

propositional version popularized by Alur and Henzinger [2]. Their seminal work has also studied its complexity over dense and discrete time [2], [3], as well as interesting decidable fragments for dense time [1]. While their work basically settled the problems for discrete time, follow-up work by other authors has extended and refined the picture for dense time, such as by studying expressive completeness [16], [17], simplifying decision procedures [20], or identifying expressive decidable fragments [5], [23], [24].

Bounded variability is a natural semantic restriction over dense time, which has been applied to various formalisms including timed automata [28], duration calculus [10], and, in our previous work, MTL [12]. Recently, we also applied it to LTL over discrete time; while it is obvious that every discrete-time model has bounded variability (given by the fixed duration associated with one discrete time unit), in [14], [15] we showed how the LTL validity problem can be simplified under the assumption that only $v < V$ change events happen every V discrete time steps. Therefore, bounded variability can be a simplifying assumption also for discrete time.

The undecidability results of Section V-A use reductions from undecidable problems of nondeterministic n -counter machines, which we introduce in Section IV. These are a kind of Minsky's counter machines [21]; their connection with MTL was first exploited by Alur and Henzinger [2].

Section VI discusses MTL fragments with lower complexity. Over discrete time, these fragments can be derived from similarly low-complexity fragment of LTL [9], which have been extensively studied by several authors [27], [7], [19].

II. TIMED WORDS AND VARIABILITY

We denote a generic time domain by \mathbb{T} . In the paper, \mathbb{T} is either the *discrete* set of the nonnegative integers \mathbb{N} , or the *dense* (and *continuous*) set of the nonnegative reals $\mathbb{R}_{\geq 0}$.

An *interval* is a convex subset of the time domain, represented by a pair $\langle a, b \rangle$, where \langle and \rangle are square or round brackets to respectively denote inclusion or exclusion of the endpoint. We use the pseudo-arithmetic expressions $> s$, $\geq s$, $< s$, $\leq s$, and $= s$ as abbreviations for the intervals (s, ∞) , $[s, \infty)$, $[0, s)$, $[0, s]$ and $[s, s]$. We assume a binary encoding of constants in the time domain unless explicitly stated otherwise.

Given a time domain \mathbb{T} and a finite alphabet set \mathcal{P} of atomic propositions, a *timed word* over \mathbb{T} is a countably infinite sequence of pairs $\omega = (\sigma_0, t_0)(\sigma_1, t_1)(\sigma_2, t_2) \cdots$ such that:

- 1) Each integer $k \geq 0$ denotes a *position* in a timed word;
- 2) For each k , σ_k is a (nonempty, w.l.o.g.) subset of \mathcal{P} denoting the propositions holding at position k ; and $t_k \in \mathbb{T}$ is a timestamp denoting the time of the occurrence at position k ;
- 3) The timestamps are strictly monotonic, that is $t_h > t_k$ iff $h > k$, and diverging, that is for all $t \in \mathbb{T}$ there exists k such that $t_k > t$ (divergence is subsumed by monotonicity in discrete time).

We also conventionally assume that $t_0 = 0$. The set of all timed words over \mathbb{T} is denoted by \mathcal{BT} .

A timed word ω has *variability bounded by v/V* , for $V \in \mathbb{T}$ and $v \in \mathbb{N}$, iff it has no more than v positions within any closed time interval of length V : for all $k \in \mathbb{N}$, $t_{k+v} - t_k > V$. The set of all timed words over \mathbb{T} with variability bounded by v/V is denoted by $\mathcal{BT}[v/V]$.

III. MTL: METRIC TEMPORAL LOGIC

We present the syntax and semantics of propositional MTL and recall some fundamental facts about its complexity.

Syntax. MTL formulas are defined by the grammar

$$\phi ::= \top \mid p \mid \neg\phi_1 \mid \phi_1 \wedge \phi_2 \mid \bigcup_J(\phi_1, \phi_2),$$

where p ranges over the alphabet \mathcal{P} , and J is an interval of the time domain with integer endpoints. We assume the standard definitions for *false*: \perp , and for the derived Boolean connectives: \vee , \Rightarrow , and \Leftrightarrow . The symbol α abbreviates the formula $\bigvee_{p \in \mathcal{P}} p$, which holds iff some proposition holds. We introduce the derived temporal operators *eventually*: $\diamond_J(\phi) = \bigcup_J(\top, \phi)$; *globally* (also, *always*): $\square_J(\phi) = \neg \diamond_J(\neg\phi)$; *action until*: $\hat{\bigcup}_J(\phi_1, \phi_2) = \bigcup_J(\alpha \Rightarrow \phi_1, \phi_2)$; and *next*: $\bigcirc_J(\phi) = \hat{\bigcup}_J(\perp, \phi)$. Operator precedence is: \neg has the highest precedence, then \wedge , then \vee , then \Rightarrow , then all temporal operators, and finally \Leftrightarrow . We may omit the parentheses around arguments when unambiguous, and drop intervals $[0, \infty)$.

