rather than polynomial time) computability.

Next, the induction rule of **CLA6** is the same (1.1), but with the weaker requirement that $F(x)$ be an exponentially (rather than polynomially) bounded formula. Here “*exponentially bounded*” means that every subformula $\Box y E(y)$ of $F(x)$ looks like $\Box y(|y| \leq t \rightarrow G(y))$ and every subformula $\sqcup y E(y)$ looks like $\sqcup y(|y| \leq t \wedge G(y))$, where t is any $0, ', +, \times$ -combination of z_1, \dots, z_n (rather than $|z_1|, \dots, |z_n|$ as earlier), where z_1, \dots, z_n are any variables different from y and not bound by \forall or \exists within $F(x)$. It was proven in [38] that **CLA6** is sound and complete in the same sense as **CLA4**, but with respect to elementary recursive time (which can be seen to coincide with elementary recursive space) computability.

Finally, the induction rule of **CLA7** is also the same (1.1), but with no restrictions on $F(x)$ whatsoever. This system was shown in [38] to be sound and complete in the same sense as **CLA4**, but with respect to primitive recursive time (= space) computability.

Efficiency-oriented systems of clarithmetic, such as **CLA4**, can be seen as programming languages where “programming” simply means theorem-proving. The soundness of the underlying system guarantees that any proof that can be written will be translatable into a program that runs efficiently and indeed is a solution of the problem expressed by the target formula of the proof. Note that the problem of verifying whether a program meets its specification, which is generally undecidable, is fully neutralized here: the “specification” is nothing but the target formula of the proof, and the proof itself, while encoding an efficient program, also automatically serves as a verification of the correctness of that program. Furthermore, every step/formula of the proof can be viewed as its own (best possible) “comment”. In a more ambitious and, at this point, somewhat fictional perspective, after developing reasonable theorem-provers, CoL-based efficiency-oriented systems can be seen as declarative programming languages in an extreme sense, where human “programming” just means writing a formula expressing the problem whose efficient solution is sought for systematic usage in the future. That is, a program simply coincides with its specification. The compiler’s job would be finding a proof (the hard part) and translating it into a machine-language code (the easy part). The process of compiling could thus take long but, once compiled, the program would run fast ever after.

Beginning from the first study of (the system $I\Delta_0$ of) *bounded arithmetic* by Parikh [41] in 1971, numerous complexity-oriented systems have been studied in the literature. A notable advantage of CoL-based systems over the other systems with similar aspirations, which typically happen to be inherently weak systems, is having actually or potentially

unlimited strength, with the latter including the full expressive and deductive power of classical logic and Peano arithmetic. In view of the above-outlined potential applications, the importance of this feature is obvious: the stronger a system, the better the chances that a proof/program will be found for a declarative, non-preprocessed, ad hoc specification of the goal. Among the other appealing features of clarithmetics is being semantically meaningful in the full generality of their languages, scalable, and easy to understand in their own rights. Syntactically they also tend to be remarkably simple, as we had a chance to see from the above description of **CLA4–CLA7**. To achieve adequacy, alternative approaches often need to do a serious amount of pushing and shoving, including extending the language of arithmetic through symbols for functions expressible in the kind old Peano arithmetic but no longer adequately expressible after the latter has been “tampered with”, and introducing a few tens of new axioms to compensate for the loss of deductive power when switching from Peano arithmetic to weaker versions of it. Needless to say that the alternative systems typically understand and deal with computational problems in the narrow sense of functions, while clarithmetics deal with a very general class of interactive computational problems.

Having said the above, the most potentially consequential result of the present paper (Theorems 10.3 and 11.1) is sharpening the earlier-mentioned preservation theorem for Logical Consequence. As we remember, the preservation theorem of [34] states that, as long as Ω is a class of functions closed under composition and containing all polynomial functions, Logical Consequence preserves Ω time and Ω space computabilities. With these conditions, the extent of applicability of **CL12** (of Logical Consequence, that is) as a basis for complexity-oriented systems was far from being fully revealed. For instance, the old form of the preservation theorem did not guarantee the soundness of **CL12**-based theories for sublinear (logarithmic, polylogarithmic etc.) space computabilities, as the class of sublinear functions does not contain all polynomial functions; or, it did not guarantee the soundness of **CL12**-based theories for exponential time (or space) complexity, because the class of exponential functions is not closed under composition; etc. The present paper strengthens the above preservation theorem by removing the “closed under composition” and “containing all polynomial functions” requirements on Ω , replacing them with much weaker and finer conditions, bringing into play amplitude complexity along with time and space complexities (as it turns out, the latter can very substantially depend on the former, so that no advanced studies of interactive complexity can avoid explicitly dealing with amplitude complexity). This new adequacy result for Logical Consequence dramatically broadens the applicability of **CL12** as a basis for complexity-oriented applied theories.

This paper is self-contained and does not assume any prior familiarity with CoL. To that end, Sections 2, 3, 4 (in part), 6 and 7 merely serve the purpose of reintroducing the relevant parts of CoL, and can be more or less safely omitted by readers already well familiar with the subject.⁴

2. CONSTANT GAMES AND “PROPOSITIONAL” OPERATIONS

In this and the subsequent few sections we present definitions of the basic relevant concepts. A reader aspiring to get additional insights is recommended to consult the first 10 sections

⁴It should be acknowledged that over 90% of the contents of those sections are borrowed from some earlier papers on CoL, mostly [34]. Already about 30 articles have been written in the course of developing CoL, and producing 30 original introductions and explanations of the same concepts — or trying to make them *look* original — is neither technically nor ethically feasible.

of [27], which provide a tutorial-style introduction to the subject. It should be however noted that the definitions of game operations given in the present section are different from — yet equivalent to — the “canonical” definitions of the same operations found in [27]. The same applies to some later definitions as well.

As we already know, computational problems in CoL are understood as games between two players: **Machine**, symbolically \top , and **Environment**, symbolically \perp . These names will not always be capitalized, and they may take articles. A **move** means any finite string over the standard keyboard alphabet. A **labeled move (labmove)** is a move prefixed with \top or \perp , with such a prefix (**label**) indicating which player has made the move. A **run** is a (finite or infinite) sequence of labmoves, and a **position** is a finite run.

We will be using the letters Φ, Γ, Δ for runs, and $\alpha, \beta, \gamma, \delta$ for moves. The letter \wp will always be a variable for players, and

$$\neg\wp$$

will mean “ \wp ’s adversary” (“the other player”). Runs will be often delimited by “ \langle ” and “ \rangle ”, with $\langle \rangle$ thus denoting the **empty run**. The meaning of an expression such as $\langle \Phi, \wp\alpha, \Gamma \rangle$ must be clear: this is the result of appending to the position $\langle \Phi \rangle$ the labmove $\langle \wp\alpha \rangle$ and then the run $\langle \Gamma \rangle$.

The following is a formal definition of constant games, combined with some less formal conventions regarding the usage of certain terminology.

Definition 2.1. A **constant game** is a pair $A = (\mathbf{Lr}^A, \mathbf{Wn}^A)$, where:

1. \mathbf{Lr}^A is a set of runs satisfying the condition that a (finite or infinite) run is in \mathbf{Lr}^A iff all of its nonempty finite initial segments are in \mathbf{Lr}^A .⁵ The elements of \mathbf{Lr}^A are said to be **legal runs** of A , and all other runs are said to be **illegal**. We say that α is a **legal move** for \wp in a position Φ of A iff $\langle \Phi, \wp\alpha \rangle \in \mathbf{Lr}^A$; otherwise α is **illegal**. When the last move of the shortest illegal initial segment of Γ is \wp -labeled, we say that Γ is a \wp -**illegal** run of A .

2. \mathbf{Wn}^A is a function that sends every run Γ to one of the players \top or \perp , satisfying the condition that if Γ is a \wp -illegal run of A , then $\mathbf{Wn}^A\langle \Gamma \rangle = \neg\wp$. When $\mathbf{Wn}^A\langle \Gamma \rangle = \wp$, we say that Γ is a \wp -**won** (or **won by** \wp) run of A ; otherwise Γ is **lost** by \wp . Thus, an illegal run is always lost by the player who has made the first illegal move in it.

A constant game A is said to be **elementary** iff $\mathbf{Lr}^A = \{\langle \rangle\}$, i.e., A does not have any nonempty legal runs. There are exactly two elementary constant games, for which we use the same symbols as for the two players. One is \top with $\mathbf{Wn}^\top\langle \rangle = \top$, and the other is \perp with $\mathbf{Wn}^\perp\langle \rangle = \perp$. Standard true sentences, such as “snow is white” or “ $0=0$ ”, are understood as the game \top , and false sentences, such as “snow is black” or “ $0=1$ ”, as the game \perp . Correspondingly, the two games \top and \perp may be referred to as **propositions**.

The operation of **prefixation** takes two arguments: a constant game A and a position Φ that must be a legal position of A (otherwise the operation is undefined), and returns the game $\langle \Phi \rangle A$. Intuitively, $\langle \Phi \rangle A$ is the game playing which means playing A starting (continuing) from position Φ . That is, $\langle \Phi \rangle A$ is the game to which A **evolves** (will be “**brought down**”) after the moves of Φ have been made. Here is a definition:

Definition 2.2. Let A be a constant game and Φ a legal position of A . The game $\langle \Phi \rangle A$ is defined by:

- $\mathbf{Lr}^{\langle \Phi \rangle A} = \{\Gamma \mid \langle \Phi, \Gamma \rangle \in \mathbf{Lr}^A\}$;

⁵This condition implies that the empty run $\langle \rangle$, having no nonempty initial segments, is always among the elements of \mathbf{Lr}^A .

- $\mathbf{Wn}^{\langle\Phi\rangle A}\langle\Gamma\rangle = \mathbf{Wn}^A\langle\Phi, \Gamma\rangle$.

Convention 2.3. A terminological convention important to remember is that we often identify a legal position Φ of a game A with the game $\langle\Phi\rangle A$. So, for instance, we may say that the move 1 by \perp brings the game $B_0 \sqcap B_1$ down to the position B_1 . Strictly speaking, B_1 is not a position but a game, and what *is* a position is $\langle\perp 1\rangle$, which we here identified with the game $B_1 = \langle\perp 1\rangle(B_0 \sqcap B_1)$.

Note that, in order to define the **Lr** component of a constant game A , it is sufficient to specify what the **initial legal (lab)moves** — i.e., the elements of $\{\wp\alpha \mid \langle\wp\alpha\rangle \in \mathbf{Lr}^A\}$ — are, and to what game the game A is brought down after such an initial legal labmove $\wp\alpha$ is made; then, the set of legal runs of A will be uniquely defined (this can be seen to hold even in recursive definitions of game operations as in clauses 1 and 3 of Definition 2.4 below). Similarly, note that defining the **Wn** component only for the legal runs of A would be sufficient, for then it uniquely extends to all runs. With these observations in mind, we define the operations $\neg, \wedge, \vee, \sqcap, \sqcup$ as follows:

Definition 2.4. Let A, B, A_0, A_1, \dots be constant games, and n a positive integer.

1. $\neg A$ (**negation**) is defined by:

(i): $\langle\wp\alpha\rangle \in \mathbf{Lr}^{\neg A}$ iff $\langle\neg\wp\alpha\rangle \in \mathbf{Lr}^A$. Such an initial legal labmove $\wp\alpha$ brings the game down to $\neg(\langle\neg\wp\alpha\rangle A)$.

(ii): Whenever Γ is a legal run of $\neg A$, $\mathbf{Wn}^{\neg A}\langle\Gamma\rangle = \top$ iff $\mathbf{Wn}^A\langle\neg\Gamma\rangle = \perp$. Here

$$\neg\Gamma$$

means Γ with each label \wp changed to $\neg\wp$.

2. $A_0 \sqcap \dots \sqcap A_n$ (**choice conjunction**) is defined by:

(i): $\langle\wp\alpha\rangle \in \mathbf{Lr}^{A_0 \sqcap \dots \sqcap A_n}$ iff $\wp = \perp$ and $\alpha = i \in \{0, \dots, n\}$.⁶ Such an initial legal labmove $\perp i$ brings the game down to A_i .

(ii): Whenever Γ is a legal run of $A_0 \sqcap \dots \sqcap A_n$, $\mathbf{Wn}^{A_0 \sqcap \dots \sqcap A_n}\langle\Gamma\rangle = \perp$ iff Γ looks like $\langle\perp i, \Delta\rangle$ ($i \in \{0, \dots, n\}$) and $\mathbf{Wn}^{A_i}\langle\Delta\rangle = \perp$.

3. $A_0 \wedge \dots \wedge A_n$ (**parallel conjunction**) is defined by:

(i): $\langle\wp\alpha\rangle \in \mathbf{Lr}^{A_0 \wedge \dots \wedge A_n}$ iff $\alpha = i.\beta$, where $i \in \{0, \dots, n\}$ and $\langle\wp\beta\rangle \in \mathbf{Lr}^{A_i}$. Such an initial legal labmove $\wp i.\beta$ brings the game down to

$$A_0 \wedge \dots \wedge A_{i-1} \wedge \langle\wp\beta\rangle A_i \wedge A_{i+1} \wedge \dots \wedge A_n.$$

(ii): Whenever Γ is a legal run of $A_0 \wedge \dots \wedge A_n$, $\mathbf{Wn}^{A_0 \wedge \dots \wedge A_n}\langle\Gamma\rangle = \top$ iff, for each $i \in \{0, \dots, n\}$, $\mathbf{Wn}^{A_i}\langle\Gamma^i\rangle = \top$. Here

$$\Gamma^i.$$

⁶Here and in the other clauses of this definition, the number i is identified with its binary representation. It should be pointed out that the binary representation of zero in the present case is meant to be the symbol “0”. This is as opposed to our later (Section 3) treatment of constants, where the “official” binary representation of zero is agreed to be the empty bit string.

means the result of removing, from Γ , all labmoves except those that look like $\wp i.\alpha$, and then further changing each such (remaining) $\wp i.\alpha$ to $\wp \alpha$.⁷

4. $A_0 \sqcup \dots \sqcup A_n$ (**choice disjunction**) and $A_0 \vee \dots \vee A_n$ (**parallel disjunction**) are defined exactly as $A_0 \sqcap \dots \sqcap A_n$ and $A_0 \wedge \dots \wedge A_n$, respectively, only with “ \top ” and “ \perp ” interchanged.

5. $A \rightarrow B$ (**strict reduction**) is treated as an abbreviation of $(\neg A) \vee B$.

We also agree that, when $k = 1$, $A_1 \sqcap \dots \sqcap A_k$ simply means A_1 , and so do $A_1 \sqcup \dots \sqcup A_k$, $A_1 \wedge \dots \wedge A_k$ and $A_1 \vee \dots \vee A_k$. We further agree that, when the set $\{A_1, \dots, A_k\}$ is empty ($k = 0$, that is), both $A_1 \sqcap \dots \sqcap A_k$ and $A_1 \wedge \dots \wedge A_k$ mean \top , while both $A_1 \sqcup \dots \sqcup A_k$ and $A_1 \vee \dots \vee A_k$ mean \perp .

Example 2.5. The game

$$(0=0 \sqcap 0=1) \rightarrow (10=11 \sqcap 10=10) ,$$

i.e.

$$\neg(0=0 \sqcap 0=1) \vee (10=11 \sqcap 10=10),$$

has thirteen legal runs, which are:

- 1: $\langle \rangle$. It is won by \top , because \top is the winner in the right \vee -disjunct (consequent).
- 2: $\langle \top 0.0 \rangle$. (The labmove of) this run brings the game down to $\neg 0=0 \vee (10=11 \sqcap 10=10)$, and \top is the winner for the same reason as in the previous case.
- 3: $\langle \top 0.1 \rangle$. It brings the game down to $\neg 0=1 \vee (10=11 \sqcap 10=10)$, and \top is the winner because it wins in both \vee -disjuncts.
- 4: $\langle \perp 1.0 \rangle$. It brings the game down to $\neg(0=0 \sqcap 0=1) \vee 10=11$. \top loses as it loses in both \vee -disjuncts.
- 5: $\langle \perp 1.1 \rangle$. It brings the game down to $\neg(0=0 \sqcap 0=1) \vee 10=10$. \top wins as it wins in the right \vee -disjunct.
- 6-7: $\langle \top 0.0, \perp 1.0 \rangle$ and $\langle \perp 1.0, \top 0.0 \rangle$. Both bring the game down to the false $\neg 0=0 \vee 10=11$, and both are therefore lost by \top .
- 8-9: $\langle \top 0.1, \perp 1.0 \rangle$ and $\langle \perp 1.0, \top 0.1 \rangle$. Both bring the game down to the true $\neg 0=1 \vee 10=11$, which makes \top the winner.
- 10-11: $\langle \top 0.0, \perp 1.1 \rangle$ and $\langle \perp 1.1, \top 0.0 \rangle$. Both bring the game down to the true $\neg 0=0 \vee 10=10$, so \top wins.
- 12-13: $\langle \top 0.1, \perp 1.1 \rangle$ and $\langle \perp 1.1, \top 0.1 \rangle$. Both bring the game down to the true $\neg 0=1 \vee 10=10$, so \top wins.

Later we will be using some relaxed informal jargon already established in CoL for describing runs and strategies, referring to moves via the intuitive meanings of their effects on the game. For instance, the initial labmove $\top 0.0$ in a play of the game $p \sqcap q \rightarrow r$ we can characterize as “ \top made the move 0.0”. Remembering the meaning of the prefix “0.” of this move, we may as well say that “ \top made the move 0 in the antecedent”. Further remembering the effect of such a move on the antecedent, we may just as well say “ \top chose p (or the left \sqcap -conjunct) in the antecedent”. We may also say something like “ \top (made the move that) brought the game down to $p \rightarrow r$ ”, or “ \top (made the move that) brought the antecedent down to p ”.

⁷Intuitively, Γ^i is the run played in the A_i component. The present condition thus means that \top wins a \wedge -conjunction of games iff it wins in each conjunct.

To define the operation \downarrow in the style of Definition 2.4, we need some preliminaries. What we call a **tree of games** is a structure defined inductively as an element of the smallest set satisfying the following conditions:

- Every constant game A is a tree of games. The one-element sequence $\langle A \rangle$ is said to be the **yield** of such a tree, and the **address** of A in this tree is the empty bit string.
- Whenever \mathcal{A} is a tree of games with yield $\langle A_1, \dots, A_m \rangle$ and \mathcal{B} is a tree of games with yield $\langle B_1, \dots, B_n \rangle$, the pair $\mathcal{A} \circ \mathcal{B}$ is a tree of games with **yield** $\langle A_1, \dots, A_m, B_1, \dots, B_n \rangle$. The **address** of each A_i in this tree is $0w$, where w is the address of A_i in \mathcal{A} . Similarly, the address of each B_i is $1w$, where w is the address of B_i in \mathcal{B} .

Example 2.6. Where A, B, C, D are constant games, $(A \circ B) \circ (C \circ (A \circ D))$ is a tree of games with yield $\langle A, B, C, A, D \rangle$. The address of the first A of the yield, to which we may as well refer as the first **leaf** of the tree, is 00 ; the address of the second leaf B is 01 ; the address of the third leaf C is 10 ; the address of the fourth leaf A is 110 ; and the address of the fifth leaf D is 111 .

Note that \circ is not an operation on games, but just a symbol used instead of the more common comma to separate the two parts of a pair. And a tree of games itself is not a game, but a collection of games arranged into a certain structure, just as a sequence of games is not a game but a collection of games arranged as a list.

For bit strings u and w , we will write $u \preceq w$ to indicate that u is a (not necessarily proper) **prefix** (initial segment) of w .

Definition 2.7. Let A_1, \dots, A_n ($n \geq 1$) be constant games, and \mathcal{T} be a tree of games with yield $\langle A_1, \dots, A_n \rangle$. Let w_1, \dots, w_n be the addresses of A_1, \dots, A_n in \mathcal{T} , respectively. The game $\downarrow \mathcal{T}$ (the **branching recurrence** of \mathcal{T}) is defined by:

(i): $\langle \wp \alpha \rangle \in \mathbf{Lr}^{\delta \mathcal{T}}$ iff one of the following conditions is satisfied:

- (1) $\wp \alpha = \wp u. \beta$, where $u \preceq w_i$ for at least one $i \in \{1, \dots, n\}$ and, for each i with $u \preceq w_i$, $\langle \wp \beta \rangle \in \mathbf{Lr}^{A_i}$. We call such a (lab)move a **nonreplicative** (lab)move. It brings the game down to

$$\downarrow \mathcal{T}' ,$$

where \mathcal{T}' is the result of replacing A_i by $\langle \wp \beta \rangle A_i$ in \mathcal{T} for each i with $u \preceq w_i$. If here u is w_i (rather than a proper prefix of such) for one of $i \in \{1, \dots, n\}$, we say that the move $\wp u. \beta$ is **focused**. Otherwise it is **unfocused**.

- (2) $\wp \alpha = \perp w_i$, where $i \in \{1, \dots, n\}$. We call such a (lab)move a **replicative** (lab)move. It brings the game down to $\downarrow \mathcal{T}'$, where \mathcal{T}' is the result of replacing A_i by $(A_i \circ A_i)$ in \mathcal{T} .

(ii): Whenever Γ is a legal run of $\downarrow \mathcal{T}$, $\mathbf{Wn}^{\delta \mathcal{T}} \langle \Gamma \rangle = \top$ iff, for each $i \in \{1, \dots, n\}$ and every infinite bit string v with $w_i \preceq v$, we have $\mathbf{Wn}^{A_i} \langle \Gamma^{\preceq v} \rangle = \top$. Here

$$\Gamma^{\preceq v}$$

means the result of deleting, from Γ , all labmoves except those that look like $\wp u. \alpha$ for some bit string u with $u \preceq v$, and then further changing each such (remaining) labmove $\wp u. \alpha$ to $\wp \alpha$.⁸

⁸Intuitively, $\Gamma^{\preceq v}$ is the run played in one of the multiple “copies” of A_i that have been generated in the play, with v acting as a (perhaps longer than necessary yet meaningful) “address” of that copy.

Example 2.8. Let

$$G = p \sqcup (q \sqcap (r \sqcap (s \sqcup t))),$$

where p, q, r, s, t are constant elementary games. And let

$$\Gamma = \langle \perp:, \top.1, \perp 0.0, \perp 1.1, \perp 1:, \perp 10.0, \perp 11.1, \top 11.0 \rangle.$$

Then Γ is a legal run of $\downarrow G$. Below we trace, step by step, the effects of its moves on $\downarrow G$.

The 1st (lab)move $\perp:$ means that \perp replicates the (only) leaf of the tree, with the address of that leaf being the empty bit string. This move brings the game down to — in the sense that $\langle \perp: \rangle \downarrow G$ is — the following game:

$$\downarrow \left((p \sqcup (q \sqcap (r \sqcap (s \sqcup t)))) \circ (p \sqcup (q \sqcap (r \sqcap (s \sqcup t)))) \right).$$

The 2nd move $\top.1$ means choosing the second \sqcup -disjunct $q \sqcap (r \sqcap (s \sqcup t))$ in *both* leaves of the tree. This is so because the addresses of those leaves are 0 and 1, and the empty bit string — seen between “ \top ” and “.1” in $\top.1$ — is an initial segment of both addresses. The effect of this unfocused move is the same as the effect of the two consecutive focused moves $\top 0.1$ and $\top 1.1$ (in whatever order) would be, but \top might have its reasons for having made an unfocused move. Among such reasons could be that \top did not notice \perp 's initial move (or the latter arrived late over the asynchronous network) and thought that the position was still G , in which case making the moves $\top 0.1$ and $\top 1.1$ would be simply illegal. Note also that the ultimate effect of the move $\top.1$ on the game would remain the same as it is now even if this move was made before the replicative move $\perp:$. It is CoL's striving to achieve this sort of flexibility and asynchronous-communication-friendliness that has determined our seemingly “strange” choice of trees rather than sequences as the underlying structures for \downarrow -games. Any attempt to deal with sequences instead of trees would encounter the problem of violating what CoL calls the *static* (speed-independent) property of games, defined later in Section 3.

Anyway, the position resulting from the second move of Γ is

$$\downarrow \left((q \sqcap (r \sqcap (s \sqcup t))) \circ (q \sqcap (r \sqcap (s \sqcup t))) \right).$$

The effect of the 3rd move $\perp 0.0$ is choosing the left \sqcap -conjunct q in the left (0-addressed) leaf of the tree, which results in

$$\downarrow \left(q \circ (q \sqcap (r \sqcap (s \sqcup t))) \right).$$

Similarly, the 4th move $\perp 1.1$ chooses the right \sqcap -conjunct in the right leaf of the tree, resulting in

$$\downarrow \left(q \circ (r \sqcap (s \sqcup t)) \right).$$

The 5th move $\perp 1:$ replicates the right leaf, bringing the game down to

$$\downarrow \left(q \circ ((r \sqcap (s \sqcup t)) \circ (r \sqcap (s \sqcup t))) \right).$$

The 6th move $\perp 10.0$ chooses the left \sqcap -conjunct in the second (00-addressed) leaf, and, similarly, the 7th move $\perp 11.1$ chooses the right \sqcap -conjunct in the third (11-addressed) leaf. These two moves bring the game down to

$$\downarrow \left(q \circ (r \circ (s \sqcup t)) \right).$$

The last, 8th move $\top 11.0$ chooses the left \sqcup -disjunct of the third leaf, and the final position is

$$\downarrow(q \circ (r \circ s)).$$

According to clause (ii) of Definition 2.7, Γ is a \top -won run of $\downarrow G$ iff, for any infinite bit string v , $\Gamma^{\preceq v}$ is a \top -won run of G . Observe that for any infinite — or, “sufficiently long” finite — bit string v , $\Gamma^{\preceq v}$ is either $\langle \top 1, \perp 0 \rangle$ (if $v = 0\dots$) or $\langle \top 1, \perp 1, \perp 0 \rangle$ (if $v = 10\dots$) or $\langle \top 1, \perp 1, \perp 1, \top 0 \rangle$ (if $v = 11\dots$). We also have $\langle \top 1, \perp 0 \rangle G = q$, $\langle \top 1, \perp 1, \perp 0 \rangle G = r$ and $\langle \top 1, \perp 1, \perp 1, \top 0 \rangle G = s$. So it is no accident that we see q, r, s at the leaves in the final position. Correspondingly, the game is won iff each one of these three propositions is true.

The cases where \perp makes infinitely many replications in a run Γ of a game $\downarrow H$ and hence the “eventual tree” is infinite are similar, with the only difference that the “addresses” of the “leaves” of such a “tree”, corresponding to different plays of H , may be infinite bit strings. But, again, the overall game $\downarrow H$ will be won by \top iff all of those plays — all $\Gamma^{\preceq v}$ where v is an infinite bit string, that is — are \top -won plays of H .

Definition 2.9. Let B, A_1, \dots, A_n ($n \geq 0$) be constant games. We define

$$A_1, \dots, A_n \circ - B$$

— let us call it the (generalized) **Turing reduction** of B to A_1, \dots, A_n — as the game

$$\downarrow A_1 \wedge \dots \wedge \downarrow A_n \rightarrow B.$$

3. NOT-NECESSARILY-CONSTANT GAMES AND “QUANTIFIER” OPERATIONS

Constant games can be seen as generalized propositions: while propositions in classical logic are just elements of $\{\top, \perp\}$, constant games are functions from runs to $\{\top, \perp\}$. Our concept of a (not-necessarily-constant) game defined in this section generalizes that of a constant game in the same sense as the classical concept of a predicate generalizes that of a proposition.

Throughout CoL, we have a fixed set VARIABLES of expressions called **variables**,

$$\text{VARIABLES} = \{var_1, var_2, var_3, \dots\}.$$

Each variable is thus the string var_i for some positive decimal numeral i . However, we seldom write variables in this form and, instead, usually use the letters (metavariables) $r, s, t, u, v, w, x, y, z$ for them. We also have another fixed set CONSTANTS of expressions called **canonical constants**, or simply **constants**. Specifically,

$$\text{CONSTANTS} = \{\epsilon, 1, 10, 11, 100, 101, 110, 111, 1000, \dots\}.$$

Constants are thus **binary numerals** — the strings matching the regular expression $\epsilon \cup 1(0 \cup 1)^*$, where ϵ is the empty string. We will be often identifying such strings — by some rather innocent abuse of concepts — with the natural numbers represented by them in the standard binary notation, and vice versa. Note that ϵ represents 0. For this reason, following tradition, we shall usually write 0 instead of ϵ , keeping in mind that, in such cases, the length $|0|$ of the string 0 should (still) be seen to be 0 rather than 1. We will be mostly using a, b, c, d as metavariables for constants.

A **universe** (of discourse) is a pair $U = (Dom, Den)$, where Dom , called the **domain** of U , is a nonempty set, and Den , called the **denotation** of U , is a (total) function of type $\text{CONSTANTS} \rightarrow Dom$. The elements of Dom will be referred to as the **individuals** of U .

The intuitive meaning of $d = Den(c)$ is that the individual d is the **denotat** of the constant c and thus c is a **name** of d . So, the function Nam from Dom to the powerset of `CONSTANTS` satisfying the condition $c \in Nam(d) \Leftrightarrow d = Den(c)$ can be called the **naming** of U . Of course, whenever convenient, a universe can be characterized in terms of its naming rather than denotation.

