

Averaging in LTL

Patricia Bouyer, Nicolas Markey, Raj mohan Matteplackel

► **To cite this version:**

Patricia Bouyer, Nicolas Markey, Raj mohan Matteplackel. Averaging in LTL. Baldan, Paolo and Gorla, Daniele. Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14), Sep 2014, Rome, Italy. Springer, 8704, pp.266-280, 2014, <10.1007/978-3-662-44584-6_19>. <hal-01062173>

HAL Id: hal-01062173

<https://hal.inria.fr/hal-01062173>

Submitted on 11 Sep 2014

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

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

Averaging in LTL

Patricia Bouyer, Nicolas Markey,
Raj Mohan Matteplackel

September 10, 2014

Research report LSV-14-02 (Version 2)



Laboratoire Spécification & Vérification

École Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

Averaging in LTL

Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel

LSV – CNRS & ENS Cachan – France

Abstract. For the accurate analysis of computerized systems, powerful quantitative formalisms have been designed, together with efficient verification algorithms. However, verification has mostly remained boolean—either a property is true, or it is false. We believe that this is too crude in a context where quantitative information and constraints are crucial: correctness should be quantified!

In a recent line of works, several authors have proposed quantitative semantics for temporal logics, using e.g. *discounting* modalities (which give less importance to distant events). In the present paper, we define and study a quantitative semantics of LTL with *averaging* modalities, either on the long run or within an until modality. This, in a way, relaxes the classical Boolean semantics of LTL, and provides a measure of certain properties of a model. We prove that computing and even approximating the value of a formula in this logic is undecidable.

1 Introduction

Formal verification of computerized systems is an important issue that aims at preventing bugs in the developed computerized systems. The model-checking approach to verification consists in automatically checking that the model of a system satisfies a correctness property. The standard approach is therefore a yes/no (that is, boolean) approach: either the system satisfies the specified property, or the system does not satisfy the property. Model-checking has been widely developed and spread over the last 35 years and is a real success story.

In many applications, quantitative information is crucial; quantities can already appear at the functional level of the system (such as timing constraints between events, or bounds on various quantities like the energy consumption, ...), and many quantitative models like timed automata [4] and their weighted extension [5,7] have therefore been proposed and studied. But quantities can even have more impact on the quality of the system: how good is a system w.r.t. a property? In that case the standard boolean approach might appear as too crude: among those systems that are incorrect (in a boolean sense), some might still be better than others. In order to take this into account, the model-checking approach to verification has to be lifted to a more quantitative perspective [18]. This would allow to *quantify* the quality of systems, and to investigate their tolerance to slight perturbations.

Partly supported by ERC Starting Grant EQualIS and EU FP7 project Cassting.

There are three classical approaches for turning standard model checking to a quantitative perspective. A first approach, building on automata-based techniques to model checking, consists in defining quantitative semantics for finite state automata. This uses weighted automata [21,16], with different possible semantics. Quantitative decision problems for this setting are addressed in [13,15]. A second approach consists in defining distances between models, or between models and specifications, that can provide an accurateness measure of the model w.r.t. the specification. This approach has been developed e.g. in [12], and then extended into the *model measuring* problem [19]. A third approach is to define quantitative specification languages. For probabilistic systems, this approach is rather standard, and quantitative logics like CSL have been defined and used for model-checking [6]. More recently, this approach has been developed for quantitative but non-stochastic systems. We give more details on those approaches in the “related work” paragraph below.

Example 1 (Jobshop scheduling). Consider a finite set of machines, on which we want to schedule finitely many jobs with possibly dependencies between jobs. Standard analysis asks for the existence of a scheduler that satisfies some scheduling policy, or for optimal such schedulers. A more quality-oriented approach could consist in evaluating the average load along a schedule, or the least machine usage, or the average idle time of a given machine. Those cannot be expressed as a standard boolean model-checking question. \triangleleft

Example 2 (Mobile-phone server). Consider a server that should acknowledge any request by some grant (representing the range of frequency—the bigger the range, the larger the grant). Then the quality of such a server could be expressed as the average over all requests of the range that is allocated in response. This cannot be expressed as a standard boolean model-checking question. \triangleleft

In this paper, we propose quantitative measures of correctness based on the linear-time temporal logic LTL. More precisely, we propose a natural extension of LTL, called **avgLTL**, with two natural averaging modalities: a new average-until operator $\psi_1 \tilde{U} \psi_2$ that computes the average value of ψ_1 along the path until ψ_2 has a high value, and where the semantics of standard modalities are extended using a min-max approach; and a long-run average operator $\tilde{G} \psi$, which computes the limit of the values of ψ in the long run along the path. Developing the two examples above, we will show that this logic can express interesting properties.

We focus on the model-checking problem, which corresponds to computing the value of a run (or a Kripke structure) w.r.t. a given property, and on the corresponding decision (comparison with a threshold) and approximation problems. We show that all variants (*i.e.*, all kinds of thresholds, and both when the model is a single path and when it is a Kripke structure) of model-checking and approximation problems are undecidable. Such a robust undecidability is rather surprising (at least to us), given the positive results of [2] for a discounted semantics for LTL, of [22] for an extension of LTL with mean-payoff constraints. Despite the undecidability result for **frequency-LTL** (a boolean extension of LTL with frequency-constrained “until” modality) and for LTL with average assertions

over weighted Kripke structures [8,10], we had hope that some variants of our problem would be decidable.

However we believe these undecidability results are interesting in several respects. (i) First, up to now (see related work below), quantitative specification languages based on LTL have always involved discounting factors, which allows to only consider a bounded horizon; this helps obtaining decidability results. In several papers though, averaging in LTL is mentioned, but left as open research directions. (ii) Also, we prove robust undecidability results, in the sense that undecidability is proven both for model-checking over a path and model-checking a Kripke structure, and for all thresholds; note that many cases require a specific proof. (iii) Finally, our proof techniques are non-trivial and may be interesting in other contexts; we were not able to get a direct encoding of two-counter machines for proving the undecidability of the model-checking problem over Kripke structures, and had to use a diagonal argument; this is due to convergence phenomena that arise in the context of quantitative model-checking, and which have mostly been omitted so far in the rest of the literature.

Related work. Several recent papers have proposed quantitative-verification frameworks based on temporal logic. The authors of [14] were the first to suggest giving temporal logics a quantitative semantics: they extend CTL with various new modalities involving a discount on the future (the later the event, the smaller the impact on the value of the formula). In that framework, model-checking is proven decidable.

As regards linear-time temporal logics, a first attempt to define a quantitative semantics has been proposed in [17]. However, no modality is really quantitative, only the models are quantitative, yielding finitely non-boolean values. Still, the authors suggest discounting and long-run averaging as possible extensions of their work. Another approach is tackled in [1], where functions f are added to the syntax of LTL, with the value of $f(\psi_1, \dots, \psi_k)$ on a path π being the result of applying f to the values of subformulas ψ_1, \dots, ψ_k on π . As explained in [1], this quantitative language is not that expressive: each formula only takes finitely many values. It follows that the verification problems are decidable.

Frequency-LTL, an extension of LTL with “frequency-until”, has been studied in [9], and even though it has a boolean semantics, the frequency modality gives a quantitative taste to the logic: $\phi_1 \mathbf{U}^c \phi_2$ holds true along a path whenever there is a position along that path at which ϕ_2 holds, and the frequency of ϕ_1 along the prefix is at least c . This paper shows the undecidability of the satisfiability problem. We discuss this approach in more details in Section 8, since it shares some techniques with ours.

Finally the recent work [2] is the closest to ours. It studies LTL extended with a discounted until modality: roughly, the values of the subformulas are multiplied by a discount factor, which decreases and tends to zero with the distance to the evaluation point. This way, the further the witness, the lower the value. An automata-based algorithm is given to decide the threshold problem. Due to discounting, whether the value of a formula is larger than some threshold on a path can be checked on a bounded prefix of the path. On the other hand, adding

local average (*i.e.*, the average of finitely many subformulas) yields undecidability (for the existence of a path with value $1/2$). We will discuss with more details this paper in Section 8.

2 Average-LTL

Let \mathcal{P} be a finite set of atomic propositions. A *quantitative Kripke structure* over \mathcal{P} is a 4-tuple $\mathcal{K} = \langle V, v_0, E, \ell \rangle$ where V is a finite set of vertices, $v_0 \in V$ is the initial vertex, $E \subseteq V \times V$ is a set of transitions (which we assume total, meaning that for each $v \in V$, there exists $v' \in V$ s.t. $(v, v') \in E$) and $\ell: V \rightarrow ([0, 1] \cap \mathbb{Q})^{\mathcal{P}}$ is a labelling function, associating with each state the value of each atomic proposition in that state. The Kripke structure \mathcal{K} is said *qualitative* whenever for every $v \in V$ and $p \in \mathcal{P}$, $(\ell(v))(p) \in \{0, 1\}$. A run or path in a Kripke structure \mathcal{K} from $v \in V$ is a finite or infinite sequence $\pi = (v_i)_{i \in I}$ (where I is a (bounded or unbounded) interval of \mathbb{N} containing 0) s.t. $v_0 = v$ and $(v_{i-1}, v_i) \in E$ for all relevant $i \in I \setminus \{0\}$. The size $|\pi|$ of π is the cardinality of I . In the sequel, we will be interested in the sequence $\ell(\pi) = (\ell(v_i))_{i \in I}$, and we will often identify a run with the sequence in $([0, 1] \cap \mathbb{Q})^{\mathcal{P} \times I}$ it defines. Given a run $\pi = (v_i)_{i \in I}$ and an integer j , we write $\pi_{\geq j}$ for the run $(v_{i+j})_{i \geq 0, i+j \in I}$.

We now introduce the logic average-LTL (**avgLTL** for short) and its interpretation over infinite runs. The syntax of **avgLTL** over \mathcal{P} is given by:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathbf{X} \varphi \mid \varphi \mathbf{U} \varphi \mid \mathbf{G} \varphi \mid \varphi \tilde{\mathbf{U}} \varphi \mid \tilde{\mathbf{G}} \varphi.$$

where $p \in \mathcal{P}$. Notice that negation is only allowed on atomic propositions. We write **LTL** for the fragment where $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{G}}$ are not allowed.

Let $\pi = (v_i)_{i \in \mathbb{N}}$ be an infinite run, and φ be an **avgLTL** formula. The valuation $\llbracket \pi, \varphi \rrbracket$ is then given as follows:

$$\begin{aligned} \llbracket \pi, p \rrbracket &= (\ell(v_0))(p) & \llbracket \pi, \neg p \rrbracket &= 1 - (\ell(v_0))(p) \\ \llbracket \pi, \psi_1 \vee \psi_2 \rrbracket &= \max\{\llbracket \pi, \psi_1 \rrbracket, \llbracket \pi, \psi_2 \rrbracket\} & \llbracket \pi, \mathbf{X} \psi \rrbracket &= \llbracket \pi_{\geq 1}, \psi \rrbracket \\ \llbracket \pi, \psi_1 \wedge \psi_2 \rrbracket &= \min\{\llbracket \pi, \psi_1 \rrbracket, \llbracket \pi, \psi_2 \rrbracket\} \\ \llbracket \pi, \mathbf{G} \psi \rrbracket &= \inf_{i \in \mathbb{N}} \llbracket \pi_{\geq i}, \psi \rrbracket \\ \llbracket \pi, \psi_1 \mathbf{U} \psi_2 \rrbracket &= \sup_{i \in \mathbb{N}} \min\{\llbracket \pi_{\geq i}, \psi_2 \rrbracket, \min_{0 \leq j < i} (\llbracket \pi_{\geq j}, \psi_1 \rrbracket)\} \\ \llbracket \pi, \tilde{\mathbf{G}} \psi \rrbracket &= \liminf_{i \rightarrow \infty} (\sum_{j=0}^{j < i} \llbracket \pi_{\geq j}, \psi \rrbracket) / i \\ \llbracket \pi, \psi_1 \tilde{\mathbf{U}} \psi_2 \rrbracket &= \sup \left(\{\llbracket \pi, \psi_2 \rrbracket\} \cup \left\{ \min\{\llbracket \pi_{\geq i}, \psi_2 \rrbracket, (\sum_{j=0}^{j < i} \llbracket \pi_{\geq j}, \psi_1 \rrbracket) / i\} \mid i > 0\} \right\} \right) \end{aligned}$$

We recover the boolean semantics for the standard operators when all atomic propositions have either value 0 (false) or value 1 (true). Note that in that case we might abusively consider that $v_i \in 2^{\mathcal{P}}$, recording the set of atomic propositions with value 1 at each position. The first five rules are standard and natural in a quantitative setting. The semantics of the **U**- and **G**-modalities are also natural: they extend the standard equivalences $\psi_1 \mathbf{U} \psi_2 \equiv \psi_2 \vee (\psi_1 \wedge \mathbf{X}(\psi_1 \mathbf{U} \psi_2))$, and $\mathbf{G} \psi \equiv \psi \wedge \mathbf{X} \mathbf{G} \psi$ to a quantitative setting. The last two modalities are specific

to our setting: formula $\psi_1 \tilde{\mathbf{U}} \psi_2$ computes the average of formula ψ_1 for the i first steps, and then compares the value with that of ψ_2 at the $(i + 1)$ -st step. The best choice of i (if it exists) is then selected, and gives the value to the formula. Formula $\tilde{\mathbf{G}} \psi$ computes the average of ψ in the long-run.

We come back to our two illustrative examples given in the introduction, to show how our logic can be used to express natural properties.

Example 3 (Jobshop scheduling). We come back to Example 1, assuming a set of n machines. Let `load` be an atomic proposition having value k/n at state s if k machines are in use in that state. Notice that we could equivalently use the local averaging operator \oplus of [2] in order to have `load` defined as the average of the atomic propositions indicating which machines are in use. Then formula $\varphi_1 = \text{load } \tilde{\mathbf{U}} \text{stop}$ evaluated on a schedule computes the average machine use along that schedule, if `stop` is a boolean atomic proposition which holds true when all jobs are finished. A schedule assigning value 1 to φ_1 could be seen as an optimal schedule, where no computation power is lost. A schedule assigning a small value to formula φ_1 is a schedule with a large loss of computation power.

On the other hand formula $\varphi_2 = \text{load } \mathbf{U} \text{stop}$ will evaluate to the smallest instantaneous machine use along a schedule. Note that syntactically it is a standard until, but it evaluates differently in our quantitative framework. \triangleleft

Example 4 (Mobile phone server). The quality of the server of Example 2 can be expressed as the average over all requests of the frequency allocated in response. We can write such a property as $\varphi_3 = \tilde{\mathbf{G}} (\neg \text{req} \vee \text{no_grant } \mathbf{U} \text{grant})$, where `req` and `no_grant` are boolean atomic propositions with the obvious meaning, and `grant` is an atomic proposition with value in $[0, 1]$ representing the quality of the allocated range of frequencies (the closer to 1, the better). Larger values of φ_3 then indicate better frequency allocation algorithms. \triangleleft

We also evaluate formulas of `avgLTL` over Kripke structures. If v is a state of the Kripke structure \mathcal{K} and $\varphi \in \text{avgLTL}$, then we define: $\llbracket (\mathcal{K}, v), \varphi \rrbracket = \sup \{ \llbracket \pi, \varphi \rrbracket \mid \pi \text{ is an infinite run of } \mathcal{K} \text{ from } v \}$. We simply write $\llbracket \mathcal{K}, \varphi \rrbracket$ when $v = v_0$ is the initial vertex of \mathcal{K} . Notice that considering the supremum here corresponds to the *existential* semantics of boolean LTL, where the aim is to find a path satisfying the formula.