Semantics. Given a timed word $\omega = (\sigma_0, t_0)(\sigma_1, t_1) \cdots$ and a position $k \in \mathbb{N}$, the *pointwise* satisfaction relation \models_p for an MTL formula ϕ is inductively defined as follows:

$$\begin{aligned} \omega, k \models_p \top; \\ \omega, k \models_p p & \quad \text{iff } p \in \sigma_k; \\ \omega, k \models_p \neg\phi_1 & \quad \text{iff } \omega, k \not\models_p \phi_1; \\ \omega, k \models_p \phi_1 \wedge \phi_2 & \quad \text{iff } \omega, k \models_p \phi_1 \text{ and } \omega, k \models_p \phi_2; \\ \omega, k \models_p \bigcup_J(\phi_1, \phi_2) & \quad \text{iff there exists } h > k \text{ such that:} \\ & \quad t_h - t_k \in J, \omega, h \models_p \phi_2, \text{ and,} \\ & \quad \text{for all } k < x < h, \omega, x \models_p \phi_1; \\ \omega \models_p \phi & \quad \text{iff } \omega, 0 \models_p \phi. \end{aligned}$$

The semantics of \bigcup and $\hat{\bigcup}$ coincide under the pointwise semantics, since formulas are only evaluated at positions, where α invariably holds. The (derived) semantics of next is: $\omega, k \models_p \bigcirc_J(\phi_1)$ iff $t_{k+1} - t_k \in J$ and $\omega, k+1 \models_p \phi_1$; that is, the next position has timestamp in J relative to the current one, and ϕ_1 holds there.

Given a timed word ω as above and a time instant $t \in \mathbb{T}$, the *continuous* satisfaction relation \models_c for an MTL formula ϕ is inductively defined as follows (we only list the cases that differ from the pointwise semantics):

$$\begin{aligned} \omega, t \models_c p & \quad \text{iff there exists } k \in \mathbb{N} \text{ such that:} \\ & \quad t_k = t \text{ and } p \in \sigma_k; \\ \omega, t \models_c \bigcup_J(\phi_1, \phi_2) & \quad \text{iff there exists } u > t \text{ such that:} \\ & \quad u - t \in J, \omega, u \models_c \phi_2, \text{ and,} \\ & \quad \text{for all } t < v < u, \omega, v \models_c \phi_1. \end{aligned}$$

Over dense time, the continuous semantics generalizes the pointwise semantics in the sense that the former is strictly more expressive [8], [5]. The semantics of next under continuous semantics is, however, analogous to that over the pointwise semantics, thanks to the usage of $\hat{\bigcup}$ in its definition.

Remark 1. In the following, we assume the pointwise semantics over discrete time \mathbb{N} , and the continuous semantics over dense time $\mathbb{R}_{\geq 0}$. Our results for dense time are also transferable to the pointwise semantics *mutatis mutandis*, provided *past operators* are available (see Section VII).

Complexity: general models. Satisfiability of MTL formulas is highly undecidable over dense time, where it is Σ_1^1 -hard [2]. It is instead decidable over discrete time, with an **EXSPACE**-complete decidability problem [2] (which translates to doubly-exponential deterministic time). Over discrete time, the high complexity is essentially due to the succinctness of the binary encoding (the expressiveness is the same as LTL).

Complexity: bounded models. Bounded variability is a semantic restriction that reduces the complexity of MTL. In fact, we proved that satisfiability of MTL over dense-time models with variability bounded by v/V , for any given v/V , is **EXSPACE**-complete [12], matching the complexity of MTL over discrete time, as well as that of other decidable dense-time logics [1], [16]. The following is a corollary of our previous results, which we use in this paper.

Corollary 2. *For any v, v' , and V , it is decidable whether an MTL formula has some model over $\mathbb{R}_{\geq 0}$ with variability bounded by v'/V but not by v/V .*

Proof: We showed that the MTL satisfiability problem over $\mathcal{BR}_{\geq 0}[v/V]$ is decidable for generic v/V [12, Corollary 1]; and that we can encode in MTL the bounded variability constrain as well as its complement [12, Section 4.3]. ■

In recent work [14], [15], we showed how the notion of bounded variability can reduce the complexity of MTL over discrete time as well. While bounded variability does not affect the exponential-space worst-case complexity, since discrete-time models have inherently bounded variability, it can reduce the complexity in practice. Precisely, when studying the variability of an arbitrary LTL formula ϕ over behaviors with variability bounded by any given v/V , with $v < V$, we can consider a simplified ϕ' whose size depends on v but not on the distances encoded in ϕ through next operators. While the results of [14], [15] target LTL, it is clear that they carry over to MTL over discrete time.

IV. COUNTER MACHINES

Counter machines [21] are powerful computational devices, widely used in formal language theory. We use a nondeterministic version of counter machines, and derive some complexity results which we use in the remainder.

Definition 3. An n -counter machine executes programs consisting of a finite list of instructions with labels ℓ_0, ℓ_1, \dots and operating on n integer counter variables v_0, \dots, v_{n-1} . An instruction is one of the following:

halt	terminate computation
if $v_k > 0$ goto ℓ_i, ℓ_j	conditional branch
inc v_k	increment counter
dec v_k	decrement counter

where the conditional branch consists in jumping to ℓ_i or ℓ_j nondeterministically if counter v_k is non-zero; and decrementing a counter with zero value is undefined. Computations start at location ℓ_0 with all counters equal to zero and proceed according to the obvious semantics of instructions. Without loss of generality, assume that instruction **halt** occurs exactly once and that the last instruction in the list is either **halt** or a branch.

For n -counter machines, with $n \geq 2$, the halting problem (deciding whether the location with **halt** is visited in some computation) is Σ_1^0 -complete (**RE**-complete: undecidable but semidecidable); the non-halting problem (deciding whether some computation does not halt) is Σ_2^0 -complete; the recurring computation problem (deciding whether location ℓ_0 is visited infinitely often in some computation) is Σ_1^1 -hard [3].²

A. Bounded and Unbounded Counters

Consider the following decision problems for n -counter machines:

- bounded counter:** given an integer β , decide whether v_0 overflows β in some computation;
- finite counter:** decide whether there exists β such that $v_0 \leq \beta$ in all computations;

Theorem 4. *The bounded counter problem is Σ_1^0 -complete; the finite counter problem is Σ_2^0 -complete.*

Proof: We prove hardness by reduction from, respectively, the halting and non-halting problems of n -counter machines. We then report the simpler corresponding completeness proofs.