The **ideal universe** is the universe whose domain is `CONSTANTS` and whose denotation is the identity function on that domain. All earlier papers on CoL, except [31, 34, 37, 38], dealt only with the ideal universe. This was for simplicity considerations, yielding no loss of generality as no results relied on the assumption that the underlying universe was ideal. Our present treatment, however, for both technical and philosophical reasons, *does* call for the above-defined, more general, concept of a universe — a universe where some individuals may have unique names, some have many names, and some have no names at all.⁹ Note that real-world universes are typically not ideal: not all people living or staying in the United States have social security numbers; most stars and planets of the Galaxy have no names at all, while some have several names (Morning Star = Evening Star = Venus); etc. A natural example of a non-ideal universe from the world of mathematics would be the one whose domain is the set of real numbers, only some of whose elements have names, such as 5, $1/3$, $\sqrt{2}$ or π . Generally, even if the set of constants was not fixed, no universe with an uncountable domain would be “ideal” for the simple reason that there can only be countably many names. This is so because names, by their very nature and purpose, have to be finite objects. Observe also that many properties of common interest, such as computability or decidability, are usually sensitive to how objects (individuals) are named, as they deal with the names of those objects rather than the objects themselves. For instance, strictly speaking, computing a function $f(x)$ means the ability to tell, after seeing a (the) name of an arbitrary object \mathbf{a} , to produce a (the) name of the object \mathbf{b} with $\mathbf{b} = f(\mathbf{a})$. Similarly, an algorithm that decides a predicate $p(x)$ on a set S , strictly speaking, takes not elements of S — which may be abstract objects such as numbers or graphs — but rather names of those elements (such as binary numerals or codes). It is not hard to come up with a nonstandard naming of the natural numbers through binary numerals where the predicate “ x is even” is undecidable. On the other hand, for any undecidable arithmetical predicate $p(x)$, one can come up with a naming such that $p(x)$ becomes decidable — for instance, one that assigns even-length names to all \mathbf{a} with $p(\mathbf{a})$ and assigns odd-length names to all \mathbf{a} with $\neg p(\mathbf{a})$. Classical logic exclusively deals with individuals of a universe without a need for also considering names for them, as it is not concerned with decidability or computability. CoL, on the other hand, with its computational semantics, inherently calls for being more careful about differentiating between individuals and their names, and hence for explicitly considering universes in the form (Dom, Den) rather than just Dom as classical logic does.

By a **valuation** on a universe $U = (Dom, Den)$, or a **U -valuation**, we mean a (total) function e of type `VARIABLES` \rightarrow Dom . For the purposes of this definition, the Den component of U is thus irrelevant and, whenever convenient, we may talk about valuations in terms of just domains instead of universes. The set of all U -valuations will be denoted by `VALUATIONS`(U). When a universe U (or its domain) is fixed, irrelevant or clear from the context, we may omit an explicit reference to it and simply say “valuation”. References

⁹Further generalizations are possible if and when a need arises. Namely, one may depart from our present assumption that the set of constants is infinite and/or fixed, as long as there is a fixed constant — say, 0 — that belongs to every possible set of constants ever considered. No results of this or any earlier papers on CoL would be in any way affected by doing so.

to U can be similarly omitted when talking about individuals, denotats, names or some later-defined concepts such as those of a game or a function.

In the above terms, a classical predicate p can be understood as a function that sends each valuation e to a proposition, i.e., to a constant predicate. Similarly, what we call a game sends valuations to constant games:

Definition 3.1. Let U be a universe. A **game on U** is a (total) function A from U -valuations to constant games. For a valuation e , we write $e[A]$ rather than $A(e)$ to denote the value of A at e . Such a constant game $e[A]$ is said to be an **instance** of A . For readability, we usually write \mathbf{Lr}_e^A and \mathbf{Wn}_e^A instead of $\mathbf{Lr}^{e[A]}$ and $\mathbf{Wn}^{e[A]}$.

Just as it is the case with propositions versus predicates, constant games in the sense of Definition 2.1 will be thought of as special, constant cases of games in the sense of Definition 3.1. Namely, in the context of a given universe U , each constant game A' is the game A (on U) such that, for every valuation e , $e[A] = A'$. From now on we will no longer distinguish between such A and A' , so that, if A is a constant game, it is its own instance, with $A = e[A]$ for every valuation e .

Where n is a natural number, we say that a game A is **n -ary** iff there are n variables such that, for any two valuations e_1 and e_2 that agree on all those variables, we have $e_1[A] = e_2[A]$. Note that, if A is n -ary, then it is also $(n + 1)$ -ary, $(n + 2)$ -ary, etc. A game that is n -ary for some n is said to be **finitary**. The present paper is going to exclusively deal with finitary games and, for this reason, we agree that, from now on, when we say “game”, we always mean “finitary game”.

For a variable x and valuations e_1, e_2 , we write $e_1 \equiv_x e_2$ to mean that the two valuations agree on all variables other than (perhaps) x .

We say that a game A **depends** on a variable x iff there are two valuations e_1, e_2 with $e_1 \equiv_x e_2$ such that $e_1[A] \neq e_2[A]$. An n -ary game thus depends on at most n variables. And constant games are nothing but 0-ary games, i.e., games that do not depend on any variables.

We say that a (not necessarily constant) game A is **elementary** iff so are all of its instances.

Just as constant games are generalized propositions, games can be seen as generalized predicates. Namely, in the context of a given universe U , we will view each predicate p on U 's domain as the elementary game such that, for every valuation e , $\mathbf{Wn}_e^p \langle \rangle = \top$ iff p is true at e . And vice versa: every elementary game p will be viewed as the predicate on U 's domain which is true at a given valuation e iff $\mathbf{Wn}_e^p \langle \rangle = \top$. Thus, for us, “predicate” and “elementary game” are synonyms. Accordingly, any standard terminological or notational conventions familiar from the literature for predicates also apply to them viewed as elementary games.

There are two different yet not always clearly differentiated understandings of predicates in the literature. One — more common — understanding is *extensional*, according to which an n -ary predicate is a set of n -tuples of individuals (that satisfy p). The other understanding — let us call it *subextensional* — sees a predicate p as a set of valuations (at which p is true). Note that, unlike the extensional understanding, the subextensional understanding is variable-sensitive: it, for instance, sees $var_1 \geq 100$ and $var_2 \geq 100$ as different predicates, because one depends on var_1 while the other depends on var_2 instead. Our understanding of predicates and, more generally, our understanding of games, is clearly subextensional. We agree on the following understanding of functions in the same subextensional style:

Definition 3.2. Let $U = (Dom, Den)$ be a universe. A **function on U** is a mapping f from $VALUATIONS(U)$ to Dom . For a valuation e , we write $e[f]$ rather than $f(e)$ to denote the value of f at e .

Much of our terminology for games naturally extends to functions as well. Namely, where n is a natural number, we say that a function f is **n -ary** iff there are $\leq n$ variables such that, for any two valuations e_1 and e_2 that agree on all those variables, we have $e_1[f] = e_2[f]$. A function that is n -ary for some n is said to be **finitary**. The present paper exclusively deals with finitary functions, and we agree that, from now on, “function” always means “finitary function”. We say that a function f **depends** on a variable x iff there are two valuations e_1, e_2 with $e_1 \equiv_x e_2$ such that $e_1[f] \neq e_2[f]$. **Constant functions** are nothing but 0-ary functions, i.e. functions that do not depend on any variables.

Convention 3.3. In the context of a given universe $U = (Dom, Den)$, all individuals, constants and variables will be (simultaneously) understood as (denoting) functions on U . Namely:

1. An individual a is the (constant) function that sends every valuation e to a .
2. A constant c is the (constant) function that sends every valuation e to $Den(c)$.
3. A variable x is the (not-necessarily-constant) function that sends every valuation e to $e(x)$.

Definition 3.4. Let W be a game or a function on a universe U , x_1, \dots, x_n be pairwise distinct variables, and f_1, \dots, f_n be functions on U . The result of **substituting f_1, \dots, f_n for x_1, \dots, x_n in W** , denoted by $W(f_1/x_1, \dots, f_n/x_n)$, is defined as the game on the same universe U such that, for every valuation e , $e[W(f_1/x_1, \dots, f_n/x_n)] = e'[W]$, where e' is the valuation that sends each $x_i \in \{x_1, \dots, x_n\}$ to (the value of) $e[f_i]$ and agrees with e on all other variables.

Following the standard readability-improving practice established in the literature for predicates and functions, we will often fix pairwise distinct variables x_1, \dots, x_n for a game or function W and write W as $W(x_1, \dots, x_n)$. Representing W in this form sets a context in which we can write $W(f_1, \dots, f_n)$ to mean the same as the more clumsy expression $W(f_1/x_1, \dots, f_n/x_n)$.

Just as the Boolean operations straightforwardly extend from propositions to all predicates, our operations $\neg, \wedge, \vee, \rightarrow, \sqcap, \sqcup, \circ, \ominus$ extend from constant games to all games, with each operation taking one or more games on a given (common-for-all-arguments) universe and returning a game on the same universe. This is done by simply stipulating that $e[\dots]$ commutes with all of those operations: $\neg A$ is the game such that, for every valuation e , $e[\neg A] = \neg e[A]$; $A \sqcap B$ is the game such that, for every e , $e[A \sqcap B] = e[A] \sqcap e[B]$; etc. So does the operation of prefixation: provided that Φ is a legal position of every instance of A , $\langle \Phi \rangle A$ is understood as the unique game such that, for every e , $e[\langle \Phi \rangle A] = \langle \Phi \rangle e[A]$.

Definition 3.5. Below $A(x)$ is any game on a universe U . On the same universe:

1. The game $\sqcap x A(x)$ (**choice universal quantification**) is defined by stipulating that, for every U -valuation e , we have:

- (i): $\langle \wp \alpha \rangle \in \mathbf{Lr}_e^{\sqcap x A(x)}$ iff $\wp = \perp$ and α is $\#c$ for some constant c . Such an initial legal labmove $\perp \#c$ brings the game $e[\sqcap x A(x)]$ down to $e[A(c)]$.
- (ii): Whenever Γ is a legal run of $\sqcap x A(x)$, $\mathbf{Wn}_e^{\sqcap x A(x)} \langle \Gamma \rangle = \perp$ iff Γ looks like $\langle \perp \#c, \Delta \rangle$ and $\mathbf{Wn}_e^{A(c)} \langle \Delta \rangle = \perp$.

2. The game $\sqcup xA(x)$ (**choice existential quantification**) is defined in exactly the same way, only with \top and \perp interchanged.

Thus, every initial legal move of $\sqcap xA(x)$ or $\sqcup xA(x)$ is the string $\#c$ for some $c \in \text{CONSTANTS}$, which in our informal language we may refer to as “the constant chosen (by the corresponding player) for x ”. Note that $\sqcap xA(x)$ is nothing but the infinite choice conjunction $A(0) \sqcap A(1) \sqcap \dots$, with the only technical difference that, in the former, the initial, “choice” move should be prefixed with the symbol “#”. Similarly, $\sqcup xA(x)$ is essentially nothing but the infinite choice disjunction $A(0) \sqcup A(1) \sqcup \dots$.

We say that a game A is **unistructural** iff, for any two valuations e_1 and e_2 , we have $\mathbf{Lr}_{e_1}^A = \mathbf{Lr}_{e_2}^A$. Of course, all constant or elementary games are unistructural. It can also be easily seen that all our game operations preserve the unistructural property of games. For the purposes of the present paper, considering only unistructural games is sufficient.

We define the remaining operations \forall and \exists only for unistructural games.

Definition 3.6. Below $A(x)$ is any unistructural game on a universe U . On the same universe:

1. The game $\forall xA(x)$ (**blind universal quantification**) is defined by stipulating that, for every U -valuation e , we have:

- (i): $\langle \wp\alpha \rangle \in \mathbf{Lr}_e^{\forall xA(x)}$ iff $\langle \wp\alpha \rangle \in \mathbf{Lr}_e^{A(x)}$. Such an initial legal labmove $\wp\alpha$ brings the game $e[\forall xA(x)]$ down to $e[\forall x\langle \wp\alpha \rangle A(x)]$.
- (ii): Whenever Γ is a legal run of $e[\forall xA(x)]$, $\mathbf{Wn}_e^{\forall xA(x)}\langle \Gamma \rangle = \top$ iff, for every valuation g with $g \equiv_x e$, $\mathbf{Wn}_g^{A(x)}\langle \Gamma \rangle = \top$.

2. The game $\exists xA(x)$ (**blind existential quantification**) is defined in exactly the same way, only with \top and \perp interchanged.

Example 3.7. Let G be the following game on the ideal universe, with the predicates *Even* and *Odd* having their expected meanings:

$$\forall y \left(\text{Even}(y) \sqcup \text{Odd}(y) \rightarrow \sqcap x (\text{Even}(x+y) \sqcup \text{Odd}(x+y)) \right).$$

Then the sequence $\langle \perp 1.\#11, \perp 0.0, \top 1.1 \rangle$ is a legal run of G , the effects of the moves of which are shown below:

$$\begin{aligned} G : & \quad \forall y \left(\text{Even}(y) \sqcup \text{Odd}(y) \rightarrow \sqcap x (\text{Even}(x+y) \sqcup \text{Odd}(x+y)) \right) \\ \langle \perp 1.\#11 \rangle G : & \quad \forall y \left(\text{Even}(y) \sqcup \text{Odd}(y) \rightarrow \text{Even}(11+y) \sqcup \text{Odd}(11+y) \right) \\ \langle \perp 1.\#11, \perp 0.0 \rangle G : & \quad \forall y \left(\text{Even}(y) \rightarrow \text{Even}(11+y) \sqcup \text{Odd}(11+y) \right) \\ \langle \perp 1.\#11, \perp 0.0, \top 1.1 \rangle G : & \quad \forall y \left(\text{Even}(y) \rightarrow \text{Odd}(11+y) \right) \end{aligned}$$

The play hits (ends as) the true proposition $\forall y (\text{Even}(y) \rightarrow \text{Odd}(11+y))$ and hence is won by \top .

Example 3.8. The sequence $\langle \top 0.1., \perp 1.\#10, \top 0.1.0.\#10, \top 0.1.0.\#10, \perp 0.1.0.\#100, \top 0.1.1.\#100, \top 0.1.1.\#10, \perp 0.1.1.\#1000, \top 1.\#1000 \rangle$ is a legal run of the game

$$\forall x (x^3 = (x \times x) \times x), \sqcap x \sqcap y \sqcup z (z = x \times y) \circ - \sqcap x \sqcup y (y = x^3). \quad (3.1)$$

Below we see how the game evolves according to the scenario of this run:

	$\forall x(x^3 = (x \times x) \times x), \Box x \Box y \Box z(z = x \times y) \circ - \Box x \Box y(y = x^3)$
$\top 0.1.:$ yields	$\forall x(x^3 = (x \times x) \times x), \Box x \Box y \Box z(z = x \times y) \circ \Box x \Box y \Box z(z = x \times y) \circ - \Box x \Box y(y = x^3)$
$\perp 1.$ #10 yields	$\forall x(x^3 = (x \times x) \times x), \Box x \Box y \Box z(z = x \times y) \circ \Box x \Box y \Box z(z = x \times y) \circ - \Box y(y = 10^3)$
$\top 0.1.0.$ #10 yields	$\forall x(x^3 = (x \times x) \times x), \Box y \Box z(z = 10 \times y) \circ \Box x \Box y \Box z(z = x \times y) \circ - \Box y(y = 10^3)$
$\top 0.1.0.$ #10 yields	$\forall x(x^3 = (x \times x) \times x), \Box z(z = 10 \times 10) \circ \Box x \Box y \Box z(z = x \times y) \circ - \Box y(y = 10^3)$
$\perp 0.1.0.$ #100 yields	$\forall x(x^3 = (x \times x) \times x), (100 = 10 \times 10) \circ \Box x \Box y \Box z(z = x \times y) \circ - \Box y(y = 10^3)$
$\top 0.1.1.$ #100 yields	$\forall x(x^3 = (x \times x) \times x), (100 = 10 \times 10) \circ \Box y \Box z(z = 100 \times y) \circ - \Box y(y = 10^3)$
$\top 0.1.1.$ #10 yields	$\forall x(x^3 = (x \times x) \times x), (100 = 10 \times 10) \circ \Box z(z = 100 \times 10) \circ - \Box y(y = 10^3)$
$\perp 0.1.1.$ #1000 yields	$\forall x(x^3 = (x \times x) \times x), (100 = 10 \times 10) \circ (1000 = 100 \times 10) \circ - \Box y(y = 10^3)$
$\top 1.$ #1000 yields	$\forall x(x^3 = (x \times x) \times x), (100 = 10 \times 10) \circ (1000 = 100 \times 10) \circ - 1000 = 10^3$

The play hits a true proposition and hence is won by \top . Note that here, unlike the case in the previous example, \top is the winner no matter what universe we consider and what the meanings of the expressions $x \times y$ and x^3 are. In fact, \top has a “purely logical” winning strategy in this game, in the sense that the strategy is successful regardless of whether things have their standard arithmetical meanings or some other meanings. This follows from the promised soundness of **CL12** and the fact — illustrated later in Example 7.2 — that (3.1) is provable in **CL12**. Such “purely logical” strategies in CoL are called *uniform solutions* (see Section 6).

We close this section by defining a couple of additional concepts that are central in CoL, even if nothing in the present paper directly or indirectly relies on them. For either player \wp , a run Φ is said to be a \wp -**delay** of a run Γ iff (1) for both players $\wp' \in \{\top, \perp\}$, the subsequence of \wp' -labeled moves of Φ is the same as that of Γ , and (2) for any $n, k \geq 1$, if the n th \wp -labeled move is made later than (is to the right of) the k th $\neg\wp$ -labeled move in Γ , then so is it in Φ . For instance, $\langle \perp \alpha_1, \top \beta_1, \perp \alpha_2, \top \beta_2 \rangle$ is a \top -delay of $\langle \top \beta_1, \perp \alpha_1, \top \beta_2, \perp \alpha_2 \rangle$. Next, we say that a run is \wp -**legal** iff it is not \wp -illegal. Now, we say that a constant game A is **static** iff, for either player \wp , whenever a run Φ is a \wp -delay of a run Γ , we have: (i) if Γ is a \wp -legal run of A , then so is Φ , and (ii) if Γ is a \wp -won run of A , then so is Φ . This concept extends to all games by stipulating that a not-necessarily-constant game is static iff all of its instances are so.

Static games form a very wide natural subclass of all games. Intuitively, such games are interactive tasks where the relative speeds of the players are irrelevant, as it never hurts a player to postpone making moves. In other words, static games are games that are contests of intellect rather than contests of speed. And one of the theses that CoL philosophically relies on is that static games present an adequate formal counterpart of our intuitive concept of “pure”, speed-independent interactive computational problems. Correspondingly, CoL restricts its attention (more specifically, possible interpretations of the atoms of its formal language) to static games. Every elementary game is trivially static, and the class of static games turns out to be closed under all game operations studied in CoL. This can be seen to immediately imply that all games expressible in the language of the later-defined logic **CL12** are static (as well as finitary and unistructural).

4. INTERACTIVE MACHINES

In traditional game-semantical approaches, including those of Lorenzen [39], Hintikka [10] and Blass’s [3], players’ strategies are understood as *functions* — typically as functions from interaction histories (positions) to moves, or sometimes (Abramsky and Jagadeesan [1]) as functions that only look at the latest move of the history. This *strategies-as-functions* approach, however, is generally inapplicable in the context of CoL, whose relaxed semantics, in striving to get rid of “bureaucratic pollutants” and only deal with the remaining true essence of games, does not impose any regulations on which player can or should move in a given situation. Here, in many cases, either player may have (legal) moves, and then it is unclear whether the next move should be the one prescribed by \top ’s strategy function or the one prescribed by the strategy function of \perp . For a game semantics whose ambition is to provide a comprehensive, natural and direct tool for modeling interaction, the strategies-as-functions approach would be less than adequate, even if technically possible. This is so for the simple reason that the strategies that real computers follow are not functions. If the strategy of your personal computer was a function from the history of interaction with you, then its performance would keep noticeably worsening due to the need to read the continuously lengthening — and, in fact, practically infinite — interaction history every time before responding. Fully ignoring that history and looking only at your latest keystroke in the spirit of [1] is also certainly not what your computer does, either. The advantages of our approach thus become especially appreciable when one tries to bring complexity theory into interactive computation: hardly (m)any really meaningful and interesting complexity-theoretic concepts can be defined for games (particularly, games that may last long) with the strategies-as-functions approach.

In CoL, (\top ’s effective) strategies are defined in terms of interactive machines, where computation is one continuous process interspersed with — and influenced by — multiple “input” (environment’s moves) and “output” (machine’s moves) events. Of several, seemingly rather different yet equivalent, machine models of interactive computation studied in CoL, this paper only employs the most basic, **HPM** (“Hard-Play Machine”) model.

An HPM is a Turing machine with the additional capability of making moves. The adversary can also move at any time, and such moves are the only nondeterministic events from the machine’s perspective. Along with one or more ordinary read/write **work tapes**,¹⁰

¹⁰This is the first time in the literature on CoL that multiple-work-tape HPMs are considered. All earlier HPMs had a single work tape. Just as in the traditional theory of computation, as long as we are willing to

the machine has an additional, read-only tape called the **run tape**.¹¹ The latter, serving as a dynamic input, at any time spells the “current position” of the play. Its role is to make the evolving run fully visible to the machine. In these terms, an algorithmic solution (\top ’s winning strategy) for a given constant game A is understood as an HPM \mathcal{M} such that, no matter how the environment acts during its interaction with \mathcal{M} (what moves it makes and when), the run incrementally spelled on the run tape is a \top -won run of A . As for \perp ’s strategies, there is no need to define them: all possible behaviors by \perp are accounted for by the different possible nondeterministic updates of the run tape of an HPM.

In the above outline, we described HPMs in a relaxed fashion, without being specific about details such as, say, how, exactly, moves are made by the machine, how many moves either player can make at once, what happens if both players attempt to move “simultaneously”, etc. As it happens, all reasonable design choices yield the same class of winnable games as long as we only consider static games, including all games expressible in the language of logic **CL12**. Such games are not necessarily constant but, due to being finitary, they can and will be thought of to be constant by identifying them with their \sqcap -closures. Correspondingly, in this paper, we use the term “**computational problem**”, or simply “**problem**”, as a synonym of “constant static game”.

While design choices are largely unimportant and “negotiable”, we still want to agree on some technical details for clarity. Just like an ordinary Turing machine, an HPM has a finite set of **states**, one of which has the special status of being the **start state**. There are no accept, reject, or halt states, but there are specially designated states called **move states**. For simplicity, we do not allow the start state to (also) be a move state. Each tape of the machine has a beginning but no end, and is divided into infinitely many **cells**, arranged in the left-to-right order. At any time, each cell contains one symbol from a certain fixed finite set of **tape symbols**. The **blank** symbol `BLANK`, as well as \top and \perp , are among the tape symbols. We also assume that these three symbols are not among the symbols that the body of any (legal or illegal) move can ever contain, i.e. not among the symbols of the keyboard alphabet of which, as we agreed in Section 2, all moves are composed. Each tape has its own **scanning head**, at any given time looking (located) at one of the cells of the tape.

For technical purposes, we additionally assume the (physical or imaginary) presence of a **buffer**.¹² The size of the latter is unlimited and, at any time, it contains some (possibly empty) finite string over the keyboard alphabet. The function of the buffer is to let the machine construct a (“large”) move piece-by-piece before officially making such a move.

A transition from one **computation step** (“**clock cycle**”, “**time**”) to another happens according to the fixed **transition function** of the machine. The latter, depending on the current state and the symbols seen by the scanning heads on the corresponding tapes, deterministically prescribes: (1) the next state that the machine should assume; (2) the tape symbol by which the old symbol should be overwritten in the current cell (the cell

ignore certain “small” polynomial time differences in computation (game-playing) efficiency, nothing depends on how many — one or more — work tapes are allowed.

¹¹Together with the work and run tapes, the HPMs from the earlier literature on CoL also had an additional tape called the *valuation tape*. The latter becomes redundant in our present treatment due to the fact that we are exclusively interested in constant games or finitary games identified with their (constant) \sqcap -closures.

¹²The earlier HPM models did not include buffers. The reason for a modification in the present style will be explained in Section 5.

currently scanned by the head), for each work tape individually; (3) the (finite, possibly empty) string over the keyboard alphabet that should be appended to the content of the buffer; and (4) the (not necessarily the same) direction — stay put, one cell to the left, or one cell to the right — in which each scanning head should move. It is stipulated that when the head of a tape is looking at the first (leftmost) cell, an attempt to move to the left results in staying put. The same happens when the head tries to move to the right while looking at a cell containing `BLANK`. We additionally assume that `BLANK` can never be written¹³ by the machine on its work tapes. Note that, in view of the above conditions, at any time, all cells of a work tape found to the right of a blank cell are also blank. The same holds for the run tape, of course.

When the machine starts working, it is in its start state, all scanning heads are looking at the leftmost cells of the corresponding tapes, all work tapes are blank (i.e., all their cells contain `BLANK`), the run tape does not contain any \top -labeled moves (but it may contain some \perp -labeled moves, signifying that Environment has made those moves “at the very beginning”), and the buffer is empty. Whenever the machine enters a move state, the move α written in the buffer by that time, in the \top -prefixed form $\top\alpha$, is automatically appended to the contents of the run tape, and the buffer is simultaneously emptied. Here we assume that α includes (as a suffix) the string γ that the machine tried to put into the buffer on that transition. For terminological and conceptual convenience, we think of this situation as that γ is actually written into the buffer before it, immediately after that (on the same transition), migrates from the buffer to the run tape together with the rest of α . Also, on every transition, any finite sequence $\perp\beta_1, \dots, \perp\beta_m$ of \perp -labeled moves may be nondeterministically appended to the contents of the run tape. If the above two events happen on the same clock cycle, then both $\top\alpha$ and $\perp\beta_1, \dots, \perp\beta_m$ will be appended to the contents of the run tape, where $\top\alpha$ can (nondeterministically) go before, after or anywhere in the middle of $\perp\beta_1 \dots \perp\beta_m$. Whenever we say that a move (by either player) was **made at time c** , or *on the transition to step c* , we mean that c is the clock cycle on which that move (first) appeared on the run tape. We agree that the count of clock cycles starts from 0 rather than 1.

In the future, when describing the work of a machine, we may use the jargon “**retire**”. What will be meant by retiring is going into an infinite loop that makes no moves, puts nothing into the buffer, and does not reposition the scanning heads. Retiring thus achieves the same effect as halting would achieve if this was technically allowed.

A **configuration** is a full description of the situation in the machine at some given computation step. It consists of records of the (“current”) contents of the work and run tapes, the content of the buffer, the location of each scanning head, and the state of the machine. A **computation branch** of the machine is an infinite sequence C_0, C_1, C_2, \dots of configurations, where C_0 is an **initial configuration** (one described at the beginning of the preceding paragraph), and every C_{i+1} is a configuration that could have legally followed (again, in the sense explained earlier) C_i according to the transition function of the machine. In less formal contexts, we may say “**play**” instead of “computation branch”. For a computation branch B , the **run spelled by B** is the run Γ incrementally spelled on the run tape in the corresponding scenario of interaction. We say that such a Γ is a **run generated by** the machine.

¹³I.e., `BLANK` can never replace a non-blank symbol.

We say that a given HPM \mathcal{M} **wins** (**computes**, **solves**) a given constant game A — and write $\mathcal{M} \models A$ — iff every run generated by \mathcal{M} is a \top -won run of A . We say that A is **computable** iff there is an HPM \mathcal{M} with $\mathcal{M} \models A$; such an HPM is said to be an (algorithmic) **solution**, or **winning strategy**, for A .

5. INTERACTIVE COMPLEXITY

At present, the theory of interactive computation is far from being well developed, and even more so is the corresponding complexity theory. The studies of interactive computation in the context of complexity, while having been going on since long ago, have been relatively scattered and ad hoc: more often than not, interaction has been used for better understanding certain complexity issues for traditional, non-interactive problems rather than being treated as an object of systematic studies in its own rights (examples would be alternating computation [7], or interactive proof systems and Arthur-Merlin games [9, 2]). As if complexity theory was not “complex” enough already, taking it to the interactive level would most certainly generate a by an order of magnitude greater diversity of species from the complexity zoo. The paper [34] made the first attempt to bring complexity concepts into CoL. The present paper significantly refines that attempt. The main novelty, as mentioned in Section 1, is throwing amplitude complexity (see below) into the mix. It should be however pointed out that our way of measuring the complexity of strategies is merely one out of a huge and interesting potential variety of complexity measures meaningful and useful in the interactive context.

For the following two definitions, we assume the presence of some fixed function that sends every move α to a natural number called the *magnitude* of α .

Definition 5.1. In the context of a given computation branch (play) of a given HPM \mathcal{M} :

- (1) By the **background** of a clock cycle c we mean the greatest of the magnitudes of Environment’s moves made before time c , or is 0 if there are no such moves.
- (2) By the **timecost** of a cycle c we mean $c - d$, where d is the greatest cycle with $d < c$ on which a move was made by Environment, or is 0 if there is no such cycle.
- (3) By the **spacecost** of a cycle c we mean the maximum number of cells in which \mathcal{M} has ever written anything on any (any *one*) work tape before time c .¹⁴
- (4) If \mathcal{M} makes a move on a cycle c , then the **background** (resp. **timecost**, resp. **spacecost**) of that move¹⁵ means the background (resp. timecost, resp. spacecost) of c .

¹⁴Here, of course, “written anything” means “written anything new”, i.e., “overwritten an old symbol by a new one”, for, technically, not writing simply means writing/repeating the old symbol of the cell. Note that we only look at the “worst-case” work tape, deeming the numbers of cells written on all other (if any) work tapes irrelevant. Also note that we only count the numbers of *cells* ever written rather the numbers of *writes* — that is, we only count the first write in a given cell, with all subsequent writes contributing nothing extra to the spacecost. Further note that the number of cells ever *visited* (rather than *written*) by the scanning head of the worst-case work tape of \mathcal{M} before time c is either the same as the spacecost of c , or exceeds the latter by 1. This is so because, as we agreed, a scanning head can never move to the right while looking at a blank cell. Finally, note that the spacecost of the first step/cycle of a computation branch is always 0, and the spacecost of every subsequent step is nothing but the number of non-blank cells (which form a contiguous block) on the worst-case work tape of \mathcal{M} at the preceding step.