Example 5. We develop a small toy example to illustrate how simple formulas can be evaluated in the (qualitative) Kripke structure depicted on Fig. 1.

Consider the `avgLTL` formulas $a \tilde{\mathbf{U}} b$ and $c \tilde{\mathbf{U}} b$. For the first formula we have $\llbracket a \cdot b \cdot c^\omega, a \tilde{\mathbf{U}} b \rrbracket = 1$ (the supremum being reached at the second position along the run), and therefore $\llbracket \mathcal{K}, a \tilde{\mathbf{U}} b \rrbracket = 1$.

Now, for the formula $c \tilde{\mathbf{U}} b$ and the same run as above, we have $\llbracket a \cdot b \cdot c^\omega, c \tilde{\mathbf{U}} b \rrbracket = 0$: indeed, the right-hand-side formula b has value zero everywhere except at position 1, but the average of c on the previous positions is zero. For the run $a \cdot (b \cdot c)^\omega$, considering all positions (but position 1) where b is non-zero, we get $\llbracket a \cdot (b \cdot c)^\omega, c \tilde{\mathbf{U}} b \rrbracket = \sup \{ n / (2n + 1) \mid n \in \mathbb{N}_{>0} \} = 1/2$. Note that the value $1/2$ is not reached by any prefix. Now consider the run $\pi'_k = a \cdot b \cdot c^k \cdot (b \cdot c)^\omega$, for some




Fig. 1: A Kripke structure \mathcal{K} .

positive integer k . Then we have $\llbracket \pi'_k, c \tilde{\mathbf{U}} b \rrbracket = \sup \{(k+n)/(k+2n+2) \mid n \in \mathbb{N}\}$. When $k \geq 3$, the supremum is $k/(k+2)$, which is reached for $n = 0$ (i.e., at the second occurrence of b). From this we get that $\llbracket \mathcal{K}, c \tilde{\mathbf{U}} b \rrbracket = 1$. However no run witnesses that value. \triangleleft

3 The problems we consider

In this paper, we consider the following two problems:

Existence problem: given a Kripke structure \mathcal{K} , an **avgLTL** formula φ , and a threshold $\bowtie c$ (with $\bowtie \in \{<, \leq, =, \geq, >\}$ and $c \in [0, 1] \cap \mathbb{Q}$), is there a path π in \mathcal{K} such that $\llbracket \pi, \varphi \rrbracket \bowtie c$?

Value problem: given a Kripke structure \mathcal{K} , an **avgLTL** formula φ , and a threshold $\bowtie c$ (with $\bowtie \in \{<, \leq, =, \geq, >\}$ and $c \in [0, 1] \cap \mathbb{Q}$), does $\llbracket \mathcal{K}, \varphi \rrbracket \bowtie c$?

Note that both problems are different since, as illustrated in Example 5, it can be the case that $\llbracket \mathcal{K}, \varphi \rrbracket = 1$ even though no path of \mathcal{K} assigns value 1 to φ .

We also consider their approximation variants, defined as follows:

Approximate existence problem: given a Kripke structure \mathcal{K} , an **avgLTL** formula φ , a value $c \in [0, 1] \cap \mathbb{Q}$ and $\varepsilon > 0$, is there a path π in \mathcal{K} such that $c - \varepsilon < \llbracket \pi, \varphi \rrbracket < c + \varepsilon$?

Approximate value problem: given a Kripke structure \mathcal{K} , an **avgLTL** formula φ , a value $c \in [0, 1] \cap \mathbb{Q}$ and $\varepsilon > 0$, does $c - \varepsilon < \llbracket \mathcal{K}, \varphi \rrbracket < c + \varepsilon$?

4 Model checking **avgLTL** is undecidable

In the sequel, we prove that **avgLTL** model-checking is *robustly* undecidable, in the sense that all the problems above are undecidable, for all threshold conditions considered. We would like to emphasize that different kinds of threshold give rise to different problems, and could have led to different decidability results. For instance, given a Kripke structure \mathcal{K} and an **avgLTL** formula φ , $\llbracket \mathcal{K}, \varphi \rrbracket > 1/2$ iff there exists an infinite run π in \mathcal{K} such that $\llbracket \pi, \varphi \rrbracket > 1/2$. On the other hand, $\llbracket \mathcal{K}, \varphi \rrbracket = 1/2$ iff there exists a sequence of infinite runs $(\pi^n)_{n \in \mathbb{N}}$ such that

$\llbracket \pi^n, \varphi \rrbracket \leq 1/2$ for every n , and $\lim_{n \rightarrow \infty} \llbracket \pi^n, \varphi \rrbracket = 1/2$. These remarks advocate for a clear and exhaustive study of the different problems with all the different thresholds.

Additionally, we believe that our original proof techniques (in particular the diagonal argument used to circumvent convergence phenomena for the model-checking of Kripke structures) are of particular interest and could be used in related settings. We discuss further these issues and related works in Section 8

We can now state the main results of the paper.

Theorem 6. *The existence problem is undecidable, for every threshold of the form $\bowtie 1/2$, with $\bowtie \in \{<, \leq, =, \geq, >\}$.*

Theorem 7. *The value problem is undecidable, for every threshold of the form $\bowtie 1/2$, with $\bowtie \in \{<, \leq, =, \geq, >\}$.*

We present these results as two distinct theorems, since proofs require very different techniques, even though a similar encoding is used.

Remark 8. Our proofs only involve qualitative Kripke structures. We present the results for $c = 1/2$, but our proofs could be adapted to handle any other rational value in $(0, 1)$ (e.g. by inserting fake actions in the encoding).

Now, if the approximate variants were decidable, then taking e.g. $c = 1$ and $\epsilon = 1/2$, we could decide e.g. whether a formula has value larger than $1/2$, contradicting the previous theorems. Hence:

Theorem 9. *The approximate existence and value problems are undecidable.*

The rest of the paper presents the main ideas of the proof. Due to lack of space, the full proofs could not be included here, but can be found in the research report [11] associated to this paper.

5 Proof of Theorem 6

We only give an explanation of the undecidability for the existence problem with threshold $\geq 1/2$ (the other types of thresholds require a twist in the construction, but no fundamental new argument).

The proof relies on an encoding of the halting problem for deterministic two-counter machines, which is well-known to be undecidable. A two-counter machine \mathcal{M} is a finite-state machine, equipped with two kinds of transitions: *update*-transitions move from one state to another one while incrementing or decrementing one of the counters; *test*-transitions keep the counters unchanged, but may lead to two different states depending on the positiveness of one of the counters. The machine has a special state, called the *halting state*, from which no transitions is possible. We assume w.l.o.g. that all the other states have exactly one outgoing transition.

A configuration of \mathcal{M} is given by the current state and the values of both counters. A run of \mathcal{M} is a sequence of consecutive configurations *which might not properly update the counters*. It is said *valid* whenever the counters are properly

updated along the run. There is a unique maximal valid run in \mathcal{M} from the initial configuration: it is either halting or infinite.

The idea of our reduction is to build a Kripke structure which generates the encodings of all (including invalid) runs of \mathcal{M} : it has to take care of the discrete structure of \mathcal{M} , but does not check that counters are properly updated along the run. Correct update of counter values will be checked using an **avgLTL** formula.

Description of the encoding. We first explain how we encode the runs of \mathcal{M} . We only give a simplified idea of the encoding. We write Q for the set of states of \mathcal{M} .

For $p \geq 2$, we write \mathbb{B}_p for the set $\{0, 1, \dots, p-1\}$. For $b \in \mathbb{B}_p$, we let $b^{+i} = b+i \pmod p$. An element of \mathbb{B}_p is abusively called a *bit*. These bits are used to distinguish between consecutive configurations. For the rest of this section, taking $p = 2$ would be sufficient, but the proof of Theorem 7 requires higher values for p . We encode configurations of \mathcal{M} using the following finite set of atomic propositions: $\mathcal{P}_p = (Q \cup \{a_0, a_1\}) \times \mathbb{B}_p \cup \{\#\}$. The symbol $\#$ will be a marker for halting computations.

Exactly one atomic proposition from \mathcal{P}_p will have value one at each position along the encoding (the other propositions having value zero). Given a bit b , a configuration $\gamma = (q, n_0, n_1)$ of \mathcal{M} is encoded as the word $\text{enc}_b(\gamma) = (q, b) \cdot (a_0, b)^{n_0} \cdot (a_1, b)^{n_1}$. For a halting configuration, we set $\text{enc}_b(\gamma) = (q_{\text{halt}}, b)$.

The bit $b \in \mathbb{B}_p$ is incremented (modulo p) from one configuration to the next one. Let $\rho = \gamma_0 \cdot \gamma_1 \cdots$ be a (not necessary valid) run in \mathcal{M} . The p -encoding of ρ is then given by:

$$p\text{-enc}(\rho) = \begin{cases} \text{enc}_{b_0}(\gamma_0) \cdot \text{enc}_{b_1}(\gamma_1) \cdot \text{enc}_{b_2}(\gamma_2) \cdots & \text{if } \rho \text{ is infinite} \\ \text{enc}_{b_0}(\gamma_0) \cdot \text{enc}_{b_1}(\gamma_1) \cdots \text{enc}_{b_{n-1}}(\gamma_{n-1}) \#^\omega & \text{if } \rho \text{ has length } n \end{cases}$$

with $b_j = j \pmod p$ for every j . We write $\text{enc}(\rho)$ if p is clear from the context.

We can easily construct a Kripke structure that generates the encodings of all possible (valid or invalid) runs of \mathcal{M} . For index p , we write $\mathcal{K}_{\mathcal{M}}^p$ for the corresponding Kripke structure. We now turn to the **avgLTL** formula, whose role is to check proper updates of the counters.

Definition of the formulas. We will define a formula $\text{consec}_{\mathcal{M}}^p$, which will be used to check that each single consecution in the run properly updates the counters. Then we define formula

$$\text{halt}_{\mathcal{M}}^p = \mathbf{F} q_{\text{halt}} \wedge \mathbf{G} \text{consec}_{\mathcal{M}}^p.$$

It is rather clear that if we can build such a formula $\text{consec}_{\mathcal{M}}^p$, then the above formula will check that the unique maximal valid run of \mathcal{M} is halting. Unfortunately, things are not that easy, and formula $\mathbf{G} \text{consec}_{\mathcal{M}}^p$ will only be able to check the validity of *finite* runs

We now focus on defining $\text{consec}_{\mathcal{M}}^p$, using the average-until modality. We only give an intuition (the full definition requires the complete encoding). Consider a

portion P of the p -encoding of a run ρ , which corresponds to a single-step of the computation of \mathcal{M} where instruction q keeps both counter values unchanged:

$$\dots (q, b) \cdot (a_0, b)^{n_0} \cdot (a_1, b)^{n_1} \cdot (q', b^{+1}) \cdot (a_0, b^{+1})^{n'_0} \cdot (a_1, b^{+1})^{n'_1} (q'', b^{+2}) \dots$$

The formula has to enforce $n'_0 = n_0$ and $n'_1 = n_1$. This is the case if, and only if, for every $\alpha \in \{1 + n_0 + n_1, 1 + n_0 + n'_1, 1 + n'_0 + n_1, 1 + n'_0 + n'_1\}$,

$$\frac{\alpha}{1 + n_0 + n_1 + 1 + n'_0 + n'_1} = \frac{1}{2}.$$

The denominator is the length of the portion from (q, b) to the position just before (q'', b^{+2}) , whereas the various values for α are the number of positions where some distinguished atomic proposition holds along this portion. For instance, $1 + n'_0 + n_1$ is the number of positions where formula $\psi = (q', b^{+1}) \vee (a_0, b^{+1}) \vee (a_1, b)$ holds along P . Computing the above quotient will be done using an $\tilde{\mathbf{U}}$ -formula: $\llbracket P, \psi \tilde{\mathbf{U}}(q'', b^{+2}) \rrbracket$ precisely equals $\frac{\alpha}{1 + n_0 + n_1 + 1 + n'_0 + n'_1}$

Using this idea, we are able to construct a formula $\text{consec}_{\mathcal{M}}^p$ (as a conjunction of several $\tilde{\mathbf{U}}$ -formulas) whose value is $1/2$ along a single step of the computation if, and only if, this step is valid (that is, it correctly updates the counters).

Correctness of the reduction. Even though formula $\text{consec}_{\mathcal{M}}^p$ properly checks the validity of a single step of the computation, it might be the case that $\llbracket p\text{-enc}(\rho), \mathbf{G} \text{consec}_{\mathcal{M}}^p \rrbracket = 1/2$, even though the whole computation is not valid: this is due to the definition of the semantics of $\tilde{\mathbf{U}}$ as the supremum over all positions of the average; in particular, a single error in the computation can be hidden in the rest of the run. Consider for instance the counter machine in Fig. 2. The unique initial and maximal valid run of \mathcal{M} halts. However, if the first transition increments counter a_0 twice, and all further transitions are properly taken, then the resulting (invalid) run will assign value $1/2$ to formula $\mathbf{G} \text{consec}_{\mathcal{M}}^p$.




Fig. 2: There is an invalid infinite run ρ such that $\llbracket p\text{-enc}(\rho), \mathbf{G} \text{consec}_{\mathcal{M}}^p \rrbracket = 1/2$

Still, we are able to prove the following classification of runs of \mathcal{M} in terms of the value of $\text{halt}_{\mathcal{M}}^p$. It proves the fact that formula $\text{consec}_{\mathcal{M}}^p$ properly checks the validity of a single step of the computation, provided the $\tilde{\mathbf{U}}$ -formulas cannot

benefit from the supremum semantics. This is the case when the run in the Kripke structure ends with $\#\omega$, which corresponds to finite runs of \mathcal{M} .

Classification 1. Fix $p \geq 2$. Let ρ be a maximal run in \mathcal{M} .

- if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^p \rrbracket = 0$;
- if ρ is finite and valid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$;
- if ρ is finite and invalid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^p \rrbracket < 1/2$.

Corollary 10. Fix $p \geq 2$. The following five statements are equivalent:

1. \mathcal{M} halts;
2. the unique initial and maximal valid run $\rho_{\mathcal{M}}$ of \mathcal{M} is such that $\llbracket p\text{-enc}(\rho_{\mathcal{M}}), \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$;
3. there exists an initial maximal run ρ in \mathcal{M} such that $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$;
4. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$;
5. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^p \rrbracket \geq 1/2$.

This corollary allows to conclude the undecidability proof of Theorem 6.

6 Proof of Theorem 7

As already mentioned, whether $\llbracket \mathcal{K}, \varphi \rrbracket > 1/2$ (and dually, $\llbracket \mathcal{K}, \varphi \rrbracket \leq 1/2$) is equivalent to the existence of a path whose value is strictly more than $1/2$, which we just proved undecidable.