Hardness of the bounded counter problem. Given a generic n -counter machine M , we reduce halting to bounded counter for $\beta = 0$ by modifying M into M' as follows. Add one counter and injectively rename all counters in the instruction list so that the new counter is called v_0 ; thus, v_0 is not mentioned in the renamed instructions. Then, replace the unique halting instruction appearing at some ℓ_h in M by two instructions: ℓ_h : **inc** v_0 followed by ℓ_h^+ : **halt**.

Since we only added deterministic instructions, there is a one-to-one correspondence between computations of M and computations of M' . A generic nondeterministic computation χ of M reaches location ℓ_h iff the unique corresponding computation χ' of M' also reaches ℓ_h . In such computations χ' , v_0 overflows β before halting at ℓ_h^+ . In all, some computation of M halts iff v_0 overflows in some computation of M' . Thus, the bounded counter problem is Σ_1^0 -hard.

Hardness of the finite counter problem. Given a generic n -counter machine M , we reduce from the non-halting problem. Create another counter machine M' with a fresh counter v_0 , which works as follows. M' simulates all computations of M deterministically: as soon as a specific computation terminates, M' backtracks the simulation and makes a different nondeterministic choice. (We omit the details of the simulation, which are straightforward.) Whenever the simulation

²These complexities easily follow from reduction of the same problems for Turing machines.

completes a halting computation of M , it increments v_0 before continuing with the next computation. If the simulation ever comes to an end (that is, if M has only finitely many computations, all halting), M' enters an infinite loop that makes v_0 diverge. Therefore, M' has only one non-halting (because either M' enters the infinite loop or M has infinitely many computations) deterministic execution.

Consider now the finite counter problem for M' . If it has answer YES, it means that the simulation eventually executes a non-halting computation of M ; from that point on, v_0 is never incremented. If it has answer NO, it means that the simulation consists of infinitely many halting computations of M , or that it reached the divergent loop and hence M had only finitely many halting computations. The answer to the non-halting problem for M is therefore the same in either case. This shows that we reduced the non-halting problem to the finite counter problem, and both are Σ_2^0 -hard.

Completeness of the bounded counter problem. We reduce the bounded counter problem (for any β) to halting, thus showing that the former is in Σ_1^0 (and hence, by combining it with the hardness result, Σ_1^0 -complete). The idea is to guard every increment to v_0 with a conditional of the form **if** $v_0 \geq \beta$ **goto** ℓ_h **else inc** v_0 , where ℓ_h is the halting location. Conversely, computations that halt in the initial programs are modified so that they run forever. Since v_0 is initially zero, a computation halts iff it overflowed in the initial program; hence, a semi-decision for halting produces a semi-decision procedure for bounded counter. The details of how to encode such modifications using standard instructions are straightforward.

Completeness of the finite counter problem. We show that the finite counter problem is in Σ_2^0 (and hence, by combining it with the hardness result, Σ_2^0 -complete) according to the definition of Σ_2^0 in the arithmetical hierarchy [26]. Let \mathcal{O}_β be the set of all counter machines where v_0 overflows β in some computation. Previously, we have shown that \mathcal{O}_β is Σ_1^0 ; hence its complement set $\overline{\mathcal{O}_\beta}$ —all counter machines where $v_0 \leq \beta$ in all computations—is Π_1^0 . The set \mathcal{F} of all counter machines for which the finite counter problem has answer YES is defined by $M \in \mathcal{F} \iff \exists \beta : \overline{\mathcal{O}_\beta}$, and hence it is Σ_2^0 . ■

B. MTL and Counter Machines

Alur and Henzinger [2] pioneered the usage of counter machines to analyze the complexity of real-time logics. Using their techniques, we show the essentials of how to encode computations of n -counter machines as MTL formulas over $\mathbb{R}_{\geq 0}$: computations are encoded as timed words; and, given a machine M , we build an MTL formula Γ_M that is satisfied precisely by the words encoding M 's computations.

Consider an n -counter machine M with $m+1$ instructions ℓ_0, \dots, ℓ_m , such that ℓ_h is the location of the unique halt instruction. We introduce the following propositions: p_k , for $0 \leq k \leq m$, which holds when M is at location ℓ_k ; and z_k , for $1 \leq k \leq n$, which we use to represent the value of counter v_k : there are as many distinct occurrences of proposition z_k over a unit interval as the value of counter

v_k in the corresponding configuration. A configuration is a tuple $\langle \ell_k, x_1, \dots, x_n \rangle$ denoting that M is at location ℓ_k and the counters store the values x_1, \dots, x_n . At each integer time instant: all propositions z_d 's are false; and exactly one of the propositions p_k 's holds, with p_0 holding initially. The p_k 's are all false everywhere else:

$$p_0 \wedge \left(\bigwedge_{1 \leq k \leq m} \square \left(p_k \Rightarrow \bigwedge_{1 \leq j \neq k \leq m} \neg p_j \wedge \bigwedge_{1 \leq d \leq n} \neg z_d \right) \wedge \bigwedge_{1 \leq k \leq m} \square \left(p_k \Rightarrow \bigvee_{1 \leq j \leq m} \bigcup_{i=1} \left(\bigwedge_{1 \leq i \leq m} \neg p_i, p_j \right) \right) \right).$$

With similar formulas, we constrain the z_k 's to occur at distinct instants: whenever z_k then $\neg z_h$ also holds simultaneously, for $h \neq k$.