¹⁵As easily understood, here and later in similar contexts, including Definition 5.2, a “move” means a move not (just) as a *string*, but (also) as an *event*, namely, the event of \mathcal{M} moving at time c .

Throughout this paper, an n -ary ($n \geq 0$) **arithmetical function** means a total function from n -tuples of natural numbers to natural numbers. As always, “unary” is a synonym of “1-ary”.

Where \mathcal{M} is an HPM and A is a constant game, a **\perp -legal play of A by \mathcal{M}** means a computation branch B of \mathcal{M} such that the run spelled by B is a \perp -legal run of A .

Definition 5.2. Let \mathcal{M} be an HPM, h a unary arithmetical function, and A a constant game. We say that:

1. \mathcal{M} **plays A in amplitude h** iff, in every \perp -legal play of A by \mathcal{M} , whenever \mathcal{M} makes a move α , the magnitude of α does not exceed $h(\ell)$, where ℓ is the background of α ;
2. \mathcal{M} **plays A in space h** iff, in every \perp -legal play of A by \mathcal{M} , the spacecost of any given clock cycle c does not exceed $h(\ell)$, where ℓ is the background of c ;
3. \mathcal{M} **plays A in time h** iff, in every \perp -legal play of A by \mathcal{M} , whenever \mathcal{M} makes a move α , the timecost of α does not exceed $h(\ell)$, where ℓ is the background of α .

When a game A is fixed or clear from the context, we will usually omit a reference to it and simply say that a given machine \mathcal{M} **plays**, or **runs**, in time (space, amplitude) h , or that h is a **bound** for the time (space, amplitude) complexity of \mathcal{M} .

Note that each clause of Definition 5.2 only looks at \perp -legal plays of A . An intuitive reason for this arrangement is that the machine is expected to “act properly” only as long as its adversary plays legally; once Environment makes an illegal move, the (confused) machine is no longer responsible for anything.

Amplitude complexity is to keep track of (set bound on) the magnitudes of \top ’s moves relative to the magnitudes of \perp ’s moves. For instance, assume “magnitude” means “size”, and consider the problem $\top x \perp y (y = 2x)$. This is about computing the function $2x$, the traditional sort of an input-output problem. A machine that solves it would play in amplitude $\lambda \ell. \ell + 1$. This is because the size of \top ’s move (“output”) will exceed the size of \perp ’s move (“input”) by 1.¹⁶ Next, our time complexity concept can be seen to be in the spirit of what is usually called *response time*. The latter generally does not and should not depend on the length of the preceding interaction history. On the other hand, it is not and should not merely be a function of the adversary’s last move, either. A similar characterization applies to our concept of space complexity. All three complexity measures are equally meaningful whether it be in the context of “short-lasting” games (such as the ones represented by the formulas of the later-defined logic **CL12**) or the context of games that may have “very long” and even infinitely long legal runs. It is worthwhile to note one substantial difference between amplitude complexity on one hand, and time and space complexities on the other hand: the former, unlike the latter, only depends on *what* runs the machine generates, regardless of *how* those runs are generated.

In [34], amplitude complexity was not explicitly considered. Also, the definitions of space and time complexities looked at all runs generated by the machine rather than only at \perp -legal runs. Also, timecost was measured as the time elapsed since the last move by either player rather than only by Environment. Finally, as mentioned, [34], just like all other earlier papers on CoL, only dealt with single-work-tape HPMs. Ignoring these minor and (mostly) inconsequential differences, and with magnitude understood as symbolwise length, our present time complexity concept is, in fact, no different from that introduced in [34]. The same is not quite the case for space complexity though. In the version of

¹⁶Unless, of course, the input is $\#0$, in which case the output is also $\#0$. But, as the other sorts of complexities, amplitude complexity is an “at most” kind of a measure.

the underlying HPM model employed in [34] and all other earlier papers on CoL, there was no buffer, and the machine had to construct the moves that it made on its work tape. We have revised that technical detail because the old arrangement did not offer sufficient flexibility. Namely, under the approach of [34], it was impossible to meaningfully talk about cases where amplitude complexity exceeds space complexity — for instance, natural and frequently occurring combinations such as logarithmic space with polynomial amplitude.¹⁷ It should be noted that, while the concept of amplitude complexity was absent in [34], due to the above-mentioned choice of the underlying HPM model, amplitude complexity was automatically bounded by space complexity there.

Definitions 5.1 and 5.2 are relative to the choice of how the magnitude of a move is exactly measured. More than one reasonable choice is possible here. The most straightforward and scalable measure would be to understand the magnitude of a move α as the “raw” length of the string α . However, we here prefer a more refined measure. Namely, we agree that:

Convention 5.3. By the **magnitude** of a move α we mean the number m defined as follows:

- if α does not contain the symbol $\#$, then $m = 0$;
- otherwise m is $|c|$ (the bitwise length of c), where c is the greatest constant such that $\#c$ is a substring of α .¹⁸

The reason for our preference of the above magnitude measure over the raw lengths of moves is purely technical, and otherwise it has no effect on asymptotic (or even much finer) analysis of complexity. The point is that, in games represented by formulas of the language of **CL12**, the only essentially varying parameter that affects move sizes is the length $|c|$ of a constant c chosen for a variable x in a $\Box xE$ or $\Box xE$ component, and we want to focus on that parameter only. For instance, consider the game (represented by) $\Box xP(x) \rightarrow (\top \wedge \Box xP(x))$. If the environment chooses c for x in the antecedent, the raw size of the corresponding move $0.\#c$ will be $|c| + 3$ rather than $|c|$, with 3 being the overhead imposed by the technical prefix “0.#”. It is natural to want to ignore that overhead. Imagine further that the machine responded to the above move by choosing the same constant c for x in (the second conjunct of) the consequent. The raw size $|c| + 5$ of the corresponding actual move $1.1.\#c$ happens to be greater than that of the adversary’s move $0.\#c$, despite the fact that the machine made “the same” — namely, “non-size-increasing” — choice. We do not want to let such necessary yet rather arbitrary “overhead bureaucracy” anyhow interfere with and annoy us in our complexity analysis.

Let A be a constant game, h a unary arithmetical function, and \mathcal{M} an HPM. We say that \mathcal{M} **wins (computes, solves) A in time h** , or that \mathcal{M} **is an h time solution for A** , iff \mathcal{M} plays A in time h and $\mathcal{M} \models A$. We say that A is **computable (solvable) in time h** iff it has an h time solution. Similarly for “**space**” or “**amplitude**” instead of “time”.

¹⁷By the way, the situation is similar with Turing machines in the traditional theory of computation. For instance, a “canonical”, single-tape Turing machine is not fit for defining sublinear space complexity, for which reason such a machine is usually modified through including, in addition to the work tape, two separate tapes for input and output. Such an input tape is reminiscent of our run tape, and output tape is reminiscent of our buffer.

¹⁸Such a c does exist. As we remember, the constant 0, as a bit string, is empty. Thus, due to the omnipresence of the empty string, it is impossible for an occurrence of $\#$ not to be “followed” by *some*, even if empty, constant.

When we say **polynomial time**, it is to be understood as “time h for some polynomial function h ”. Similarly for **polynomial space**, **polynomial amplitude**, **logarithmic space**, etc. More generally, using the asymptotic “Big-O” notation, where g is a unary arithmetical function, “time (space, amplitude) $O(g)$ ” should be understood as “time (space, amplitude) h for some function h with $h \in O(g)$ ”.

6. LANGUAGE OF **CL12** AND ITS SEMANTICS

Logic **CL12** will be axiomatically constructed in Section 7. The present section is merely devoted to its *language*. The building blocks of the formulas of the latter are:

- **Nonlogical predicate letters**, for which we use p, q as metavariables. With each predicate letter is associated a fixed nonnegative integer called its **arity**. We assume that, for any n , there are countably infinitely many n -ary predicate letters.
- **Function letters**, for which we use f, g as metavariables. Again, each function letter comes with a fixed **arity**, and we assume that, for any n , there are countably infinitely many n -ary function letters.
- The binary **logical predicate letter** =.
- **Variables** and **constants**. These are the same as the ones fixed in Section 3 — that is, the elements of VARIABLES and CONSTANTS, respectively.
- Logical connectives and quantifiers $\top, \perp, \neg, \wedge, \vee, \forall, \exists, \sqcap, \sqcup, \sqcap, \sqcup$.

Terms, for which in this paper we use τ as a metavariable, are built from variables, constants and function letters in the standard way. An **atomic formula** is $p(\tau_1, \dots, \tau_n)$, where p is an n -ary predicate letter and the τ_i are terms. When the arity of p is 0, we write p instead of $p()$. Also, we write $\tau_1 = \tau_2$ instead of $=(\tau_1, \tau_2)$, and $\tau_1 \neq \tau_2$ instead of $\neg(\tau_1 = \tau_2)$. **Formulas** are built from atomic formulas, propositional connectives \top, \perp (nullary), \neg (unary), $\wedge, \vee, \sqcap, \sqcup$ (binary), variables and quantifiers $\forall, \exists, \sqcap, \sqcup$ in the standard way, with the exception that, officially, \neg is only allowed to be applied to atomic formulas. The definitions of *free* and *bound* occurrences of variables are standard (with \sqcap, \sqcup acting as quantifiers along with \forall, \exists). A formula with no free occurrences of variables is said to be **closed**.

Note that, conceptually, \top and \perp do not count as atoms. For us, atoms are formulas containing no logical operators. The formulas \top and \perp do not qualify because they *are* (0-ary) logical operators themselves.

$\neg E$, where E is not atomic, will be understood as a standard abbreviation: $\neg\top = \perp$, $\neg\neg E = E$, $\neg(A \wedge B) = \neg A \vee \neg B$, $\neg\sqcap x E = \sqcup x \neg E$, etc. And $E \rightarrow F$ will be understood as an abbreviation of $\neg E \vee F$.

Parentheses will often be omitted — as we just did in the preceding paragraph — if there is no danger of ambiguity. When omitting parentheses, we assume that \neg and the quantifiers have the highest precedence, and \rightarrow has the lowest precedence. An expression $E_1 \wedge \dots \wedge E_n$, where $n \geq 2$, is to be understood as $E_1 \wedge (E_2 \wedge (\dots \wedge (E_{n-1} \wedge E_n) \dots))$. Sometimes we can write this expression for an unspecified $n \geq 0$ (rather than $n \geq 2$). Such a formula, in the case of $n = 1$, should be understood as simply E_1 . Similarly for \vee, \sqcap, \sqcup . As for the case of $n = 0$, the nullary \wedge and \sqcap should be understood as \top while the nullary \vee and \sqcup as \perp .

Sometimes a formula F will be represented as $F(s_1, \dots, s_n)$, where the s_i are variables. When doing so, we do not necessarily mean that each s_i has a free occurrence in F , or that every variable occurring free in F is among s_1, \dots, s_n . However, it *will* always be assumed (usually only implicitly) that the s_i are pairwise distinct, and have no bound occurrences

in F . In the context set by the above representation, $F(\tau_1, \dots, \tau_n)$ will mean the result of replacing, in F , each occurrence of each s_i by term τ_i . When writing $F(\tau_1, \dots, \tau_n)$, it will always be assumed (again, usually only implicitly) that the terms τ_1, \dots, τ_n contain no variables that have bound occurrences in F , so that there are no unsafe collisions of variables when doing replacements.

Similar — well established in the literature — notational conventions apply to terms.

A **sequent** is an expression $E_1, \dots, E_n \circ - F$, where E_1, \dots, E_n ($n \geq 0$) and F are formulas. Here E_1, \dots, E_n is said to be the **antecedent** of the sequent, and F said to be the **succedent**.

By a **free** (resp. **bound**) **variable** of a sequent we shall mean a variable that has a free (resp. bound) occurrence in one of the formulas of the sequent. A **closed sequent** is one with no free variables. For safety and simplicity, throughout the rest of this paper we assume that the sets of all free and bound variables of any sequent that we ever consider — unless strictly implied otherwise by the context — are disjoint. This restriction, of course, does not yield any loss of expressive power as variables can always be renamed so as to satisfy this condition.

An **interpretation**¹⁹ is a function $*$ that:

- Sends the word “*Universe*” to a universe $Universe^*$, called the **universe** of $*$. The **domain** and **denotation** of $*$ are understood as those of $Universe^*$, and are denoted by $Domain^*$ and $Denotation^*$, respectively.
- Sends every n -ary function letter f to an n -ary function $f^*(var_1, \dots, var_n)$ on $Universe^*$.
- Sends every nonlogical n -ary predicate letter p to an n -ary predicate $p^*(var_1, \dots, var_n)$ on $Universe^*$.

We uniquely extend the above $*$ to a mapping that sends each term τ to a function τ^* , and each formula or sequent S to a game S^* over $Universe^*$, by stipulating that:

- (1) Where c is a constant, c^* is $Denotation^*(c)$ (with the individual $Denotation^*(c)$ understood as a function according to Convention 3.3).
- (2) Where s is a variable, s^* is s (with s understood as a function according to Convention 3.3).
- (3) Where f is an n -ary function letter and τ_1, \dots, τ_n are terms, $(f(\tau_1, \dots, \tau_n))^*$ is $f^*(\tau_1^*, \dots, \tau_n^*)$.
- (4) Where τ_1 and τ_2 are terms, $(\tau_1 = \tau_2)^*$ is $\tau_1^* = \tau_2^*$.
- (5) Where p is an n -ary nonlogical predicate letter and τ_1, \dots, τ_n are terms, $(p(\tau_1, \dots, \tau_n))^*$ is $p^*(\tau_1^*, \dots, \tau_n^*)$.
- (6) $*$ commutes with all logical operators, seeing them as the corresponding game operations: \perp^* is \perp , $(E_1 \wedge \dots \wedge E_n)^*$ is $E_1^* \wedge \dots \wedge E_n^*$, $(\Box x E)^*$ is $\Box x(E^*)$, etc.
- (7) Similarly, $*$ sees the sequent symbol $\circ -$ as the same-name game operation: $(E_1, \dots, E_n \circ - F)^*$ is $E_1^*, \dots, E_n^* \circ - F^*$.

When O is a function letter, a predicate letter, a term, a formula or a sequent and $O^* = W$, we say that $*$ **interprets** O as W . We can also refer to such a W as “ O **under interpretation** $*$ ”.

When a given formula is represented as $F(x_1, \dots, x_n)$, we will typically write $F^*(x_1, \dots, x_n)$ instead of $(F(x_1, \dots, x_n))^*$. A similar practice will be used for terms as well.

¹⁹The concept of an interpretation in CoL is usually more general than the present one. Interpretations in our present sense are called **perfect**. But here we omit the word “perfect” as we do not consider any nonperfect interpretations, anyway.

We agree that, for a sequent or formula S , an interpretation $*$ and an HPM \mathcal{M} , whenever we say that \mathcal{M} is a **solution** of S^* or write $\mathcal{M} \models S^*$, we mean that \mathcal{M} is a solution of the (constant) game $\Box x_1 \dots \Box x_n (S^*)$, where x_1, \dots, x_n are exactly the free variables of S , listed according to their lexicographic order. We call the above game the \Box -**closure** of S^* , and denote it by $\Box S^*$.

Note that, for any given sequent or formula S , the **Lr** component of the game $\Box S^*$ does not depend on the interpretation $*$. Hence we can safely say “legal run of $\Box S$ ” — or even just “legal run of S ” — without indicating an interpretation applied to S .

We say that an HPM \mathcal{M} is a (polynomial time, h space, etc.) **uniform** (or **logical**) **solution** of a sequent X iff, for any interpretation $*$, \mathcal{M} is a (polynomial time, h space, etc.) solution of X^* .

Intuitively, a uniform solution is a “purely logical” solution. “Logical” in the sense that it does not depend on the universe and the meanings of the nonlogical symbols (predicate and function letters) — does not depend on a (the) interpretation $*$, that is. It is exactly these kinds of solutions that we are interested in when seeing CoL as a logical basis for applied theories or knowledge base systems. As a universal-utility tool, CoL (or a CoL-based compiler) would have no knowledge of the meanings of those nonlogical symbols (the meanings that will be changing from application to application and from theory to theory), other than what is explicitly given by the target formula and the axioms or the knowledge base of the system.

7. AXIOMATICS OF **CL12**

The purpose of the system **CL12** that we deductively construct in this section is to axiomatize the set of sequents with logical solutions. Our formulation of the system relies on the terminology and notation explained below.

- (1) A **surface occurrence** of a subformula is an occurrence that is not in the scope of any choice operators (\Box , \sqcup , \Box and/or \sqcup).
- (2) A formula not containing choice operators — i.e., a formula of the language of classical first order logic — is said to be **elementary**.
- (3) A sequent is **elementary** iff all of its formulas are so.
- (4) The **elementarization**

$$\|F\|$$

of a formula F is the result of replacing in F all \sqcup - and \sqcup -subformulas by \perp , and all \Box - and \Box -subformulas by \top . Note that $\|F\|$ is (indeed) an elementary formula.

- (5) The **elementarization** $\|G_1, \dots, G_n \circ - F\|$ of a sequent $G_1, \dots, G_n \circ - F$ is the elementary formula

$$\|G_1\| \wedge \dots \wedge \|G_n\| \rightarrow \|F\|.$$

- (6) A sequent is said to be **stable** iff its elementarization is classically valid; otherwise it is **unstable**. By “classical validity”, in view of Gödel’s completeness theorem, we mean provability in some fixed standard version of classical first-order calculus with constants, function letters and $=$, where $=$ is treated as the logical *identity* predicate (so that, say, $x=x$, $x=y \rightarrow (E(x) \rightarrow E(y))$, etc. are provable).
- (7) We will be using the notation

$$F[E]$$

to mean a formula F together with some (single) fixed surface occurrence of a subformula E . Using this notation sets a context, in which $F[H]$ will mean the result of replacing in $F[E]$ the (fixed) occurrence of E by H . Note that here we are talking about some *occurrence* of E . Only that occurrence gets replaced when moving from $F[E]$ to $F[H]$, even if the formula also had some other occurrences of E .

- (8) By a **rule** (of inference) in this section we mean a binary relation $\mathbb{Y}\mathcal{R}X$, where $\mathbb{Y} = \langle Y_1, \dots, Y_n \rangle$ is a finite sequence of sequents and X is a sequent. **Instances** of such a relation are schematically written as

$$\frac{Y_1, \dots, Y_n}{X},$$

where Y_1, \dots, Y_n are called the **premises**, and X is called the **conclusion**. Whenever $\mathbb{Y}\mathcal{R}X$ holds, we say that X **follows** from \mathbb{Y} by \mathcal{R} .

- (9) Expressions such as \vec{G}, \vec{K}, \dots will usually stand for finite sequences of formulas. The standard meaning of an expression such as \vec{G}, F, \vec{K} should also be clear.

THE RULES OF CL12

CL12 has the six rules listed below, with the following additional conditions/explanations:

- (1) In \sqcup -Choose and \sqcap -Choose, $i \in \{0, 1\}$.
- (2) In \sqcup -Choose and \sqcap -Choose, t is either a constant or a variable with no bound occurrences in the premise, and $H(t)$ is the result of replacing by t all free occurrences of x in $H(x)$ (rather than vice versa).

$$\frac{\text{\textbf{\(\(\sqcup\)-Choose}}\)} \\ \vec{G} \circ- F[H_i]}{\vec{G} \circ- F[H_0 \sqcup H_1]}$$

$$\frac{\text{\textbf{\(\(\sqcap\)-Choose}}\)} \\ \vec{G}, E[H_i], \vec{K} \circ- F}{\vec{G}, E[H_0 \sqcap H_1], \vec{K} \circ- F}$$

$$\frac{\text{\textbf{\(\(\sqcup\)-Choose}}\)} \\ \vec{G} \circ- F[H(t)]}{\vec{G} \circ- F[\sqcup x H(x)]}$$

$$\frac{\text{\textbf{\(\(\sqcap\)-Choose}}\)} \\ \vec{G}, E[H(t)], \vec{K} \circ- F}{\vec{G}, E[\sqcap x H(x)], \vec{K} \circ- F}$$

Replicate

$$\frac{\vec{G}, E, \vec{K}, E \circ- F}{\vec{G}, E, \vec{K} \circ- F}$$

Wait

$$\frac{Y_1, \dots, Y_n}{X} \quad (n \geq 0), \text{ where all of the following five conditions are satisfied:}$$

- (1) **\sqcap -Condition:** Whenever X has the form $\vec{G} \circ - F[H_0 \sqcap H_1]$, both of the sequents $\vec{G} \circ - F[H_0]$ and $\vec{G} \circ - F[H_1]$ are among Y_1, \dots, Y_n .
- (2) **\sqcup -Condition:** Whenever X has the form $\vec{G}, E[H_0 \sqcup H_1], \vec{K} \circ - F$, both of the sequents $\vec{G}, E[H_0], \vec{K} \circ - F$ and $\vec{G}, E[H_1], \vec{K} \circ - F$ are among Y_1, \dots, Y_n .
- (3) **\sqcap -Condition:** Whenever X has the form $\vec{G} \circ - F[\sqcap x H(x)]$, for some variable y not occurring in X , the sequent $\vec{G} \circ - F[H(y)]$ is among Y_1, \dots, Y_n . Here and below, $H(y)$ is the result of replacing by y all free occurrences of x in $H(x)$ (rather than vice versa).
- (4) **\sqcup -Condition:** Whenever X has the form $\vec{G}, E[\sqcup x H(x)], \vec{K} \circ - F$, for some variable y not occurring in X , the sequent $\vec{G}, E[H(y)], \vec{K} \circ - F$ is among Y_1, \dots, Y_n .
- (5) **Stability condition:** X is stable.

As will be seen in Section 8, each rule — seen bottom-up — encodes an action that a winning strategy should take in a corresponding situation, and the name of each rule is suggestive of that action. For instance, Wait (indeed) prescribes the strategy to wait till the adversary moves. This explains why we have called “Replicate” the rule which otherwise is nothing but what is commonly known as Contraction.

A **CL12-proof** of a sequent X is a sequence X_1, \dots, X_n of sequents, with $X_n = X$, such that, each X_i follows by one of the rules of **CL12** from some (possibly empty in the case of Wait, and certainly empty in the case of $i = 1$) set \mathcal{P} of premises such that $\mathcal{P} \subseteq \{X_1, \dots, X_{i-1}\}$. When a **CL12-proof** of X exists, we say that X is **provable** in **CL12**, and write **CL12** $\vdash X$.

A **CL12-proof** of a formula F will be understood as a **CL12-proof** of the empty-antecedent sequent $\circ - F$. Accordingly, **CL12** $\vdash F$ means **CL12** $\vdash \circ - F$.

Fact 7.1. **CL12** is a conservative extension of classical logic. That is, an elementary sequent $E_1, \dots, E_n \circ - F$ is provable in **CL12** iff the formula $E_1 \wedge \dots \wedge E_n \rightarrow F$ is valid in the classical sense.

Proof. Assume E_1, \dots, E_n, F are elementary formulas. If $E_1 \wedge \dots \wedge E_n \rightarrow F$ is classically valid, then $E_1, \dots, E_n \circ - F$ follows from the empty set of premises by Wait. And if $E_1 \wedge \dots \wedge E_n \rightarrow F$ is not classically valid, then $E_1, \dots, E_n \circ - F$ cannot be the conclusion of any of the rules of **CL12** except Replicate. However, applying (bottom-up) Replicate does not take us any closer to finding a proof of the sequent, as the premise still remains an unstable elementary sequent. \square

CL12 can also be seen to be a conservative extension of the earlier known logic **CL3** studied in [19].²⁰ The latter is nothing but the empty-antecedent fragment of **CL12** without function letters and identity.

Example 7.2. In this example, \times is a binary function letter and 3 is a unary function letter. We write $x \times y$ and x^3 instead of $\times(x, y)$ and $^3(x)$, respectively. The following sequence of sequents is a **CL12-proof** of the sequent (3.1) from Example 3.8. It may be worth observing that the strategy used by \top in that example, in a sense, “follows” our present proof step-by-step in the bottom-up direction. And this is no accident: as we are going to see in the course of proving the soundness of **CL12**, every **CL12-proof** rather directly encodes a winning strategy.

1. $\forall x(x^3 = (x \times x) \times x), t = s \times s, r = t \times s \circ - r = s^3$ Wait: (no premises)

²⁰Essentially the same logic, called **L**, was in fact known as early as in [13].

2. $\forall x(x^3=(x \times x) \times x), t=s \times s, r=t \times s \multimap \sqcup y(y=s^3)$ \sqcup -Choose: 1
3. $\forall x(x^3=(x \times x) \times x), t=s \times s, \sqcup z(z=t \times s) \multimap \sqcup y(y=s^3)$ Wait: 2
4. $\forall x(x^3=(x \times x) \times x), t=s \times s, \sqcap y \sqcup z(z=t \times y) \multimap \sqcup y(y=s^3)$ \sqcap -Choose: 3
5. $\forall x(x^3=(x \times x) \times x), t=s \times s, \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcup y(y=s^3)$ \sqcap -Choose: 4
6. $\forall x(x^3=(x \times x) \times x), \sqcup z(z=s \times s), \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcup y(y=s^3)$ Wait: 5
7. $\forall x(x^3=(x \times x) \times x), \sqcap y \sqcup z(z=s \times y), \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcup y(y=s^3)$ \sqcap -Choose: 6
8. $\forall x(x^3=(x \times x) \times x), \sqcap x \sqcap y \sqcup z(z=x \times y), \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcup y(y=s^3)$ \sqcap -Choose: 7
9. $\forall x(x^3=(x \times x) \times x), \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcup y(y=s^3)$ Replicate: 8
10. $\forall x(x^3=(x \times x) \times x), \sqcap x \sqcap y \sqcup z(z=x \times y) \multimap \sqcap x \sqcup y(y=x^3)$ Wait: 9

Example 7.3. The formula $\forall x p(x) \rightarrow \sqcap x p(x)$ is provable in **CL12**. It follows from $\forall x p(x) \rightarrow p(y)$ by Wait. The latter, in turn, follows by Wait from the empty set of premises.

On the other hand, the formula $\sqcap x p(x) \rightarrow \forall x p(x)$, i.e. $\sqcup x \neg p(x) \vee \forall x p(x)$, is not provable. Indeed, its elementarization is $\perp \vee \forall x p(x)$, which is not classically valid. Hence $\sqcup x \neg p(x) \vee \forall x p(x)$ cannot be derived by Wait. Replicate can also be dismissed for obvious reasons. This leaves us with \sqcup -Choose. But if $\sqcup x \neg p(x) \vee \forall x p(x)$ is derived by \sqcup -Choose, then the premise should be $\neg p(t) \vee \forall x p(x)$ for some variable or constant t . The latter, however, is a classically non-valid elementary formula and hence, by Fact 7.1, is not provable.

Example 7.4. The formula $\sqcap x \sqcup y(p(x) \rightarrow p(y))$ is provable in **CL12** as follows:

1. $p(s) \rightarrow p(s)$ Wait:
2. $\sqcup y(p(s) \rightarrow p(y))$ \sqcup -Choose: 1
3. $\sqcap x \sqcup y(p(x) \rightarrow p(y))$ Wait: 2

On the other hand, the formula $\sqcup y \sqcap x(p(x) \rightarrow p(y))$ can be seen to be unprovable, even though its classical counterpart $\exists y \forall x(p(x) \rightarrow p(y))$ is a classically valid elementary formula and hence provable in **CL12**.

Example 7.5. While the formula $\forall x \exists y(y=f(x))$ is classically valid and hence provable in **CL12**, its constructive counterpart $\sqcap x \sqcup y(y=f(x))$ can be easily seen to be unprovable. This is no surprise. In view of the expected soundness of **CL12**, provability of $\sqcap x \sqcup y(y=f(x))$ would imply that every function f is computable (furthermore, computable in a “uniform” way), which, of course, is not the case.

Exercise 7.6. To see the resource-consciousness of **CL12**, show that it does not prove $p \sqcap q \rightarrow (p \sqcap q) \wedge (p \sqcap q)$, even though this formula has the form $F \rightarrow F \wedge F$ of a classical tautology. Then show that, in contrast, **CL12** proves the *sequent* $p \sqcap q \multimap (p \sqcap q) \wedge (p \sqcap q)$ because, unlike the antecedent of a \rightarrow -combination, the antecedent of a \multimap -combination is reusable (through Replicate).

Exercise 7.7. Show that **CL12** $\vdash \sqcup x \sqcap y p(x, y) \multimap \sqcup x(\sqcap y p(x, y) \wedge \sqcap y p(x, y))$. Then observe that, on the other hand, **CL12** does not prove any of the formulas

$$\begin{aligned}
 \sqcup x \sqcap y p(x, y) &\rightarrow \sqcup x(\sqcap y p(x, y) \wedge \sqcap y p(x, y)); \\
 \sqcup x \sqcap y p(x, y) \wedge \sqcup x \sqcap y p(x, y) &\rightarrow \sqcup x(\sqcap y p(x, y) \wedge \sqcap y p(x, y)); \\
 \sqcup x \sqcap y p(x, y) \wedge \sqcup x \sqcap y p(x, y) \wedge \sqcup x \sqcap y p(x, y) &\rightarrow \sqcup x(\sqcap y p(x, y) \wedge \sqcap y p(x, y)); \\
 &\dots
 \end{aligned}$$

Intuitively, this contrast is due to the fact that, even though both $\downarrow A$ and $\wedge A = A \wedge A \wedge \dots$ are resources allowing to reuse A any number of times, the “branching” form of reusage offered by $\downarrow A$ is substantially stronger than the “parallel” form of reusage offered by $\wedge A$. $\downarrow \sqcup x \sqcap y p(x, y) \rightarrow \sqcup x (\sqcap y p(x, y) \wedge \sqcap y p(x, y))$ is a valid principle of CoL while $\wedge \sqcup x \sqcap y p(x, y) \rightarrow \sqcup x (\sqcap y p(x, y) \wedge \sqcap y p(x, y))$ is not.