We now turn to the more interesting cases of $=$ (the result for \geq and $<$ directly follows, as we explain at the end of this proof). We were not able to write a direct proof as previously, because we could not distinguish between counter machines that have a halting computation (whose encoding has value $1/2$ against formula $\text{halt}_{\mathcal{M}}^p$ above) and counter machines that have sequences of computations whose encodings have values *converging* to $1/2$.




Fig. 3: A non-halting two-counter machine for which $\llbracket \mathcal{K}_{\mathcal{M}}^p, \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$

Example 11. We consider the deterministic two-counter machine \mathcal{M} of Fig. 3, having q_0 as its initial state. The unique initial and maximal valid run of \mathcal{M} is infinite (it loops in $q_1 \rightleftharpoons q_2$). A single error can make the transition from q_1 to q_3 available, from which valid consecutions lead to q_{halt} . The *weight* of this error can be arbitrarily small, as it can occur with an arbitrarily large value of a_0 . It is not difficult to check that $\llbracket \mathcal{K}_{\mathcal{M}}^p, \text{halt}_{\mathcal{M}}^p \rrbracket = 1/2$ (for any $p \geq 2$). \triangleleft

Analysis of a non-halting two-counter machine. We consider a deterministic accept/reject two-counter machine \mathcal{M} : such machines have two halting states, now named q_{accept} and q_{reject} . Their computations may still be infinite. We consider formula $\mathbf{consec}_{\mathcal{M}}^p$ again, and define $\mathbf{accept}_{\mathcal{M}}^p = \mathbf{F} q_{accept} \wedge \mathbf{G} \mathbf{consec}_{\mathcal{M}}^p$.

We first analyse the impact of the first error along a finite run ρ of \mathcal{M} onto the value of $\mathbf{G} \mathbf{consec}_{\mathcal{M}}^p$, and we are able to show the following surprising but crucial lemma (remember the example of Fig. 2) whose proof requires long technical developments. The condition imposed on p is a sufficient condition for “detecting” invalid consecutions along finite runs. The computation leading to this value is explained in the long version [11] of this work.

Lemma 12. *Fix $p \geq 927$. Let ρ be a finite invalid run of \mathcal{M} . Assume $\rho_i \rho_{i+1}$ is the first invalid consecution along ρ , and write \mathbf{step}_i for the portion of $p\text{-enc}(\rho)$ corresponding to that consecution. Pick $n \geq 30$ such that $\llbracket \mathbf{step}_i, \mathbf{consec}_{\mathcal{M}}^p \rrbracket \leq 1/2 - 1/n$. Then $\llbracket p\text{-enc}(\rho), \mathbf{G} \mathbf{consec}_{\mathcal{M}}^p \rrbracket \leq 1/2 - 1/n$.*

This allows to prove the next fundamental result:

Lemma 13. *Fix $p \geq 927$, and assume that $\llbracket \mathcal{K}_{\mathcal{M}}^p, \mathbf{accept}_{\mathcal{M}}^p \rrbracket = 1/2$, but that no run ρ of \mathcal{M} has $\llbracket p\text{-enc}(\rho), \mathbf{accept}_{\mathcal{M}}^p \rrbracket = 1/2$. Then the unique initial and maximal valid run of \mathcal{M} is infinite.*

We sketch the proof of this lemma, since it contains an interesting argument.

Sketch of proof. Let ρ be the unique initial and maximal valid run of $\mathcal{K}_{\mathcal{M}}^p$. Let $(\rho^n)_{n \in \mathbb{N}}$ be a sequence of initial and maximal runs such that $\llbracket p\text{-enc}(\rho^n), \mathbf{accept}_{\mathcal{M}}^p \rrbracket > 1/2 - 1/n$ (such a sequence exists by hypothesis, but runs ρ^n might be invalid). Pick $n \geq 30$, and let $\rho_{i_n}^n \rho_{i_n+1}^n$ be the first invalid consecution of ρ^n . Write \mathbf{step}_{i_n} for the portion of $p\text{-enc}(\rho^n)$ corresponding to that consecution. Applying Lemma 12, we get that $\llbracket \mathbf{step}_{i_n}, \mathbf{consec}_{\mathcal{M}}^p \rrbracket > 1/2 - 1/n$. Since $\rho_{i_n}^n \rho_{i_n+1}^n$ is an invalid consecution, we also have that $\llbracket \mathbf{step}_{i_n}, \mathbf{consec}_{\mathcal{M}}^p \rrbracket < 1/2$. It follows that $1/(|\mathbf{step}_{i_n}| - 1) < 1/n$, which implies that $|\mathbf{step}_{i_n}| > n$. Now, the prefix of ρ of size i_n coincides with that of ρ^n , since $\rho_{i_n}^n \rho_{i_n+1}^n$ is the first invalid consecution. We conclude that ρ contains configurations of arbitrarily large size, so that the sum of the two counters is unbounded along ρ . Hence ρ is infinite. \square

A diagonal argument. Any deterministic Turing machine can be simulated by a deterministic two-counter machine [20]. In particular, given a deterministic Turing machine B , we can build a deterministic two-counter machine $\mathcal{M}(B)$ whose computation mimics the run of B on input B . Then $\mathcal{M}(B)$ accepts (resp. rejects, does not halt) if, and only if, B accepts (resp. rejects, does not halt on) input B .

We fix $p \geq 927$, and define the following function \mathcal{H} , which takes as input a deterministic Turing machine B :

$$\mathcal{H}(B) = \begin{cases} \text{accept} & \text{if } \llbracket \mathcal{K}_{\mathcal{M}(B)}^p, \mathbf{accept}_{\mathcal{M}(B)}^p \rrbracket = 1/2 \\ \text{reject} & \text{otherwise} \end{cases}$$

Proposition 14. *The function \mathcal{H} is not computable.*

Proof. Towards a contradiction, assume \mathcal{H} is computable. Let $\mathcal{T}_{\mathcal{H}}$ be a deterministic Turing machine that computes \mathcal{H} . Notice in particular that $\mathcal{T}_{\mathcal{H}}$ halts on all its inputs; we assume that it ends in its state $q_{accept}^{\mathcal{T}}$ when \mathcal{H} accepts the input, and in $q_{reject}^{\mathcal{T}}$ when \mathcal{H} returns *reject*.

We now define the following deterministic Turing machine \mathcal{C} , which takes as input a deterministic Turing machine B :

$\mathcal{C}(B)$: Simulate $\mathcal{T}_{\mathcal{H}}$ on B ;
 If the simulation ends in $q_{accept}^{\mathcal{T}}$ then goto $q_{reject}^{\mathcal{C}}$, otherwise goto $q_{accept}^{\mathcal{C}}$.

The Turing machine \mathcal{C} terminates on all its inputs, since so does $\mathcal{T}_{\mathcal{H}}$; also, \mathcal{C} is deterministic, and we can therefore run \mathcal{C} on input \mathcal{C} itself.

Assume \mathcal{C} accepts input \mathcal{C} . This means that $\mathcal{H}(\mathcal{C})$ rejects, which means that $\llbracket \mathcal{K}_{\mathcal{M}(\mathcal{C})}^p, \text{accept}_{\mathcal{M}(\mathcal{C})}^p \rrbracket < 1/2$. This means that $\mathcal{M}(\mathcal{C})$ does not accept (by a straightforward extension of Corollary 10 to accept/reject two-counter machines), and therefore \mathcal{C} does not accept \mathcal{C} , contradicting our hypothesis.

Hence \mathcal{C} rejects input \mathcal{C} , so that $\llbracket \mathcal{K}_{\mathcal{M}(\mathcal{C})}^p, \text{accept}_{\mathcal{M}(\mathcal{C})}^p \rrbracket = 1/2$. However, since \mathcal{C} does not accept \mathcal{C} , the unique initial and maximal valid run of $\mathcal{M}_{\mathcal{C}}$ is either infinite or rejecting. Applying Lemma 13 to $\mathcal{M}_{\mathcal{C}}$, we get that it is actually infinite. This means that the simulation of $\mathcal{T}_{\mathcal{H}}$ on input \mathcal{C} does not terminate. This contradicts the fact that $\mathcal{T}_{\mathcal{H}}$ terminates on every input. Therefore \mathcal{H} is not computable. \square

Theorem 7 is a direct consequence of this lemma for threshold = 1/2. Now, using Classification 1, for a deterministic two-counter machine \mathcal{M} , it holds that $\llbracket \mathcal{K}_{\mathcal{M}}^p, \text{accept}_{\mathcal{M}}^p \rrbracket = 1/2$ iff $\llbracket \mathcal{K}_{\mathcal{M}}^p, \text{accept}_{\mathcal{M}}^p \rrbracket \geq 1/2$. Hence the above proof applies to threshold $\geq 1/2$ as well. The case of $< 1/2$ is the dual of $\geq 1/2$: if \mathcal{K} is a Kripke structure and φ an *avgLTL* formula, $\llbracket \mathcal{K}, \varphi \rrbracket < 1/2$ iff it is not the case that $\llbracket \mathcal{K}, \varphi \rrbracket \geq 1/2$, which proves the result for threshold $< 1/2$ as well.

7 Proof of Theorem 9

We now discuss the undecidability of the approximate variants. It relies on the same encoding as that for the existence problem and threshold $> 1/2$. For that threshold, we have a classification of the runs similar to Classification 1, for formula $\text{halt}_{\mathcal{M}}^{p, >}$: for every maximal run ρ in \mathcal{M} :

- if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket = 0$;
- if ρ is finite and valid, then $1/2 < \llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket < 3/4$;
- if ρ is finite and invalid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket \leq 1/2$.

We deduce that \mathcal{M} halts iff there exists an initial and maximal valid run π in $\mathcal{K}_{\mathcal{M}}^p$ with $1/2 < \llbracket \pi, \text{halt}_{\mathcal{M}}^{p, >} \rrbracket < 3/4$. This shows undecidability of the approximate existence problem.

Now, we also have in this case the equivalence with $1/2 < \llbracket \mathcal{K}_{\mathcal{M}}^p, \text{halt}_{\mathcal{M}}^{p, >} \rrbracket < 7/8$ (not $3/4$ since there might be some convergence phenomenon towards value $3/4$), which also shows the undecidability of the approximate value problem.

8 Discussion on related works

In this section, we would like to illustrate the difficulty of lifting temporal-logic model checking from the qualitative to the quantitative setting. As we saw in this paper, several new convergence phenomena do appear, which make the problem complex, but also make the proofs difficult. Our undecidability proofs in this paper involve difficult techniques to properly handle the convergence phenomena that appear in the semantics of the logic. This difficulty has led to several wrong arguments in the related literature, as we now illustrate.

We first discuss the logic **frequency-LTL** of [9]. This logic has a boolean semantics, but extends **LTL** with a **frequency-U** modality, which gives it a quantitative taste: formula $\phi_1 \mathbf{U}^c \phi_2$ holds true along a path π whenever there is a position n along π at which ϕ_2 holds, and the number of previous positions where ϕ_1 holds is larger than or equal to $c \cdot n$ (hence c is a lower bound on the frequency of ϕ_1 on the prefix before ϕ_2 holds). Note that it need not be the case that the position n is the first position where ϕ_2 holds: for instance *abbcaaac* satisfies formula $a \mathbf{U}^{\frac{1}{2}} c$, but at the first occurrence of c , the frequency of a on the prefix is $1/3$, which is less than $1/2$; the correct witnessing position for $a \mathbf{U}^{\frac{1}{2}} c$ is the second occurrence of c , where the frequency of a becomes $4/7$. In **frequency-LTL**, there is no convergence phenomena, but some possibly unbounded search for some witnessing position. Then evaluating $b \mathbf{U}^{\frac{1}{2}} c$ on a path π is not equivalent to comparing formula $b \tilde{\mathbf{U}} c$ to value $1/2$ on path π : first because of convergence phenomena (as illustrated in Example 5), and because in our quantitative setting, the value of the right-hand-side subformula could be less than $1/2$.

It is shown in [9] that the validity problem for **frequency-LTL** is undecidable, and our reduction shares similarities with that reduction (but we believe that our reduction is simpler, and the result stronger, since it uses no nested $\tilde{\mathbf{U}}$). However the undecidability proof (as written in [9]) has a flaw: it relies on the claim that “[t]he formula $b \mathbf{U}^{\frac{1}{2}} l \wedge \hat{b} \mathbf{U}^{\frac{1}{2}} l$ enforces the pattern $b^m \hat{b}^m l \dots$ ” (the order of b ’s and \hat{b} ’s is enforced by another **LTL** formula). This claim is wrong in general, since the $\mathbf{U}^{\frac{1}{2}}$ -formulas might not refer to the same occurrences of l . The proof can be patched¹, and one way is to restrict to paths that end with $\#\omega$ for some marker $\#$; in that way a backward argument can be used to check proper encoding of the execution of the two-counter machine (this is actually what we do in the proof of Theorem 6).

We now discuss the logic **discounted-LTL** of [2]. This logic gives a quantitative semantics to an extension of **LTL**, with a new **discounted-U** modality: given a discount function η , the value of formula $\phi_1 \mathbf{U}_\eta \phi_2$ along a path π is the supremum over all positions n along π of the minimum of the value of ϕ_2 at that position, discounted by $\eta(n)$, and of the values of ϕ_1 at every earlier position i , discounted by $\eta(i)$. Satisfiability is proven decidable; it is shown undecidable when adding the local average operator \oplus , which computes the average of two formulas.

¹ Personal communication with the authors.

Those results are then extended to the model-checking problem². While the first result extends properly for threshold $< c$ (since the infimum over all paths is smaller than c if, and only if, there is a path that evaluates to a value smaller than c ; hence convergence phenomena are avoided), it is not valid for Theorem 3 of [2] (which is stated with threshold $> c$). Also, undecidability of the model-checking problem with local-average operator (Theorem 6 of [2]) is not correct since it does not take convergence phenomena into account. A corrected version of the proof is available in [3]; while it does not use a diagonal argument as we do, the undecidability proof is not a direct encoding of a two-counter machine, but requires computing the value of two different formulas in order to encode the halting problem.

This all shows that extending temporal logics to a quantitative setting is more than a simple exercise: complex convergence phenomena come into play, which have to be understood and handled with extreme care. We hope that our work will provide new insights about these problems, and believe that our techniques can be useful for handling them.

9 Conclusion and future work

We believe that our logic `avgLTL` is a very relevant logic in many applications. It provides a way of *measuring* some properties, such as the average load of the CPUs in scheduling applications. We proved that the value of a formula can not be computed—and not even approximated. For the interesting case however (deciding whether $\llbracket \mathcal{K}, \phi \rrbracket \geq \eta$), we had to resort to an original diagonal argument to get around convergence phenomena.

Our negative results certainly echo back the fact, mentioned e.g. in [17], that averaging does not fit well with classical automata-based approaches for temporal logics. Indeed, averaging gives rise to new values that are not present in the original automaton. Discounting LTL instead of averaging has the same difficulty, but this is compensated by the fact that when discounting, the value of a formula can be approximated by considering only a finite prefix of a run [2].