Each time interval $[t, t+1)$, for $t \in \mathbb{N}$, encodes the $(t+1)$ -th configuration reached during a valid computation: p_k holding at t means that M is at location ℓ_k ; and, for $1 \leq j \leq n$, z_j holds over $[t, t+1)$ exactly as many times as the integer value stored in counter v_j . The initial configuration $\langle \ell_0, 0, \dots, 0 \rangle$ is encoded by

$$\bigwedge_{1 \leq j \leq n} \square_{[0,1]}(\neg z_j).$$

The encoding of any instruction refers to a current time $t \in \mathbb{N}$ and defines the state over $[t+1, t+2)$ as a modification of the state over $[t, t+1)$. The most significant operation is the increment: ℓ_k : **inc** v_c , whose MTL encoding declares that the state in the next interval has exactly one more occurrence of z_c than it has in the current interval:

$$\square \left(p_k \Rightarrow \left(\begin{array}{l} \diamond_{=1} p_{k+1} \\ \wedge \bigwedge_{1 \leq d \neq c \leq n} \square_{(0,1)}(z_d \Leftrightarrow \diamond_{=1} z_d) \\ \wedge \square_{(0,1)}(z_c \Rightarrow \diamond_{=1} z_c) \\ \wedge \bigcup_{(0,1)} \left(\begin{array}{l} \diamond_{=1} z_c \Rightarrow z_c, \\ \neg z_c \wedge \diamond_{=1} z_c \wedge \\ \bigcup_{>0}(\neg z_c \wedge \diamond_{=1}(\neg z_c), p_{k+1}) \end{array} \right) \end{array} \right) \right). \quad (1)$$

In (1)'s consequent, the first conjunct states that ℓ_{k+1} is the next location visited (since this is not a branch instruction). The second conjunct states that the values of all counters other than v_c are unchanged: for every occurrence of some z_d in the current interval, there is an occurrence exactly one time unit later in the next interval and vice versa; hence occurrences of z_d are “copied” from the current to the next interval. Similarly, the third conjunct declares that v_c does not decrease (z_c 's occurrences in the current interval are copied into the next one). The fourth conjunct asserts that there exists an instant, after the last occurrence of z_c in the current interval and before the next occurrence of p_{k+1} at the beginning of the next interval, such that z_c occurs exactly once at the corresponding instant in the next interval. This new distinct occurrence of z_c is always possible thanks to the density of the temporal domain; thus any value of counters can be stored in a unit time interval. The encoding of other instructions is similar, with the halting instruction determining an indefinite repetition of the final configuration in the future.

Remark 5. Over pointwise semantics, we can express a behavior analogous to (1) using past operators. The key

observation [23] is that the “copy” of a counter v_d can be expressed as $\square_{(0,1)}(z_d \Rightarrow \diamond_{=1} z_d)$ and $\square_{(1,2)}(z_d \Rightarrow \overleftarrow{\diamond}_{=1} z_d)$, where $\overleftarrow{\diamond}_{=1}(\phi)$ holds iff its arguments held one time unit in the past.

V. THE COMPLEXITY OF BOUNDED VARIABILITY

Given a time domain \mathbb{T} and a formula ϕ , we define two decision problems—the second is a generalization of the first—that deal with ϕ ’s bounded variability. We write $\mathcal{B}(\phi)$ to denote the subset of a set \mathcal{B} of timed words that satisfy ϕ .

$BV_{\mathbb{T}}(v, V)$: Determine whether every model of ϕ over \mathbb{T} has variability bounded by v/V : does $\mathcal{B}\mathbb{T}(\phi) \subseteq \mathcal{B}\mathbb{T}[v/V](\phi)$?

$\overline{BV}_{\mathbb{T}}$: Determine whether there exist v, V such that the answer to $BV_{\mathbb{T}}(v, V)$ is YES: does $\exists v, V : \mathcal{B}\mathbb{T}(\phi) \subseteq \mathcal{B}\mathbb{T}[v/V](\phi)$?

A bar denotes the corresponding *complement* problems: $\overline{BV}_{\mathbb{T}}(v, V)$ asks whether some model of ϕ has variability not bounded by v/V (bounded by v'/V for some $v' > v$, or unbounded); $\overline{BV}_{\mathbb{T}}$ asks whether, for every v, V , some model of ϕ has variability not bounded by v/V . Notice that the latter is not the same as asking if some model of ϕ has unbounded variability: it may as well be that every model of ϕ has bounded variability, but no variability bounds all of the models.

This section establishes the complexity of the decision problems for $\mathbb{T} = \mathbb{R}_{\geq 0}$ (Section V-A) and $\mathbb{T} = \mathbb{N}$ (Section V-B).

A. Complexity of Bounded Variability over Continuous Time

Both variants of the bounded variability problems just introduced are undecidable over continuous time, but with different undecidability degrees in the arithmetical hierarchy; in both cases, however, the undecidability degree is lesser than MTL satisfiability, which is highly undecidable (Σ_1^1 -hard [2]).

Theorem 6. $BV_{\mathbb{R}_{\geq 0}}(v, V)$ is $\Pi_1^0 = \mathbf{coRE}$ -complete; $\overline{BV}_{\mathbb{R}_{\geq 0}}$ is Σ_2^0 -complete.

Lemma 7. $BV_{\mathbb{R}_{\geq 0}}(v, V)$ is in $\Pi_1^0 = \mathbf{coRE}$.

Proof: We give a procedure to semi-decide $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V)$; this establishes that $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V) \in \mathbf{RE}$ and thus $BV_{\mathbb{R}_{\geq 0}}(v, V) \in \mathbf{coRE}$ by complement.