8. SOUNDNESS AND COMPLETENESS OF **CL12**

Throughout this paper, the letter ℓ is used as a variable for the *background parameter* in terms of which complexity is measured. When writing $f(\ell)$, for whatever arithmetical function f , in some cases it is to be understood as the quantity $f(\ell)$, while in some other cases as the function $\lambda \ell. f(\ell)$. We typically do not explicitly use the lambda operator, and whether $f(\ell)$ is meant to stand for the function $\lambda \ell. f(\ell)$ or for the value of that function at the argument ℓ should be determined based on the context.

We say that a computation branch (play) of a given HPM \mathcal{M} is **provident** iff there are infinitely many configurations in it with empty buffer contents. Intuitively, this means that there are no “unfinished moves”: every (nonempty) move that \mathcal{M} has started to construct in its buffer was eventually made in the play.

We define the **native magnitude** of a sequent X as the smallest integer $\mathfrak{c} \geq 0$ such that X contains no constant c with $|c| > \mathfrak{c}$. Thus, if X contains no constants, or no constants other than 0, its native magnitude is 0 (remember from Section 3 that, when denoting a constant, “0” stands for the empty bit string ϵ and hence its size is 0 rather than 1).

By a **minimal amplitude logical solution** of a sequent X we shall mean a $\max\{\ell, \mathfrak{c}\}$ amplitude logical solution of X , where \mathfrak{c} is the native magnitude of X . Of course, $\max\{\ell, \mathfrak{c}\} = O(\ell)$, simply meaning linear amplitude. However, unlike space or time complexities, we may not always be willing to settle for merely asymptotic analysis when it comes to amplitude complexity.

Definition 8.1. We say that a logical solution \mathcal{M} of a sequent X is **well behaved** iff the following five conditions are satisfied:

- (1) There is an integer \mathfrak{r} such that, in every \perp -legal play of X^{21} by \mathcal{M} , at most \mathfrak{r} replicative moves are made in the antecedent.
- (2) There is an integer \mathfrak{k} such that, in every \perp -legal play of X by \mathcal{M} , the run-tape head of \mathcal{M} revisits any given labmove $\wp\alpha$ at most \mathfrak{k} times. Here by **revisiting** such a $\wp\alpha$ we mean entering the cell containing the last symbol of this string in the right-to-left direction; that is, coming back to the labmove after it has already been fully scanned.
- (3) In every \perp -legal play of X by \mathcal{M} , every move that \mathcal{M} makes in one of the components of the antecedent of X is focused.
- (4) Every \perp -legal play of X by \mathcal{M} is provident.
- (5) \mathcal{M} has only one work tape.

Theorem 8.2. *For any sequent X , the following conditions are equivalent:*

- (i) **CL12** $\vdash X$.
- (ii) X has a logical solution.

²¹Remember our convention (page 28), according to which we may unambiguously say “legal play/run of X ” even if, strictly speaking, X is not a game, but rather turns into one only after some interpretation* is applied to it.

(iii) X has a well behaved, minimal amplitude, constant space and linear time logical solution \mathcal{M} .

Furthermore, the implication (i) \Rightarrow (iii) holds in the strong sense that there is an effective procedure that takes an arbitrary **CL12**-proof \mathbb{P} of an arbitrary sequent X and constructs an HPM \mathcal{M} satisfying (iii).

Remark 8.3. In the “Furthermore” clause of the above theorem, “effective” can be strengthened to “efficient”. Indeed, let \mathcal{M}' be an HPM that, using the procedure whose existence is claimed in the “Furthermore” clause, constructs \mathcal{M} and then, simulating \mathcal{M} , plays just as \mathcal{M} would play. It can be seen with some thought that such an \mathcal{M}' (as opposed to \mathcal{M} itself) is constructed from \mathbb{P} in linear time. It inherits, from \mathcal{M} , the property of being a well-behaved, minimal amplitude logical solution of X . With properly arranged details of simulation, the space and time complexities of \mathcal{M}' , asymptotically, can also be easily achieved to be the same as those of \mathcal{M} . Thus \mathcal{M}' , just like \mathcal{M} , satisfies condition of (iii) of the theorem.

A similar comment applies to the later-stated Theorems 10.3 and 11.1.

We begin our proof of Theorem 8.2 by noting that the implication (iii) \Rightarrow (ii) is trivial. The implication (ii) \Rightarrow (i) was proven in [34], and that proof is reproduced in Appendix C of this paper.²² So, here we only need to verify the implication (i) \Rightarrow (iii), in the strong sense of the “Furthermore” clause of the theorem. This is what the entire rest of the present section is devoted to, for the exception of its last subsection.

Our proof of (i) \Rightarrow (iii) proceeds by induction on the length of (the number of sequents involved in) a **CL12**-proof of X and, as such, is nothing but a combination of six cases, corresponding to the six rules of **CL12** by which the final sequent X could have been derived from its premises (if any).

In each case, our efforts will be focused on showing how to construct a single-work-tape HPM \mathcal{M} — a logical solution of the conclusion — from an arbitrary instance

$$\frac{Y_1, \dots, Y_n}{X}$$

($n \geq 0$) of the rule and arbitrary HPMs $\mathcal{N}_1, \dots, \mathcal{N}_n$ — well behaved, minimal amplitude, constant space and linear time logical solutions of the premises, that exist according to the induction hypothesis. It will be immediately clear from our description of \mathcal{M} that it is constructed effectively (thus taking care of the “Furthermore” claim of the theorem), and that the work of \mathcal{M} in no way depends on an interpretation $*$ applied to the sequents involved, so that the solution is logical.

Since an interpretation $*$ is typically irrelevant in such proofs, we will usually omit mentioning it and write simply S where, strictly speaking, S^* is meant. That is, we identify formulas or sequents with the games into which they turn once an interpretation is applied to them. Accordingly, in contexts where S^* has to be understood as $\Box S^*$ anyway (e.g., when talking about computability of S^*), we may further omit “ \Box ” and write S instead of $\Box S$.

In all cases, whether it be the description of \mathcal{M} ’s work or our further analysis of its play of the target sequent X , we shall implicitly rely on what in [19, 27] is called the **Clean Environment Assumption**, according to which \mathcal{M} ’s adversary never makes illegal moves.

²²As mentioned, the underlying HPM model dealt with in [34] was not exactly the same as the present one; yet the above proof goes through in either case without a need for any readjustments whatsoever.

That is, all plays of X by \mathcal{M} are \perp -legal. Such an assumption is legitimate because, once Environment makes an illegal move in a given play, \mathcal{M} is an automatic winner regardless of how it works afterwards, trivially satisfying the required or relevant well-behavedness and complexity conditions because the latter only look at \perp -legal runs.

8.1. \sqcup -Choose.

$$\frac{\vec{G} \circ - F[H_i]}{\vec{G} \circ - F[H_0 \sqcup H_1]}$$

Assume (induction hypothesis) that \mathcal{N} is a well behaved, minimal amplitude, constant space and linear time logical solution of the premise $\vec{G} \circ - F[H_i]$ ($i \in \{0, 1\}$). We want to construct a well behaved, minimal amplitude, constant space and linear time logical solution \mathcal{M} of the conclusion $\vec{G} \circ - F[H_0 \sqcup H_1]$.

The basic idea for the strategy followed by \mathcal{M} is very simple: \sqcup -Choose most directly encodes an action that \mathcal{M} should perform in order to successfully solve the conclusion. Namely, \mathcal{M} should choose H_i in the $H_0 \sqcup H_1$ component and then continue playing like \mathcal{N} . \mathcal{M} wins because the above initial move — call it $\beta_{initial}$ — brings the conclusion down to the premise, and \mathcal{N} wins the latter.

In more detail, this is how \mathcal{M} works. Let k be the number of free variables in $\vec{G} \circ - F[H_0 \sqcup H_1]$. At the beginning of the play, \mathcal{M} waits till Environment chooses constants for all of those free variables. That is, until k \perp -labeled moves appear on the run tape. If Environment never makes k moves, \mathcal{M} wins. Besides, in this case, \mathcal{M} will not be billed for any time, because it did not make any moves. Nor will it be billed for any space, because it does not write anything on the work tape while waiting. Suppose now Environment does make k moves. Then \mathcal{M} goes ahead and makes the above-mentioned move $\beta_{initial}$ signifying selecting H_i . By this time, \mathcal{M} 's space consumption remains 0 as in the preceding case. As for the time bill, with a moment's thought we can see that it will be at most proportional to $k\ell$, where ℓ is the background of $\beta_{initial}$.²³ The time consumption so far is linear, that is.

After \mathcal{M} chooses H_i , the main part of its work consists in simulating \mathcal{N} forever. To do this, at any time \mathcal{M} needs to remember the “current” configuration of \mathcal{N} , which at the beginning is \mathcal{N} 's start configuration in the scenario where \mathcal{N} 's adversary has made all of the initial k moves right at the beginning of the play, and which will be continuously updated to the “next configuration” by the simulation routine. Maintaining this information requires remembering \mathcal{N} 's “current” state, work-tape content, work-tape head location, run-tape content, run-tape head location, and the content of \mathcal{N} 's buffer. The first piece of this information — \mathcal{N} 's state — \mathcal{M} remembers through its own state memory without utilizing any work-tape space. Next, \mathcal{M} remembers \mathcal{N} 's work-tape content and work-tape head location by fully synchronizing those with its own work-tape content and work-tape head location. That is, at any time, the work-tape content of \mathcal{M} is an exact copy of that of \mathcal{N} , and the work-tape head location of \mathcal{M} also coincides with that of \mathcal{N} .²⁴ There is no need

²³Note that, if Environment was slow in making moves, \mathcal{M} is not billed for waiting, because \mathcal{M} 's time counter is reset to 0 every time Environment moves; thus, \mathcal{M} 's time bill will be the biggest in the scenario where Environment made all k moves at once, right during the first clock cycle. In this case \mathcal{M} will have to read all k moves, plus spend a constant amount of time on making the move $\beta_{initial}$.

²⁴In fact, since \mathcal{N} runs in constant space, \mathcal{M} could remember these two pieces of information (just like the information on \mathcal{N} 's state) in its state memory rather than on its work tape.

for \mathcal{M} to remember \mathcal{N} 's run-tape content (which would require linear rather than constant space) on its work tape. That is because all information about the content of the imaginary run tape of \mathcal{N} resides on \mathcal{M} 's own run tape. \mathcal{M} does not need to remember on its work tape \mathcal{N} 's run-tape head location either (which otherwise would require a logarithmic rather than constant amount of space). Rather, just as in the case of the work-tape head location, \mathcal{M} keeps track of this information through correspondingly positioning its own run-tape head, usually (with an exception discussed below) letting it scan the same symbol as \mathcal{N} 's imaginary run-tape head is scanning. Finally, \mathcal{M} keeps the content of its buffer identical to that of \mathcal{N} 's imaginary buffer, and thus consumes no work-tape space on remembering the corresponding piece of information. Of course, in addition to remembering the “current” configuration of \mathcal{N} , \mathcal{M} also needs to maintain (the never-changing) full information on \mathcal{N} 's transition function. But, again, there is no need to keep this information on the work tape — it is understood that, instead, it is simply “hardwired” into \mathcal{M} 's finite control.

The purpose of simulation is to play as \mathcal{N} would play. Namely, every time \mathcal{N} goes from one configuration to another,²⁵ \mathcal{M} correspondingly updates its representation of \mathcal{N} 's “current” configuration using the built-in knowledge of \mathcal{N} 's transition function. Among other things, this means that, every time \mathcal{N} adds some nonempty string to the content of its buffer, \mathcal{M} puts the same string into its own buffer. Every time \mathcal{N} makes a move (enters a move state), \mathcal{M} does the same. And every time \mathcal{N} looks up a symbol on its run tape, \mathcal{M} looks up the same symbol on its own run tape and feeds the corresponding information back to the simulation routine. Since the magnitude of $\beta_{initial}$ is 0 and all other moves that \mathcal{M} makes are also made in the corresponding scenario by \mathcal{N} , \mathcal{M} runs in precisely the same amplitude as \mathcal{N} does. Next, note that \mathcal{M} utilizes the same amount of work-tape space as \mathcal{N} does. And the time that \mathcal{M} consumes is also the same (or, rather, “essentially the same” as we are going to see below) as the time consumed by \mathcal{N} . So, \mathcal{M} inherits \mathcal{N} 's constant space and linear time complexities. Observe that the native magnitude \mathfrak{c} of the premise $\vec{G} \circ - F[H_i]$ does not exceed the native magnitude of the conclusion $\vec{G} \circ - F[H_0 \sqcup H_1]$. Hence, the fact that \mathcal{M} runs in the same amplitude $\max\{\ell, \mathfrak{c}\}$ as \mathcal{N} does additionally implies that, as desired, \mathcal{M} is a minimal amplitude logical solution of $\vec{G} \circ - F[H_0 \sqcup H_1]$.

One detail of simulation suppressed so far requires further attention. In the above explanation we pretended that the content of \mathcal{M} 's run tape (more precisely, of its so-far-scanned portion) was the same as that of the simulated \mathcal{N} 's imaginary run tape. This, however, is not exactly so. The difference between the two contents is that \mathcal{M} 's run tape will have $\top\beta_{initial}$ as an additional labmove. This is not a serious problem though. Let us call the portion of \mathcal{M} 's run tape on which $\top\beta_{initial}$ is written the *special zone*. \mathcal{M} handles the complication caused by the presence of $\top\beta_{initial}$ by remembering (through its state memory), at any time, the integer D indicating “how many labmoves away its run-tape head is from the special zone”. Namely, if the run-tape head is scanning a cell in the special zone, then $D = 0$; if it is scanning the labmove written immediately after (resp. before) the special zone, then $D = 1$ (resp. $D = -1$); if it is scanning the labmove immediately after (resp. before) the previously mentioned labmove, then $D = 2$ (resp. $D = -2$); and so on. As long as D is not 0, the simulation proceeds in the “normal” way described in the preceding two paragraphs, with \mathcal{M} 's run-tape head scanning precisely the same symbol as \mathcal{N} 's imaginary run-tape head is scanning, every time moving that head in the same direction

²⁵And, of course, this happens “all the time”.

as the direction in which \mathcal{N} 's run-tape head moves. An exception occurs when \mathcal{M} 's run-tape head enters the special zone. When this happens, \mathcal{M} temporarily stops simulation, and continues moving its run-tape head in the same direction (left-to-right or right-to-left) in which it entered the special zone. This continues until \mathcal{M} 's run-tape head exits the special zone, after which the normal simulation resumes. Thus, every time \mathcal{M} crosses the special zone from one end to the other, \mathcal{M} spends a constant amount of additional time. Remembering that \mathcal{N} is well behaved,²⁶ this event (crossing the special zone) will only happen a constant number of times. Thus, overall, the additional time consumption of \mathcal{M} (additional to the quantity established in the preceding paragraph) is also constant, meaning that, asymptotically, \mathcal{M} 's time complexity is not affected. Finally, it is immediately obvious that \mathcal{M} inherits the well-behavedness of \mathcal{N} . To summarize, \mathcal{M} is a well behaved, minimal amplitude, constant space and linear time logical solution of $\vec{G} \circ - F[H_0 \sqcup H_1]$, as desired.

8.2. \sqcap -Choose. This case is similar to the preceding one.

8.3. \sqcup -Choose.

$$\frac{\vec{G} \circ - F[H(\mathfrak{t})]}{\vec{G} \circ - F[\sqcup x H(x)]}$$

Taking into account that the choice existential quantifier is nothing but a “long” choice disjunction, this case is also rather similar to the case of \sqcup -Choose. Assume that \mathcal{N} is a well behaved, minimal amplitude, constant space and linear time logical solution of the premise. We want to construct a same kind of a solution \mathcal{M} for the conclusion.

First, consider the case of \mathfrak{t} being a constant that occurs (not only in the premise but also) in the conclusion. We let \mathcal{M} be a machine that works as follows. At the beginning, \mathcal{M} waits till Environment specifies some constants for all free variables of $\vec{G} \circ - F[\sqcup x H(x)]$. For readability, we continue referring to the resulting game as $\vec{G} \circ - F[\sqcup x H(x)]$, even though, strictly speaking, it is $e[\vec{G} \circ - F[\sqcup x H(x)]]$, where e is a valuation that agrees with the choices that Environment has just made for the free variables of the sequent. Now \mathcal{M} makes the move $\beta_{initial}$ that brings $\vec{G} \circ - F[\sqcup x H(x)]$ down to $\vec{G} \circ - F[H(\mathfrak{t})]$. For instance, if $\vec{G} \circ - F[\sqcup x H(x)]$ is $\vec{G} \circ - E \wedge (K \vee \sqcup x H(x))$ and thus $\vec{G} \circ - F[H(\mathfrak{t})]$ is $\vec{G} \circ - E \wedge (K \vee H(\mathfrak{t}))$, then $\beta_{initial}$ is 1.1.1.# \mathfrak{t} . After this move, \mathcal{M} “turns itself into \mathcal{N} ” in the same fashion as in the proof of the case of \sqcup -Choose. The only difference between the present case and the case of \sqcup -Choose worth pointing out is that the magnitude of $\beta_{initial}$ is no longer 0 — rather, it is $|\mathfrak{t}|$. But note that, since \mathfrak{t} occurs in the conclusion, $|\mathfrak{t}|$ does not exceed the native magnitude \mathfrak{c} of the conclusion. This guarantees that \mathcal{M} runs in amplitude $\max\{\ell, \mathfrak{c}\}$, as desired. As for the bounds on time and space complexities, their being as desired is guaranteed by the same reasons as in the case of \sqcup -Choose. So is \mathcal{M} 's being well behaved.

Next, consider the case of \mathfrak{t} being a variable that is among the free variables of the conclusion. It can be handled in a similar way to the above, with the only difference that now $\beta_{initial}$ is a move signifying choosing, for x , the constant chosen by Environment for \mathfrak{t} . To make such a move, \mathcal{M} finds on its run tape the corresponding \perp -labeled move, and then copies its symbols into its buffer in a single pass (after first putting into the buffer the necessary technical/addressing prefix for the move, of course). This only takes a linear amount of time and a constant (if any) amount of space. In analyzing the time overhead

²⁶What is relevant here is condition 2 of Definition 8.1.

imposed on simulation by the presence of the labmove $\top\beta_{initial}$ on \mathcal{M} 's run tape, it is worth pointing out that a difference between the present case and the case of \sqcup -Choose is that, while the length of the “special zone” in the former case was constant, in the present case it is linear. This is however fine, because, as before, the special zone will be crossed only a constant number of times, so that the time overhead remains linear. As for the amplitude complexity of \mathcal{M} , it also remains as desired, because the magnitude of $\beta_{initial}$ does not exceed its background ℓ .

The remaining possibility to consider is that of \mathfrak{t} being either a constant with no occurrence in the conclusion, or a variable with no (free) occurrence in the conclusion. By induction on the length of the proof of the premise, it is very easy to see that such a proof remains a proof with \mathfrak{t} replaced by 0 (or any other constant for that matter) everywhere in it. Thus, we may assume that the premise of $\vec{G} \circ - F[\sqcup x H(x)]$ is simply $\vec{G} \circ - F[H(0)]$, i.e., that $\mathfrak{t} = 0$. Now, this case will be handled exactly as the earlier case of \mathfrak{t} being a constant that occurs in the conclusion. When dealing with that case, the assumption that \mathfrak{t} had an occurrence in the conclusion was only needed to guarantee that $|\mathfrak{t}|$ did not exceed the native magnitude of the conclusion. In the present case, i.e. when $\mathfrak{t} = 0$, we can no longer assume that \mathfrak{t} occurs in the conclusion. However, $|\mathfrak{t}|$ still does not exceed the native magnitude of the conclusion because $|\mathfrak{t}| = 0$.

8.4. \sqcap -Choose. This case is similar to the preceding one.

8.5. Replicate.

$$\frac{\vec{G}, E, \vec{K}, E \circ - F}{\vec{G}, E, \vec{K} \circ - F}$$

Remembering that we agreed to see no distinction between sequents and the games they represent, and disabbreviating \vec{G} , \vec{K} and $\circ -$, the premise and the conclusion of this rule can be rewritten as the following two games, respectively:

$$\circlearrowleft G_1 \wedge \dots \wedge \circlearrowleft G_m \wedge \circlearrowleft E \wedge \circlearrowleft K_1 \wedge \dots \wedge \circlearrowleft K_n \wedge \circlearrowleft E \rightarrow F; \quad (8.1)$$

$$\circlearrowleft G_1 \wedge \dots \wedge \circlearrowleft G_m \wedge \circlearrowleft E \wedge \circlearrowleft K_1 \wedge \dots \wedge \circlearrowleft K_n \rightarrow F. \quad (8.2)$$

Assume \mathcal{N} is a well behaved, minimal amplitude, constant space and linear time logical solution of the premise. We let a logical solution \mathcal{M} of the conclusion be a machine that works as follows.

After Environment chooses some constants for all free variables of (8.2), \mathcal{M} makes a replicative move in the $\circlearrowleft E$ component of the latter, thus bringing the game down to

$$\circlearrowleft G_1 \wedge \dots \wedge \circlearrowleft G_m \wedge \circlearrowleft (E \circ E) \wedge \circlearrowleft K_1 \wedge \dots \wedge \circlearrowleft K_n \rightarrow F \quad (8.3)$$

(more precisely, it will be not (8.3) but $e[(8.3)]$, where e is a valuation that agrees with Environment's choices for the free variables of the sequent. As we did earlier, however, notationally we ignore this difference). Later the section on \mathcal{M} 's run tape containing this replicative move will be treated as a “special zone” in the same way as in the case of \sqcup -Choose, so we may safely pretend that it is simply not there.

Now we need to observe that (8.3) is “essentially the same as” (8.1), so that \mathcal{M} , through simulation, can continue playing (8.3) “essentially as” \mathcal{N} would play (8.1) in the scenario where the adversary of \mathcal{N} chose the same constants for free variables as the adversary of

\mathcal{M} just did. All that \mathcal{M} needs to do to account for the minor technical differences between (8.3) and (8.1) is to make a very simple “reinterpretation” of moves. Namely:

- Any move made within any of the $\downarrow G_i$ or $\downarrow K_i$ components of (8.3) \mathcal{M} sees exactly as \mathcal{N} would see the same move in the same component of (8.1), and vice versa.
- Any (focused) move of the form $0.\alpha$ (resp. $1.\alpha$) made in the $\downarrow(E \circ E)$ component of (8.3) \mathcal{M} sees as \mathcal{N} would see the move α as if it was made in the first (resp. second) $\downarrow E$ component of (8.1), and vice versa.
- Any (unfocused) move of the form $\epsilon.\alpha$ made by Environment in the $\downarrow(E \circ E)$ component of (8.3) \mathcal{M} sees as \mathcal{N} would see the move α made twice (but on the same clock cycle) by its adversary: once in the first $\downarrow E$ component of (8.1), and once in the second $\downarrow E$ component of (8.1).

Due to the third clause above, the difference between the contents of the run tapes of \mathcal{M} and the simulated \mathcal{N} in the present case is greater than in the case of \sqcup -Choose. This means that the run-tape head of \mathcal{M} may need to make certain additional back-and-forth journeys to properly maintain simulation. Without going into details, we simply point out that this can be handled in a way similar to the way we handled the presence of the “special zone” when discussing \sqcup -Choose. In view of the fact that \mathcal{N} is well behaved, with some thought it can be seen that \mathcal{M} only needs a constant amount of extra space and a linear amount of extra time for this, so that \mathcal{M} , just like \mathcal{N} , runs in constant space and linear time. \mathcal{M} ’s amplitude complexity obviously remains the same as that of \mathcal{N} , which implies that \mathcal{M} is a minimal amplitude logical solution as required. Finally, it can be seen that \mathcal{M} inherits \mathcal{N} ’s well-behavedness as well.

8.6. Wait.

$$\frac{Y_1, \dots, Y_n}{X}$$

(where $n \geq 0$ and the \sqcap -, \sqcup -, \sqcap -, \sqcup - and Stability conditions are satisfied).

We shall rely on the following lemma. It can be verified by a straightforward induction on the complexity of Z , which we omit. Remember that $\langle \rangle$ stands for the empty run, and $\|Z\|$ for the elementarization of Z .

Lemma 8.4. *For any sequent Z , valuation e and interpretation $*$, $\mathbf{Wn}_e^{Z*} \langle \rangle = \mathbf{Wn}_e^{\|Z\|*} \langle \rangle$.*

Assume $\mathcal{N}_1, \dots, \mathcal{N}_n$ are well behaved, minimal amplitude, constant space and linear time logical solutions of Y_1, \dots, Y_n , respectively. We let \mathcal{M} , a logical solution of X , be a machine that works as follows.

At the beginning, as always, \mathcal{M} waits till Environment chooses some constants for all free variables of the conclusion. Let e be a (the) valuation that agrees with the choices just made by Environment (in the previous cases, we have suppressed the e parameter, but now we prefer to deal with it explicitly). So, the conclusion is now brought down to $e[X]$. After this event, \mathcal{M} continues waiting until Environment makes one more move. If such a move is never made, then the run of (the \sqcap -closure of) X generated in the play can be simply seen as the empty run of $e[X]$. Due to the Stability condition, $\|X\|$ is classically valid, meaning that $\mathbf{Wn}_e^{\|X\|} \langle \rangle = \top$. But then, in view of Lemma 8.4, $\mathbf{Wn}_e^X \langle \rangle = \top$. This makes \mathcal{M} the winner without having consumed any space or any move-preceding time for which it can be billed.

Suppose now Environment makes a move α . With a little thought, one can see that any (legal) move α by Environment brings the game $e[X]$ down to $g[Y_i]$ for a certain valuation g and one of the premises Y_i of the rule. For example, if X is $P \circ - (E \sqcap F) \vee \sqcap x G(x)$ where P is atomic, then a legal move α by Environment should be either 1.0.0 or 1.0.1 or 1.1.# c for some constant c . In the case $\alpha = 1.0.0$, the above-mentioned premise Y_i will be $P \circ - E \vee \sqcap x G(x)$, and g will be the same as e . In the case $\alpha = 1.0.1$, Y_i will be $P \circ - F \vee \sqcap x G(x)$, and g , again, will be the same as e . Finally, in the case $\alpha = 1.1.\#c$, Y_i will be $P \circ - (E \sqcap F) \vee G(y)$ for a variable y not occurring in X , and g will be the valuation that sends y to c 's denotat and agrees with e on all other variables, so that $g[P \circ - (E \sqcap F) \vee G(y)]$ is $e[P \circ - (E \sqcap F) \vee G(c)]$, with the latter being the game to which $e[X]$ is brought down by the labmove $\perp 1.1.\#c$.

After the above event, \mathcal{M} does the usual trick of turning itself into — and continuing playing as — \mathcal{N}_i , with the only difference that, if $g \neq e$, the behavior of \mathcal{N}_i should be followed for the scenario where the adversary of the latter, at the very beginning of the play, chose constants for the free variables of Y_i in accordance with g rather than e .

It is left to the reader to convince himself or herself that, with adequately arranged details of simulation, \mathcal{M} is as desired.

8.7. On the completeness of CL12. While CoL takes no interest in nonalgorithmic “solutions” of problems, it would still be a pity to let one fact go unobserved. As noted, the completeness proof for **CL12** from [34] (namely, the $(ii) \Rightarrow (i)$ part of our present Theorem 8.2) will be reproduced in Appendix C. An analysis of that proof reveals that virtually nothing in it relies on the fact that a purported HPM solving the problem represented by a **CL12**-unprovable sequent, whose non-existence is proven, follows an algorithmic strategy. So, the completeness result can be strengthened by saying that, if **CL12** does not prove a sequent X , then X does not even have a nonalgorithmic logical solution. Precisely defining the meaning of a “nonalgorithmic”, or rather “not-necessarily-algorithmic” logical solution, is not hard. The most straightforward way to do so would be to simply take our present definition of a logical solution but generalize its underlying model of computation by allowing HPMs to have oracles — in the standard sense — for whatever functions.

As an aside, among the virtues of CoL is that it eliminates the need for many ad hoc inventions such as the just-mentioned oracles. Namely, observe that a problem A is computable by an HPM with an oracle for a function $f(x)$ if and only if the problem $\sqcap x \sqcup y (y = f(x)) \circ - A$ is computable in the ordinary sense (i.e., computable by an ordinary HPM without any oracles). So, a CoL-literate person, regardless of his or her aspirations, would never really have to speak in terms of oracles or nonalgorithmic strategies. This explains why ‘*CoL takes no interest in nonalgorithmic “solutions” of problems*’.

9. LOGICAL CONSEQUENCE

Definition 9.1. We say that a formula F is a **logical consequence** of formulas E_1, \dots, E_n ($n \geq 0$) iff the sequent $E_1, \dots, E_n \circ - F$, in the role of X , satisfies any of the (equivalent) conditions (i)-(iii) of Theorem 8.2.

As noted in Section 1, the following rule, which we (also) call **Logical Consequence**, will be the only logical rule of inference in **CL12**-based applied systems:

From E_1, \dots, E_n conclude F as long as F is a logical consequence of E_1, \dots, E_n .

A reader familiar with earlier essays on CoL would remember that, philosophically speaking, computational *resources* are symmetric to computational problems: what is a problem for one player to solve is a resource that the other player can use. Namely, having a problem A as a computational resource intuitively means having the (perhaps externally provided) ability to successfully solve/win A . For instance, as a resource, $\Box x \sqcup y (y = x^2)$ means the ability to tell the square of any number.