We are currently investigating two directions in order to get decidability results: first by adding discounting on the right-hand-side formula (while keeping averaging on the left-hand-side); second, by considering the *qualitative* cases of `avgLTL`, namely whether a formula has value 0 or 1. One difficulty here is that in some cases the witnesses are a family of paths, instead of just a single path.

References

1. S. Almagor, U. Boker, and O. Kupferman. Formalizing and reasoning about quality. In ICALP'13, LNCS 7966, p. 15–27. Springer, 2013.

² Note that the value of a Kripke structure is defined there as the infimum over all paths, and not as the supremum as we do here; this corresponds to the duality between existential *vs* universal quantification in standard LTL

2. S. Almagor, U. Boker, and O. Kupferman. Discounting in LTL. In TACAS'14, LNCS, Lecture Notes in Computer Science. Springer, 2014. To appear.
3. S. Almagor, U. Boker, and O. Kupferman. Discounting in LTL. Research Report 1406.4249, arXiv, 2014. 21 pages.
4. R. Alur and D. L. Dill. A theory of timed automata. *Theor. Computer Science*, 126(2):183–235, 1994.
5. R. Alur, S. La Torre, and G. J. Pappas. Optimal paths in weighted timed automata. In HSCC'01, LNCS 2034, p. 49–62. Springer, 2001.
6. A. Aziz, K. Sanwal, V. Singhal, and R. K. Brayton. Model-checking continuous-time Markov chains. *ACM Transactions on Computational Logic*, 1(1):162–170, 2000.
7. G. Behrmann, A. Fehnker, T. Hune, K. G. Larsen, P. Pettersson, J. Romijn, and F. Vaandrager. Minimum-cost reachability for priced timed automata. In HSCC'01, LNCS 2034, p. 147–161. Springer, 2001.
8. U. Boker, K. Chatterjee, T. A. Henzinger, and O. Kupferman. Temporal specifications with accumulative values. In LICS'11, p. 43–52. IEEE Comp. Soc. Press, 2011.
9. B. Bollig, N. Decker, and M. Leucker. Frequency linear-time temporal logic. In TASE'12, p. 85–92. IEEE Comp. Soc. Press, 2012.
10. P. Bouyer, P. Gardy, and N. Markey. Quantitative verification of weighted kripke structures. Research Report LSV-14-08, Laboratoire Spécification et Vérification, ENS Cachan, France, 2014. 26 pages.
11. P. Bouyer, N. Markey, and R. M. Matteplackel. Quantitative verification of weighted kripke structures. Research Report LSV-14-02, Laboratoire Spécification et Vérification, ENS Cachan, France, 2014. 35 pages.
12. P. Černý, T. A. Henzinger, and A. Radhakrishna. Simulation distances. *Theor. Computer Science*, 413(1):21–35, 2012.
13. K. Chatterjee, L. Doyen, and T. A. Henzinger. Quantitative languages. *ACM Transactions on Computational Logic*, 11(4), 2010.
14. L. de Alfaro, M. Faella, T. A. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. *Theor. Computer Science*, 345(1):139–170, 2005.
15. L. Doyen. *Games and Automata: From Boolean to Quantitative Verification*. Mémoire d'habilitation, ENS Cachan, France, 2012.
16. M. Droste, W. Kuich, and W. Vogler, editors. *Handbook of Weighted Automata*. Springer, 2009.
17. M. Faella, A. Legay, and M. Stoelinga. Model checking quantitative linear time logic. In QAPL'08, ENTCS 220, p. 61–77. Elsevier Science, 2008.
18. T. A. Henzinger. Quantitative reactive models. In MODELS'12, LNCS 7590, p. 1–2. Springer, 2012.
19. T. A. Henzinger and J. Otop. From model checking to model measuring. In CONCUR'13, LNCS 8052, p. 273–287. Springer, 2013.
20. M. L. Minsky. *Computation: Finite and Infinite Machines*. Prentice Hall, Inc., 1967.
21. M.-P. Schützenberger. On the definition of a family of automata. *Inf. & Cont.*, 4(2-3):245–270, 1961.
22. T. Tomita, S. Hiura, S. Hagihara, and N. Yonezaki. A temporal logic with mean-payoff constraints. In ICFEM'12, LNCS 7635, p. 249–265. Springer, 2012.

Technical appendix

In this appendix, we develop the full proofs of our results. For the sake of readability, this appendix repeats the arguments already developed in the main part of the paper, together with detailed proofs.

Because we sometimes evaluate **avgLTL** on finite portions of runs, we extend the semantics of **avgLTL** to both finite and infinite runs. We extend the grammar of **avgLTL** with the dual-**X** operator, written $\overline{\mathbf{X}}\varphi$.

Let $\pi \in ([0, 1]^{\mathcal{P}})^I$ be a finite or infinite run, and φ be an **avgLTL** formula. Then, writing $\pi = (v_i)_{i \in I}$, the valuation $\llbracket \pi, \varphi \rrbracket$ is given as follows:

$$\begin{aligned}
\llbracket \pi, p \rrbracket &= v_0(p) \\
\llbracket \pi, \neg p \rrbracket &= 1 - v_0(p) \\
\llbracket \pi, \psi_1 \vee \psi_2 \rrbracket &= \max\{\llbracket \pi, \psi_1 \rrbracket, \llbracket \pi, \psi_2 \rrbracket\} \\
\llbracket \pi, \psi_1 \wedge \psi_2 \rrbracket &= \min\{\llbracket \pi, \psi_1 \rrbracket, \llbracket \pi, \psi_2 \rrbracket\} \\
\llbracket \pi, \mathbf{X} \varphi \rrbracket &= \begin{cases} 0 & \text{if } |\pi| = 1 \\ \llbracket \pi_{\geq 1}, \varphi \rrbracket & \text{otherwise} \end{cases} \\
\llbracket \pi, \overline{\mathbf{X}} \varphi \rrbracket &= \begin{cases} 1 & \text{if } |\pi| = 1 \\ \llbracket \pi_{\geq 1}, \varphi \rrbracket & \text{otherwise} \end{cases} \\
\llbracket \pi, \mathbf{G} \varphi \rrbracket &= \inf_{i \in I} \llbracket \pi_{\geq i}, \varphi \rrbracket \\
\llbracket \pi, \psi \mathbf{U} \varphi \rrbracket &= \sup_{i \in I} \min\{\llbracket \pi_{\geq i}, \varphi \rrbracket, \min_{0 \leq j < i} (\llbracket \pi_{\geq j}, \psi \rrbracket)\} \\
\llbracket \pi, \tilde{\mathbf{G}} \varphi \rrbracket &= \liminf_{i \in I, i \rightarrow \infty} (\sum_{j=0}^{j < i} \llbracket \pi_{\geq j}, \varphi \rrbracket) / i \\
\llbracket \pi, \psi \tilde{\mathbf{U}} \varphi \rrbracket &= \sup \left(\{\llbracket \pi, \varphi \rrbracket\} \cup \{\min\{\llbracket \pi_{\geq i}, \varphi \rrbracket, (\sum_{j=0}^{j < i} \llbracket \pi_{\geq j}, \psi \rrbracket) / i\} \mid i \in I \setminus \{0\}\} \right)
\end{aligned}$$

Notice that both semantics coincide on infinite runs.

A Proof of Theorem 6

We begin with proving undecidability for the existence problem. The proof relies on an encoding of deterministic two-counter machines. For completeness of the presentation, and for fixing notations, we first define two-counter machines.

A.1 Two-counter machines

A two-counter machine \mathcal{M} is a tuple $\langle Q, \delta, \{a_0, a_1\}, q_0, q_{halt} \rangle$, where Q is a finite set of states, $q_0, q_{halt} \in Q$ are the initial and halting states, a_0 and a_1 are the two counter names, and δ is the transition relation, defined as a subset of

$$\left(Q' \times \{a_0, a_1\} \times \{++, --\} \times Q' \right) \cup \left(Q' \times \{a_0, a_1\} \times Q' \times Q' \right) \cup \left(Q' \times \{q_{halt}\} \right)$$

where $Q' = Q \setminus \{q_{halt}\}$. The first kind of transitions increment and decrement the counters, whereas the second kind of transitions is for conditional jumps and zero-tests. The last kind of transitions is for halting. The semantics of \mathcal{M} is

given as an (infinite-state) transition system, where configurations are elements of $Q \times \mathbb{N} \times \mathbb{N}$, and there is a move $(q, n_0, n_1) \rightarrow_{\mathcal{M}} (q', n'_0, n'_1)$ whenever one of the following conditions holds:

- (*increment*) there is a transition $(q, a_i ++, q') \in \delta$ such that $n'_i = n_i + 1$ and $n'_{1-i} = n_{1-i}$;
- (*decrement*) there is a transition $(q, a_i --, q') \in \delta$ such that $n'_i = n_i - 1$ and $n'_{1-i} = n_{1-i}$;
- (*conditional jump*) there is a transition $(q, a_i = 0, q_{=0}, q_{>0}) \in \delta$ such that if $n_i = 0$, then $q' = q_{=0}$, else $q' = q_{>0}$; furthermore, $(n'_0, n'_1) = (n_0, n_1)$.

Finally there are moves $(q, n_0, n_1) \rightarrow_{\mathcal{M}} q_{halt}$ as soon as $(q, q_{halt}) \in \delta$. In our two-counter machines, we impose that each decrement be preceded by an appropriate conditional jump checking that the counter to be decremented has positive value. A two-counter machine is said *deterministic* whenever for every $q \neq q_{halt}$, there is *exactly* one element $(q, a_i ++, q')$ or $(q, a_i --, q')$ or $(q, a_i = 0, q_{=0}, q_{>0})$ or (q, q_{halt}) in δ .

Given two configurations $\gamma = (q, n_0, n_1)$ and $\gamma' = (q', n'_0, n'_1)$, the sequence $\gamma\gamma'$ is called a *valid consecution* of \mathcal{M} whenever $\gamma \rightarrow_{\mathcal{M}} \gamma'$; it is only a *consecution* whenever there exist $n''_0, n''_1 \in \mathbb{N}$ such that $\gamma \rightarrow_{\mathcal{M}} (q', n''_0, n''_1)$ (counter values can be incorrectly updated).

A finite or infinite sequence $(\gamma_i)_{i \in I}$ of configurations is called a run (resp. a valid run) whenever for every $i \in I$, $\gamma_i\gamma_{i+1}$ is a consecution (resp. valid consecution). It is said initial whenever $\gamma_0 = (q_0, 0, 0)$. If \mathcal{M} is deterministic, there is a unique maximal valid initial run, which is either finite (if it ends in q_{halt}) or infinite. The halting problem for \mathcal{M} asks whether the maximal valid initial run of \mathcal{M} is finite, and ends with counter values zero.

Theorem 15 ([20]). *The halting problem for deterministic two-counter machines is undecidable.*

We now show that for every $\bowtie \in \{<, \leq, =, \geq, >\}$, given a deterministic two-counter machine \mathcal{M} , we can construct a Kripke structure $\mathcal{K}_{\mathcal{M}}^p$ and an avgLTL formulas $\mathbf{halt}_{\mathcal{M}}^{p, \bowtie}$, such that \mathcal{M} halts if, and only if, there is a path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \mathbf{halt}_{\mathcal{M}}^{p, \bowtie} \rrbracket \bowtie 1/2$.

For convenience in our encoding, we assume w.l.o.g. that two-counter machines increment and decrement counters by 2. We fix for the rest of this section a deterministic two-counter machine $\mathcal{M} = \langle Q, \delta, \{a_0, a_1\}, q_0, q_{halt} \rangle$.

A.2 Encoding the runs of \mathcal{M}

If $p \in \mathbb{N}$, we write \mathbb{B}_p for the set $\{0, 1, \dots, p-1\}$. For $b \in \mathbb{B}_p$, we let $b^{+i} = b + i \bmod p$. An element of \mathbb{B}_p will abusively be called a *bit*. We encode configurations of \mathcal{M} using the following finite set of atomic propositions:

$$\mathcal{P}_p = \left(Q \cup (Q \times \{a_0, a_1\} \times \mathbb{B}_2^2) \right) \times \left(\mathbb{B}_p \times (Q \cup \{\perp\}) \right) \cup \{\#\}$$

where $p \geq 2$ will be defined later on. Exactly one atomic proposition from this subset will be true at each position along the encoding. A configuration $\gamma = (q, n_0, n_1)$ is encoded as a sequence of $1 + n_0 + n_1$ propositions in \mathcal{P}_p as follows:

$$\mathbf{enc}_{b,q'}(\gamma) = [q, (b, q')] \cdot [(q, a_0, C_0), (b, q')]^{n_0} \cdot [(q, a_1, C_1), (b, q')]^{n_1}.$$

All propositions encoding the same configuration share the same second part $(b, q') \in \mathbb{B}_p \times (Q \cup \{\perp\})$, whose role will be explained later. The first part of the propositions in this encoding respectively represent the current state q , n_0 copies of (q, a_0, C_0) , where $C_0[0]$ encodes the truth value of “ $n_0 > 0$ ” and $C_0[1] = 0$, and n_1 copies of (q, a_1, C_1) , with $C_1[0] = C_0[0]$ and $C_1[1]$ is the truth value of “ $n_1 > 0$ ”.

For an halting configuration, we set $\mathbf{enc}_{b,q'}(\gamma) = [q_{halt}, (b, q')]$ (which we might simply write q_{halt} in the sequel, as no ambiguity may arise). The bit $b \in \mathbb{B}_p$ is incremented (modulo p) from one configuration to the next one. The value q' records the name of the previous instruction (for technical reasons).

Let $\rho = \gamma_0 \cdot \gamma_1 \cdots$ be a finite or infinite (not necessary valid) run in \mathcal{M} , and write q_i for the state of γ_i . The p -encoding of ρ is then given by:

$$p\text{-enc}(\rho) = \mathbf{enc}_{b_0, \perp}(\gamma_0) \cdot \mathbf{enc}_{b_1, q_0}(\gamma_1) \cdot \mathbf{enc}_{b_2, q_1}(\gamma_2) \cdots \mathbf{enc}_{b_i, q_{i-1}}(\gamma_i) \cdots$$

with $b_j = j \bmod p$ for every j , if ρ is infinite and

$$p\text{-enc}(\rho) = \mathbf{enc}_{b_0, \perp}(\gamma_0) \cdot \mathbf{enc}_{b_1, q_0}(\gamma_1) \cdot \mathbf{enc}_{b_2, q_1}(\gamma_2) \cdots \mathbf{enc}_{b_n, q_{n-1}}(\gamma_n) \#^\omega$$

with $b_j = j \bmod p$ for every j , if ρ is finite. We simply write $\mathbf{enc}(\rho)$ if p is clear in the context. For such a sequence, if $i \in \mathbb{N}$, we write

$$\mathbf{step}_i(\mathbf{enc}(\rho)) = \mathbf{enc}_{b_i, q_{i-1}}(\gamma_i) \cdot \mathbf{enc}_{b_{i+1}, q_i}(\gamma_{i+1}) \cdot [q_{i+2}, (b_{i+2}, q_{i+1})],$$

with the convention that $q_{-1} = \perp$. Roughly this is the encoding of the i -th step of the computation. We also write $\mathbf{from}_i(\mathbf{enc}(\rho))$ for the suffix of $\mathbf{enc}(\rho)$ starting at $\mathbf{enc}_{b_i, q_{i-1}}(\gamma_i)$.