Consider a generic MTL formula ϕ . Some model of ϕ has variability not bounded by v/V iff: (a) some model of ϕ has variability bounded by v'/V but not by v/V , for some $v' > v$; or (b) some model of ϕ has unbounded variability. Since we are dealing with divergent models only (see Section II), (b) can only occur with models where the variability is bounded up to any finite time t , but the variability bound increases indefinitely over time.³

For any finite time T , let $\phi[T]$ denote the MTL formula which restricts the evaluation of ϕ to the finite time interval $[0, T]$. This can be constructed as follows: add a fresh proposition e constrained by $\phi_e = \bigcup_{=T}(e, e \wedge \square_{>0} \neg e)$. Rewrite ϕ

in negation normal form, and replace every atom q by $e \Rightarrow q$. Postulate that, if e is false, all other propositions in \mathcal{P} are false as well: $\phi_{\mathcal{P}} = \square(\neg e \Rightarrow \bigwedge_{p \in \mathcal{P}} \neg p)$. Finally, $\phi[T]$ is $\phi_e \wedge \phi \wedge \phi_{\mathcal{P}}$. Since no event occurs after finite time T , all models of $\phi[T]$ have variability bounded by x/T , for some finite (possibly very large) x . Therefore, some model of $\phi[T]$ has variability not bounded by t/T iff some model of $\phi[T]$ has variability bounded by t'/T but not by t/T , for some $x \geq t' > t$.

We can now describe a procedure P_1 that semi-decides $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V)$; it consists of the following steps:

- 1) Initially, $\delta := v + 1$ and $\Delta := V + 1$;
- 2) Using Corollary 2, decide whether $\phi[\Delta]$ has some model with variability bounded by δ/V but not by v/V ;
- 3) If it does, stop and return YES;
- 4) Otherwise $\delta := \delta + 1$, $\Delta := \Delta + 1$, and go to (2).

If the answer to $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V)$ is YES, then either (a) or (b) above holds; let us show that, in both cases, P_1 terminates with the correct answer.

If (a) is the case, let ω_a be a model with variability bounded by v'/V but not by v/V for some $v' > v$; that is, ω_a has \bar{v} events, for $v < \bar{v} \leq v'$, over some time interval $[x, x + V]$. In this case, P_1 terminates with YES as soon as $\delta \geq \bar{v}$ and $\Delta \geq x + V$.

If (b) is the case, let ω_b be a model with unbounded variability; since variability is unbounded, there exists a time T such that: ω_b has $v' > v$ events over some time window $[x, x + V]$, for $0 \leq x < x + V \leq T$. In this case, P_1 terminates with YES as soon as $\delta \geq v'$ and $\Delta \geq T$. ■

Lemma 8. $BV_{\mathbb{R}_{\geq 0}}(v, V)$ is **coRE-hard**.

Proof: We reduce the bounded counter problem (Section IV-A) of 2-counter machines to $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V)$; the lemma follows by Theorem 4 through complement problems.

Consider a generic 2-counter machine M with counters v_0 and v_1 . We construct an MTL formula Γ_M that encodes the computations of M along the lines of Section IV-B, but with some modifications. For $t \in \mathbb{N}$, the $(t + 1)$ -th configuration $\langle \ell_k, x_0, x_1 \rangle$ is encoded over the time interval $[4t, 4t + 4)$ as follows: p_k holds at $4t$, z_0 holds x_0 times over $(4t + 1, 4t + 2)$, z_1 holds x_1 times over $(4t + 3, 4t + 4)$, and no propositions hold elsewhere over the whole $[4t, 4t + 4)$. With this spacing of counter events, we can see that the models of Γ_M are such that any interval of length 1 includes at most as many events as the largest value held by a counter during some computation. Thus, Γ_M has some model with variability not bounded by $\beta/1$, which is an instance of $\overline{BV}_{\mathbb{R}_{\geq 0}}(v, V)$, iff a counter overflows β in some computation of M .

Now we have only established whether *some* counter overflows in M , whereas the bounded counter problem specifically targets overflows of v_0 . To close the gap, we encode the overflowing of v_0 in M as an MTL formula $\Xi_{\beta}^{v_0}$:

$$\diamond \left(\left(\bigvee_{0 \leq k \leq m} p_k \right) \wedge \overbrace{\diamond_{(0,1)} \left(z_0 \wedge \diamond_{(0,1)} \left(z_0 \wedge \dots \right) \right)}^{\beta+1 \text{ nested diamonds}} \right).$$

³In related work, we called similar behaviors “Berkeley” [13], [11].

Thanks to the padding, the nested diamonds evaluate to true iff there are at least $\beta + 1$ distinct occurrences of z_0 in the slot corresponding to one configuration. Thus, v_0 overflows β in M iff $\Gamma_M \wedge \Xi_\beta^{v_0}$ has some model with variability not bounded by $\beta/1$. ■

Lemma 9. $BV_{\mathbb{R}_{\geq 0}}$ is in Σ_2^0 .

Proof: Given the definition of Σ_2^0 in the arithmetical hierarchy [26], it is sufficient to provide an enumeration of all MTL formulas ϕ for which the answer to $BV_{\mathbb{R}_{\geq 0}}$ is YES, relative to an oracle for $BV_{\mathbb{R}_{\geq 0}}(v, V)$, which is in Π_1^0 by Lemma 7. To this end, we dovetail [25, Chap. 3] through all pairs (v, ϕ) of nonnegative integers $v \in \mathbb{N}$ and MTL formulas ϕ . For each pair, if the answer to $BV_{\mathbb{R}_{\geq 0}}(v, 1)$ is YES for ϕ , then the answer to $BV_{\mathbb{R}_{\geq 0}}$ also is YES for ϕ . It is clear that this enumeration eventually finds all formulas for which the answer to $BV_{\mathbb{R}_{\geq 0}}$ is YES. ■

Lemma 10. $BV_{\mathbb{R}_{\geq 0}}$ is Σ_2^0 -hard.