According to the following thesis, logical consequence lives up to its name. A justification for it, as well as an outline of its significance, was provided in Section 1:

Thesis 9.2. *Assume E_1, \dots, E_n and F are formulas such that there is a *-independent (whatever interpretation *) intuitive description and justification of a winning strategy for F^* , which relies on the availability and “recyclability” — in the strongest sense possible — of E_1^*, \dots, E_n^* as computational resources. Then F is a logical consequence of E_1, \dots, E_n .*

Example 9.3. Imagine a **CL12**-based applied formal theory, in which we have already proven two facts: $\forall x (x^3 = (x \times x) \times x)$ (the meaning of “cube” in terms of multiplication) and $\Box x \Box y \sqcup z (z = x \times y)$ (the computability of multiplication), and now we want to derive $\Box x \sqcup y (y = x^3)$ (the computability of “cube”). This is how we can reason to justify $\Box x \sqcup y (y = x^3)$:

Consider any s (selected by Environment for x in $\Box x \sqcup y (y = x^3)$). We need to find s^3 . Using the resource $\Box x \Box y \sqcup z (z = x \times y)$, we first find the value t of $s \times s$, and then the value r of $t \times s$. According to $\forall x (x^3 = (x \times x) \times x)$, such an r is the sought s^3 .

Thesis 9.2 promises that the above intuitive argument will be translatable into a **CL12**-proof of

$$\forall x (x^3 = (x \times x) \times x), \Box x \Box y \sqcup z (z = x \times y) \circ - \Box x \sqcup y (y = x^3)$$

(and hence the succedent will be derivable in the theory by Logical Consequence as the formulas of the antecedent are already proven). Such a proof indeed exists — see Example 7.2.

While Thesis 9.2 is about the *completeness* of Logical Consequence, the following Theorems 10.3 and 11.1 are about *soundness*, establishing that Logical Consequence preserves computability, and does so in a certain very strong sense. They are the most important results of the present paper: as noted in Section 1, it is these two theorems that dramatically broaden the applicability of **CL12** (of Logical Consequence, that is) as a logical basis for complexity-oriented applied theories. Theorem 10.3 carries good news for the cases where time efficiency is of main concern, and otherwise we are willing to settle for at least linear space. Theorem 11.1 does the same but for the cases where the primary concern is space efficiency — namely, when we want to keep the latter sublinear.

10. FIRST PRESERVATION THEOREM

We say that a given HPM \mathcal{M} plays a given constant game A **providently** iff all \perp -legal plays of A by \mathcal{M} are provident. Intuitively, this means that, as long as the adversary plays A legally, \mathcal{M} never starts constructing a (nonempty) move in its buffer unless it is going to actually make such a move later. A **provident solution** of a given constant game A means an HPM \mathcal{M} with $\mathcal{M} \models A$ that plays A providently.

Due to being technical and probably not particularly interesting for most readers, proofs of the following two Lemmas 10.1 and 10.2 are postponed to Appendix A and Appendix B, respectively. Here we only present brief outlines of the ideas underlying those proofs.

Lemma 10.1. *There is an effective procedure that takes an arbitrary HPM \mathcal{N} , together with an arbitrary formula E , and constructs an HPM \mathcal{Q} such that, for any interpretation $*$ and any unary arithmetical functions $\mathbf{a}, \mathbf{s}, \mathbf{t}$, if \mathcal{N} is an \mathbf{a} amplitude, \mathbf{s} space and \mathbf{t} time solution of E^* , then \mathcal{Q} is a provident, \mathbf{a} amplitude, $O(\mathbf{s})$ space and $O(\mathbf{t})$ time solution of E^* . Such a \mathcal{Q} has the same number of work tapes as \mathcal{N} does.*

Proof idea. We let \mathcal{Q} simulate \mathcal{N} and play exactly as \mathcal{N} does yet in a way that guarantees providence. Namely, every time \mathcal{N} starts constructing a new move α in its buffer, \mathcal{Q} waits — without hurrying to put any symbols of α into its own buffer — to see if \mathcal{N} eventually makes such a move; if and when the simulation shows that \mathcal{N} moves, \mathcal{Q} resimulates \mathcal{N} 's work, this time actually copying the evolving content of α into its own buffer and entering a move state once that content is mature. Through an appropriate arrangement of the details of \mathcal{Q} 's work, one can achieve the satisfaction of all conditions of the lemma. ■

Lemma 10.2.

1. *Consider any formula E . There is a number c such that, for any HPM \mathcal{Q} , interpretation $*$ and unary arithmetical function \mathbf{a} , if \mathcal{Q} is a provident and \mathbf{a} -amplitude solution of E^* , then, at any step t of any \perp -legal play of E by \mathcal{Q} , the number of symbols in \mathcal{Q} 's buffer does not exceed $\mathbf{a}(\ell) + c$, where ℓ is the background of t .*
2. *Consider any sequent X with a well behaved and minimal amplitude logical solution \mathcal{Q} . There is a number c such that, at any step t of any \perp -legal play of X by \mathcal{Q} , the number of symbols in \mathcal{Q} 's buffer does not exceed $\ell + c$, where ℓ is the background of t .*

Proof idea. Assuming all the conditions of clause 1 of the lemma, \mathcal{Q} never constructs an “oversized” — longer than $\mathbf{a}(\ell)$ plus a constant — move in its buffer because, due to providence, it would have to actually make such a move sooner or later, in which case it would violate its amplitude constraints (unless, of course, by good luck, Environment makes an illegal or “very long” move meanwhile). Clause 2 is similar, taking into account that providence is one of the necessary conditions for being well behaved. ■

Below and later, when f is a unary arithmetical function and m a natural number, $f^m(x)$ means the m -fold composition of f with itself applied to x , i.e. $f(f(\dots(x)\dots))$ with f repeated m times. On the other hand, $(f(x))^m$ has its usual meaning of the m th power of $f(x)$, i.e. $f(x)$ multiplied by itself m times.

Theorem 10.3. *There is an effective procedure that takes an arbitrary CL12-proof \mathbb{P} of an arbitrary sequent $E_1, \dots, E_n \circ - F$, arbitrary HPMs $\mathcal{N}_1, \dots, \mathcal{N}_n$ and constructs an HPM \mathcal{M} satisfying the condition below:*

Assume $$ is an interpretation and $\mathbf{a}, \mathbf{s}, \mathbf{t}$ are unary arithmetical functions such that:*

- (i): *For each $i \in \{1, \dots, n\}$, \mathcal{N}_i is an \mathbf{a} amplitude, \mathbf{s} space and \mathbf{t} time solution of E_i^* .*
- (ii): *For any x , $\mathbf{a}(x) \geq \max\{x, \mathbf{c}\}$, where \mathbf{c} is the native magnitude of $E_1, \dots, E_n \circ - F$.*
- (iii): *For any x , $\mathbf{s}(x) \geq x$.*

(iv): For any x , $\mathfrak{t}(x) \geq x$.

Then there is a number \mathfrak{b} which only depends on \mathbb{P} such that, with \mathfrak{R} abbreviating $\mathfrak{a}^{\mathfrak{b}}(\ell)$, \mathcal{M} is a solution of F^* that runs in amplitude \mathfrak{R} , space $O(\mathfrak{s}(\mathfrak{R}))$ and time $O(\mathfrak{t}(\mathfrak{R}))$.

The rest of the present section is devoted to a proof of the above theorem.

Consider an arbitrary sequent $E_1, \dots, E_n \circ - F$ together with a **CL12**-proof \mathbb{P} of it, and let \mathfrak{c} be the native magnitude of $E_1, \dots, E_n \circ - F$. Let $\mathfrak{a}, \mathfrak{s}, \mathfrak{t}$ be functions satisfying conditions (ii)-(iv) of the theorem. By Theorem 8.2, there is a well behaved logical solution \mathcal{K} of $E_1, \dots, E_n \circ - F$ which runs in $\max\{\ell, \mathfrak{c}\}$ amplitude, constant space and linear time, and such a \mathcal{K} — fix it — can be effectively constructed from \mathbb{P} . Note that, in view of conditions (ii)-(iv) of the present theorem, \mathcal{K} automatically also runs in amplitude \mathfrak{a} , space $O(\mathfrak{s})$ and time $O(\mathfrak{t})$. Consider an arbitrary interpretation $*$ (which, as done before, we shall notationally suppress) and arbitrary HPMs $\mathcal{N}_1, \dots, \mathcal{N}_n$ satisfying condition (i) of the theorem. Below we describe an HPM \mathcal{M} such that \mathcal{M} wins F under the same interpretation $*$. It is important to note that our *construction* of \mathcal{M} does not depend on $*$, $\mathfrak{a}, \mathfrak{s}, \mathfrak{t}$ and hence on our assumption that conditions (i)-(iv) are satisfied; only our *claim* that \mathcal{M} wins F , and our further claims about its complexities, do.

As always, we let our machine \mathcal{M} , at the beginning of the play, wait till Environment selects a constant — call it c_x — for each free variable x of F . Let us fix e as the valuation that sends every free variable x of F to the denotat of the corresponding constant c_x , and (arbitrarily) sends all other variables to 0.

We describe the work of \mathcal{M} afterwards at a high level. A more detailed description and analysis would be neither feasible (since it would be prohibitively long and technical) nor necessary. From our description it will be immediately clear that \mathcal{M} is constructed effectively, so this issue will not be explicitly addressed.

To understand the idea, let us first consider the simple case where \mathcal{K} never makes any replicative moves in the antecedent of $E_1, \dots, E_n \circ - F$. The main part of the work of \mathcal{M} consists in continuously monitoring its run tape to see if Environment has made any new moves, combined with simulating, in parallel, a play of $E_1, \dots, E_n \circ - F$ by the machine \mathcal{K} and — for each $i \in \{1, \dots, n\}$ — a play of E_i by the machine \mathcal{N}_i . During simulation, \mathcal{M} “imagines” that, at the beginning of the play, for each free variable x of the corresponding formula or sequent, the adversary of each machine has chosen the same constant c_x as \mathcal{M} ’s adversary did. After the above initial moves by the real and imaginary adversaries, each of the $n + 2$ games

$$G \in \{E_1, \dots, E_n, E_1, \dots, E_n \circ - F, F\}$$

that we consider here (the first $n + 1$ games imaginary and the last one real) will be brought down to $e[G]$ but, for readability and because e is fixed, we shall usually omit e and write simply G instead of $e[G]$.

Since we here assume that \mathcal{K} never makes any replications in the antecedent of $E_1, \dots, E_n \circ - F$, playing this game essentially means simply playing

$$E_1 \wedge \dots \wedge E_n \rightarrow F, \tag{10.1}$$

so, in what follows, we will talk about (10.1) instead of $E_1, \dots, E_n \circ - F$.

As earlier in similar situations, we assume that, in the real play of F , Environment does not make any illegal moves (the Clean Environment Assumption). We can also safely

assume that the simulated machines do not make any illegal moves in the corresponding games, or else our assumptions about their winning those games would be wrong.²⁷

If so, what \mathcal{M} does in the above mixture of the real and simulated plays is that it applies copycat between $n + 1$ pairs of (sub)games, real or imaginary. Namely, it mimics, in (the real play of) F , \mathcal{K} 's moves made in the consequent of (the imaginary play of) (10.1), and vice versa: uses Environment's moves made in the real play of F as \mathcal{K} 's (imaginary) adversary's moves in the consequent of (10.1). Further, for each $i \in \{1, \dots, n\}$, \mathcal{M} uses the moves made by \mathcal{N}_i in E_i as \mathcal{K} 's imaginary adversary's moves in the E_i component of (10.1), and vice versa: uses the moves made by \mathcal{K} in that component as \mathcal{N}_i 's imaginary adversary's moves in E_i .

Therefore, the final positions hit by the $n + 2$ imaginary and real plays of

$$E_1, \dots, E_n, E_1 \wedge \dots \wedge E_n \rightarrow F \text{ and } F$$

will retain the above forms, i.e., will be

$$E'_1, \dots, E'_n, E'_1 \wedge \dots \wedge E'_n \rightarrow F' \text{ and } F'$$

for some E'_1, \dots, E'_n, F' . Our assumption that the machines $\mathcal{N}_1, \dots, \mathcal{N}_n$ and \mathcal{K} win the games E_1, \dots, E_n and $F_1 \wedge \dots \wedge F_n \rightarrow F$ implies that each $G \in \{E'_1, \dots, E'_n, E'_1 \wedge \dots \wedge E'_n \rightarrow F'\}$ is \top -won, in the sense that $\mathbf{Wn}^G \langle \rangle = \top$. It is then obvious that so should be F' . Thus, the (real) play of F brings it down to the \top -won F' , meaning that \mathcal{M} wins F .

Let us fix \mathbf{b} as the total number of occurrences of choice operators $\sqcap, \sqcup, \sqcap, \sqcup$ in the \sqcap -closure $\sqcap(10.1)$ of (10.1). Note that no legal run of $\sqcap(10.1)$ or of any of the games $\sqcap E_1, \dots, \sqcap E_n, \sqcap F$ will contain more than \mathbf{b} labmoves.

Consider an arbitrary play (computation branch) of \mathcal{M} , and an arbitrary clock cycle c in the context of that play. Let ℓ be the background of c . In the simulations of \mathcal{K} and $\mathcal{N}_1, \dots, \mathcal{N}_n$, every move made by the imaginary adversary of one of these machines is a copy²⁸ of either a move made by Environment in the real play, or a move made by one of the machines $\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n$ during simulation. Let β_1, \dots, β_m be the moves by simulated machines that \mathcal{M} detects by time c , arranged according to the times of their detections. Let

$$\mathcal{G}_1, \dots, \mathcal{G}_m \in \{\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n\}$$

be the machines that made these moves, respectively. Remembering that all machines run in amplitude \mathbf{a} , the magnitude of β_1 cannot exceed $\mathbf{a}(\ell)$. That is because, by the time when \mathcal{G}_1 made the move β_1 , all (if any) moves by \mathcal{G}_1 's imaginary adversary were copies of moves made by Environment in the real play rather than moves made by some other simulated machines, and hence the background of β_1 in the simulated play of \mathcal{G}_1 did not exceed ℓ . For similar reasons, with $\mathbf{a}(\ell)$ now acting in the role of ℓ and \mathcal{G}_2 in the role of \mathcal{G}_1 , the magnitude of β_2 cannot exceed $\mathbf{a}(\mathbf{a}(\ell))$. Similarly, the magnitude of β_3 cannot exceed $\mathbf{a}(\mathbf{a}(\mathbf{a}(\ell)))$, etc. Also, by our selection of the constant \mathbf{b} , at most \mathbf{b} moves can be made altogether by the

²⁷Since we need to construct \mathcal{M} no matter whether those assumptions are true or not, we can let \mathcal{M} simply retire as soon as some illegal/unexpected behavior is detected.

²⁸Here and later in similar contexts, "copy" should be understood in a generous sense, referring only the "actual meanings" of moves, with their "addressing" prefixes otherwise possibly modified/readjusted. For instance, if Environment made the move α in the real play of F , the corresponding "copy" move made by \mathcal{M} in the imaginary play of (10.1) for the imaginary adversary of \mathcal{K} will be $1.\alpha$ rather than α , with the prefix "1." merely indicating that the move α is made in the consequent. It is important to note that such prefix readjustments in no way affect the magnitudes of the "copied" moves. This is one of the reasons why we chose to understand magnitude as in Convention 5.3 rather than as raw size.

simulated machines, so $m \leq \mathfrak{b}$. Thus, by time c , the magnitude of no move made in this mixture of real and simulated plays can exceed $\mathfrak{a}^{\mathfrak{b}}(\ell)$. This means that \mathcal{M} , which merely mimics some of \mathcal{K} 's moves, runs in amplitude $\mathfrak{R} = \mathfrak{a}^{\mathfrak{b}}(\ell)$ as promised.

The next thing to clarify is why the space and time complexities of \mathcal{M} are also as promised. Here we need to look into more details of the work of \mathcal{M} . First of all, let us agree for the rest of this section — or, rather, reiterate our earlier general convention — that, in the (often only implicit) context of a given computation step of \mathcal{M} , as in the preceding paragraph, ℓ stands for the background of that step. As we saw in the preceding paragraph, the magnitude of no real or imaginary move will ever exceed $\mathfrak{R} = \mathfrak{a}^{\mathfrak{b}}(\ell)$ for that very ℓ . So, while ℓ is the “current” background in the context of \mathcal{M} , the “current” background in the context of a simulated machine can instead be as great as \mathfrak{R} . Therefore, for both simplicity and safety, in what follows we prefer to talk in terms of (the implicitly “current” value of) \mathfrak{R} rather than ℓ .

A legitimate question that we have never addressed so far is the one about the number of work tapes that \mathcal{M} has. This is our arrangement. Per each simulated machine $\mathcal{G} \in \{\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n\}$, \mathcal{M} has:

- As many work tapes as \mathcal{G} has. At any time during simulation, \mathcal{M} keeps the content of each such tape identical to the imaginary content of the corresponding work tape of \mathcal{G} . Furthermore, \mathcal{M} positions the scanning head of such a tape also exactly where the scanning head of the corresponding work tape of \mathcal{G} would be. In this way, the two work tapes (including their scanning heads) — one real and one imaginary — are perfectly synchronized.
- One work tape to maintain the imaginary content of \mathcal{G} 's run tape. Again, \mathcal{M} keeps the contents and the scanning head locations of the two tapes identical; the synchronization is perfect just like in the preceding case, except for occasional interrupts during which \mathcal{M} is updating the content of this tape through adding a new labmove to it.
- One work tape to maintain the imaginary content of \mathcal{G} 's buffer. As we may guess, again, the contents of such a tape and such a buffer are synchronized, with the scanning head of \mathcal{M} 's work tape always located at the leftmost blank cell. As in the preceding case, the synchronization is disrupted only occasionally. Namely, once \mathcal{G} enters a move state and thus instantaneously clears its buffer, \mathcal{M} takes a number of steps to correspondingly blank out its work tape.²⁹

Of course, \mathcal{M} also needs to keep track of certain other pieces of information. Namely, it needs to remember the state of each simulated machine, as well as remember the (never changing) transition functions of those machines. This information, due to its amount's being bounded (not exceeding a certain constant), can be and is kept track of via \mathcal{M} 's state memory rather than work tape memory, thus contributing nothing to \mathcal{M} 's space complexity.

Thus, at any time, \mathcal{M} 's overall work-tape space consumption consists of three components: (1) the space needed to remember the work tape contents of the simulated machines; (2) the space needed to remember the run tape contents of the simulated machines; and (3) the space needed to remember the buffer contents of the simulated machines.

²⁹Since we agreed in Section 4 that HPMS never write BLANK on their work tapes, deleting or “blinking out” here could be understood as changing the content of each cell to some special and never-used-elsewhere (“pseudo-blank”) symbol.

Each simulated machine contributes to component (1) the amount that is equal to its own space consumption $O(\mathfrak{s}(\mathfrak{R}))$. Since there is only a constant number $(n + 1)$ of such machines, the overall cost of component (1) remains $O(\mathfrak{s}(\mathfrak{R}))$.

The run tape of each of the $n + 1$ simulated machines can contain at most \mathfrak{b} labmoves, and the size of each such labmove is bounded by $O(\mathfrak{R})$. So, the overall cost of component (2) is $O((n + 1) \times \mathfrak{R})$, i.e. $O(\mathfrak{R})$, which, in view of condition (iii) of our theorem, can be generously rewritten as $O(\mathfrak{s}(\mathfrak{R}))$.

In view of Lemma 10.1, we may assume that the machines $\mathcal{N}_1, \dots, \mathcal{N}_n$ are provident solutions of the corresponding games. Therefore, applying clause 1 of Lemma 10.2 to $\mathcal{N}_1, \dots, \mathcal{N}_n$ and applying clause 2 of the same lemma to (the well behaved and minimal amplitude) \mathcal{K} , we find that representing the buffer contents of any of the $n + 1$ simulated machines takes at most $O(\mathfrak{a}(\mathfrak{R}))$ space. In fact, with a more careful analysis, this quantity can be lowered to $O(\mathfrak{R})$, which then, as in the preceding paragraph, can be rewritten as the target $O(\mathfrak{s}(\mathfrak{R}))$. This is so because, as it can be seen with a little thought, due to providence (or well-behavedness which implies providence), whenever the buffer of a given simulated machine is nonempty on a given cycle (meaning that the machine is going to move at some later time), the maximum possible background of that cycle is $\mathfrak{a}^{\mathfrak{b}-1}(\ell)$ rather than $\mathfrak{a}^{\mathfrak{b}}(\ell)$ established earlier for the general case; and if so, Lemma 10.2 tells us that the buffer content can be recorded in space $O(\mathfrak{a}(\mathfrak{a}^{\mathfrak{b}-1}(\ell)))$, i.e. $O(\mathfrak{a}^{\mathfrak{b}}(\ell))$, i.e. $O(\mathfrak{R})$, rather than the earlier claimed $O(\mathfrak{a}(\mathfrak{R}))$. From here, as in the cases of components (1) and (2), we conclude that the overall cost of component (3) does not exceed $O(\mathfrak{s}(\mathfrak{R}))$.

Summarizing the above four paragraphs, the overall space consumption of \mathcal{M} — and thus (“even more so”) the space consumption of any particular work tape of \mathcal{M} — is indeed $O(\mathfrak{s}(\mathfrak{R}))$, as desired.

\mathcal{M} initializes its simulation routine by copying Environment’s initial moves (in the “ $\perp 1$.”-prefixed form) onto its work tape that keeps track of the content of \mathcal{K} ’s run tape. After that it keeps its run-tape head at the leftmost blank cell and, every time a new move α is detected made by either player \wp , further adds the string $\wp 1.\alpha$ to the content of the above-mentioned work tape. \mathcal{M} similarly copies the moves made by the simulated machines onto its work tapes representing the run tapes of those machines and their adversaries. An exception here is the case of “1.”-prefixed moves made by \mathcal{K} ; each such move α , after (or simultaneously with) copying it on its work tape representing \mathcal{K} ’s run tape, \mathcal{M} further copies into its buffer, and then makes α as a move in the real play.³⁰ Since the number of moves that will be made in the mixture of the real and simulated plays does not exceed the constant \mathfrak{b} and the magnitude of each move does not exceed \mathfrak{R} , the above move-copying work, altogether, contributes only $O(\mathfrak{R})$ to \mathcal{M} ’s overall time consumption. In view of

³⁰At this point one may notice that, despite our having qualified \mathcal{M} ’s strategy as applying copycat, the run Γ generated in the real play of F may be not exactly the same as the run Δ generated in the consequent of (10.1). For instance, if Environment made a move β while \mathcal{M} was constructing the above α in its buffer, the labmove $\perp\beta$ will appear *before* $\top\alpha$ in Γ but *after* $\top\alpha$ in Δ . Is this a problem? Not at all. It is not hard to see that, even if the orders of labmoves in Γ and Δ are not exactly the same, both runs are legal and they bring F down to the same position, in the sense that $\langle \Gamma \rangle F = \langle \Delta \rangle F$, and this is all that eventually matters. Using the terminology of the end of Section 3, Γ is a \top -delay of Δ , which, taking into account that the game F is static, is just as good as if the two runs were identical. Namely, Δ ’s being \top -won implies that Γ is also \top -won. Similarly, by the Clean Environment Assumption, Γ is \perp -legal; but Γ ’s being a \top -delay of Δ implies that Δ is a \perp -delay of Γ ; this, by the definition of static games, further implies that Δ is \perp -legal. Thus, the assumption that Δ was a \perp -legal run of F , on which (whether the reader noticed or not) we have implicitly relied in our argument, was legitimate.

condition (iv) of our theorem, asymptotically this quantity adds nothing to our target time $O(\mathfrak{s}(\mathfrak{R}))$ and hence can be ignored.

The above move-copying events can be viewed as exceptions occurring occasionally during the otherwise “normal” simulation process. The latter consists in following the $n + 1$ simulated machines, in parallel and step by step, via keeping correspondingly updating (within the same, single, step) the corresponding work tape contents and head locations. That is, during the “normal” simulation process, \mathcal{M} works at exactly the same pace as the simulated machines, spending a single clock cycle of its own on tracing one step of all simulated machines simultaneously. Taking into account that each of the simulated machines runs in time $O(\mathfrak{t}(\mathfrak{R}))$ and hence consumes at most $O(\mathfrak{t}(\mathfrak{R}))$ steps before responding with a move, it is not hard to see that the number of steps performed by \mathcal{M} during the “normal” simulation process since Environment’s last move (or since step 0 if there is no such move) each time before making a new move of its own is $O(\mathfrak{t}(\mathfrak{R}))$. Thus, \mathcal{M} indeed runs in time $O(\mathfrak{t}(\mathfrak{R}))$, as desired.

Whatever we have said so far was about the simple case where \mathcal{K} makes no replicative moves in the antecedent of $E_1, \dots, E_n \circ - F$. How different is the general case, where \mathcal{K} can make replications? Not very different. The overall work of \mathcal{M} remains the same, with the only difference that, every time \mathcal{K} replicates one of E_i (more precisely, to whatever game a given copy of E_i has evolved by that time), \mathcal{M} splits the corresponding simulation of \mathcal{N}_i into two identical copies (using two different sets of its work tapes), with the same past but possibly diverging futures. This increases the number of simulated plays and the corresponding number of to-be-synchronized (by the copycat routine) pairs of games, as well as the number of work tapes that \mathcal{M} utilizes, but otherwise \mathcal{M} continues working as in the earlier described scenario. \mathcal{M} is guaranteed to win for the same reasons as before. Furthermore, the complexity analysis that we provided earlier still remains valid as long as we appropriately readjust (increase) the value of the parameter \mathfrak{b} . As we remember, \mathfrak{b} was chosen to be the total number of occurrences of choice (“move-inducing”) operators in the \sqcap -closure of the sequent. The only reason for this choice of \mathfrak{b} was to make sure that it was a bound on the lengths of (numbers of moves in) legal runs of the game represented by the sequent. If, however, \mathcal{K} makes some replicative moves in the antecedent, then each occurrence of a choice operator in the antecedent of the original sequent may eventually give rise to many moves rather than only one move. But how many moves? Luckily, as we remember, \mathcal{K} is well behaved, so that, even if it makes replicative moves, it does so only a certain bounded (constant) number of times that only depends on \mathbb{P} . This means that the lengths of the runs of the \sqcap -closure of $E_1, \dots, E_n \circ - F$ generated by \mathcal{K} still remain bounded. Setting the new value of \mathfrak{b} to that bound, \mathfrak{b} thus remains constant. And this is all that is necessary for our earlier complexity analysis to remain valid.

11. SECOND PRESERVATION THEOREM

The conditions of the following theorem essentially differ from those of Theorem 10.3 only in that they replace the “at least linear” requirement (condition (iii)) on the space complexity of the premise-solving machines by the weaker “at least logarithmic” condition. The price to pay is an increase of the time complexity of the conclusion-solving machine relative to the time complexities of the premise-solving machines.

Theorem 11.1. *There is an effective procedure that takes an arbitrary **CL12**-proof \mathbb{P} of an arbitrary sequent $E_1, \dots, E_n \circ - F$, arbitrary HPMs $\mathcal{N}_1, \dots, \mathcal{N}_n$ and constructs an HPM \mathcal{M} satisfying the condition below:*

Assume $$ is an interpretation and $\mathbf{a}, \mathbf{s}, \mathbf{t}$ are unary arithmetical functions such that:*

(i): *For each $i \in \{1, \dots, n\}$, \mathcal{N}_i is an \mathbf{a} amplitude, \mathbf{s} space and \mathbf{t} time solution of E_i^* .*

(ii): *For any x , $\mathbf{a}(x) \geq \max\{x, \mathbf{c}\}$, where \mathbf{c} is the native magnitude of $E_1, \dots, E_n \circ - F$.*

(iii): *For any x , $\mathbf{s}(x) \geq \log(x)$.*

(iv): *For any x , $\mathbf{t}(x) \geq x$ and $\mathbf{t}(x) \geq \mathbf{s}(x)$.³¹*

Then there are numbers \mathbf{b}, \mathbf{d} which only depend on \mathbb{P} such that, with \mathfrak{R} abbreviating $\mathbf{a}^{\mathbf{b}}(\ell)$, \mathcal{M} is a solution of F^ that runs in amplitude \mathfrak{R} , space $O(\mathbf{s}(\mathfrak{R}))$ and time $O((\mathbf{t}(\mathfrak{R}))^{\mathbf{d}})$.*

The rest of the present section is exclusively devoted to a proof of the above theorem.

Consider an arbitrary sequent $E_1, \dots, E_n \circ - F$ together with a **CL12**-proof \mathbb{P} of it. Let \mathbf{c} be the native magnitude of $E_1, \dots, E_n \circ - F$, and let $\mathbf{a}, \mathbf{s}, \mathbf{t}$ be functions satisfying conditions (ii), (iii) and (iv) of the theorem. By Theorem 8.2, there is a well behaved logical solution \mathcal{K} of $E_1, \dots, E_n \circ - F$ which runs in $\max\{\ell, \mathbf{c}\}$ amplitude, constant space and linear time, and such a \mathcal{K} — fix it — can be effectively constructed from \mathbb{P} . In view of conditions (ii)-(iv), \mathcal{K} automatically also runs in amplitude $\mathbf{a}(\ell)$, space $O(\mathbf{s}(\ell))$ and time $O(\mathbf{t}(\ell))$. Assume that, in the context of a given arbitrary interpretation (which we notationally suppress as before), $\mathcal{N}_1, \dots, \mathcal{N}_n$ are HPMs satisfying condition (i) of our theorem. Below we describe an HPM \mathcal{M} such that \mathcal{M} wins F . From the description it will be immediately clear that \mathcal{M} is constructed effectively.