Example 16. Consider the two-counter machine \mathcal{M} in Fig. 4 and $p \geq 2$. Consider the run $\rho = (q_0, 0, 0) \cdot (q_1, 1, 0) \cdot (q_3, 1, 0) \cdot (q_4, 1, 1)$. Then

$$\begin{aligned} p\text{-enc}(\rho) &= [q_0, (b_0, \perp)] \cdot [q_1, (b_1, q_0)] \cdot [(q_1, a_0, \frac{1}{0}), (b_1, q_0)] \cdot \\ &\quad [q_3, (b_2, q_1)] \cdot [(q_3, a_0, \frac{1}{0}), (b_2, q_1)] \cdot \\ &\quad [q_4, (b_3, q_3)] \cdot [(q_4, a_0, \frac{1}{0}), (b_3, q_3)] \cdot [(q_4, a_1, \frac{1}{1}), (b_3, q_3)] \end{aligned}$$

where $b_i = i \bmod p$. ◁

In our reduction, the encoding of (valid or invalid) runs of our two-counter machines are generated by a Kripke structure, and validity is expressed as an \mathbf{avgLTL} formula. We first define the Kripke structure, then the \mathbf{avgLTL} formula, and finally prove the correctness of the reduction.

```

 $q_0$  : inc  $a_0$ ; goto  $q_1$ 
 $q_1$  : if ( $a_1 = 0$ ) goto  $q_3$  else goto  $q_2$ 
 $q_2$  : inc  $a_0$ ; goto  $q_1$ 
 $q_3$  : inc  $a_1$ ; goto  $q_4$ 
 $q_4 = q_{accept}$ 

```




Fig. 4: A two-counter machine \mathcal{M} (textual and graphical representations)

A.3 Definition of the Kripke structure

The Kripke structure $\mathcal{K}_{\mathcal{M}}^p = (V, v_0, E, \ell)$ is defined over the set \mathcal{P}_p of atomic propositions. The set of vertices V is the set of atomic propositions \mathcal{P}_p , and ℓ is the identity function. The initial vertex is $v_0 = [q_0, (0, \perp)]$. The set E of edges is defined as follows (for clarity we forget about quantifications):

$$\begin{aligned}
E = & \{ [q, (b, q'')] \rightarrow [q', (b^{+1}, q)] \mid (q, = 0, = 0) \rightsquigarrow_{\mathcal{M}} q' \} \\
& \cup \{ [q, (b, q'')] \rightarrow [(q, a_0, \frac{1}{0}), (b, q'')] \} \\
& \cup \{ [q, (b, q'')] \rightarrow [(q, a_1, \frac{0}{1}), (b, q'')] \} \\
& \cup \{ [(q, a_0, \frac{1}{0}), (b, q'')] \rightarrow [(q, a_1, \frac{1}{1}), (b, q'')] \} \\
& \cup \{ [(q, a_j, C), (b, q'')] \rightarrow [(q, a_j, C), (b, q'')] \} \\
& \cup \{ [(q, a_1, \frac{0}{1}), (b, q'')] \rightarrow [q', (b^{+1}, q)] \mid (q, = 0, > 0) \rightsquigarrow_{\mathcal{M}} q' \} \\
& \cup \{ [(q, a_0, \frac{0}{0}), (b, q'')] \rightarrow [q', (b^{+1}, q)] \mid (q, > 0, = 0) \rightsquigarrow_{\mathcal{M}} q' \} \\
& \cup \{ [(q, a_1, \frac{1}{1}), (b, q'')] \rightarrow [q', (b^{+1}, q)] \mid (q, > 0, > 0) \rightsquigarrow_{\mathcal{M}} q' \} \\
& \cup \{ [q_{halt}, (b, q'')] \rightarrow \# \} \cup \{ \# \rightarrow \# \}
\end{aligned}$$

where q'' ranges over $Q \cup \{\perp\}$ and $\rightsquigarrow_{\mathcal{M}}$ is defined by $(q, \bowtie_0 0, \bowtie_1 0) \rightsquigarrow_{\mathcal{M}} q'$ when there exists naturals $n_0 \bowtie_0 0$ and $n_1 \bowtie_1 0$ such that $(q, n_0, n_1) \rightarrow_{\mathcal{M}} (q', n'_0, n'_1)$ for some $n'_0, n'_1 \in \mathbb{N}$.

The construction is illustrated on Fig. 5. Note the role of the boolean vector $C \in \mathbb{B}_2^2$, which ensure that the **if** constraint on the transition to the next configuration is satisfied by the current one.

With this construction, it is not difficult to argue that:

Proposition 17 (Paths in $\mathcal{K}_{\mathcal{M}}^p$). *Let ρ be a sequence of configurations of \mathcal{M} . Then ρ is a run of \mathcal{M} iff $p\text{-enc}(\rho)$ is a path in $\mathcal{K}_{\mathcal{M}}^p$.*

A.4 Definition of the formulas

We now build formulas to characterize the path in $\mathcal{K}_{\mathcal{M}}^p$ which encodes the valid run of \mathcal{M} . We first develop formulas and proofs for the equality and non-strict inequality problems. The case of strict inequalities is handled separately.

We use the following two formulas to express the halting property for \mathcal{M} in two different ways:

$$\text{halt}_{\mathcal{M}}^{p, \geq} = \mathbf{F} q_{halt} \wedge \mathbf{G} \text{consec}_{\mathcal{M}}^{p, \geq} \qquad \text{halt}_{\mathcal{M}}^{p, \leq} = \mathbf{G} \neg q_{halt} \vee \mathbf{G} \text{consec}_{\mathcal{M}}^{p, \leq}$$




Fig. 5: A small part of the Kripke structure $\mathcal{K}_{\mathcal{M}}^p$. The constraint “if $(q, \bowtie_0 0, \bowtie_1 0) \rightsquigarrow_{\mathcal{M}} q'$ ” means that there is a move $(q, n_0, n_1) \rightarrow_{\mathcal{M}} (q', n'_0, n'_1)$ in \mathcal{M} , for some naturals $n_0 \bowtie_0 0$ and $n_1 \bowtie_1 0$.

where $\text{consec}_{\mathcal{M}}^{p, \leq}$ and $\text{consec}_{\mathcal{M}}^{p, \geq}$ are formulas that we will now define. For sake of readability we now simply write consec^{\leq} and consec^{\geq} , and even simply consec^{\bowtie} if we want to speak of those two formulas.

Formulas consec^{\bowtie} aim at checking that counter values are properly updated after an instruction. We first give an intuition of how it works, before defining it. Consider a portion of the encoding of a run ρ :

$$\begin{aligned} & \dots [q', (b, q)] \cdot [(q', a_0, C_0), (b, q)]^{n_0} \cdot [(q', a_1, C_1), (b, q)]^{n_1} \cdot \\ & [q'', (b^{+1}, q')] \cdot [(q'', a_0, C'_0), (b^{+1}, q')]^{n'_0} \cdot [(q'', a_1, C'_1), (b^{+1}, q')]^{n'_1} \cdot \\ & [q''', (b^{+2}, q'')] \dots \end{aligned}$$

where instruction q' keeps both counter values unchanged. The formula has to enforce $n'_0 = n_0$ and $n'_1 = n_1$. This is the case if, and only if, for every $\alpha \in \{1 + n_0 + n_1, 1 + n_0 + n'_1, 1 + n'_0 + n_1, 1 + n'_0 + n'_1\}$,

$$\frac{\alpha}{1 + n_0 + n_1 + 1 + n'_0 + n'_1} = \frac{1}{2}.$$

The denominator is the length of the portion from $[q', (b, q)]$ until $[q''', (b^{+2}, q'')]$, whereas the various values for α are the number of positions where some distinguished atomic proposition holds along this portion. For instance, $1 + n'_0 + n_1$ is

the number of positions where formula

$$(q'', (b^{+1}, q')) \vee [(q'', a_0, C'_0), (b^{+1}, q')] \vee [(q', a_1, C_1), (b, q)]$$

holds true. Then computing the quotient will be made thanks to an $\tilde{\mathbf{U}}$ -formula.

We now define the $\mathbf{consec}^{\bowtie}$ formulas. Let us first define the following convenient notations: if $S \subseteq \{0, 1\}$ we set $\bar{S} = \{0, 1\} \setminus S$; furthermore we will use the symbol \star_T as a disjunction on “any possible value” in T . We might omit to mention T when the range is clear from the context. For instance, $[(q, a_s, \star), (b, \star)]$ is a shorthand for $\bigvee_{C \in \mathbb{B}_2^2} \bigvee_{q' \in Q \cup \{\perp\}} [(q, a_s, C), (b, q')]$, and $[q, (b, \star)]$ represents $\bigvee_{q' \in Q \cup \{\perp\}} [q, (b, q')]$. We let $\Psi_b = q_{halt} \vee [\star_Q, (b, \star)]$ characterize the first proposition in the encoding of configurations having bit b .

Then, for every $b \in \mathbb{B}_p$, for every $S \subseteq \{0, 1\}$ and for every $q \in Q$, we define a formula $\Phi_{b,S}^q$, which depends on the type of instruction q , as follows:

- If q is a conditional jump condition (which keeps counter values unchanged), then for every $S \subseteq \{0, 1\}$, we define

$$\Phi_{b,S}^q = \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)].$$

For instance, when $S = \{0\}$, formula $\Phi_{b,S}^q$ catches the n_0 states of the encoding of counter a_0 in all configurations having bit b , as well as the n_1 states encoding counter a_1 in configurations b^{+1} , and the first state of configuration b^{+1} .

- If q increments counter a_r then we define

$$\Phi_{b,S}^q = \begin{cases} [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)] & \text{if } r \in S \\ \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{otherwise} \end{cases}$$

- If q decrements counter a_r then we define

$$\Phi_{b,S}^q = \begin{cases} \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{if } r \in S \\ [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)] & \text{otherwise} \end{cases}$$

Let $\varphi_{b,S}^q$ be the formula³ $[q, (b, \star)] \Rightarrow (\Phi_{b,S}^q \tilde{\mathbf{U}} \Psi_{b+2})$. We then define:

$$\mathbf{consec}^{\geq} = \bigwedge_{q \in Q} \bigwedge_{b \in \mathbb{B}_p} \bigwedge_{S \subseteq \{0,1\}} \varphi_{b,S}^q \quad \mathbf{consec}^{\leq} = \bigwedge_{q \in Q} \bigwedge_{b \in \mathbb{B}_p} \bigvee_{S \subseteq \{0,1\}} \varphi_{b,S}^q.$$

³ $\alpha \Rightarrow \beta$ is defined in the usual way when α is an atomic proposition.

A.5 Correctness of the reduction

In the following series of lemmas, we establish a strong link between valid runs in \mathcal{M} and paths of $\mathcal{K}_{\mathcal{M}}^p$ which have value $1/2$ on formulas $\mathbf{consec}^{\boxtimes}$. For the rest of this section we assume that $p \geq 2$.

We begin with a simple, generic lemma:

Lemma 18. *Let $\varphi \in \mathbf{avgLTL} \setminus \{\mathbf{G}, \tilde{\mathbf{G}}\}$. If π' is a prefix of a path π , then $\llbracket \pi', \varphi \rrbracket \leq \llbracket \pi, \varphi \rrbracket$. Furthermore if π is infinite, then $\llbracket \pi, \varphi \rrbracket = \limsup_{n \rightarrow \infty} \llbracket \pi_{<n}, \varphi \rrbracket$, where $\pi_{<n}$ is the prefix of length n of π .*

Proof. We prove the result by induction on φ . The arguments for atomic propositions and their negations, as well as for boolean operators, are trivial. The case where $\varphi = \mathbf{X}\psi$ is also easy, with the case $|\pi'| = 1$ being handled separately.

When $\varphi = \varphi_1 \mathbf{U} \varphi_2$, we use the induction hypothesis on φ_1 and φ_2 to get

$$\min\{\llbracket \pi'_{\geq i}, \varphi_2 \rrbracket, \min\{\llbracket \pi'_{\geq j}, \varphi_1 \rrbracket \mid 0 \leq j < i\}\} \leq \min\{\llbracket \pi_{\geq i}, \varphi_2 \rrbracket, \min\{\llbracket \pi_{\geq j}, \varphi_2 \rrbracket \mid 0 \leq j < i\}\}$$

for all $0 \leq i \leq |\pi'|$. The result follows. The argument for $\tilde{\mathbf{U}}$ is similar. \square

We fix the following notations for our technical lemmas. We let $\rho = \rho_0 \rho_1 \cdots$ be a finite or infinite run of \mathcal{M} , and $\pi = \mathbf{enc}(\rho)$. We write $\pi = \pi_0 \pi_1 \cdots$ where π_i is the encoding of ρ_i (with bit b_i and state q_{i-1} given by the prefix up to ρ_i). We note

$$\pi_j = [q_j, (b_j, q_{j-1})] \cdot [(q_j, a_0, C_0^j), (b_j, q_{j-1})]^{n_0^j} \cdot [(q_j, a_1, C_1^j), (b_j, q_{j-1})]^{n_1^j}.$$

We define $\eta_i^S = \Phi_{b_i, S}^{q_i} \tilde{\mathbf{U}} \Psi_{b_i^+}$. In particular, $\llbracket \mathbf{from}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket = \llbracket \mathbf{from}_i(\pi), \eta_i^S \rrbracket$ and $\llbracket \mathbf{step}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket = \llbracket \mathbf{step}_i(\pi), \eta_i^S \rrbracket$. We write $\llbracket \mathbf{step}_i(\pi), \eta_i^S \rrbracket = p_i^S / r_i$ ($\mathbf{step}_i(\pi)$ is finite, so we know that the value of η_i^S is a rational), before reduction (*i.e.*, r_i is the size of $\pi_i \pi_{i+1}$). The following results directly follow from the definitions of the formulas.

Lemma 19. *If $(q_i, b_i) \neq (q, b)$, then $\llbracket \mathbf{step}_i(\pi), \varphi_{b, S}^q \rrbracket = \llbracket \mathbf{from}_i(\pi), \varphi_{b, S}^q \rrbracket = 1$. For every $S \subseteq \{0, 1\}$, $\llbracket \mathbf{step}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket + \llbracket \mathbf{step}_i(\pi), \varphi_{b_i, \bar{S}}^{q_i} \rrbracket = 1$.*

Lemma 20. *The following assertions are equivalent:*

1. *the consecution $\rho_i \rho_{i+1}$ is valid;*
2. *for every $S \subseteq \{0, 1\}$ $\llbracket \mathbf{step}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket = 1/2$;*
3. *$\llbracket \mathbf{step}_i(\pi), \mathbf{consec}^{\geq} \rrbracket = 1/2$;*
4. *$\llbracket \mathbf{step}_i(\pi), \mathbf{consec}^{\leq} \rrbracket = 1/2$.*

Proof. Assume $\rho_i \rho_{i+1}$ is a valid consecution. We prove that $\llbracket \mathbf{step}_i(\pi), \eta_i^S \rrbracket = 1/2$ for every $S \subseteq \{0, 1\}$, which implies the other three statements (thanks to the remarks above). We assume q_i increments counter a_0 . The other cases can be handled similarly.