Proof: We reduce the finite counter problem (Section IV-A) of n -counter machines to $BV_{\mathbb{R}_{\geq 0}}$; the lemma follows by Theorem 4.

This reduction is the trickiest among those in this paper. The difficulty lies in the fact that, while the finite counter problem refers a specific counter v_0 , $BV_{\mathbb{R}_{\geq 0}}$ considers variability of all propositions; while it is easy to reduce from general to specific, here we need to build a reduction in the opposite direction. The proof of Lemma 8 involves a similar mismatch, but things are simpler there, thanks to the existence of a known bound β , which we can monitor explicitly; now, instead, the bound is existentially quantified. An easy solution would be to change the definition of $BV_{\mathbb{R}_{\geq 0}}$ to refer a specific proposition that varies, but that would weaken the result proved. Instead, we leverage nondeterminism to “guess” the bound.

Build a counter machine M_x which simulates computations of M as follows. Every computation of M_x starts by nondeterministically storing a positive integer x in a fresh counter v_x . This is achieved by the instructions:

$$\begin{aligned} \ell_0 &: \text{inc } v_x \\ \ell_1 &: \text{if } v_x > 0 \text{ goto } \ell_0, \ell_2 \end{aligned} \quad (2)$$

If ℓ_1 's nondeterministic branch eventually jumps to ℓ_2 , the rest of M_x 's program simulates all computations of M by dovetailing [25, Chap. 3], so that the simulation does not get stuck in non-terminating computations of M . Additionally, whenever v_0 overflows $x - 1$, the simulation halts; and if M had only finitely many computations, the simulation concludes with an infinite idle loop (unless it has previously halted upon v_0 overflowing).

Consider now the MTL formula $\Gamma = \Gamma_{M_x} \wedge \diamond p_h$, where Γ_{M_x} encodes M_x 's computations as in Section IV-B, and ℓ_h is the unique halting location of M_x . Thus, the models of Γ describe all valid computations of M_x that halt (and hence, in particular, that do not get stuck forever in the initial loop (2) that increments v_x —something we get for free given that we are reducing between undecidable problems).

Consider now $BV_{\mathbb{R}_{\geq 0}}$ for Γ . If its answer is YES, then there must be only finitely many models that satisfy Γ ; otherwise, they would include simulations for all values of x , which would entail that, for every possible bound x , there exists a model where v_0 overflows x , against the hypothesis that the answer is YES (i.e., all models are bounded). Therefore, there is a finite bound on v_0 in all computations of M , given by one plus the maximum of values reached by v_0 in all finitely many models. Conversely, if the answer to $BV_{\mathbb{R}_{\geq 0}}$ for Γ is NO, then there must be infinitely many models that satisfy Γ , that is one for every value of x ; in fact, all such models are halting, and hence if they are finitely many the maximum of all counters in all such models would be well defined and finite, against the hypothesis that the answer is NO (i.e., there always is an unbounded model). The existence of halting models for all values of x entails that v_0 overflows any finite value in some computation. In summary, the answer to $BV_{\mathbb{R}_{\geq 0}}$ for Γ is YES iff the answer to the finite counter problem for M is YES. This concludes the reduction. ■

B. Complexity of Bounded Variability over Discrete Time

The complexity of bounded variability over discrete time paints a picture quite different from that over continuous time. Lemmas 11–12 prove that $BV_{\mathbb{N}}(v, V)$ is **EXSPACE**-complete, which is the same complexity as MTL satisfiability over \mathbb{N} . On the other hand, it is clear that $BV_{\mathbb{N}}$ is decidable in constant time, since every discrete-time MTL formula has variability bounded by x/x for any integer $x > 0$, precisely because the time domain is discrete, and hence there is a hard upper bound on the variability of events.

Lemma 11. $BV_{\mathbb{N}}(v, V)$ is **EXSPACE**-hard.

Proof: We polynomial-time reduce MTL satisfiability to $\overline{BV_{\mathbb{N}}}(v, V)$; the lemma follows since **EXSPACE** is closed under complement.

The decision procedure for discrete-time TPTL is based on the following fundamental property [3, Lemma 5]: a TPTL formula ψ is satisfiable iff it has a model where the difference between any pair of consecutive timestamps is always less than or equal to the product δ_ψ of all constants appearing in ψ . The same property holds of MTL formulas over the integers [2].

Since we are considering timed ω -words, which have infinitely many events, the property entails that an MTL formula ϕ is satisfiable iff there exists a timed word ω such that: ω has at least one event with timestamp $t \leq \delta_\phi$ and $\omega \models \phi$. Therefore, a generic MTL formula ϕ is satisfiable iff some of its models have variability not bounded by $0/\delta_\phi$, that is iff the answer to $\overline{BV_{\mathbb{N}}}(0, \delta_\phi)$ is YES. Assuming a binary encoding of constants, as it is customary, δ_ϕ is polynomial in the size of ϕ (because the product of n constants has size $O(n^2)$ in binary), thus the reduction is done in polynomial time. ■

Lemma 12. $BV_{\mathbb{N}}(v, V)$ is in **EXSPACE**.

Proof: We show how to encode the requirement that a model has variability bounded by v/V as an MTL formula $B_{v, V}$.