Since the present case is considerably more challenging than the case dealt with in Section 10, we will make a simplifying assumption that the formulas of $E_1, \dots, E_n \circ - F$ have no free occurrences of variables. This special case rather easily (in the style of Section 10) generalizes to all cases, and we leave making such a generalization to the reader if necessary. Furthermore, as we did at the beginning of the proof of Theorem 10.3, we will adopt another simplifying assumption, according to which \mathcal{K} never makes any replicative moves in the antecedent of the sequent, so that the game it plays is (essentially)

$$E_1 \wedge \dots \wedge E_n \rightarrow F. \quad (11.1)$$

Again, this case then generalizes to all cases in the same way and for similar reasons as at the end of Section 10, so the issue of such a generalization will no longer be explicitly addressed in the present proof. Finally, as before, we assume that neither \mathcal{M} 's environment nor any of the machines that it simulates ever make illegal moves.

As in the proof of Theorem 10.3, the idea for the strategy of \mathcal{M} is to let \mathcal{K} (through simulation) play against $\mathcal{N}_1, \dots, \mathcal{N}_n$ in the corresponding n components of the antecedent of (11.1), and mimic, in the real play of F , the play of \mathcal{K} in the consequent of (11.1). Then, for reasons pointed out in the proof of Theorem 10.3, \mathcal{M} is guaranteed to win F . Furthermore, for exactly the same reasons as before, \mathcal{M} is also guaranteed to run in amplitude $\mathfrak{R} = \mathbf{a}^{\mathbf{b}}(\ell)$

³¹While this condition is always automatically satisfied in the traditional complexity theory, in our case this is not so. After all, think of a scenario where the machine consumes a lot of space but makes no moves.

as desired, where, as in Section 10, \mathfrak{b} is the number of occurrences of choice operators in the \sqcap -closure of (11.1).³²

So, we only need to understand how to achieve — through appropriately (re)arranging the details of \mathcal{M} 's work — that the space and time complexities of \mathcal{M} are also as promised in the present theorem. In the context of a given computation step of \mathcal{M} , as before, ℓ will stand for the background of that (“current”) step. As noted in Section 10, then the “current” background of any simulated machine can be at most \mathfrak{R} . Therefore, as before, in our complexity analysis we will typically think and talk in terms of (the implicitly “current” value of) \mathfrak{R} rather than ℓ .

In our present case \mathcal{M} can no longer afford to perform simulation in the same way as it did in the preceding section, because the amount $\mathfrak{s}(\mathfrak{R}) \geq \log(\mathfrak{R})$ of work-tape space available (asymptotically) to it may be less than the previously available amount $\mathfrak{s}(\mathfrak{R}) \geq \mathfrak{R}$. What caused \mathcal{M} 's high space consumption in Section 10 was the fact that it had represented the contents of the run tapes and the buffers of the simulated machines on its work tapes, which generally takes $O(\mathfrak{R})$ rather than $O(\mathfrak{s}(\mathfrak{R}))$ space. In the present case, when dealing with a j th computation step of a machine $\mathcal{G} \in \{\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_1\}$, we let \mathcal{M} in its work-tape memory only keep representations of the other (and some additional, previously redundant) components of the corresponding configuration of \mathcal{G} . Namely, with “current” below referring to an arbitrary given j th computation step of \mathcal{G} , \mathcal{M} maintains the following pieces of information — call them together the **sketch** of the j th configuration (computation step) of \mathcal{G} :

- 1st component::** The current state of \mathcal{G} . This only takes a constant amount of space.³³
- 2nd component::** The current contents of the work tapes of \mathcal{G} . This piece of information can be represented with $O(\mathfrak{s}(\mathfrak{R}))$ space, because \mathcal{G} runs in space \mathfrak{s} and the magnitude of no (real or imaginary) move exceeds \mathfrak{R} .
- 3rd component::** The current locations of the work-tape heads of \mathcal{G} . The amount of space needed for this obviously does not exceed the preceding amount $O(\mathfrak{s}(\mathfrak{R}))$ — in fact, it is $O(\log(\mathfrak{s}(\mathfrak{R})))$.
- 4th component::** The current location of the run-tape head of \mathcal{G} . This, in view of the fact that the magnitude of no move exceeds \mathfrak{R} and that there is a constant bound (namely, \mathfrak{b}) on the maximum number of moves that can emerge on the run tape of \mathcal{G} , obviously takes $O(\log(\mathfrak{R}))$ space. But, in view of condition (iii) of the theorem, $\mathfrak{s}(\mathfrak{R}) \geq \log(\mathfrak{R})$. So, the present component can be remembered with $O(\mathfrak{s}(\mathfrak{R}))$ space.
- 5th component::** The number of moves that \mathcal{G} has made so far (at steps $\leq j$) in the play. This number can never exceed \mathfrak{b} , so holding it in memory only takes a constant amount of space.
- 6th component::** The current number of symbols in the buffer of \mathcal{G} . For the same reasons as those relied upon on page 45 when proving Theorem 10.3, we find that the size of the buffer content of \mathcal{G} is $O(\mathfrak{R})$. Hence, recording this quantity only takes $O(\log(\mathfrak{R}))$ space and therefore, as in the case of the 4th component, $O(\mathfrak{s}(\mathfrak{R}))$ space.

³²In the general case, i.e., in the case where \mathcal{K} may make replications in the antecedent, the number \mathfrak{b} will be greater, but will still remain a constant, depending on the **CL12**-proof \mathbb{P} of the sequent but not on any moves made in the play.

³³Of course, as in the preceding section, \mathcal{M} can just as well keep track of this piece of information through its state memory rather than work-tape memory.

7th component:: The (possibly empty) string α that has been added to the buffer of \mathcal{G} when it made a transition to the j th step from the preceding, $(j - 1)$ th, step; here we assume that if $j = 0$, i.e., if there is no preceding step, then the string α is empty. In either case, recording such an α , of course, only takes a constant amount of space.

8th component:: In addition, if \mathcal{G} is \mathcal{K} , the sketch has a record of the first few (bounded number of) symbols from the current content of \mathcal{G} 's buffer sufficient to determine in which of the $n + 1$ components of (11.1) the move that is being constructed in that buffer is meant. For instance, if $n = 2$, then remembering the first four or even three symbols would be sufficient. That is because every (legal) move that \mathcal{K} makes will have either the prefix 0.0. indicating that the move is being made in E_1 , or the prefix 0.1. indicating that the move is being made in E_2 , or the prefix 1. indicating that the move is being made in F .

Summing up all of the above quantities, we find that maintaining the sketch for any step of any given simulated machine \mathcal{G} takes $O(\mathfrak{s}(\mathfrak{R}))$ space.

Unfortunately, the sketch of a given computation step j of \mathcal{G} alone is not sufficient to fully trace the subsequent steps of \mathcal{G} and thus successfully conduct simulation. One reason is that, in order to compute (the sketch of) the $(j + 1)$ th step of \mathcal{G} , one needs to know the content of the cell scanned by the run-tape head of \mathcal{G} . However, sketches do not keep track of what is on the run tape, and that information — unless residing on the run tape of \mathcal{M} itself by good luck³⁴ — is generally forgotten. Another reason is that, when \mathcal{G} makes a move on a given step, \mathcal{M} is unable to act accordingly (copy that move in the real play or use it as another machine's imaginary adversary's move) because it does not remember the content of the buffer of \mathcal{G} — this information, just like information on the content of the run tape, is not present in the sketch.

We handle the above difficulty by letting the simulation routine recompute the missing information every time such information is needed. This is done through recursive calls to the routine itself. Required space efficiency here is achieved at the expense of time: the same computations will generally be performed many times over and over because the procedure, to save space, keeps forgetting certain crucial and reusable results of its previous computations.

Properly materializing the above general idea requires quite some care though, namely when it comes to the details of how the simulations of different machines should be synchronized. Among the crucial conditions for our recursive procedure to work within the required complexity limits is to make sure that the depth of the recursion stack never exceeds a certain constant bound. For simplicity, we design \mathcal{M} as a single-work-tape machine, even though with many work tapes \mathcal{M} could have achieved a “somewhat” better time efficiency.

Below, by an **agent** we shall mean either Environment (that is, \mathcal{M} 's real adversary in its play of F) or any one of the $n + 1$ machines $\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n$.

We let \mathcal{M} , in addition to the sketches for the simulated steps of the simulated machines, maintain (one single/common copy of) what we call the **global history**. The latter is a list of all moves made by all $n + 2$ agents throughout the mixture of the real and imaginary plays “so far”. More precisely, this is not a list of moves themselves, but rather entries with certain partial information on those moves. Namely, the entry for each move α does not indicate the actual content of α (which would require $O(\mathfrak{R})$ space, thus possibly exceeding

³⁴Namely, the “lucky” case is when \mathcal{G} is \mathcal{K} and it is trying to read some move made by its adversary in the consequent of (11.1).

the available amount $O(\mathfrak{s}(\mathfrak{A}))$), but rather only the size of α , i.e. the number of symbols in α . Recording this quantity only takes $O(\log(\mathfrak{A}))$ and — in view of condition (iii) of the theorem — $O(\mathfrak{s}(\mathfrak{A}))$ space. The entry for α also indicates which of the $n + 2$ agents has made (is the **author** of) the move; in addition, if the author of the move is \mathcal{K} , there is a record of in which of the $n + 1$ components of (11.1) the move was made. Since representing the last two pieces of information only takes a constant amount of space, it does not add anything to the asymptotic size of an entry, so the entry for each move α in the global history takes $O(\mathfrak{s}(\mathfrak{A}))$ space. Further, remembering where the number \mathfrak{b} came from, we see that the number of entries in the global history can never exceed \mathfrak{b} . Since \mathfrak{b} is constant, we find that \mathcal{M} only spends the same $O(\mathfrak{s}(\mathfrak{A}))$ amount of space on maintaining the overall global history. While a move α is not the same as the entry for it in the global history, in the sequel we will often terminologically identify these two.

What do we need the global history for? As noted earlier, during its work, \mathcal{M} will often have to resimulate some already simulated portions of the work of one or another machine \mathcal{G} . To make such a resimulation possible, it is necessary to have information on the times at which \mathcal{G} 's adversary has made its moves in the overall scenario that we are considering and re-constructing. Recording the actual move *times* as they were detected during the initial simulation, however, could take us beyond our space limits (think of a situation where, say, \mathcal{K} waits “very long” before its environment makes a move). So, instead, we only keep track — via the global history — of the *order* of moves. Then we neutralize the problem of not remembering the “actual” times of \mathcal{G} 's adversary's moves by simply assuming that \mathcal{G} 's adversary always makes its moves instantaneously in response to \mathcal{G} 's moves. The point is that, if \mathcal{G} wins its game, it does so in all scenarios, including the above scenario of an instantaneously moving adversary. This is our preliminary and very rough attempt to explain the use of the global history, of course, and more can be understood after seeing the details of \mathcal{M} 's work below.

What follows is a high-level yet relatively detailed description of the work of \mathcal{M} . The latter relies on the three subprocedures called UPDATE SKETCH, FETCH SYMBOL and MAKE HISTORY. We start with UPDATE SKETCH.

11.1. Procedure UPDATE SKETCH. In the context of a given global history \mathbb{H} , this procedure takes the sketch $\mathbb{S}_j^{\mathcal{G}}$ of a given computation step j of a given machine $\mathcal{G} \in \{\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n\}$, and returns the sketch $\mathbb{S}_{j+1}^{\mathcal{G}}$ of the next computation step $j + 1$ of the same machine.

Let m be the 5th component of $\mathbb{S}_j^{\mathcal{G}}$. The number m tells us how many moves \mathcal{G} had made by time j . In most cases, UPDATE SKETCH will be used while re-constructing some past episode of \mathcal{G} 's work. It is then possible that the global history contains an $(m + 1)$ th move by \mathcal{G} . If so, then such a move, together with all subsequent moves by whichever agents, are “future moves” from the perspective of the j th step of \mathcal{G} that UPDATE SKETCH is currently dealing with. This means that, when “imagining” the situation at the j th step of \mathcal{G} , those moves should be discarded. So, let \mathbb{H}' be the result of deleting from the global history \mathbb{H} the $(m + 1)$ th move of \mathcal{G} and all subsequent moves (if there are no such moves, then simply $\mathbb{H}' = \mathbb{H}$). Next, *relevant* to \mathcal{G} are only the moves from \mathbb{H}' that are made either by \mathcal{G} , or by \mathcal{G} 's imaginary adversary. For instance, if \mathcal{G} is \mathcal{N}_1 , then the only relevant moves are those made either by \mathcal{N}_1 or by \mathcal{K} in the E_1 component of (11.1). Correspondingly, let \mathbb{H}'' be the result of further deleting from \mathbb{H}' all moves that are not relevant to \mathcal{G} .³⁵ Thus,

³⁵Note that, if $\mathcal{G} = \mathcal{K}$, then all moves are relevant to it, so $\mathbb{H}' = \mathbb{H}''$.

the moves of \mathbb{H}'' prefixed with corresponding labels are exactly the labmoves that \mathcal{G} would see — in the same order as they appear in \mathbb{H}'' — on its run tape at step j , except that, if \mathcal{G} is \mathcal{K} , each move of \mathbb{H}'' should additionally take an *addressing prefix* — a string indicating in which of the $n + 1$ components of (11.1) the move is made. For instance, if the author of the move is \mathcal{N}_1 , then on its run tape \mathcal{K} will see $\perp 0.0.\alpha$ rather than just $\perp\alpha$. Here “0.0.” is the addressing prefix of the move, indicating that α is made in the first conjunct of the antecedent of (11.1).

The information contained in $\mathbb{S}_j^{\mathcal{G}}$ is “almost” sufficient for UPDATE SKETCH to calculate the sought value of $\mathbb{S}_{j+1}^{\mathcal{G}}$. The only missing piece of information is the symbol s scanned by the run-tape head of \mathcal{G} on step j . UPDATE SKETCH thus needs, first of all, to figure out what that symbol s is. To do this, using the move sizes recorded in the global history, UPDATE SKETCH computes the length p of the “active” content of \mathcal{G} ’s run tape; that is, p is the ordinal number of the rightmost non-blank cell of the run tape of \mathcal{G} on its j th step. Next, let q (found in the 4th component of $\mathbb{S}_j^{\mathcal{G}}$) be the number indicating the location of the run-tape head of \mathcal{G} on step j . UPDATE SKETCH compares q with p . If $q > p$, it concludes that s is BLANK. Otherwise, s should be a symbol of some labmove λ residing on \mathcal{G} ’s run tape. Below, by the *preamble* of λ we shall mean λ ’s label and addressing prefix,³⁶ and we shall refer to the rest of λ as (the move) α .

In a rather similar manner, based on the above q and using the relevant information contained in \mathbb{H} , UPDATE SKETCH figures out whether s is in λ ’s preamble or in λ ’s (“main”) α part. If s resides in the preamble, UPDATE SKETCH directly figures out what particular symbol s is. Otherwise, if s resides within α , UPDATE SKETCH finds two integers X and Y . Here X is the number of moves made by α ’s author \mathcal{A} before this agent made the move α . And Y is the ordinal number of the sought symbol s within α . If \mathcal{A} is (the real) Environment,³⁷ using X and Y , UPDATE SKETCH finds s on the run tape of \mathcal{M} . In any other case, UPDATE SKETCH calls the below-described procedure FETCH SYMBOL on (\mathcal{A}, X, Y) . As will be seen later, FETCH SYMBOL then returns the sought symbol s . Thus, in any case, UPDATE SKETCH now knows the symbol s read by the run-tape head of \mathcal{G} on step j . Omitting details, we just want to once again point out that doing all of the above — except executing the FETCH SYMBOL subroutine — only requires some straightforward logarithmic space arithmetic in the style seen in the preceding paragraph.

Keeping the above s in mind, UPDATE SKETCH now additionally consults $\mathbb{S}_j^{\mathcal{G}}$ (namely, the 2nd and the 3rd components of it) for the symbols scanned by the work-tape heads of \mathcal{G} on step j , as well as for the state of \mathcal{G} on that step (the 1st component of $\mathbb{S}_j^{\mathcal{G}}$). Using this information and its knowledge of \mathcal{G} ’s transition function, UPDATE SKETCH is now able to find the state of \mathcal{G} on step $j + 1$, the directions in which its scanning heads moved when transitioning from step j to step $j + 1$, the symbols by which the old scanned symbols were replaced on the work tapes, the string added to the buffer on the transition, the new size of (number of symbols in) the buffer, the new number of \mathcal{G} -authored moves (the latter will be m if the state on step $j + 1$ is not a move state, and $m + 1$ otherwise), and the first few symbols of the buffer if required. To summarize omitting the straightforward details of these calculations, UPDATE SKETCH finds (all 8 components of) the sought sketch $\mathbb{S}_{j+1}^{\mathcal{G}}$.

³⁶ λ has a (nonempty) addressing prefix if and only if $\mathcal{G} = \mathcal{K}$.

³⁷And, hence, \mathcal{G} is \mathcal{K} .

11.2. Procedure FETCH SYMBOL. In the context of a given global history \mathbb{H} , this procedure takes a machine $\mathcal{G} \in \{\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n\}$ and two numbers X, Y , where X is smaller than the number of \mathcal{G} 's moves in the global history, and Y is a number not exceeding the length of \mathcal{G} 's $(X + 1)$ th move there. The goal of FETCH SYMBOL is to return, through rerunning \mathcal{G} , the Y th symbol of the $(X + 1)$ th move of \mathcal{G} .³⁸

To achieve the above goal, FETCH SYMBOL creates a variable \mathbb{S} to hold a sketch of \mathcal{G} , and sets the initial value of \mathbb{S} to the **initial sketch** of \mathcal{G} . By the latter we mean the sketch of the initial configuration of \mathcal{G} , i.e. the configuration where \mathcal{G} is in its start state, the buffer and the work tapes are empty,³⁹ and all scanning heads are looking at the leftmost cells of their tapes.

After the above initialization step, FETCH SYMBOL performs the following subprocedure:

- (1) Remember the 5th and 6th components — call them Z and T , respectively — of \mathbb{S} , and then perform UPDATE SKETCH on \mathbb{S} . Let \mathbb{S}' be the resulting sketch, and let σ be the 7th component of \mathbb{S}' . Below, as always, $|\sigma|$ means the length of (number of symbols in) σ . Also, let τ be the length of the addressing prefix of the $(X + 1)$ th \mathcal{G} -authored move recorded in \mathbb{H} ; it is understood here that $\tau = 0$ if $\mathcal{G} \neq \mathcal{K}$.
- (2) If $Z = X$ and $T < Y + \tau \leq T + |\sigma|$, then return the $(Y + \tau - T)$ th symbol of σ . Otherwise, update (the value of) \mathbb{S} to \mathbb{S}' , destroy your memory of Z and T to recycle space, and go back to step 1.

11.3. Procedure MAKE HISTORY. For the convenience of description, below we assume that the number n from (11.1) is at least 1. The case of $n = 0$ is simpler if not trivial.

The procedure MAKE HISTORY takes a global history \mathbb{H} as an argument and, treating \mathbb{H} as a variable that may undergo updates, acts according to the following prescriptions:

Stage 1:: Create variables $\mathbb{S}^{\mathcal{K}}, \mathbb{S}^{\mathcal{N}_1}, \dots, \mathbb{S}^{\mathcal{N}_n}$, each one to hold a sketch of the corresponding simulated machine from the list $\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n$. Initialize these variables to the initial sketches (see Subsection 11.2) of the corresponding machines. Proceed to Stage 2.

Stage 2:: See if Environment has made a new move (this can be done, say, by counting Environment's moves on \mathcal{M} 's run tape, and comparing their number with the number of Environment-authored moves recorded in the global history). If yes, update the global history \mathbb{H} by adding to it a record for that move, and repeat MAKE HISTORY. If not, go to Stage 3.

Stage 3::

(a): Perform UPDATE SKETCH on $\mathbb{S}^{\mathcal{K}}$. Let $\mathbb{T}^{\mathcal{K}}$ be the resulting sketch.

(b): If \mathcal{K} did not make a **globally new** move on its transition from $\mathbb{S}^{\mathcal{K}}$ to $\mathbb{T}^{\mathcal{K}}$,⁴⁰ change the value of the variable $\mathbb{S}^{\mathcal{K}}$ to $\mathbb{T}^{\mathcal{K}}$, and proceed to Stage 4. Here and later in similar contexts, by a “globally new” move we mean a move not recorded in the global history \mathbb{H} . Figuring out whether \mathcal{K} made a globally new move is easy. Technically, \mathcal{K} made a globally new move if and only if, firstly, it *did* make a move, i.e., the 1st component

³⁸Here, if $\mathcal{G} = \mathcal{K}$, the addressing prefix of the move should be ignored when counting its symbols.

³⁹As for the run tape, what is on it is irrelevant because a sketch has no record of the run-tape content anyway.

⁴⁰Here and later in similar contexts, we terminologically identify sketches with the corresponding steps of the corresponding machines.

of $\mathbb{T}^{\mathcal{K}}$ is a move state; and secondly, such a move is not recorded in \mathbb{H} , meaning that the 5th component of $\mathbb{T}^{\mathcal{K}}$ exceeds the total number of \mathcal{K} 's moves recorded in \mathbb{H} .

- (c): If \mathcal{K} made a globally new move and the move was in one of the E_i components of (11.1), update \mathbb{H} by adding to it a record for that move, and repeat MAKE HISTORY.
- (d): Suppose now \mathcal{K} made a globally new move — call it α — in the F component of (11.1). Let X be the 5th component of $\mathbb{S}^{\mathcal{K}}$, and Y be the 6th component of $\mathbb{S}^{\mathcal{K}}$. Thus, Y is the size of the string $1.\alpha$, and X is the number of moves that \mathcal{K} had made before it made the move α . In this case, call FETCH SYMBOL $Y - 2$ times: first on $(\mathcal{K}, X, 3)$, then on $(\mathcal{K}, X, 4)$, \dots , finally on (\mathcal{K}, X, Y) . Copy to the buffer (of \mathcal{M}) each of the $Y - 2$ symbols returned by these calls. Once all calls have been made, go to a move state. Notice that this results in \mathcal{M} making the move α in the real play. Now update the global history \mathbb{H} by adding to it a record for the move α made by \mathcal{K} , and repeat MAKE HISTORY.

Stage $(3 + i)$ ($1 \leq i \leq n$)::

- (a): Perform UPDATE SKETCH on $\mathbb{S}^{\mathcal{N}_i}$. Let $\mathbb{T}^{\mathcal{N}_i}$ be the resulting sketch.
- (b): If \mathcal{N}_i did not make a globally new move (in the sense explained for \mathcal{K} in Stage 3) on its transition from $\mathbb{S}^{\mathcal{N}_i}$ to $\mathbb{T}^{\mathcal{N}_i}$, then update the value of $\mathbb{S}^{\mathcal{N}_i}$ to $\mathbb{T}^{\mathcal{N}_i}$. After that, if $i < n$, go to Stage $3 + i + 1$, and if $i = n$, go to Stage 2.
- (c): If \mathcal{N}_i made a globally new move, then update the global history \mathbb{H} by adding to it a record for that move, and repeat MAKE HISTORY.

11.4. The overall strategy and an example of its run. The overall strategy followed by \mathcal{M} consists in creating the variable \mathbb{H} , initializing it to the empty global history, and then (forever) running MAKE HISTORY.

Note that MAKE HISTORY will be iterated at most \mathfrak{b} times, because every iteration increases the number of moves in \mathbb{H} , and that number, as already observed, cannot exceed \mathfrak{b} . Since MAKE HISTORY is restarted only finitely many times, the last iteration of it never terminates.

Let us look at an example scenario to make sure we understand the work of \mathcal{M} . For simplicity, assume $n = 1$, so that (11.1) is $E_1 \rightarrow F$. At the very beginning of its work, \mathcal{M} creates the variable \mathbb{H} and sets its value to the empty global history $\langle \rangle$. The rest of its work just consists in running MAKE HISTORY. So, in what follows, we can use “ \mathcal{M} ” and “MAKE HISTORY” as synonyms.

During Stage 1, MAKE HISTORY creates two variables $\mathbb{S}^{\mathcal{K}}$ and $\mathbb{S}^{\mathcal{N}_1}$, and sets their values to the initial sketches of \mathcal{K} and \mathcal{N}_1 , respectively. The result of this step reflects the start situation, where “nothing has happened yet” in the mixture of the real play of F by \mathcal{M} and the simulated plays of $E_1 \rightarrow F$ and E_1 by the machines \mathcal{K} and \mathcal{N}_1 .

Now MAKE HISTORY starts performing, over and over, Stages 2 through 4. The work in those stages can be characterized as “global simulation”. This is a routine that keeps updating, in turn and one step at a time, the two sketches $\mathbb{S}^{\mathcal{K}}$ (Stage 3) and $\mathbb{S}_1^{\mathcal{N}}$ (Stage 4) to the sketches of the “next configurations” of the corresponding machines in the scenario where the adversaries of those machines have made no moves; simultaneously, MAKE HISTORY keeps checking (Stage 2) the run tape of \mathcal{M} to see if Environment has made a move. This will continue until either Environment or one of the two simulated machines is detected to make a move. In our example, let us imagine that Environment was the first of the three agents to make a move, and such a move was α_1 . What happens in this case?

MAKE HISTORY simply restarts the global simulation by resetting (in Stage 1) the two sketches $\mathbb{S}^{\mathcal{K}}$ and $\mathbb{S}^{\mathcal{N}_1}$ to the initial sketches of \mathcal{K} and \mathcal{N}_1 . The earlier-described “Stage 2 through Stage 4 over and over” routine will be repeated, with the only difference that the global history \mathbb{H} is now showing the presence of the Environment-authored α_1 . This means that the simulation of the machine \mathcal{K} will now proceed in the scenario where, at the very beginning of the play, \mathcal{K} ’s adversary had made the move α_1 in the F component. So, every time the simulated \mathcal{K} tries to read one of the symbols of α_1 on its imaginary run tape, MAKE HISTORY— \mathcal{M} , that is — looks that symbol up on its own run tape. As for \mathcal{N}_1 , its (re)simulation will proceed exactly as before (during the first iteration of MAKE HISTORY), for its imaginary adversary has not yet made any moves. By switching to this new scenario, MAKE HISTORY, in fact, deems the previous scenario (the scenario where none of the agents had made any moves) invalid, and simply forgets about it. This new, 2nd attempt of global simulation (the second iteration of MAKE HISTORY, that is) will continue until one of the three agents, again, is detected to make a move.

Let us say it is again Environment, which makes move α_2 . Then MAKE HISTORY again correspondingly updates \mathbb{H} , deems the previous global simulation scenario invalid, forgets it and restarts global simulation for the scenario where, from the very beginning, \mathcal{K} ’s adversary had made the two moves α_1 and α_2 (and the adversary of \mathcal{N}_1 , again, had made no moves). This 3rd attempt of global simulation will continue until, again, a move is detected by one of the three agents.

Let us say this time it is \mathcal{K} , which makes move β in the consequent of $E_1 \rightarrow F$. In this event, through invoking FETCH SYMBOL as many times as the length of β , MAKE HISTORY — \mathcal{M} , that is — assembles β in its buffer symbol-by-symbol, and then makes the move β in the real play. After that, as always when a move by one of the agents is detected, the global simulation restarts. Now the global history \mathbb{H} is showing (records for) the sequence $\langle \alpha_1, \alpha_2, \beta \rangle$ of three moves. In the present, 4th attempt of global simulation, \mathcal{N}_1 is (re)simulated again in the scenario where no moves had been made by \mathcal{N}_1 or its adversary. Similarly, as in the previous case, \mathcal{K} is resimulated in the scenario where, at the very beginning, its adversary had made the two moves α_1 and α_2 . The only difference between the present attempt of global simulation and the previous one is that, once \mathcal{K} is detected to make the expected move β , nothing special happens. Namely, the global history is not updated (as β is already there), the move β is not made in the real play (because it has already been made), and the global simulation continues in the ordinary fashion rather than restarts. The present attempt of global simulation, again, will be interrupted if and when one of the agents is detected to make a globally new move, i.e. a move not recorded in the global history.

Let us say it is again Environment, which makes move α_3 . As always, a record for α_3 is added to \mathbb{H} and the global simulation restarts. For \mathcal{N}_1 everything will be the same as before. As for \mathcal{K} , its resimulation will start in the scenario where, at the beginning of the play, its adversary had made the moves α_1 and α_2 . We already know that, in this scenario, sooner or later, \mathcal{K} will make its previously detected move β . Once this event is detected, \mathcal{K} ’s simulation continues for the scenario where its adversary responded by the move α_3 *immediately* after \mathcal{K} made the move β .

Let us now imagine that the above, fifth attempt of global simulation detects that \mathcal{K} has made a move γ in the E_1 component of $E_1 \rightarrow F$. As always, \mathbb{H} is correspondingly updated and the global simulation restarts. For \mathcal{K} , it will be in the same scenario as before (the adversary had made the moves α_1 and α_2 in the F component right at the beginning of the

play, and made the third move α_3 there right after \mathcal{K} moved β), with the only difference that detecting \mathcal{K} 's expected move γ does not cause MAKE HISTORY to restart. As for \mathcal{N}_1 , this time it will be simulated in the scenario where, at the very beginning of the play, its adversary had made the move γ . So, every time the simulated \mathcal{N}_1 tries to read a symbol of γ from its imaginary run tape, FETCH SYMBOL is called to get that symbol and feed it to the simulation.