Since $\rho_i \rho_{i+1}$ is valid, we have $n_0^{i+1} = n_0^i + 2$ (remember that \mathcal{M} increments by 2), and $n_1^{i+1} = n_1^i$. Fix $S = \{0\}$ and consider formula η_i^S . Then

$$\begin{aligned} \llbracket \text{step}_i(\pi), \eta_i^S \rrbracket &= (2 + n_0^i + n_1^{i+1}) / (2 + n_0^i + n_1^i + n_0^{i+1} + n_1^{i+1}) \\ &= (2 + n_0^i + n_1^i) / (2 + n_0^i + n_1^i + 2 + n_0^i + n_1^i) \\ &= 1/2. \end{aligned}$$

Similarly we can show that other formulas also evaluate to 1/2 and therefore $\llbracket \text{step}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket = 1/2$ for every $S \subseteq \{0, 1\}$.

Conversely, if $\llbracket \text{step}_i(\pi), \text{consec}^{\geq} \rrbracket = 1/2$, applying the second property of Lemma 19, we get that for every $S \subseteq \{0, 1\}$, $\llbracket \text{step}_i(\pi), \varphi_{b_i, S}^{q_i} \rrbracket = 1/2$. From there we easily get that all the counters are correctly updated. We conclude that $\rho_i \rho_{i+1}$ is a valid consecution. The case when $\llbracket \text{step}_i(\pi), \text{consec}^{\leq} \rrbracket = 1/2$ is similar. \square

Lemma 21. *If ρ is valid, then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\bowtie} \rrbracket = 1/2$, for $\bowtie \in \{\leq, \geq\}$.*

Proof. From Lemmas 20 and 18, for every i , we have $\llbracket \text{from}_i(\pi), \text{consec}^{\bowtie} \rrbracket \geq 1/2$. Fix $S \subseteq \{0, 1\}$, and consider formula η_i^S . Let $J_i = \{j \geq i \mid (q_j, b_j) = (q_i, b_i)\}$. Formula $\varphi_{b_i, S}^{q_i}$ may have value one only on states belonging to some $\pi_j \pi_{j+1}$ with $j \in J_i$; it has value zero anywhere else. As consecution $\rho_j \rho_{j+1}$ is valid, Lemma 20 entails $\llbracket \text{step}_j(\pi), \eta_j^S \rrbracket = \llbracket \text{step}_j(\pi), \eta_i^S \rrbracket = 1/2$. Then we write:

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \sup_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\ &\leq \sup_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k} \\ &= \sup_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} 1/2 \cdot r_k}{\sum_{k \in J_i, i \leq k \leq j} r_k} \\ &= 1/2 \end{aligned}$$

The first equality holds because $\pi_k p_{i_{k+1}}$ and $\pi_{k'} \pi_{k'+1}$ do not overlap for any two $k, k' \in J_i$ with $k \neq k'$ (thanks to bits $b \in \mathbb{B}_p$, assuming $p \geq 2$). We conclude that $\llbracket \text{from}_i(\pi), \text{consec}^{\bowtie} \rrbracket = 1/2$, and finally $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\bowtie} \rrbracket = 1/2$. \square

The following lemma deals with invalid finite runs.

Lemma 22. *Assume ρ is finite and invalid. Then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket < 1/2$ and $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\leq} \rrbracket > 1/2$.*

Proof. Let $\rho_i \rho_{i+1}$ be the last invalid consecution along ρ . We first focus on formula consec^{\geq} . From Lemma 20, $\llbracket \text{step}_i(\pi), \text{consec} \rrbracket \neq 1/2$. We also know from Lemma 19 that for every $S \subseteq \{0, 1\}$,

$$\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket + \llbracket \text{from}_i(\pi), \eta_i^{\bar{S}} \rrbracket = 1$$

So for some $S \subseteq \{0, 1\}$, $\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket < 1/2$. We pick such a set S . Reusing the notations of the previous proof, it must be the case that for every $j \in J_i \setminus \{i\}$ (if any), $p_j^S/r_j = 1/2$ (by choice of i), and there is a positive integer e such that $(p_i^S + e)/r_i = 1/2$. We get:

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \max_{j \in J_i} \frac{\sum_{k \in J_i, i < k \leq j} p_k^S}{\sum_{k \in J_i, i < k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\ &\leq \max_{j \in J_i} \left(\frac{(p_i^S + e) + \sum_{k \in J_i, i < k \leq j} p_k^S}{r_i + \sum_{k \in J_i, i < k \leq j} r_k} \right. \\ &\quad \left. - \frac{e}{r_i + \sum_{k \in J_i, i < k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \\ &= 1/2 - \frac{e}{|\text{from}_i(\pi)|} \end{aligned}$$

We conclude that $\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket < 1/2$.

The proof for formula consec^{\leq} is dual. We pick S such that $\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket > 1/2$ and write $e > 0$ such that $\frac{p_i^S - e}{r_i} = 1/2$. We then get:

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \max_{j \in J_i} \frac{\sum_{k \in J_i, i < k \leq j} p_k^S}{\sum_{k \in J_i, i < k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\ &\geq \frac{(p_i^S - e)}{r_i} + \frac{e}{r_i} \quad (\text{taking } j = i) \end{aligned}$$

We conclude that $\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket > 1/2$. \square

As a corollary we get the following correspondence between value-1/2 assigned to formula $\text{halt}_{\mathcal{M}}^{\times}$ and valid halting runs in \mathcal{M} .

Corollary 23. *Let ρ be a maximal run of \mathcal{M} . If $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{\times} \rrbracket = 1/2$, then ρ is a finite valid run ending in the halting state, witnessing the fact that \mathcal{M} is halting.*

Proof. We begin with $\text{halt}_{\mathcal{M}}^{\geq} = \mathbf{F} q_{\text{halt}} \wedge \mathbf{G} \text{consec}^{\geq}$. If $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{\geq} \rrbracket = 1/2$, then ρ is finite and ends in the halting state (otherwise $\llbracket \text{enc}(\rho), \mathbf{F} q_{\text{halt}} \rrbracket = 0$, and the conjunction with $\mathbf{G} \text{consec}^{\geq}$ would also be zero) and $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{\geq} \rrbracket = \llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket$. Assume that ρ is invalid; applying Lemma 22, we get that $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket < 1/2$ (or $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\leq} \rrbracket > 1/2$), which is a contradiction. Therefore ρ is valid.

Now consider $\text{halt}_{\mathcal{M}}^{\leq} = \mathbf{G} \neg q_{\text{halt}} \vee \mathbf{G} \text{consec}^{\geq}$. Again, $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{\leq} \rrbracket = 1/2$ entails that ρ reaches the halting state (otherwise $\llbracket \text{enc}(\rho), \mathbf{G} \neg q_{\text{halt}} \rrbracket = 1$, and the disjunction with $\mathbf{G} \text{consec}^{\leq}$ would also be one). Hence $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{\leq} \rrbracket = \llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\leq} \rrbracket$. If this has value 1/2, then by Lemma 22, ρ is valid. \square

We can notice here that it is important to consider halting runs in the above proof. Indeed, if ρ is infinite, it may be the case that $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket = 1/2$

whereas ρ is non-valid (a small mistake at the beginning of the computation might be compensated later along the run).

Gathering all these results, we obtain the following classification of runs in \mathcal{M} (or equivalently paths in $\mathcal{K}_{\mathcal{M}}^p$), from which Corollary 24 and Theorem 6 (for $\bowtie \in \{\leq, =, \geq\}$) follow.

Classification 2 (w.r.t. formula $\text{halt}_{\mathcal{M}}^{p;\geq}$ and $\text{halt}_{\mathcal{M}}^{p;\leq}$). Fix $p \geq 2$. Let ρ be a maximal run in \mathcal{M} .

- if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\geq} \rrbracket = 0$ and $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\leq} \rrbracket = 1$;
- if ρ is finite and valid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\geq} \rrbracket = \llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\leq} \rrbracket = 1/2$;
- otherwise, $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\geq} \rrbracket < 1/2$ and $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\leq} \rrbracket > 1/2$.

Corollary 24. Fix $p \geq 2$. Given $\bowtie \in \{\leq, \geq\}$, the following four statements are equivalent:

1. \mathcal{M} halts;
2. the unique initial and maximal valid run $\rho_{\mathcal{M}}$ of \mathcal{M} is such that $\llbracket \text{enc}(\rho_{\mathcal{M}}), \text{halt}_{\mathcal{M}}^{p;\bowtie} \rrbracket = 1/2$;
3. there exists an initial maximal run ρ in \mathcal{M} such that $\llbracket \text{enc}(\rho), \text{halt}_{\mathcal{M}}^{p;\bowtie} \rrbracket = 1/2$;
4. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^{p;\bowtie} \rrbracket = 1/2$;
5. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^{p;\geq} \rrbracket \geq 1/2$;
6. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^{p;\leq} \rrbracket \leq 1/2$.

A.6 Handling strict inequalities

Here we explain how we extend the above proof to handle strict inequalities.

Case $\llbracket \pi, \varphi \rrbracket > 1/2$. We use the same encoding as above, but bits b now range over \mathbb{B}_p with $p \geq 6$. We twist formulas $\Phi_{b,S}^q$ to ensure that, reusing notations of Section A.5:

$$(i) \frac{p_i^S - 1}{r_i} + \frac{p_i^{\bar{S}} - 1}{r_i} = 1;$$

$$(ii) \text{consecution } \rho_i \rho_{i+1} \text{ is valid if, and only if, for every } S \subseteq \{0, 1\}, \frac{p_i^S - 1}{r_i} = \frac{1}{2}.$$

The formulas are defined as follows. If q is a conditional jump:

$$\Phi_{b,S}^q = [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)]$$

If q increments counter a_r :

$$\Phi_{b,S}^q = \begin{cases} \left[[q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)] \right. \\ \quad \left. \vee \left([(\star, a_r, \star), (b^{+1}, q)] \wedge \mathbf{X} \neg [(\star, a_r, \star), (b^{+1}, q)] \right) \right] & \text{if } r \in S \\ \left[[q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \right] & \text{otherwise} \end{cases}$$

If q decrements counter a_r :

$$\Phi_{b,S}^q = \begin{cases} [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{if } r \in S \\ [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] \vee [\star_Q, (b^{+1}, q)] \\ \vee \left([(q, a_r, \star), (b, \star)] \wedge \mathbf{X} \neg [(q, a_r, \star), (b, \star)] \right) & \text{otherwise} \end{cases}$$

One easily checks that these formulas fulfill conditions (i) and (ii) above. We then let $\varphi_{b,S}^q$ be the formula $[q, (s, \star)] \Rightarrow (\Phi_{b,S}^q \tilde{\mathbf{U}} \Psi_{b+2})$, and

$$\mathbf{consec}_{\mathcal{M}}^{p,>} = \bigwedge_{q \in Q} \bigwedge_{b \in \mathbb{B}_p} \bigwedge_{S \subseteq \{0,1\}} \varphi_{b,S}^q \quad \mathbf{halt}_{\mathcal{M}}^{p,>} = \mathbf{F} q_{\mathbf{halt}} \wedge \mathbf{G} \mathbf{consec}_{\mathcal{M}}^{p,>}$$

Later we will write simply $\mathbf{consec}^>$.

We now prove correctness of this new reduction. Lemma 21 extends as follows:

Lemma 25. *If ρ is valid and counters are bounded along ρ (which includes the case when ρ is finite), then $\llbracket \mathbf{enc}(\rho), \mathbf{G} \mathbf{consec}^> \rrbracket > 1/2$.*

Proof. We show that $\llbracket \mathbf{from}_i(\pi), \eta_i^S \rrbracket = 1/2 + 1/r_i$. Reusing the notations of the proof of Lemma 21, we let

$$\beta_j = \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}$$

so that $\llbracket \mathbf{from}_i(\pi), \eta_i^S \rrbracket = \sup_{j \in J_i} \beta_j$. Pick $j \in J_i \setminus \{i\}$. Then:

$$\beta_j = \frac{\sum_{k \in J_i, i \leq k \leq j} (p_k^S - 1) + \sum_{k \in J_i, i \leq k \leq j} 1}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}$$

Because bits b range over \mathbb{B}_p with $p \geq 6$, it holds that

$$2 \sum_{k \in J_i, i \leq k \leq j} 1 \leq \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|$$

We deduce that $\beta_j \leq 1/2$ for all $j \in J \setminus \{i\}$. Now, $\beta_i = p_i^S / r_i = 1/2 + 1/r_i > 1/2$, so that $\llbracket \mathbf{from}_i(\pi), \eta_i^S \rrbracket = 1/2 + 1/r_i$. If counters are bounded along ρ , then r_i is bounded, hence the expected result. \square

Remark 26. Note that in the above proof, the value of $\llbracket \mathbf{enc}(\rho), \mathbf{G} \mathbf{consec}^> \rrbracket$ can be made as close as we want to $1/2$ by simply increasing a bit the counters before halting (we can add finitely many transitions before really going to $q_{\mathbf{halt}}$ that increase one of the counters – so that value r_n (n is the length of the execution) is made as big as we want, hence $1/2 + 1/r_n$ as close to $1/2$ as we want). We write $\mathcal{K}_{\mathcal{M},\delta}^p$ for the Kripke structure which corresponds to counter machine \mathcal{M} and which ensures that counters are large enough when reaching the halting state, so that $1/r_n < \delta$

We now analyse invalid finite runs, rephrasing Lemma 22 as follows:

Lemma 27. *Assume ρ is finite and invalid. Then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec} \rrbracket \leq 1/2$.*

Proof. First notice that if $\rho_i \rho_{i+1}$ is the last invalid consecution along ρ , then there is a set S such that $\frac{p_i^S - 1}{r_i} \neq \frac{1}{2}$. Using condition (i), there is a set S such that $\frac{p_i^S - 1}{r_i} < \frac{1}{2}$. Let $e \geq 1$ be such that $\frac{p_i^S - 1 + e}{r_i} = \frac{1}{2}$. By choice of i , for any $k \in J_i \setminus \{i\}$, it holds $\frac{p_k^S - 1}{r_k} = \frac{1}{2}$. Then:

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \max_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\ &= \max_{j \in J_i} \left(\frac{(p_i^S - 1 + e) + \sum_{k \in J_i, i < k \leq j} (p_k^S - 1) + \sum_{k \in J_i, i < k \leq j} 1}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} + \right. \\ &\quad \left. \frac{1 - e}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \end{aligned}$$

Now, notice that

$$2 \sum_{k \in J_i, i < k \leq j} 1 \leq \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|$$

because bits b range over \mathbb{B}_p with $p \geq 6$ (note that here, $p \geq 4$ would be sufficient). Therefore:

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &\leq \max_{j \in J_i} \left(1/2 + \frac{1 - e}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \\ &\leq 1/2 \quad (\text{since } e \geq 1) \end{aligned}$$

In particular, $\llbracket \text{from}_i(\pi), \text{consec} \rrbracket \leq 1/2$. \square

We conclude with the following classification of runs in \mathcal{M} , from which Corollary 28 and Theorem 6 (for $>$) follow.