If $v = 0$, then $B_{v,V} = \Box \perp$. Otherwise, we can adapt the techniques we introduced for LTL [14]. Consider $v > 0$ fresh propositions p_i , for $i = 1, \dots, v$. Proposition p_1 holds initially, followed by p_2, \dots, p_v in sequence; the sequence repeats indefinitely:

$$B_v = p_1 \wedge \bigwedge_{1 \leq k \leq v} \left(\Box (p_1 \Leftrightarrow \bigcirc p_{k \oplus 1}) \wedge \Box \left(p_k \Rightarrow \bigwedge_{1 \leq h \neq k \leq v} \neg p_h \right) \right)$$

where $a \oplus b$ is a shorthand for $1 + ((a + b) \bmod v)$. Since every p_k holds in a different position, we can express bounded variability by requiring that the timestamp of the next $(v + 1)$ -th position in the future be greater than V with respect to the current position's (and note that $k \oplus v = k$):

$$B_{v,V} = B_v \wedge \bigwedge_{1 \leq k \leq v} \Box (p_k \Rightarrow \mathsf{U}_{>V}(\neg p_k, p_k)).$$

Thus, $\phi \Rightarrow B_{v,V}$ is valid iff the answer to $BV_{\mathbb{N}}(v, V)$ for ϕ is YES.

The only problem with this reduction is that $B_{v,V}$ has size exponential in the size of the instance of $BV_{\mathbb{N}}(v, V)$ assuming a binary encoding of constants. Precisely, the blow-up occurs because B_v has size polynomial in v , which is exponential in the size of a binary encoding of v . Encoding the modulo- v counter in binary (using $n = \lceil \log_2 v \rceil + 1$ propositions) would not help: while updates to the counter itself can be done with formulas of size polynomial in n , there is no easy way to express in MTL the fact that the timestamp of the “next” occurrence is greater than V (with respect to the current position's) without enumerating all $2^n = v$ values for the counter.

The blow-up is, however, inessential and only due to the fact that MTL operators do not include compact “counting” modalities. We omit the details for brevity, but it is clear that one can extend the standard decision procedures for MTL [2] to handle counting modalities without affecting the complexity of the logic. Specifically, we could introduce an operator $K_{\mathcal{J}}^n \psi$ with the semantics: $\omega, k \models K_{\mathcal{J}}^n \psi$ iff $t_{k+n} - t_k \in \mathcal{J}$ and $\omega, k + n \models \psi$. $B_{v,V}$ for $v > 0$ is then equivalent to $\Box (K_{>V}^v \top)$; assuming a binary encoding of constants, this has size linear in the size of the encodings of v and V . ■

VI. BOUNDED VARIABILITY IN SIMPLE CASES

The complexity results of Section V pose some major limitations to deciding bounded variability for *generic* MTL formulas. However, the outlook may be better if we target *fragments* of MTL that are still sufficiently expressive but for which reasoning about bounded variability is simpler than in the general case. We call such fragments “bounded friendly”. We give two examples of non-trivial bounded-friendly fragments, one for discrete and one for dense time.

A. Simpler Bounded Variability over Discrete Time

Over discrete time, MTL essentially boils down to an exponentially succinct version of LTL. Therefore, we can try to lift some complexity results about simpler fragments of

LTL [27], [7], [4] to MTL over \mathbb{N} , and use them to identify bounded-friendly fragments.

Consider the two dual MTL fragments $\mathcal{F}_{\diamond, \bigcirc}^+$ and $\mathcal{F}_{\Box, \bigcirc}^+$: $\mathcal{F}_{\diamond, \bigcirc}^+$ (respectively, $\mathcal{F}_{\Box, \bigcirc}^+$) denotes the MTL fragment using only the $\diamond_{\mathcal{J}}$ (respectively, $\Box_{\mathcal{J}}$) and $\bigcirc_{\mathcal{J}}$ modalities (which we now regard as primitive), the propositional connectives \wedge and \vee , and where negations only appear on atomic propositions. Satisfiability for these fragments is decidable in exponential time.

Lemma 13. *Satisfiability of $\mathcal{F}_{\diamond, \bigcirc}^+$ and of $\mathcal{F}_{\Box, \bigcirc}^+$ over \mathbb{N} is EXP-complete.*

Proof: Consider the LTL fragment $\mathcal{L}_{F,X}^+$ which only uses the eventually and next LTL modalities, the propositional connectives \wedge and \vee , and where negations only appear on atomic propositions; [27, Th. 3.7] proves that satisfiability for $\mathcal{L}_{F,X}^+$ is NP-complete. We outline how to transform a generic $\mu \in \mathcal{F}_{\diamond, \bigcirc}^+$ into a $\lambda \in \mathcal{L}_{F,X}^+$ such that μ and λ are equisatisfiable; the converse transformation is also derivable along the same lines. In general, the size of λ is exponential in the size of μ due to the fact that metric constraints are encoded in binary in μ . The lemma follows as a manifestation of the “succinctness phenomenon” [25, Chap. 20].

The models of $\mathcal{L}_{F,X}^+$ are denumerable sequences $w = w_0 w_1 \dots$ such that w_k is the set of propositions that hold at step k . We conventionally assume that a step in w corresponds to one discrete time instant; thus, a generic \mathbb{N} -timed word $\omega = (\sigma_0, t_0) (\sigma_1, t_1) \dots$ uniquely corresponds to a sequence $w = w_0 w_1 \dots$ such that: for all $k \in \mathbb{N}$, $w_{t_k} = \sigma_k$; and, for all h 's that are not timestamps of ω , $w_h = \{\epsilon\}$, where ϵ is a special proposition denoting absence of a reading.

We define a translation τ from $\mathcal{F}_{\diamond, \bigcirc}^+$ to $\mathcal{L}_{F,X}^+$ inductively as follows, for $a, b \in \mathbb{N}$ and $c \in \mathbb{N} \cup \{\infty\}$:

$$\begin{aligned} \tau \left(\diamond_{[a,b]}(\pi) \right) &= X_a(\pi_{\tau} \vee \overbrace{X(\pi_{\tau} \vee \dots)}^{b-a \text{ nested } X\text{s}}) \\ \tau \left(\diamond_{[a,\infty)}(\pi) \right) &= X_a F(\pi_{\tau}) \\ \tau \left(\bigcirc_{[a,c]}(\pi) \right) &= \overbrace{X(\epsilon \wedge X(\epsilon \wedge \dots))}^{a-1 \text{ nested } X\text{s}} \wedge \tau \left(\diamond_{[a,c]}(\pi) \right) \end{aligned}$$

where $\tau(\pi) = \pi_{\tau}$ and X_k is a shorthand for k nested applications of X . τ does not otherwise change the propositional structure of formulas.