Imagine that the final globally new move detected was one by \mathcal{N}_1 . Call that move δ .⁴¹ The global simulation again restarts with the correspondingly updated \mathbb{H} . \mathcal{N}_1 is simulated exactly as in the previous scenario, with the only difference that nothing happens (namely, the global simulation does not restart) when \mathcal{N}_1 makes its (expected) move δ . As for \mathcal{K} , it is also simulated as in the previous case, with the only difference that, after \mathcal{K} makes its move γ , the simulation continues its work “imagining” that the adversary immediately responded with the move δ in the E_1 component. So, every time \mathcal{K} is reading a symbol of δ on its imaginary run tape, that symbol is found calling FETCH SYMBOL. In contrast, when reading a symbol of α_1 , α_2 or α_3 , that symbol is found on the run tape of \mathcal{M} .

The last attempt of global simulation (the one that never got discarded/reconsidered) corresponds to the “ultimate” scenario that determined \mathcal{M} 's real play. Namely, in our present example, the “ultimate” scenario in which \mathcal{N}_1 was simulated was that, at the very beginning of the play, \mathcal{N}_1 's adversary had made the move γ , to which \mathcal{N}_1 later responded by δ , and no moves were made ever after. As for \mathcal{K} , the ultimate scenario for it was that, at the very beginning, \mathcal{K} 's adversary had made the moves α_1 and α_2 in the F component of $E_1 \rightarrow F$, to which \mathcal{K} later responded by the move β in the same component; to β , the adversary instantaneously responded by the move α_3 in the F component; after that, at some point, \mathcal{K} made the move γ in the E_1 component, to which the adversary instantaneously responded by the move δ in the same component, and no further moves were ever made in the game. Since it is our assumption that \mathcal{N}_1 wins E_1 , the two-move run (consisting of γ and δ) that took place in the antecedent of $E_1 \rightarrow F$ is won by \mathcal{N}_1 and hence lost by \mathcal{K} . So, since \mathcal{K} wins the overall game $E_1 \rightarrow F$, the four-move run (consisting of $\alpha_1, \alpha_2, \beta, \alpha_3$) that took place in the consequent of $E_1 \rightarrow F$ is won by \mathcal{K} . But the same run was generated in the real play of F , which makes \mathcal{M} the winner.

Why do we need to restart the global simulation every time a globally new move is detected? The reason is that otherwise we generally would not be able to rely on calls to FETCH SYMBOL for obtaining required symbols. Going back to our example, imagine we did not restart the global simulation (MAKE HISTORY) after the moves α_1 , α_2 and α_3 were made by Environment. Perhaps (but not necessarily), as before, \mathcal{K} would still make its move β sometime between α_2 and α_3 , and then the move γ after α_3 . Fine so far. But the trouble starts when, after that event, \mathcal{N}_1 tries to read some symbol of γ from its imaginary run tape. A way to provide such a symbol is to invoke FETCH SYMBOL, which will resimulate \mathcal{K} to find that symbol. However, in order to properly resimulate \mathcal{K} up to the moment when it made the move γ (or, at least, put the sought symbol of the latter into its buffer), we need to know when (on which computation steps of \mathcal{K}), exactly, the moves α_1 , α_2 and α_3 emerged on \mathcal{K} 's run tape. Unfortunately, we do not remember this piece of information, because, as noted earlier, remembering the exact times (as opposed to merely remembering the order) of moves may require more space than we possess. So, instead, we assume that the moves α_1 and α_2 were made right at the beginning of \mathcal{K} 's play, and then α_3 was made

⁴¹It does not matter whether the event of \mathcal{N}_1 making the move δ was detected before or after the simulated \mathcal{K} already made both of its expected moves β and γ .

right after \mathcal{K} made its move β . This assumption, however, disagrees with the scenario of the original simulation, where α_1 was perhaps only made at step 100, α_2 at step 200, and α_3 perhaps 9999 steps after β . Therefore, there is no guarantee that \mathcal{K} will still generate the same move γ in response to α_3 . Restarting the global simulation — as we did — right after α_1 was made, then restarting it again after α_2 was detected, and then restarting it again after Environment made the move α_3 in response to β , neutralizes this problem. If \mathcal{K} made its move γ after α_3 in this new scenario (the scenario where its imaginary adversary always acted instantaneously), then every later resimulation, no matter how many times MAKE HISTORY is restarted, will again take us to the same move γ made after α_3 , because the global history, which “guides” resimulations, will always be showing the first four moves in the order $\alpha_1, \alpha_2, \beta, \alpha_3$. To see this, note that all updates of the global history only add some moves to it, and otherwise do not affect the already recorded moves or their order.

We also want to understand one remaining issue. As we should have noticed, FETCH SYMBOL always calls UPDATE SKETCH, and the latter, in turn, may again call FETCH SYMBOL. Where is a guarantee that infinitely many or “too many” nested calls will not occur? Let us again appeal to our present example, and imagine we (UPDATE SKETCH, that is) are currently simulating a step of \mathcal{K} sometime after \mathcal{N}_1 already has made its move δ . Whenever \mathcal{K} tries to read a symbol of δ , FETCH SYMBOL is called to resimulate \mathcal{N}_1 and find that symbol. While resimulating \mathcal{N}_1 , however, we may find that, at some point, its run-tape head is trying to read a symbol of γ . To get that symbol, FETCH SYMBOL will be again called to resimulate \mathcal{K} (the author of γ) and find that symbol. Can this process of mutual resimulations go on forever? Not really. Notice that, when FETCH SYMBOL is called on \mathcal{K} to find the sought symbol of γ , FETCH SYMBOL, guided by the global history, will resimulate \mathcal{K} only up to the moment when it made the move γ . But during that episode of \mathcal{K} 's work, the move δ was not yet on its run tape. So, FETCH SYMBOL will not have to be further called on \mathcal{N}_1 . Generally, as we are going to see in the next subsection, at any time there can be at most \mathfrak{b} nested invocations of FETCH SYMBOL or UPDATE SKETCH.

11.5. Complexity analysis. Let \mathbb{H} be a global history, m a natural number, and \mathcal{G} one of the machines $\mathcal{K}, \mathcal{N}_1, \dots, \mathcal{N}_n$. We define the \mathbb{H} -**index** of the pair (\mathcal{G}, m) as the number of moves in \mathbb{H}' , where \mathbb{H}' is the result of deleting from \mathbb{H} the $(m + 1)$ th move of \mathcal{G} and all subsequent moves by any agents; if here \mathbb{H} does not contain more than m moves authored by \mathcal{G} , then \mathbb{H}' is simply \mathbb{H} . Next, where $\mathbb{S}^{\mathcal{G}}$ is a sketch of \mathcal{G} , we define the \mathbb{H} -**index** of $\mathbb{S}^{\mathcal{G}}$ as the \mathbb{H} -index of (\mathcal{G}, m) , where m is (the value of) the 5th component of $\mathbb{S}^{\mathcal{G}}$. We extend the concept of \mathbb{H} -index to particular runs/iterations of UPDATE SKETCH and FETCH SYMBOL in the process of performing MAKE HISTORY. Namely, UPDATE SKETCH is always run on a sketch $\mathbb{S}^{\mathcal{G}}$ of a machine \mathcal{G} , and we define the \mathbb{H} -**index** of that run of UPDATE SKETCH as the \mathbb{H} -index of $\mathbb{S}^{\mathcal{G}}$. Similarly, FETCH SYMBOL is always called on a triple (\mathcal{G}, X, Y) for some machine \mathcal{G} and numbers X and Y , and we define the \mathbb{H} -**index** of such a call/run of UPDATE SKETCH as the \mathbb{H} -index of the pair (\mathcal{G}, X) (Y is thus irrelevant here). If \mathbb{H} is fixed or clear from the context, as it always is when we talk about a given iteration of MAKE HISTORY,⁴² we may omit “ \mathbb{H} -” and simply say “**index**”.

⁴²Remember that each iteration of MAKE HISTORY deals with one single value of \mathbb{H} given at the very beginning of the iteration; the value of \mathbb{H} only changes on a transition from one iteration of MAKE HISTORY to another.

Lemma 11.2. *In the process of any given iteration of MAKE HISTORY with argument \mathbb{H} , we have:*

1. *The \mathbb{H} -index of any run of UPDATE SKETCH does not exceed \mathfrak{b} .*
2. *Whenever a given run of UPDATE SKETCH calls FETCH SYMBOL, the \mathbb{H} -index of the callee is strictly smaller than that of the caller.*
3. *Whenever given run of FETCH SYMBOL calls UPDATE SKETCH, the \mathbb{H} -index of the callee does not exceed that of the caller.*

Proof. Clause 1 is immediate from the obvious fact that an index can never exceed the number of moves in the global history, and the latter, in view of our assumption that no agents ever make illegal moves, cannot exceed \mathfrak{b} . Clauses 2 and 3 can be verified through a rather straightforward (albeit perhaps somewhat long) analysis of the two procedures UPDATE SKETCH and FETCH SYMBOL; details of such an analysis are left to the reader. \square

That the amplitude complexity of \mathcal{M} is as required was observed earlier in the present proof on page 47.

Next, we examine space complexity. The space consumption of \mathcal{M} comes from the need to simultaneously maintain the global history and various sketches of the simulated machines. As observed earlier, maintaining the global history consumes $O(\mathfrak{s}(\mathfrak{R}))$ space (page 50), and each sketch also consume $O(\mathfrak{s}(\mathfrak{R}))$ space (page 49). At any given time, the global history is kept in memory in a single copy. So, to show that the overall space consumption is $O(\mathfrak{s}(\mathfrak{R}))$, we need to show that, at any given time, the number of sketches simultaneously kept in the memory of \mathcal{M} does not exceed a certain constant. But this is indeed so. Looking back at the work of MAKE HISTORY, we see that, at any time, its top level simultaneously maintains the constant number $n + 1$ of sketches: $\mathbb{S}^{\mathcal{K}}, \mathbb{S}^{\mathcal{N}_1}, \dots, \mathbb{S}^{\mathcal{N}_n}$, one for each simulated machine. It also keeps going through these sketches and updating them through UPDATE SKETCH, one after one and one step at a time. Since updates are done sequentially rather than in parallel, space used for them can be recycled, so that space consumptions for updating different sketches (this includes not only the top-level $n + 1$ sketches of MAKE HISTORY, but also many additional sketches that will emerge during calls to FETCH SYMBOL when updating each individual sketch) do not add together unless those sketches happen to be on a same branch of nested recursive calls that UPDATE SKETCH and FETCH SYMBOL make to each other. In view of Lemma 11.2, however, the depth of recursion (the height of the recursion stack at any time) is bounded, because the index of UPDATE SKETCH in the topmost level of recursion does not exceed \mathfrak{b} , and every pair of successor levels of recursion strictly decreases the index of the corresponding call of UPDATE SKETCH.

Finally, we look at time complexity. Remember that our goal here is to show that \mathcal{M} runs in time polynomial in $\mathfrak{t}(\mathfrak{R})$ — namely, time $O((\mathfrak{t}(\mathfrak{R}))^\mathfrak{d})$ for some constant number \mathfrak{d} that only depends on the proof \mathbb{P} of the sequent. Our discourse throughout the rest of this section should be understood in the context of some arbitrary but fixed computation branch of \mathcal{M} .

During the entire work of \mathcal{M} , MAKE HISTORY will be iterated at most \mathfrak{b} times. This is so because each iteration strictly increases the number of moves in the global history, which, as already observed, can never exceed \mathfrak{b} . The last iteration of MAKE HISTORY will run forever, but \mathcal{M} will not be billed for that time because it makes no moves during that period. Likewise, \mathcal{M} will not be billed for the time spent on an iteration of MAKE HISTORY that was interrupted at Stage 2, because a move by Environment resets \mathcal{M} 's time

counter to 0. Call all other (other than the above two) sorts of iterations of MAKE HISTORY **time-billable**. So, it is sufficient for us to understand how much time a single time-billable iteration of MAKE HISTORY takes. Pick any such iteration and fix it throughout the context of the rest of this subsection, including the forthcoming Lemmas 11.3 and 11.4. We let ℓ denote the background of the last clock cycle of that iteration and, as before, use \mathfrak{R} as an abbreviation of $\mathfrak{a}^{\mathfrak{b}}(\ell)$.

Lemma 11.3.

- (1) *The time consumed by any single run of UPDATE SKETCH of index i is $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$.*
- (2) *The time consumed by any single run of FETCH SYMBOL of index i is $O((\mathfrak{t}(\mathfrak{R}))^{i+2})$.*

Proof. We verify this lemma by induction on the index of the corresponding call/run of UPDATE SKETCH or FETCH SYMBOL.

Assume the index of a given run of UPDATE SKETCH is $i \geq 0$. Looking back at our description of UPDATE SKETCH, we see that this routine makes at most one call of FETCH SYMBOL. First, assume no such call is made. With some analysis it is clear that UPDATE SKETCH in this case spends $O(\mathfrak{s}(\mathfrak{R})) + O(\mathfrak{R})$ time. In view of condition (iv) of the theorem, this quantity does not exceed $O(\mathfrak{t}(\mathfrak{R}))$. So, the time consumption is $O(\mathfrak{t}(\mathfrak{R}))$ and hence, of course, also $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$. Now assume UPDATE SKETCH *does* call FETCH SYMBOL. By clause 2 of Lemma 11.2, the index j of such a call is less than i . Hence, by the induction hypothesis, the time taken by the call is $O((\mathfrak{t}(\mathfrak{R}))^{j+2})$. Since $j < i$, we may just as well say $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$ instead. In addition to this, UPDATE SKETCH only spends the same amount $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$ of time to complete its work as in the preceding case. Thus, in either case, the time consumption of UPDATE SKETCH is $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$.

Now consider a run of FETCH SYMBOL, and let $i \geq 0$ be its index. By clause 3 of Lemma 11.2, the index of any call of UPDATE SKETCH that the given run of FETCH SYMBOL makes is at most i . By the induction hypothesis, each such call of UPDATE SKETCH consumes at most $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$ time. Processing any such call (doing additional work related to it), in turn, obviously takes at most $O(\mathfrak{t}(\mathfrak{R}))$ time. So, each call of UPDATE SKETCH costs our run of FETCH SYMBOL at most $O((\mathfrak{t}(\mathfrak{R}))^{i+1})$ time. How many such calls of UPDATE SKETCH will FETCH SYMBOL make? Remembering our assumption that each machine runs in time \mathfrak{t} , with a little thought one can see that the number of calls of UPDATE SKETCH is at most $O(\mathfrak{t}(\mathfrak{R}))$. So, the overall time cost of the run of FETCH SYMBOL is $O(\mathfrak{t}(\mathfrak{R})) \times O((\mathfrak{t}(\mathfrak{R}))^{i+1}) = O((\mathfrak{t}(\mathfrak{R}))^{i+2})$. \square

Lemma 11.4. *No single run of UPDATE SKETCH or FETCH SYMBOL consumes more than $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$ time.*

Proof. As we already know from clause 1 of Lemma 11.2, the index of no run of UPDATE SKETCH may ever exceed \mathfrak{b} . So, by clause 1 of Lemma 11.3, no run of UPDATE SKETCH consumes more than $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$ time. As for FETCH SYMBOL, this procedure is only run when UPDATE SKETCH calls it. The maximum possible index of the caller, as just noted, is \mathfrak{b} . Therefore, by clause 2 of Lemma 11.2, the maximum possible index of any given run of FETCH SYMBOL is $\mathfrak{b} - 1$. Then, by clause 2 of Lemma 11.3, such a run consumes at most $O((\mathfrak{t}(\mathfrak{R}))^{(\mathfrak{b}-1)+2}) = O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$ time. \square

We are now ready to look at the time consumption of the single time-billable iteration of MAKE HISTORY fixed earlier.

Stage 1 of MAKE HISTORY obviously takes a constant amount of time, and this stage is iterated only once. So, asymptotically, it contributes nothing to the overall time consumption of the procedure.

Stage 2 of a time-billable iteration of MAKE HISTORY obviously takes $O(\mathfrak{R})$ time and hence, in view of condition (iv) of the theorem, not more than $O(\mathfrak{t}(\mathfrak{R}))$ time. Namely, checking out the run tape may require moving the run-tape head of \mathcal{M} (at most) from one end of (the non-trivial, i.e. non-blank part of) the tape to the other end, and the length of that segment of the tape is at most \mathfrak{b} times \mathfrak{R} .

Stage 3 starts with performing UPDATE SKETCH (substage (a)), and this, by Lemma 11.4, takes $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$ time. With a little thought, the time taken by substages (b) and (c) of Stage 3 can be seen to be at most linear in \mathfrak{R} and hence in $\mathfrak{t}(\mathfrak{R})$. So is the time taken by substage (d) without counting the calls of FETCH SYMBOL that it performs. Each call of FETCH SYMBOL, by Lemma 11.4, additionally takes $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$ time, and altogether substage (d) performs at most $O(\mathfrak{R})$ and thus $O(\mathfrak{t}(\mathfrak{R}))$ calls. So, the overall time consumption of substage (d) is $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+2})$.

A similar but simpler analysis applies to each stage $3 + i$ ($1 \leq i \leq n$), after which we find that the time consumption of such a stage is $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+1})$.

To summarize, none of the $3 + n$ stages of the iteration of MAKE HISTORY takes more than $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+2})$ time. Stage 1 is repeated only once, and the remaining stages are repeated at most $O(\mathfrak{t}(\mathfrak{R}))$ times as can be seen with a little thought, keeping in mind that the iteration of MAKE HISTORY that we are dealing with is a time-billable one and that each simulated machine runs in time \mathfrak{t} , including the machine whose move interrupted the iteration. If so, the overall time consumption is $O((\mathfrak{t}(\mathfrak{R}))^{\mathfrak{b}+3})$. Taking $\mathfrak{d} = \mathfrak{b} + 3$ completes our proof of the claim of the theorem regarding the time complexity of \mathcal{M} .

12. FINAL REMARKS

The following is an immediate corollary of Theorem 11.1:

Corollary 12.1. *Whenever a formula F is a logical consequence of formulas E_1, \dots, E_n and the latter have polynomial amplitude, logarithmic space and polynomial time solutions under a given interpretation $*$, so does the former. Such a solution can be effectively constructed from a **CL12**-proof of $E_1, \dots, E_n \circ - F$ and the solutions of E_1^*, \dots, E_n^* .*

To see the import of the above corollary, imagine we want to construct a clarithmetical theory S in the style of systems **CLA4-CLA7** discussed in Section 1, such that the number-theoretic problem expressed by any formula F provable in S has a polynomial amplitude, logarithmic space and polynomial time solution. Call this property of formulas the *Polynomial-Logarithmic-Polynomial* property. To achieve this, it would be sufficient to just make sure that every nonelementary (containing some choice operators) axiom of S enjoys the Polynomial-Logarithmic-Polynomial property, and that every nonlogical inference rule of S — if such are present — preserves the Polynomial-Logarithmic-Polynomial property of formulas. As for the elementary (choice-operator-free) axioms of S , they just need to be true, in which case they also automatically enjoy the Polynomial-Logarithmic-Polynomial property because the problems they express are “solved” by an HPM that does nothing and hence “runs” in 0 amplitude, 0 space and 0 time. Then we have a guarantee that every theorem of S also enjoys the Polynomial-Logarithmic-Polynomial property as desired. This is so because the only logical inference rule of S (just as of any other **CL12**-based

applied theory) is Logical Consequence, which, according to Corollary 12.1, preserves the Polynomial-Logarithmic-Polynomial property. As an aside, note that we would have full flexibility in selecting or varying the set of elementary axioms of S as long as those axioms are true, with such variations having no effect on the soundness of S with respect to the Polynomial-Logarithmic-Polynomial property. Typically we would want the elementary axioms of S to be Peano axioms, but nothing holds us from choosing any stronger or weaker collection of elementary axioms if and when a need arises.

It is easily understood that Corollary 12.1 is but one of a series of nice corollaries of Theorems 10.3 and 11.1. Various corollaries in the same style will be dealing with other natural and important complexity triples, such as Linear-Logarithmic-Polynomial, Polynomial-Polylogarithmic-Polynomial, Linear-Polynomial-Exponential, Polynomial-Polynomial-Exponential, etc. — you name it! This means that **CL12** is an adequate and very scalable common logical basis for a wide class of complexity-oriented or complexity-sensitive applied systems. **CL12** — more precisely, the associated rule of Logical Consequence — is adequate because, on one hand, by Theorems 10.3, 11.1 and their corollaries in the style of Corollary 12.1, it is sound for such systems, and, on the other hand, by Theorem C.1 and/or Thesis 9.2 (feel free to also throw the discussion of Subsection 8.7 into the mix), it is as strong as a logical rule of inference could possibly be.

We want to close this paper with a couple of terminological conventions for future use. By the **unarification** of an n -ary ($n \geq 0$) arithmetical function h we shall mean the unary arithmetical function h' such that, for any number ℓ , $h'(\ell) = h(\ell, \dots, \ell)$, with n occurrences ℓ in “ (ℓ, \dots, ℓ) ”.

The following convention conservatively generalizes Definition 5.2 from unary functions to functions of any arities. Such a generalization is of no conceptual interest, but it may offer some technical convenience in many treatments.

Convention 12.2. Assume h is a (not-necessarily-unary) arithmetical function, \mathcal{M} is an HPM, and A is a constant game. By saying that \mathcal{M} plays A in time h we shall mean that \mathcal{M} plays A (in the sense of Definition 5.2) in time h' , where h' is the unarification of h . Similarly for space and amplitude.

Convention 12.3. Assume H is a set of arithmetical functions, \mathcal{M} is an HPM, and A is a constant game. By saying that \mathcal{M} plays A in time H we shall mean that \mathcal{M} plays A (in the sense of convention 12.2) in time h for some $h \in H$. Similarly for “ \mathcal{M} is an H time solution of A ”, “ \mathcal{M} runs in time H ”, “ \mathcal{M} is an H time machine”, etc. Similarly for space and amplitude.

Finally, for brevity’s sake, we also agree on the following:

Convention 12.4. Assume h_1, h_2, h_3 are arithmetical functions, \mathcal{M} is an HPM, and A is a constant game. By saying that \mathcal{M} plays A in **tricomplexity** (h_1, h_2, h_3) we shall mean that \mathcal{M} plays A (in the sense of Convention 12.2) in amplitude h_1 , space h_2 and time h_3 . Similarly for “ \mathcal{M} is an (h_1, h_2, h_3) tricomplexity solution of A ”, “ \mathcal{M} runs in (h_1, h_2, h_3) tricomplexity”, “ \mathcal{M} is an (h_1, h_2, h_3) tricomplexity machine”, etc. Similarly for (H_1, H_2, H_3) instead of (h_1, h_2, h_3) , where H_1, H_2, H_3 are sets of arithmetical functions.

REFERENCES

- [1] S. Abramsky and R. Jagadeesan. Games and full completeness for multiplicative linear logic. *Journal of Symbolic Logic*, 59:543-574, 1994.

- [2] L. Babai and M. Shlomo. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36:254-276, 1988.
- [3] A. Blass. A game semantics for linear logic. *Annals of Pure and Applied Logic*, 56:183-220, 1992.
- [4] M. Bauer. A PSPACE-complete first order fragment of computability logic. *ACM Transactions on Computational Logic*, volume 15, no 1, paper 1, pages 1-12, 2014.
- [5] M. Bauer. The computational complexity of propositional cirquent calculus. *Logical Methods in Computer Science*, volume 11, issue 1, paper 12, pages 1-16, 2015.
- [6] S. Buss. *Bounded Arithmetic* (revised version of Ph. D. thesis). Bibliopolis, 1986.
- [7] A. Chandra, D. Kozen and L. Stockmeyer. Alternation. *Journal of the ACM*, 28:114-113, 1981.
- [8] J.Y. Girard. Linear logic. *Theoretical computer science*, 50:1-102, 1987.
- [9] S. Goldwasser, S. Micali and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18:186-208, 1989.
- [10] J. Hintikka. *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic*. Clarendon Press 1973.
- [11] G. Japaridze. A constructive game semantics for the language of linear logic. *Annals of Pure and Applied Logic*, 85:87-156, 1997.
- [12] G. Japaridze. The propositional logic of elementary tasks. *Notre Dame Journal of Formal Logic*, 41:171-183, 2000.
- [13] G. Japaridze. The logic of tasks. *Annals of Pure and Applied Logic* 117:263-295, 2002.
- [14] G. Japaridze. Introduction to computability logic. *Annals of Pure and Applied Logic*, 123:1-99, 2003.
- [15] G. Japaridze. Propositional computability logic I. *ACM Transactions on Computational Logic*, 7:302-330, 2006.
- [16] G. Japaridze. Propositional computability logic II. *ACM Transactions on Computational Logic*, 7:331-362, 2006.
- [17] G. Japaridze. Introduction to cirquent calculus and abstract resource semantics. *Journal of Logic and Computation*, 16:489-532, 2006.
- [18] G. Japaridze. Computability logic: a formal theory of interaction. In *Interactive Computation: The New Paradigm*. D. Goldin, S. Smolka and P. Wegner, eds., Springer, pages 183-223, 2006.
- [19] G. Japaridze. From truth to computability I. *Theoretical Computer Science*, 357:100-135, 2006.
- [20] G. Japaridze. From truth to computability II. *Theoretical Computer Science*, 379:20-52, 2007.
- [21] G. Japaridze. The logic of interactive Turing reduction. *Journal of Symbolic Logic*, 72:243-276, 2007.
- [22] G. Japaridze. Intuitionistic computability logic. *Acta Cybernetica*, 18:77-113, 2007.
- [23] G. Japaridze. The intuitionistic fragment of computability logic at the propositional level. *Annals of Pure and Applied Logic*, 147:187-227, 2007.
- [24] G. Japaridze. Cirquent calculus deepened. *Journal of Logic and Computation*, 18:983-1028, 2008.
- [25] G. Japaridze. Sequential operators in computability logic. *Information and Computation*, 206:1443-1475, 2008.
- [26] G. Japaridze. Many concepts and two logics of algorithmic reduction. *Studia Logica*, 91:1-24, 2009.
- [27] G. Japaridze. In the beginning was game semantics. In *Games: Unifying Logic, Language, and Philosophy*. O. Majer, A.-V. Pietarinen and T. Tulenheimo, eds, Springer, pages 249-350, 2009.
- [28] G. Japaridze. Towards applied theories based on computability logic. *Journal of Symbolic Logic*, 75:565-601, 2010.
- [29] G. Japaridze. Toggling operators in computability logic. *Theoretical Computer Science*, 412:971-1004, 2011.
- [30] G. Japaridze. From formulas to cirquents in computability logic. *Logical Methods in Computer Science*, volume 7, issue 2, paper 1, pages 1-55, 2011.
- [31] G. Japaridze. Introduction to clarithmetic I. *Information and Computation*, 209:1312-1354, 2011.
- [32] G. Japaridze. Separating the basic logics of the basic recurrences. *Annals of Pure and Applied Logic*, 163:377-389, 2012.
- [33] G. Japaridze. A new face of the branching recurrence of computability logic. *Applied Mathematics Letters*, 25:1585-1589, 2012.
- [34] G. Japaridze. A logical basis for constructive systems. *Journal of Logic and Computation*, 22:605-642, 2012.
- [35] G. Japaridze. The taming of recurrences in computability logic through cirquent calculus, Part I. *Archive for Mathematical Logic*, 52:173-212, 2013.

- [36] G. Japaridze. The taming of recurrences in computability logic through cirquent calculus, Part II. *Archive for Mathematical Logic*, 52:213-259, 2013.
- [37] G. Japaridze. Introduction to clarithmetic III. *Annals of Pure and Applied Logic*, 165:241-252, 2014.
- [38] G. Japaridze. Introduction to clarithmetic II. Manuscript at <http://arxiv.org/abs/1004.3236>
- [39] P. Lorenzen. Ein dialogisches Konstruktivitätskriterium. In *Infinitistic Methods*, in PWN, Proc. Symp. *Foundations of Mathematics*, Warsaw, pages 193-200, 1961.
- [40] I. Mezhirov and N. Vereshchagin. On abstract resource semantics and computability logic. *Journal of Computer and System Sciences*, 76:356-372, 2010.
- [41] R. Parikh. Existence and feasibility in arithmetic. *Journal of Symbolic Logic*, 36:494-508, 1971.
- [42] M. Qu, J. Luan, D. Zhu and M. Du. On the toggling-branching recurrence of computability logic. *Journal of Computer Science and Technology*, 28:278-284, 2013.
- [43] W. Xu and S. Liu. Soundness and completeness of the cirquent calculus system CL6 for computability logic. *Logic Journal of the IGPL*, 20:317-330, 2012.
- [44] W. Xu and S. Liu. The countable versus uncountable branching recurrences in computability logic. *Journal of Applied Logic*, 10:431-446, 2012.
- [45] W. Xu and S. Liu. The parallel versus branching recurrences in computability logic. *Notre Dame Journal of Formal Logic*, 54:61-78, 2013.
- [46] W. Xu. A propositional system induced by Japaridze’s approach to IF logic. *Logic Journal of the IGPL*, 22:982-991, 2014.

APPENDIX A. PROOF OF LEMMA 10.1

Pick an arbitrary HPM \mathcal{N} together with an arbitrary formula E . We may assume that E is closed, or else consider $\Box E$ instead. Let us fix \mathfrak{d} as the greatest possible number of labmoves in any legal run of E .⁴³ We want to (see how to) construct an HPM \mathcal{Q} which is as promised in the lemma. As usual, in both our construction of such a machine \mathcal{Q} and our further analysis of its work, we — mostly implicitly — rely on the Clean Environment Assumption.