Classification 3 (w.r.t. formula $\text{halt}_{\mathcal{M}}^{p, >}$). *Fix $p \geq 2$. Let ρ be a maximal run in \mathcal{M} .*

- if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket = 0$;
- if ρ is finite and valid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket > 1/2$;
- if ρ is finite and invalid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket \leq 1/2$.

Note that thanks to Remark 26, the above classification can be refined, and actually for every $\delta > 0$, we can build a Kripke structure $\mathcal{K}_{\mathcal{M}, \delta}^p$ such that:

- if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket = 0$;
- if ρ is finite and valid, then $1/2 < \llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket < 1/2 + \delta$;
- if ρ is finite and invalid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket \leq 1/2$.

Corollary 28. Fix $p \geq 2$. The following four statements are equivalent:

1. \mathcal{M} halts;
2. the unique initial and maximal valid run $\rho_{\mathcal{M}}$ of \mathcal{M} is such that $\llbracket p\text{-enc}(\rho_{\mathcal{M}}), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket > 1/2$;
3. there exists an initial maximal run ρ in \mathcal{M} such that $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, >} \rrbracket > 1/2$;
4. there exists an initial path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^{p, >} \rrbracket > 1/2$.

Case $\llbracket \pi, \varphi \rrbracket < 1/2$. We use the same encoding as before, but bits range over \mathbb{B}_p with $p \geq 4$. We again modify formulas $\Phi_{b,S}^q$ to ensure that:

- (i) $\frac{p_i^S+1}{r_i} + \frac{p_i^{\bar{S}}+1}{r_i} = 1$;
- (ii) if consecution $\rho_i \rho_{i+1}$ is valid, then for every $S \subseteq \{0, 1\}$, $\frac{p_i^S+1}{r_i} = 1/2$;

If q is a conditional jump:

$$\Phi_{b,S}^q = \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)]$$

If q increments counter a_r :

$$\Phi_{b,S}^q = \begin{cases} [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{if } r \in S \\ \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} \left([(\star, a_s, \star), (b^{+1}, q)] \wedge \mathbf{X} [(\star, a_s, \star), (b^{+1}, q)] \right) & \text{otherwise} \end{cases}$$

If q decrements counter a_r :

$$\Phi_{b,S}^q = \begin{cases} \bigvee_{s \in S} \left([(q, a_s, \star), (b, \star)] \wedge \mathbf{X} [(q, a_s, \star), (b, \star)] \right) \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{if } r \in S \\ [q, (b, \star)] \vee \bigvee_{s \in S} [(q, a_s, \star), (b, \star)] \vee \bigvee_{s \in \bar{S}} [(\star, a_s, \star), (b^{+1}, q)] & \text{otherwise} \end{cases}$$

One easily checks that conditions (i) and (ii) above are fulfilled. We then let $\varphi_{b,S}^q$ be the formula $[q, (s, \star)] \Rightarrow (\Phi_{b,S}^q \tilde{\mathbf{U}} \Psi_{b+2})$, and

$$\text{consec}_{\mathcal{M}}^{p, >} = \bigwedge_{q \in Q} \bigwedge_{b \in \mathbb{B}_p} \bigvee_{S \subseteq \{0, 1\}} \varphi_{b,S}^q \quad \text{halt}_{\mathcal{M}}^{p, >} = \mathbf{G} \neg q_{\text{halt}} \vee \mathbf{G} \text{consec}_{\mathcal{M}}^{p, >}$$

We will simply write $\text{consec}^{>}$ now.

Lemma 21 then writes as follows:

Lemma 29. If ρ is finite and valid, then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{<} \rrbracket < 1/2$.

Proof. We take the same notations as in the proof of Lemma 21. Then:

$$\begin{aligned}
\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \max_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\
&= \max_{j \in J_i} \left(\frac{\sum_{k \in J_i, i \leq k \leq j} (p_k^S + 1) + 1/2 \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right. \\
&\quad \left. - \frac{\sum_{k \in J_i, i \leq k \leq j} 1 + 1/2 \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \\
&= \max_{j \in J_i} \left(1/2 - \frac{\sum_{k \in J_i, i \leq k \leq j} 1 + 1/2 \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \\
&= 1/2 - \min_{j \in J_i} \left(\frac{\sum_{k \in J_i, i \leq k \leq j} 1 + 1/2 \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \right) \\
&< 1/2
\end{aligned}$$

Since ρ is finite, we conclude that $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^< \rrbracket < 1/2$. \square

The case of invalid finite runs is a bit more difficult to handle (and this is where we use the fact that $b \in \mathbb{B}_p$ with $p \geq 4$). Lemma 22 rephrases as:

Lemma 30. *If ρ is invalid, then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^< \rrbracket \geq 1/2$.*

Proof. First assume that ρ is finite (the case where ρ is infinite follows by considering its finite prefixes). We notice that if $\rho_i \rho_{i+1}$ is the last invalid consecution along ρ , then there is a set S such that $\frac{p_i^S + 1}{r_i} \neq \frac{1}{2}$. Using property (i), there is a set S such that $\frac{p_i^S + 1}{r_i} > \frac{1}{2}$. Let $e \geq 1$ be such that $\frac{p_i^S + 1 - e}{r_i} = \frac{1}{2}$. For $k \in J_i \setminus \{i\}$, it holds that $\frac{p_k^S + 1}{r_k} = \frac{1}{2}$.

Then we can compute:

$$\begin{aligned}
\llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \max_{j \in J_i} \frac{\sum_{k \in J_i, i \leq k \leq j} p_k^S}{\sum_{k \in J_i, i \leq k \leq j} r_k + \sum_{i < h < j, h \notin J_i, h-1 \notin J_i} |\pi_h|} \\
&\geq \left(\frac{(p_i^S + 1 - e)}{r_i} + \frac{e - 1}{r_i} \right) \quad (\text{taking } j = i) \\
&\geq 1/2
\end{aligned}$$

In particular, $\llbracket \text{from}_i(\pi), \text{consec}^< \rrbracket \geq 1/2$. \square

We end up with the following classification of runs in \mathcal{M} , which entails Corollary 31 and conclude the proof of Theorem 6.

Classification 4 (w.r.t. formula $\text{halt}_{\mathcal{M}}^{p, <}$). *Fix $p \geq 2$. Let ρ be a maximal run in \mathcal{M} .*

– if ρ is infinite, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, <} \rrbracket = 1$;

- if ρ is finite and valid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, <} \rrbracket < 1/2$;
- if ρ is finite and invalid, then $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, <} \rrbracket \geq 1/2$.

Corollary 31. Fix $p \geq 2$. The following four statements are equivalent:

1. \mathcal{M} halts;
2. the unique initial valid and maximal run $\rho_{\mathcal{M}}$ of \mathcal{M} is such that $\llbracket p\text{-enc}(\rho_{\mathcal{M}}), \text{halt}_{\mathcal{M}}^{p, <} \rrbracket < 1/2$;
3. there exists an initial maximal run ρ in \mathcal{M} such that $\llbracket p\text{-enc}(\rho), \text{halt}_{\mathcal{M}}^{p, <} \rrbracket < 1/2$;
4. there exists an initial maximal path π in $\mathcal{K}_{\mathcal{M}}^p$ such that $\llbracket \pi, \text{halt}_{\mathcal{M}}^{p, <} \rrbracket < 1/2$.

B Proof of Theorem 7

We now turn to the proof of Theorem 7. We first rule out the easy case of deciding whether $\llbracket \mathcal{K}, \varphi \rrbracket > 1/2$ (and dually, $\llbracket \mathcal{K}, \varphi \rrbracket \leq 1/2$). Indeed, since $\llbracket \mathcal{K}, \varphi \rrbracket$ is the supremum of the value of φ over all initial infinite paths in \mathcal{K} , it is strictly larger than $1/2$ if, and only if, there is a path whose value is strictly more than $1/2$. This we proved undecidable in the previous section.

We now turn to the more interesting cases of $=$ (the result for \geq and $<$ directly follows, as we explain at the end of this proof). We cannot do a direct proof as previously, since we cannot distinguish between counter machines that have a halting computation (whose encoding has value $1/2$ against a given formula similar to $\text{halt}_{\mathcal{M}}^{p, \geq}$ above) and counter machines that have sequences of computations whose encodings have values *converging* to $1/2$.

Example 32. We consider the deterministic two-counter machine \mathcal{M} of Fig. 6, having q_0 as its initial state. The unique maximal valid run of \mathcal{M} is infinite (it loops in $q_1 \rightleftharpoons q_2$). A single error can make the transition from q_1 to q_3 available, from which valid consecutions lead to q_{halt} . The *weight* of this error can be arbitrarily small, as it can occur with an arbitrarily large value of a_0 . It is not difficult to check that $\llbracket \mathcal{K}_{\mathcal{M}}, \text{halt}_{\mathcal{M}}^{p, \geq} \rrbracket = 1/2$. \triangleleft




Fig. 6: A non-halting two-counter machine for which $\llbracket \mathcal{K}_{\mathcal{M}}, \text{halt}_{\mathcal{M}}^{p, \geq} \rrbracket = 1/2$

Analysis of a non-halting two-counter machine. We consider a two-counter deterministic accept/reject two-counter machine \mathcal{M} : such machines have two halting states, now named q_{accept} and q_{reject} . Their computations may still be infinite. We consider formula $\text{consec}_{\mathcal{M}}^{p;\geq}$, and formula

$$\text{accept}_{\mathcal{M}}^{p;\geq} = \mathbf{F} q_{accept} \wedge \mathbf{G} \text{consec}_{\mathcal{M}}^{p;\geq}$$

For sake of readability we write next consec^{\geq} and accept^{\geq} .

We reuse notations of the previous proofs, and rely on the same encoding as before, with bits b now ranging over \mathbb{B}_p with $p \geq 927$. This bound appears in the long and technical proof of the following lemma, which we postpone to Appendix C below. The lemma evaluates the impact of an error at one step of the computation on the value of formula consec^{\geq} .

Lemma 33. *Let ρ be a run of \mathcal{M} , and $\pi = \text{enc}(\rho)$. Fix $n \geq 30$, $p \geq 927$, and assume that: (i) $\llbracket \text{step}_i(\pi), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$ and (ii) for every $k > i$, $\llbracket \text{step}_k(\pi), \text{consec}^{\geq} \rrbracket \geq 1/2 - 1/n$. Then $\llbracket \text{from}_i(\pi), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$.*

As a consequence, we show the impact of the error at the first invalid consecution on the value of $\mathbf{G} \text{consec}^{\geq}$ along the whole execution:

Lemma 34. *Fix $p \geq 927$. Let ρ be a finite invalid run of \mathcal{M} . Assume $\rho_i \rho_{i+1}$ is the first invalid consecution along ρ . Pick $n \geq 30$ such that $\llbracket \text{step}_i(\text{enc}(\rho)), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$. Then $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$.*

Proof. If $\llbracket \text{from}_i(\text{enc}(\rho)), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$, then we are done. Assume then that $\llbracket \text{from}_i(\text{enc}(\rho)), \text{consec}^{\geq} \rrbracket > 1/2 - 1/n$. Then applying Lemma 33, there exists some $k > i$ such that $\llbracket \text{step}_k(\text{enc}(\rho)), \text{consec}^{\geq} \rrbracket < 1/2 - 1/n$.

We apply the argument iteratively; since ρ is finite, we must eventually reach an index h such that $\llbracket \text{from}_h(\text{enc}(\rho)), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$. We conclude that $\llbracket \text{enc}(\rho), \mathbf{G} \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$. \square

Lemma 35. *Fix $p \geq 927$. Assume that $\llbracket \mathcal{K}_{\mathcal{M}}^p, \text{accept}^{\geq} \rrbracket = 1/2$ but that there is no run ρ of \mathcal{M} with $\llbracket \text{enc}(\rho), \text{accept}^{\geq} \rrbracket = 1/2$. Then the unique valid run $\rho_{\mathcal{M}}$ of \mathcal{M} is infinite.*

Proof. Let ρ the maximal valid run of $\mathcal{K}_{\mathcal{M}}^p$. Let $(\rho^n)_{n \in \mathbb{N}}$ be a sequence of maximal runs such that $\llbracket \text{enc}(\rho^n), \text{accept}^{\geq} \rrbracket > 1/2 - 1/n$ (such a sequence exists by hypothesis). Pick $n \geq 30$, and let $\rho_{i_n}^n \rho_{i_n+1}^n$ be the first invalid consecution of ρ^n . Applying Lemma 34, we get that $\llbracket \text{step}_{i_n}(\text{enc}(\rho^n)), \text{consec}^{\geq} \rrbracket > 1/2 - 1/n$. Since $\text{step}_{i_n}(\text{enc}(\rho^n))$ is an invalid consecution, we also have $\llbracket \text{step}_{i_n}(\text{enc}(\rho^n)), \text{consec}^{\geq} \rrbracket < 1/2$. It follows that $1/(\llbracket \text{step}_{i_n}(\text{enc}(\rho^n)) \rrbracket - 1) < 1/n$, which implies that $|\text{step}_{i_n}(\text{enc}(\rho^n))| > n$. Now, the prefix of ρ of size i_n coincides with that of ρ^n , since $\rho_{i_n}^n \rho_{i_n+1}^n$ is the first invalid consecution. We conclude that ρ contains configurations of arbitrarily large size, so that the sum of the two counters is unbounded along ρ . Hence ρ is infinite. \square

A diagonal argument. Any deterministic Turing machine can be simulated by a deterministic two-counter machine [20]. In particular, given a deterministic Turing machine B , we can build a deterministic two-counter machine $\mathcal{M}(B)$ whose computation mimics the run of B on input B . In particular, $\mathcal{M}(B)$ accepts (resp. rejects, does not halt) if, and only if, B accepts (resp. rejects, does not halt on) input B .

We define the following function \mathcal{H} , which takes as input a deterministic Turing machine B :

$$\mathcal{H}(B) = \begin{cases} \text{accept} & \text{if } \llbracket \mathcal{K}_{\mathcal{M}(B)}^p, \text{accept}^\geq \rrbracket = 1/2 \\ \text{reject} & \text{otherwise} \end{cases}$$

where $\mathcal{M}(B)$ is the deterministic two-counter machine simulating B on input string B .

Proposition 36. *The function \mathcal{H} is not computable.*

Proof. Towards a contradiction, assume \mathcal{H} is computable. Let $\mathcal{T}_{\mathcal{H}}$ be a deterministic Turing machine that computes \mathcal{H} . Notice in particular that $\mathcal{T}_{\mathcal{H}}$ halts on all its inputs; we assume that it ends in its state $q_{\text{accept}}^{\mathcal{T}}$ when \mathcal{H} accepts the input, and in $q_{\text{reject}}^{\mathcal{T}}$ when \mathcal{H} returns *reject*.

We now define the following deterministic Turing machine \mathcal{C} , which takes as input a deterministic Turing machine B :

$\mathcal{C}(B)$: Simulate $\mathcal{T}_{\mathcal{H}}$ on B ;
 If the simulation ends in $q_{\text{accept}}^{\mathcal{T}}$ then goto $q_{\text{reject}}^{\mathcal{C}}$ otherwise goto $q_{\text{accept}}^{\mathcal{C}}$.

The Turing machine \mathcal{C} terminates on all its inputs, since so does $\mathcal{T}_{\mathcal{H}}$; also, \mathcal{C} is deterministic, and we can therefore run \mathcal{C} on input \mathcal{C} itself.

Assume \mathcal{C} accepts input \mathcal{C} . This means that $\mathcal{H}(\mathcal{C})$ rejects, which means that $\llbracket \mathcal{K}_{\mathcal{M}(\mathcal{C})}, \text{accept}^\geq \rrbracket < 1/2$. This means that $\mathcal{M}(\mathcal{C})$ does not accept (by a straightforward extension of Corollary 24 to accept/reject two-counter machines), and therefore \mathcal{C} does not accept \mathcal{C} , contradicting our hypothesis.

Hence \mathcal{C} rejects input \mathcal{C} , which means that $\llbracket \mathcal{K}_{\mathcal{M}(\mathcal{C})}, \text{accept}^\geq \rrbracket = 1/2$. However, since \mathcal{C} does not accept \mathcal{C} , the unique valid maximal run of $\mathcal{M}_{\mathcal{C}}$ is either infinite or rejecting. Applying Lemma 35 to $\mathcal{M}_{\mathcal{C}}$, we get that it is actually infinite. This means that the simulation of $\mathcal{T}_{\mathcal{H}}$ on input \mathcal{C} does not terminate. This contradicts the fact that $\mathcal{T}_{\mathcal{H}}$ terminates on all its inputs. Therefore \mathcal{H} is not computable. \square

Corollary 37. *Given a Kripke structure \mathcal{K} and a formula φ , we cannot decide whether $\llbracket \mathcal{K}, \varphi \rrbracket = 1/2$ (resp. $\llbracket \mathcal{K}, \varphi \rrbracket \geq 1/2$, $\llbracket \mathcal{K}, \varphi \rrbracket < 1/2$).*

C Proof of Lemma 33

We detail the proof of Lemma 33, explaining why we impose $n \geq 30$ and $p \geq 927$ (remember that p is the number of bits used in the encoding of the successive configurations along a run of a deterministic two-counter machine).

Lemma 33. *Let ρ be a run of \mathcal{M} , and $\pi = \text{enc}(\rho)$. Fix $n \geq 30$, $p \geq 927$, and assume that: (i) $\llbracket \text{step}_i(\pi), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$ and (ii) for every $k > i$, $\llbracket \text{step}_k(\pi), \text{consec}^{\geq} \rrbracket \geq 1/2 - 1/n$. Then $\llbracket \text{from}_i(\pi), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$.*

The proof of this lemma uses the result of Lemma 38 below. Before stating and proving Lemma 38, we first explain how it entails the result of Lemma 33: let $J_i \subseteq \mathbb{N}$ be the set $\{j \geq i \mid (q_j, b_j) = (q_i, b_i)\}$, and enumerate its items as $J_i = \{j_0 < j_1 < j_2 < \dots\}$. Obviously $j_0 = i$. Pick $S \subseteq \{0, 1\}$ such that $\llbracket \text{step}_i(\pi), \eta_i^S \rrbracket \leq 1/2 - 1/n$ (which exists because $\llbracket \text{step}_i(\pi), \text{consec}^{\geq} \rrbracket \leq 1/2 - 1/n$). Write

$$\frac{a}{b} = \llbracket \text{step}_i(\pi), \eta_i^S \rrbracket \text{ and } \frac{a_h}{b_h} = \llbracket \pi_{j_{h-1}+3}\pi_{j_{h-1}+4} \dots \pi_{j_h-1} \text{step}_{j_h}(\pi), \eta_i^S \rrbracket$$

with $b = |\text{step}_i(\pi)| - 1$ and $b_h = |\pi_{j_{h-1}+3}\pi_{j_{h-1}+4} \dots \pi_{j_h-1} \text{step}_{j_h}(\pi)| - 1$ (i.e., without simplifying fractions). Then

$$\begin{aligned} \llbracket \text{from}_i(\pi), \eta_i^S \rrbracket &= \sup_h \llbracket \pi_i \pi_{i+1} \pi_{i+2} \dots \pi_{j_h-1} \text{step}_{j_h}(\pi), \eta_i^S \rrbracket \\ &\leq \sup_h \frac{a + a_1 + \dots + a_h}{b + b_1 + \dots + b_h} \end{aligned}$$

The inequality holds because $a + a_1 + \dots + a_h$ counts all positions where $\Phi_{b_i, S}^{q_i}$ holds along $\pi_i \pi_{i+1} \pi_{i+2} \dots \pi_{j_h-1} \text{step}_{j_h}(\pi)$, while $b + b_1 + \dots + b_h$ only counts the positions in $\pi_{j_{l-1}+3}\pi_{j_{l-1}+4} \dots \pi_{j_l-1} \text{step}_{j_l}(\pi)$ for some $l \leq h$.

Remark now that, by Lemma 38, $a_h/b_h \leq 1/2 - 1/n$ for every h . By hypothesis, $a/b \leq 1/2 - 1/n$. Hence for every h :

$$\frac{a + a_1 + \dots + a_h}{b + b_1 + \dots + b_h} \leq \frac{1}{2} - \frac{1}{n}$$

which implies the expected result.

Lemma 38. *Pick $p \geq 927$, and $n \geq 30$. Let ρ be a run of \mathcal{M} , and $\pi = \text{enc}(\rho)$ (with bits ranging over \mathbb{B}_p). Let j be the least index larger than i such that $(q_i, b_i) = (q_j, b_j)$ along π . Assume that for every $i < k \leq j$, $\llbracket \text{step}_k(\pi), \text{consec}^{\geq} \rrbracket \geq 1/2 - 1/n$. Then for every $S \subseteq \{0, 1\}$, it holds $\llbracket \pi_{i+1}\pi_{i+2} \dots \pi_{j-1} \text{step}_j(\pi), \eta_i^S \rrbracket \leq 1/2 - 1/n$.*

Proof. Define the set

$$\begin{aligned} P_n &= \{\tilde{\pi} = \text{enc}(\tilde{\rho}) \mid \tilde{\pi}_{j+1} = \pi_{j+1}, \text{states}(\tilde{\pi}) = \text{states}(\pi) \\ &\text{ and } \llbracket \text{step}_k(\tilde{\pi}), \text{consec}^{\geq} \rrbracket \geq 1/2 - 1/n \text{ for every } i < k \leq j\} \end{aligned}$$

where $\text{states}(\pi)$ denotes the sequence of states of the two-counter machine visited along π (forgetting about counter values). We fix $S \subseteq \{0, 1\}$ and define

$$\tau(n) = \max_{S \subseteq \{0, 1\}} \sup_{\tilde{\pi} \in P_n} \llbracket \tilde{\pi}_{i+1}\tilde{\pi}_{i+2} \dots \tilde{\pi}_{j-1} \text{step}_j(\tilde{\pi}), \eta_i^S \rrbracket$$

Obviously, $\tau(n) \geq \tau(n+1)$ since $P_{n+1} \subseteq P_n$. We prove that $\tau(n) \leq 1/2 - 1/n$ (whenever $n \geq 30$). To prove this, it suffices to show that $\tau(30) \leq 1/2 - 1/30$, since τ is decreasing while $1/2 - 1/n$ is increasing. In the sequel, we prove this result, but keep symbol n to make the computations easier to follow.

Lemma 39. *For any $\tilde{\pi} \in P_n$ and any $i < k \leq j$, it holds*

$$\frac{n-2}{n+2} \cdot |\tilde{\pi}_k| - \frac{4n}{n+2} \leq |\tilde{\pi}_{k+1}| \leq \frac{n+2}{n-2} \cdot |\tilde{\pi}_k| + \frac{4n}{n-2}$$

Proof. This is due to the fact that $\llbracket \text{step}_k(\tilde{\pi}), \text{consec}^{\geq} \rrbracket \geq 1/2 - 1/n$. Consider $S = \emptyset$, and assume that state q_k decrements a counter (which can be proved to give the weakest constraint). Then the value of $\Phi_{b_k, S}^{q_k}$ on $\text{step}_k(\tilde{\pi})$ is $(|\tilde{\pi}_{k+1}| + 2)/(|\tilde{\pi}_{k+1}| + |\tilde{\pi}_k|)$. It follows

$$\frac{n-2}{n+2} |\tilde{\pi}_k| - \frac{4n}{n+2} \leq |\tilde{\pi}_{k+1}|$$

The other inequality is obtained by considering $S = \{0, 1\}$ and the case where q_k increments a counter. \square

Now, we fix $\tilde{\pi} \in P_n$. As in Section A.5, we write $\llbracket \text{step}_j(\tilde{\pi}), \eta_i^S \rrbracket = p_j^S / r_j$, with $r_j = |\tilde{\pi}_j| + |\tilde{\pi}_{j+1}|$. We give a first upper bound on $\llbracket \tilde{\pi}_{i+1} \tilde{\pi}_{i+2} \dots \tilde{\pi}_{j-1} \text{step}_j(\tilde{\pi}), \eta_i^S \rrbracket$:

$$\begin{aligned} \llbracket \tilde{\pi}_{i+1} \tilde{\pi}_{i+2} \dots \tilde{\pi}_{j-1} \text{step}_j(\tilde{\pi}), \eta_i^S \rrbracket &= \frac{p_j^S}{\sum_{k=i+1}^{j+1} |\tilde{\pi}_k|} \\ &\leq \frac{|\tilde{\pi}_j| + |\tilde{\pi}_{j+1}|}{\sum_{k=i+1}^{j+1} |\tilde{\pi}_k|} = \frac{|\tilde{\pi}_j| + |\tilde{\pi}_{j+1}|}{\sum_{h=0}^{j-i} |\tilde{\pi}_{j+1-h}|} \end{aligned} \quad (1)$$

Letting $\alpha_n = \frac{n-2}{n+2}$ (hence $1 - \alpha_n = \frac{4}{n+2}$) and applying Lemma 39 inductively, we get that:

$$\begin{aligned} |\tilde{\pi}_{j+1-h}| &\geq \alpha_n^h \cdot |\tilde{\pi}_{j+1}| - \sum_{g=0}^{h-1} \alpha_n^g \cdot \frac{4n}{n+2} \\ &= \alpha_n^h \cdot |\tilde{\pi}_{j+1}| - \frac{1 - \alpha_n^h}{1 - \alpha_n} \cdot \frac{4n}{n+2} \\ &= \alpha_n^h \cdot (|\tilde{\pi}_{j+1}| + n) - n \end{aligned}$$

Let h_n be the largest index $0 \leq h \leq j-i$ such that $\alpha_n^h (|\tilde{\pi}_{j+1}| + n) \geq n+1$. Notice that the inequality holds when $h=0$ (because $|\tilde{\pi}_{j+1}| \geq 1$), so that h_n is well-defined. Then (again because $|\tilde{\pi}_{j+1-h}| \geq 1$ for any h):

$$\begin{aligned} \sum_{h=0}^{j-i} |\tilde{\pi}_{j+1-h}| &\geq \sum_{h=0}^{h_n} [\alpha_n^h (|\tilde{\pi}_{j+1}| + n) - n] + \sum_{h=h_n+1}^{j-i} 1 \\ &= (|\tilde{\pi}_{j+1}| + n) \cdot \frac{1 - \alpha_n^{h_n+1}}{1 - \alpha_n} - n \cdot (h_n + 1) + (j - i - h_n) \end{aligned}$$

Using inequality $\alpha_n^{h_n} (|\tilde{\pi}_{j+1}| + n) - 1 \geq n$, we get

$$\begin{aligned} \sum_{h=0}^{j-i} |\tilde{\pi}_{j+1-h}| &\geq (|\tilde{\pi}_{j+1}| + n) \cdot \frac{1 - \alpha_n^{h_n+1}}{1 - \alpha_n} - (h_n + 1) \cdot [\alpha_n^{h_n} (|\tilde{\pi}_{j+1}| + n) - 1] \\ &\quad + (j - i - h_n) \\ &= \frac{|\tilde{\pi}_{j+1}| + n}{1 - \alpha_n} \cdot \left(1 - \alpha_n^{h_n+1} - (h_n + 1)\alpha_n^{h_n}(1 - \alpha_n)\right) + (j - i + 1) \\ &\geq \frac{|\tilde{\pi}_{j+1}| + n}{1 - \alpha_n} \cdot (1 + h_n\alpha_n^{h_n+1} - (h_n + 1)\alpha_n^{h_n}) \end{aligned}$$

Also, from Lemma 39,

$$|\tilde{\pi}_j| + |\tilde{\pi}_{j+1}| \leq \frac{|\tilde{\pi}_{j+1}|}{\alpha_n} + \frac{4n}{n-2} + |\tilde{\pi}_{j+1}| = (|\tilde{\pi}_{j+1}| + 2) \cdot \frac{2n}{n-2} \quad (2)$$

Using $n > 2$, we get

$$\begin{aligned} \llbracket \tilde{\pi}_{i+1}\tilde{\pi}_{i+2}\dots\tilde{\pi}_{j-1}\text{step}_j(\tilde{\pi}), \eta_i^S \rrbracket &\leq \frac{\frac{2n}{n-2} \cdot (|\tilde{\pi}_{j+1}| + 2)}{\frac{|\tilde{\pi}_{j+1}| + n}{1 - \alpha_n} \cdot (1 + h_n\alpha_n^{h_n+1} - (h_n + 1)\alpha_n^{h_n})} \\ &\leq \frac{8n}{n^2 - 4} \cdot \frac{1}{1 + h_n\alpha_n^{h_n+1} - (h_n + 1)\alpha_n^{h_n}}. \end{aligned}$$

We now prove that $\tau(n) \leq \frac{1}{2} - \frac{1}{n}$ when $n = 30$. First consider the function $f: h \mapsto (1 + h\alpha_n^{h+1} - (h+1)\alpha_n^h)^{-1}$. This function is easily proved to be decreasing when $h > 0$. In particular, if $h \geq 15$, $f(h) \leq f(15) \leq 1.64$. It follows that if $h_n \geq 15$, then $\tau(n) \leq 0.44 \leq 1/2 - 1/n$.

It remains to handle the case where $h_n < 15$. First, by definition of h_n , we have

$$\alpha_n^{h_n+1} \cdot (|\tilde{\pi}_{j+1}| + n) < n + 1$$

so that

$$|\tilde{\pi}_{j+1}| < \frac{n+1}{\alpha_n^{h_n+1}} - n \leq \frac{31}{\left(\frac{28}{32}\right)^{15}} - 30 \leq 200.$$

Using Equations (1) and (2) and the fact that $|\tilde{\pi}_{j+1-h}| \geq 1$ for all h , we end up with

$$\tau(30) \leq \frac{(|\tilde{\pi}_{j+1}| + 2) \cdot \frac{2 \cdot 30}{30-2}}{j-i+1} \leq \frac{202 * 60}{28} \cdot \frac{1}{p+1}$$

where the last inequality follows from the fact that $j - i$ is a multiple of p . Requiring $\tau(30) \leq 1/2 - 1/30$, we end up with the requirement that $p \geq 927$.