We let the machine \mathcal{Q} have the same number of work tapes as \mathcal{N} has. As mentioned in the proof idea given on page 41, \mathcal{Q} works through simulating \mathcal{N} . During this process, \mathcal{Q} maintains a list $L = \langle n_1, \dots, n_k, m \rangle$ of integers, where $0 \leq n_1 < \dots < n_k \leq m \leq \mathfrak{d}$. Intuitively, k is the number of moves made by \mathcal{Q} “so far”; each $n_i \in \{n_1, \dots, n_k\}$ is the number of labmoves residing on \mathcal{N} ’s run tape throughout the time period during which \mathcal{N} was “thinking on” (constructing) its i th move; and m is the number of labmoves residing on \mathcal{N} ’s run tape at the beginning of the “final” episode of simulation — the episode following \mathcal{N} ’s k th move. It is important to point out that \mathcal{Q} remembers the list L through its *state memory* rather than work-tape memory. Namely, since there is only a bounded number b of all possible values L_1, \dots, L_b of L , we can think of each state of \mathcal{Q} as being indexed by one of such values. Then, \mathcal{Q} ’s being in an L_j -indexed state can be understood as that L_j is the “current” value of L in \mathcal{Q} ’s state memory.

In a similar fashion, we let \mathcal{Q} , through its state memory, be always conscious of the ordinal number of the labmove that it is currently scanning on the run tape. We further let \mathcal{Q} keep track of the greatest value that such an ordinal number has ever achieved — that is, keep track of the overall number of labmoves it has ever seen “so far” on its run tape. This quantity will be referred to as the **historical move count**.

After initializing the values of both k and m to 0, the simulation routine of \mathcal{Q} consists of the two subroutines described below, starting from Subroutine 1. In our description of those subroutines, whenever we talk about simulating \mathcal{N} , it is to be understood in the sense of an

⁴³Such a number \mathfrak{d} is said to be the *depth* of E .

“almost perfect simulation” — namely, a simulation during which the work-tape contents of \mathcal{Q} fully coincide with the (corresponding) work-tape contents of the simulated configuration of \mathcal{N} , and so do the locations of the work-tape and run-tape scanning heads. Furthermore, the position spelled on the imaginary run tape of \mathcal{N} at any time is also either the same as or, at least, an initial segment of, the position spelled on \mathcal{Q} ’s run tape. This synchronization, however, does not extend to the buffer contents of the two machines: unless (as in stage (b) of Subroutine 2) explicitly specified otherwise, \mathcal{Q} puts no nonempty strings into its buffer during simulation, regardless of what \mathcal{N} does in this respect. In fact, for the exception of stage (b) of Subroutine 2, \mathcal{Q} ’s buffer remains empty throughout the entire process.

Subroutine 1: With the current value of the list L being $\langle n_1, \dots, n_k, m \rangle$, assume⁴⁴ there are at least m labmoves on \mathcal{Q} ’s run tape, exactly k of those labmoves are \top -labeled, and \mathcal{Q} ’s buffer is empty. If $m = \mathfrak{d}$, \mathcal{Q} retires. Otherwise:

(a) \mathcal{Q} simulates \mathcal{N} in the scenario where, from the very beginning of the play and up to (not including) the step t_1 at which \mathcal{N} makes its 1st move α_1 ,⁴⁵ the position spelled on \mathcal{N} ’s run tape consists of exactly the first n_1 labmoves (order respecting) found on \mathcal{Q} ’s run tape; then, starting from t_1 and up to the step t_2 at which \mathcal{N} makes its 2nd move α_2 , the position spelled on \mathcal{N} ’s run tape consists of exactly the first n_2 labmoves found on \mathcal{Q} ’s run tape; and so on; finally, starting from t_{k-1} and up to the step t_k at which \mathcal{N} makes its k th move α_k , the position spelled on \mathcal{N} ’s run tape consists of exactly the first n_k labmoves found on \mathcal{Q} ’s run tape. Here, if $k = 0$, there is no work to be done during the present stage (a) and, for the purposes of the following stage (b), we let t_k be 0.

(b) \mathcal{Q} continues the above simulation⁴⁶ in the scenario where, starting from step t_k , at any (“current”) time, the content of \mathcal{N} ’s run tape is fully identical to the (“current”) content of \mathcal{Q} ’s own run tape. If and when, at some time t_{k+1} during this process, \mathcal{Q} is detected to make its $(k + 1)$ th move α_{k+1} , \mathcal{Q} acts as follows, with h below standing for the historical move count at time t_{k+1} :

(b.1) If $m < h$, \mathcal{Q} sets the value of m to be equal to h , and repeats Subroutine 1.

(b.2) If $m = h$, \mathcal{Q} goes to Subroutine 2.

(Technical note: It is understood that, in either one of the above cases (b.1) or (b.2), as well as in stage (b) of Subroutine 2 below, before exiting a given iteration of the subroutine, \mathcal{Q} makes sure to delete everything on its work tapes, and move all scanning heads to the leftmost cells of the corresponding tapes.)

Subroutine 2 (with the list $L = \langle n_1, \dots, n_k, m \rangle$ inherited from the preceding iteration of Subroutine 1):

(a) \mathcal{Q} resimulates \mathcal{N} exactly as in stage (a) of Subroutine 1, from the very beginning of the play up to the step t_k at which \mathcal{N} is found to make its k th move α_k . As before, if $k = 0$, there is no work to be done during the present stage, and we simply declare the value of t_k to be 0.

(b) \mathcal{Q} continues simulation from the above step t_k up to the step t_{k+1} at which \mathcal{Q} makes its (expected) $(k + 1)$ th move α_{k+1} , in the scenario where, throughout the entire episode, \mathcal{N} ’s run tape contains precisely the first m labmoves of \mathcal{Q} ’s run tape. Also, unlike the

⁴⁴This assumption/condition will always be automatically satisfied on each actual iteration of Subroutine 1.

⁴⁵Here and later: if and when this event actually occurs, of course; however, as it happens, the expected event will indeed always occur.

⁴⁶Or, rather, *starts* simulation if $k = t_k = 0$. The same comment applies to stage (b) of Subroutine 2.

earlier simulation modes, now \mathcal{Q} puts into its buffer every string that \mathcal{N} is putting into its imaginary buffer. When, at time t_{k+1} , \mathcal{N} eventually makes its $(k+1)$ th move α_{k+1} , \mathcal{Q} enters a move state, making the same move α_{k+1} in the real play. After that, \mathcal{Q} updates the list L from $\langle n_1, \dots, n_k, m \rangle$ to $\langle n_1, \dots, n_k, m, m \rangle$ (meaning that the old value of k is now incremented by 1, with the value of the “new n_k ” set to m while the values of all other — old — elements of the list left unchanged), and goes to Subroutine 1.

Consider an arbitrary interpretation $*$ and unary arithmetical functions $\mathfrak{a}, \mathfrak{s}, \mathfrak{t}$. Assume that \mathcal{N} is an \mathfrak{a} amplitude, \mathfrak{s} space and \mathfrak{t} time solution of E^* . We want to show that, as promised in the lemma, \mathcal{Q} is a provident, \mathfrak{a} amplitude, $O(\mathfrak{s})$ space and $O(\mathfrak{t})$ time solution of E^* .

Looking back at our description of the work of \mathcal{Q} , we see that \mathcal{Q} constructs a nonempty move in its buffer only during Subroutine 2. Next, \mathcal{Q} 's work within Subroutine 2 always terminates (\mathcal{Q} returns to Subroutine 1). This is so because, during the preceding iteration of Subroutine 1, \mathcal{N} was detected to make a new, $(k+1)$ th move α_{k+1} . But Subroutine 2 resimulates \mathcal{N} 's work in exactly the same scenario in which it was simulated during the last iteration of Subroutine 1. So, while (re)simulating \mathcal{N} within Subroutine 2, \mathcal{Q} has a guarantee that, sooner or later, \mathcal{N} will be (re)detected to make its $(k+1)$ th move α_{k+1} . But, in this case, as prescribed in the (b) part of Subroutine 2, \mathcal{Q} also enters a move state, thus making a move and emptying its buffer. To summarize, every time some nonempty string appears (during Subroutine 2) in \mathcal{Q} 's buffer, sooner or later the buffer will be emptied. This means nothing but that, as desired, \mathcal{Q} plays E^* providently.

With some further analysis, which is left to the reader, one can see that every run generated by \mathcal{Q} is also a run generated by \mathcal{N} . This means that \mathcal{Q} inherits, from \mathcal{N} , the property of being an \mathfrak{a} amplitude solution of E^* .

It remains to analyze the running space and time of \mathcal{Q} . Before that, it is worth noting that, in view of the Clean Environment Assumption, all runs generated by \mathcal{Q} are (not only \perp -legal but also simply) legal. This is so because, as pointed out in the preceding paragraph, all runs generated by \mathcal{Q} are also runs generated by \mathcal{N} , and the latter, as a winning strategy for E^* , never makes illegal moves.

In view of simulating \mathcal{N} in the earlier noted “almost perfect fashion”, on each work tape \mathcal{Q} obviously consumes exactly the same amount of space as \mathcal{N} does on the corresponding work tape. For the time complexity of \mathcal{Q} , let us first try to understand how many steps a single iteration of Subroutine 2 takes. With the value of L being $\langle n_1, \dots, n_k, m \rangle$, an iteration of Subroutine 2 simulates \mathcal{N} from the very beginning of the play up to the point when it makes its $(k+1)$ th move in the scenario where \mathcal{N} 's imaginary adversary always acted instantaneously, that is, never took any time to “think” before making its moves. Under these circumstances the simulated \mathcal{N} , which runs in time \mathfrak{t} , spends at most $k\mathfrak{t}(\ell)^{47}$ time altogether. Then so does \mathcal{Q} while “almost perfectly” simulating \mathcal{N} within Subroutine 2 because, as we already know, its tape contents and scanning head locations are synchronized with those of \mathcal{N} , meaning (or, at least, making it possible) that the simulation proceeds at exactly the same pace as \mathcal{N} works. Of course, at the end of Subroutine 2, \mathcal{Q} additionally needs to clean up its work tapes and move all of its scanning heads to the leftmost cells of the corresponding tapes, but asymptotically this does not add anything to its time consumption. Thus, the time cost of each iteration of Subroutine 2 is $O(\mathfrak{t}(\ell))$.

⁴⁷Where, according to our already established notational practice, ℓ is the background of the then-current computation step of \mathcal{Q} .

A very similar analysis reveals that every terminating iteration of Subroutine 1 also costs \mathcal{Q} at most $O(\mathfrak{t}(\ell))$ steps as long as Environment makes no moves while \mathcal{Q} is performing that iteration. If Environment moves while \mathcal{Q} is performing a terminating iteration of Subroutine 2, the duration of the iteration could be greater than $O(\mathfrak{t}(\ell))$;⁴⁸ however, the portion of that duration following the time t of Environment’s (last) move remains $O(\mathfrak{t}(\ell))$ for the same reasons as in the above, “normal” cases; as for the (“excessively long”) portion of the iteration time preceding step t , its duration is irrelevant because, at step t , \mathcal{Q} ’s time counter is reset to 0.

Unlike Subroutine 2, one of the iterations of Subroutine 1 — the last iteration — never terminates. However, \mathcal{Q} will not be billed for any time spent within that iteration because, after it reaches such an iteration, it no longer makes any moves.

Putting it all together, the time bill of \mathcal{Q} for each move it makes consists of (at most) the time spent within the iteration of Subroutine 2 in which the move was made (\mathcal{Q} only moves within Subroutine 2), plus the time spent within the preceding — and hence terminating — iterations of Subroutines 1 and 2 starting from the time of Environment’s last move (or time 0 if there is no such move). We have observed above that this quantity does not exceed $O(\mathfrak{t}(\ell))$ steps per iteration. So, how many (terminating) iterations of Subroutines 1 and 2 may occur before any given iteration of Subroutine 2? Note that every such iteration — or, at least, every pair of consecutive iterations — increases the value of either the k or the m parameter of the list $L = \langle n_1, \dots, n_k, m \rangle$. And neither value can ever exceed \mathfrak{d} . Thus, there is a constant bound on the number of iterations, implying that the running time of \mathcal{Q} is $O(\mathfrak{t}(\ell))$ as desired.

Remark A.1. Looking back at our proof of Lemma 10.1, we see that it goes through not only for games of the form $\Box E^*$ where E is a formula, but, in fact, for all constant games G satisfying the following condition for some integer \mathfrak{d} :

$$\text{No legal run of } G \text{ contains more than } \mathfrak{d} \text{ labmoves.} \tag{A.1}$$

Furthermore, the machine \mathcal{Q} constructed from \mathcal{N} not only wins the game(s) G won by \mathcal{N} , but in fact generates the same runs (of relevance) as \mathcal{N} does. Specifically, Lemma 10.1 holds in the following, strong, form:

There is an effective procedure that takes an arbitrary HPM \mathcal{N} , together with an arbitrary integer \mathfrak{d} , and constructs an HPM \mathcal{Q} such that, for any constant game G satisfying condition (A.1) and any unary arithmetical functions $\mathfrak{a}, \mathfrak{s}, \mathfrak{t}$, if \mathcal{N} is an \mathfrak{a} amplitude, \mathfrak{s} space and \mathfrak{t} time solution of G , then \mathcal{Q} is a provident, \mathfrak{a} amplitude, $O(\mathfrak{s})$ space and $O(\mathfrak{t})$ time solution of G . Such a \mathcal{Q} has the same number of work tapes as \mathcal{N} does. Furthermore, for any run Γ containing at most \mathfrak{d} labmoves, Γ is a run generated by \mathcal{Q} if and only if Γ is a run generated by \mathcal{N} .

APPENDIX B. PROOF OF LEMMA 10.2

For clause 1, pick any unary arithmetical function \mathfrak{a} , formula E , interpretation $*$ and HPM \mathcal{Q} such that \mathcal{Q} is a provident, \mathfrak{a} amplitude solution of E^* . Consider any \perp -legal play of E by \mathcal{Q} and any clock cycle t .

⁴⁸Namely, this would happen due to Environment taking “too long” before moving.

Let ℓ be the background of t , Φ the position spelled on \mathcal{Q} 's run tape at time t , and α the string spelled in \mathcal{Q} 's buffer at time t . Assume α is nonempty (otherwise its size does not exceed 0 and the case is trivial). We may further assume that, in the play that we are considering, \mathcal{Q} 's adversary makes no moves at any time $\geq t$. Since \mathcal{Q} is provident, there is a (smallest) clock cycle $t' > t$ such that \mathcal{Q} makes a move α' at time t' , where α is a prefix of α' . In view of our assumption that \mathcal{Q} wins E^* , such an α' should be a legal move of E by \top in position Φ . Let m be the magnitude of α' .

Obviously there are two (smallest) finite sets A and B of strings such that any move of any legal run of E , including α' , satisfies one of the following two conditions:

- (i): α' does not contain the symbol “#”, and α' is an element of A ;
- (ii): α' is $\beta\#d$ for some element $\beta\#$ of B and some constant d .

For instance, if E is $(G_0 \sqcup G_1) \wedge \sqcup xH(x)$, then $A = \{0.0, 0.1\}$ and $B = \{1.\#\}$. Let c be the length of a longest string of $A \cup B$.⁴⁹ Note that $(A, B$ and hence) c only depends on E , not on $*$, \mathcal{Q} or \mathfrak{a} . Obviously the length of α' does not exceed $m + c$. But, by our assumption, \mathcal{Q} runs in amplitude \mathfrak{a} , meaning that $m \leq \mathfrak{a}(\ell)$. Thus, the length of α' does not exceed $\mathfrak{a}(\ell) + c$. Then the same holds for α , because α is a prefix of α' .

Clause 2, with a sequent X instead of a formula E , is handled in a rather similar way. Indeed, by assumption, \mathcal{Q} is a well behaved solution of X . This, by condition 4 of Definition 8.1, implies that \mathcal{Q} is also a provident solution of X . So, here we can rely on the providence of \mathcal{Q} in the same way as we did when dealing with clause 1. Also, by condition 1 of Definition 8.1, \mathcal{Q} makes only a bounded number of replications in X 's antecedent, which, for our purposes, makes X essentially behave as if it was a formula rather than a sequent. Namely, as before, we have two finite sets A and B such that, in the context of any legal play of X by \mathcal{Q} , any move on \mathcal{Q} 's run tape satisfies the above conditions (i) and (ii). This said, the only minor difference between the present case and the previous one is that c now should be the length of the longest string of $A \cup B$ *plus* the native magnitude of X , rather than just the length of the longest string of $A \cup B$. On the assumption that \mathcal{Q} 's runs in minimal amplitude, we then have a guarantee that the length of (α' and hence of) α does not exceed $\ell + c$, as desired.

APPENDIX C. PROOF OF THE COMPLETENESS OF **CL12**

In this appendix, as promised, we reproduce the proof of the completeness theorem for **CL12** from [34].

Theorem C.1. *Every sequent with a logical solution is provable in **CL12**.*

Proof. Assume X is a sequent not provable in **CL12**. Our goal is to show that X has no logical solution (let alone an efficient and/or well behaved logical solution).

Here we describe a *counterstrategy*, i.e., Environment's strategy, against which any particular HPM (in the usual role of \top) loses X^* for an appropriately selected interpretation $*$. In precise terms, as a mathematical object, our counterstrategy — let us call it \mathcal{C} — is a (not necessarily effective) function that prescribes, for each possible content of the run tape that may arise during the process of playing the game, a (possibly empty) sequence of moves that Environment should make on the corresponding clock cycle. In what follows,

⁴⁹In the pathological case of $A \cup B = \emptyset$, c can be taken to be 0.

whenever we say that \mathcal{C} wins or loses, we mean that so does \perp when it acts according to such prescriptions. \mathcal{C} and \perp will be used interchangeably, that is.

By a *variables-to-constants mapping* — or **vc-mapping** for short — for a sequent Y we shall mean a function whose domain is some finite set of variables that contains all (but not necessarily only) the free variables of Y and whose range is some set of constants not occurring in Y , such that to any two (graphically) different variables are assigned (graphically) different constants. When e is a vc-mapping for Y , by $e[Y]$ we shall mean the result of replacing in Y each free occurrence of every variable with the constant assigned to that variable by e .

At the beginning of the play, \mathcal{C} chooses different constants for (all) different free variables of X , also making sure that none of these constants are among the ones that occur in X . Let g be the corresponding vc-mapping for X . This initial series of moves brings X (under whatever interpretation) down to the constant game $g[X]$ (under the same interpretation).

The way \mathcal{C} works after that can be described recursively. At any time, \mathcal{C} deals with a pair (Y, e) , where Y is a **CL12**-unprovable sequent and e is a vc-mapping for Y , such that $e[Y]$ is the game to which the initial $\Box X$ has been brought down “by now”. The initial value of Y is X , and the initial value of e is the above vc-mapping g . How \mathcal{C} acts on (Y, e) depends on whether Y is stable or not.

CASE 1: Y is stable. Then there should be a **CL12**-unprovable sequent Z satisfying one of the following conditions, for otherwise Y would be derivable by Wait. \mathcal{C} selects one such Z (say, lexicographically the smallest one), and acts according to the corresponding prescription as given below.

Subcase 1.1: Y has the form $\vec{E} \circ - F[G_0 \Box G_1]$, and Z is $\vec{E} \circ - F[G_i]$ ($i = 0$ or $i = 1$). In this case, \mathcal{C} makes the move that brings Y down to Z (more precisely, $e[Y]$ down to $e[Z]$), and calls itself on (Z, e) .

Subcase 1.2: Y has the form $\vec{E}, F[G_0 \Box G_1], \vec{K} \circ - H$, and Z is $\vec{E}, F[G_i], \vec{K} \circ - H$. This subcase is similar to the previous one.

Subcase 1.3: Y has the form $\vec{E} \circ - F[\Box x G(x)]$, and Z is $\vec{E} \circ - F[G(y)]$, where y is a variable not occurring in Y . In this case, \mathcal{C} makes a move that brings Y down to $\vec{E} \circ - F[G(c)]$ for some (say, the smallest) constant c such that c is different from any constant occurring in $e[Y]$. After this move, \mathcal{C} calls itself on (Z, e') , where e' is the vc-mapping for Z that sends y to c and agrees with e on all other variables.

Subcase 1.4: Y has the form $\vec{E}, F[\Box x G(x)], \vec{K} \circ - H$, and Z is $\vec{E}, F[G(y)], \vec{K} \circ - H$, where y is a variable not occurring in Y . This subcase is similar to the previous one.

\mathcal{C} repeats the above until (the continuously updated) Y becomes unstable. This results in some finite series of moves made by \mathcal{C} . We assume that all these moves are made during a single clock cycle (remember that there are no restrictions in the HPM model on how many moves Environment can make during a single cycle).

CASE 2: Y is unstable. \mathcal{C} does not make any moves, but rather waits until its adversary makes a move.

Subcase 2.1: The adversary never makes a move. Then the run of $e[Y]$ that is generated is empty. As Y is unstable, $\|Y\|$ and hence $\|e[Y]\|$ is not classically valid. That is, $\|e[Y]\|$ is false in some classical model. But classical models are nothing but our interpretations restricted to elementary formulas. So, $\|e[Y]\|$ is false under some interpretation *. This, in view of Lemma 8.4, implies that $\mathbf{Wn}^{(e[Y])^*} \langle \rangle = \perp$ and hence \mathcal{C} is the winner.

Subcase 2.2: The adversary makes a move α . We may assume that such a move is legal, or else \mathcal{C} immediately wins. There are two further subcases to consider here:

Subsubcase 2.2.1: α is a move in the succedent, or a nonreplicative move in one of the components of the antecedent, of Y . With a little thought, it can be seen that then α brings $e[Y]$ down to $e'[Z]$, where Z is a sequent from which Y follows by one of the four Choose rules, and e' is a certain vc-mapping for Z . In this case, \mathcal{C} calls itself on (Z, e') .

Subsubcase 2.2.2: α is a replicative move in one of the components of the antecedent of Y . Namely, assume Y (after disabbreviating $\circ-$) is the game (8.2) of Subsection 8.5, and the replicative move is made in its $\downarrow E$ component. This brings $e[Y]$ down to $e[(8.3)]$. The latter, however, is “essentially the same as” $e[Z]$, where Z abbreviates the game (8.1). So, \mathcal{C} can pretend that $e[Y]$ has been brought down to $e[Z]$, and call itself on (Z, e) . The exact meaning of “pretend” here is that, after calling itself on (Z, e) , \mathcal{C} modifies its behavior — by “reinterpreting” moves — in the same style as machine \mathcal{M} modified \mathcal{N} ’s behavior in Subsection 8.5.

This completes our description of the work of \mathcal{C} .

Assume a situation corresponding to Subsubcase 2.2.2 occurs only finitely many times. Note that all other cases, except Subcase 2.1, strictly decrease the complexity of Y . So, the play finally stabilizes in a situation corresponding to Subcase 2.1 and, as was seen when discussing that subcase, \mathcal{C} wins.

Now, assume a situation corresponding to Subsubcase 2.2.2 occurs infinitely many times, that is, \mathcal{C} ’s adversary makes infinitely many replications in the antecedent. And, for a contradiction, assume that

$$\mathcal{C} \text{ loses the play of } \Box X^* \text{ for every interpretation } *. \quad (\text{C.1})$$

Let F be the (constant/closed) game/formula to which the succedent of the original $g[X]$ is eventually brought down. Similarly, let \mathcal{A} be the set of all (closed) formulas to which various copies of various formulas of the antecedent of $g[X]$ are eventually brought down. With a little thought and with Lemma 8.4 in mind, it can be seen that (C.1) implies the following:

$$\text{The set } \{\|E\| \mid E \in \mathcal{A}\} \cup \{\|\neg F\|\} \text{ is unsatisfiable (in the classical sense)}. \quad (\text{C.2})$$

By the compactness theorem for classical logic, (C.2) implies that, for some *finite* subset \mathcal{A}' of \mathcal{A} , we have:

$$\text{The set } \{\|E\| \mid E \in \mathcal{A}'\} \cup \{\|\neg F\|\} \text{ is unsatisfiable (in the classical sense)}. \quad (\text{C.3})$$

Consider a step t in the work of \mathcal{C} such that, beginning from t and at every subsequent step, the antecedent of (the then current) $e[Y]$ contains all formulas of \mathcal{A}' . It follows easily from (C.3) that, beginning from t , (the continuously updated) Y remains stable. This means that \mathcal{C} deals only with CASE 1. But, after making a certain finite number of moves as prescribed by CASE 1, Y is brought down to a stable sequent that contains no surface occurrences of \Box, \square in the succedent and no surface occurrences of \sqcup, \sqcap in the antecedent. Every such sequent follows from the empty set of premises by Wait, which is a contradiction because, as we know, the sequent Y at any step of the work of \mathcal{C} remains **CL12**-unprovable. \square

INDEX

- address 12
- addressing prefix 51
- agent 49
- amplitude (complexity) 5,24
- antecedent (of sequent) 27
- arithmetical function 24
- arity of function 17
- arity of game 16
- arity of letter 26
- atomic formula 26
- author (of move) 50

- background 23
- background parameter 32
- Big-O notation 26
- binary numeral 14
- BLANK 21
- blank symbol 21
- blind existential quantification 4,18
- blind universal quantification 4,18
- bound (for complexity) 24
- branching recurrence 12
- buffer 21

- canonical constant 14
- cell (of tape) 21
- choice conjunction 3,10
- choice disjunction 3,11
- choice existential quantification 4,18
- choice universal quantification 4,17
- Choose (\sqcup -, \sqcap -, \sqcup -, \sqcap - \sim) 29
- cirquent 3
- CL12** 3,26,28
- CLA5 - CLA10** 6
- clarithmetic 6
- Clean Environment Assumption 33
- clock cycle 21
- closed formula 26
- closed sequent 27
- closure: \sqcap - \sim 28
- CoL 2
- computable game 23
- computation branch 22
- computation step 21
- computational problem 21
- computational resource 5,40

- compute (game) 23
- conclusion (of rule) 29
- configuration 22
- constant 14,26
- constant function 17
- constant game 9
- CONSTANTS 14
- counterstrategy 66

- delay 19
- denotat (of constant) 15
- Denotation** 27
- denotation (of interpretation) 27
- denotation (of universe) 14
- depend (on variable) 16,17
- Domain** 27
- domain (of interpretation) 27
- domain (of universe) 14

- ϵ 14
- elementarization 28,28
- elementary game 2,9,16
- elementary formula 28
- elementary sequent 28
- empty run 9
- Environment 2,9
- extensional understanding of predicates 16

- FETCH SYMBOL 52
- finitary function 17
- finitary game 16
- focused 12
- formula 26
- function (on universe) 17
- function letter 26

- game (on universe) 2,9,16
- global history 49
- globally new move 52

- head (scanning \sim) 21
- historical move count 62
- HPM 20

- ideal universe 15
- illegal move 9
- illegal run 9
 - \wp -illegal 9
- index (\mathbb{H} - \sim) 56

- individual (of universe) 14
- initial configuration 22
- initial legal (lab)move 10
- initial sketch 52
- instance (of game) 16
- instance (of rule) 29
- interpret 27
- interpretation 27
- ℓ 32
- label 9
- labmove (labeled move) 9
- leaf 12
- legal move 9
- legal play: $\perp\sim$ of ... by ... 24
- legal run 9
 - \emptyset -legal 19
- logical consequence (relation) 6,39
- Logical Consequence (rule) 6,39
- logical solution 28
- lost run 9
- Lr^A** 9
- Lr_e^A** 16
- Machine 2,9
- MAKE HISTORY 52
- magnitude 23,25
- minimal amplitude logical solution 32
- move 9
- move state 21
- name (of individual) 15
- naming (of universe) 15
- native magnitude 32
- negation 3,10
- nonreplicative move 12
- \emptyset 9
- parallel conjunction 3,10
- parallel disjunction 3,11
- play 22
- position 9,10
- preamble 51
- predicate letter:
 - nonlogical 26; logical 26
- prefixation 9
- premise (of rule) 29
- proposition 9
- provident branch 32
- provident solution 40
- providently (play \sim) 40
- reduction 4
- Replicate 29
- replicative move 12
- retire 22
- revisit (labmove) 32
- rule (of inference) 29
- run 9
- run generated by machine 22
- run spelled by computation branch 22
- run tape 21
- sequent 3,27
- sketch 48
- solution 2,23
- solve (problem, game) 23
- space (complexity) 24
- spacecost 24
- stable formula 28
- start state 21
- state (of HPM) 21
- static game 19
- strict reduction 11
- subextensional understanding of predicates
 - or functions 16
- substitution of variables 17
- succedent 27
- surface occurrence 28
- tape symbol 21
- term 26
- time (complexity) 24
- time-billable 58
- timecost 23
- transition function 21
- tree of games 12
- tricomplexity 60
- Turing reduction 14
- unarification 60
- unfocused 12
- uniform solution 28
- unistructural game 18
- Universe** 27
- universe (of discourse) 14

- universe (of interpretation) 27
- unstable formula 28
- UPDATE SKETCH 50
- valuation 15
- VALUATIONS 15
- var_i 14
- variable 14,26
- VARIABLES 14
- vc-mapping 67
- Wait 29
- well behaved logical solution 32
- win 23
- \mathbf{Wn}^A 9
- \mathbf{Wn}_e^A 16
- winning strategy 23
- won run 9
- work tape 20
- yield 12
- \vdash 30
- \top (as game) 9
- \top (as player) 2,9
- \perp (as game) 9
- \perp (as player) 2,9
- \neg (as operation on games) 3,10
- \neg (as operation on players) 9
- \neg (as operation on runs) 10
- \wedge 3,10
- \vee 3,11
- \rightarrow 4,11
- \forall 4,18
- \exists 4,18
- \square 3,10
- \sqcup 3,11
- \square 4,17,28
- \sqcup 4,18
- \wedge 32
- \circ 12
- $\circ-$ 4,14
- \circ 12
- $\#$ 17
- \equiv 16
- $\langle \rangle$ 9
- \sqsubset 12
- \vDash 23
- $|\dots|$ 6,25
- $\|\dots\|$ 28,28
- $F[E]$ 28
- $e[A]$ 16
- $\langle \Phi \rangle A$ 9
- $A(f_1/x_1, \dots, f_n/x_n)$ 17
- Γ^α 10
- $\Gamma^{\preceq v}$ 12