It should be clear that a generic μ is satisfiable over timed words over \mathbb{N} iff $\tau(\mu)$ is satisfiable: models of $\tau(\mu)$ are obtained from models of μ according to the mapping of timed words described above. Furthermore, the size of $\tau(\mu)$ is $O(2^{|\mu|})$, since τ “unrolls” the constants succinctly represented in μ , which results in a possible exponential blow-up. This establishes the lemma for $\mathcal{F}_{\diamond, \bigcirc}^+$. The same complexity result for $\mathcal{F}_{\Box, \bigcirc}^+$ follows by duality of \Box and \diamond . ■

We can leverage Lemma 13 to show that $\mathcal{F}_{\diamond, \bigcirc}^+$ is *bounded friendly*. Let v, V be any variability bounds, with $v > 0$ w.l.o.g. First, note that the equivalent $\mathcal{F}_{\Box, \bigcirc}^+$ formulas $\Box \Box_{(0,\nu]}(\perp)$ and $\Box \bigcirc_{>\nu}(\top)$, for $\nu = \lceil V/v \rceil$, hold only for models with

variability bounded by v/V (specifically, they are stricter than the definition of bounded variability). Consider now a generic MTL formula ϕ written as $\phi' \wedge \psi$, where $\psi \in \mathcal{F}_{\diamond, \square}^+$. The implication $\psi \Rightarrow B_{v,V} \equiv \neg\psi \vee B_{v,V}$, where $B_{v,V}$ is one of the two just defined $\mathcal{F}_{\square, \diamond}^+$ formulas implying bounded variability, is an $\mathcal{F}_{\square, \diamond}^+$ formula: push in the outermost negation $\neg\psi$, and use the duality between \diamond and \square . The validity of $\psi \Rightarrow B_{v,V}$ can thus be decided in singly exponential time (Lemma 13), as opposed to $BV_{\mathbb{R}}(v, V)$ or the validity of ϕ which are **EXPSpace**-complete: solving them for a generic ϕ takes *time* doubly exponential in $|\phi|$. We can thus decide whether $\psi \Rightarrow B_{v,V}$ is valid in singly exponential time; if it is, ϕ has bounded variability a fortiori, and hence we can study its validity with the simplified algorithms [14], [15].

We can show by duality that $\mathcal{F}_{\diamond, \square}^+$ is also bounded friendly; for example, instead of the validity of $\psi \Rightarrow B_{v,V}$, we equivalently consider the unsatisfiability of $\psi \wedge \neg B_{v,V}$.

B. Simpler Bounded Variability over Continuous Time

While MTL is highly undecidable over dense time, a number of expressive yet decidable fragments thereof have been identified. MITL is the fragment of MTL where intervals are non-punctual, that is non-singular; MITL is fully decidable with **EXPSpace**-complete complexity [1], [16]. More recently, other decidable expressive fragments have been identified that allow singular intervals [24]; BMTL and SMTL, in particular, are interesting because their expressive power is incomparable with MITL's.

From the point of view of deciding bounded variability, however, MITL remains the most suitable choice. SMTL validity has a non-elementary decision problem; while this is still better than the undecidable $BV_{\mathbb{R}_{\geq 0}}(v, V)$, it makes it intractable in practice. BMTL validity, in contrast, is decidable in **EXPSpace**; however, BMTL cannot express invariance properties since only finite intervals are allowed, and it is clear that bounded variability is a form of invariance property since it involves whole timed words.

Let us show that MITL is bounded friendly. Let v, V be any variability bounds, with $v > 0$ w.l.o.g. (the limit case $v = 0$ can be handled separately). Note that the MITL formula $B_{v,V} = \square\square_{(0,v]}(\perp)$, for $\nu = \lceil V/v \rceil$, subsumes variability bounded by v/V . Consider now a generic MTL formula ϕ written as $\phi' \wedge \psi$, where ψ is an MITL formula. The implication $\psi \Rightarrow B_{v,V}$ obviously also is an MITL formula. The validity of $\psi \Rightarrow B_{v,V}$ is thus decidable; if $\psi \Rightarrow B_{v,V}$ is valid, ϕ has bounded variability a fortiori, and hence its validity is decidable [12].

As a final remark, notice how leveraging MITL's bounded friendliness can still be useful to determine the satisfiability of MTL specifications not entirely expressed in MITL: as long as the part ψ expressible in MITL entails bounded variability, the rest ϕ' of the actual specification can use any MTL operator, including singular intervals.

VII. DISCUSSION: OTHER MTL SEMANTICS

Remark 1 clarified that the results of this paper assume infinite timed words, and the continuous semantics over dense time. While these are perfectly common assumptions (and the naturalness of the pointwise semantics over dense time has been questioned [16]), it is still interesting to get an idea of how our results would change under a different semantics.

Over discrete time, it is straightforward to notice that all complexity results proved in V-B carry over to the finite-word semantics (where decidability also has the same **EXPSpace** complexity). Over dense time, it is possible to extend the results of Section V-A to the *signal* semantics of [1], [13]—which can be seen as a variant of the continuous semantics—by reusing some of the constructions and definitions of [12].

The situation over dense timed words (both finite and infinite) under the pointwise semantics (and no past operators) is different. For both finite and infinite words, $BV_{\mathbb{R}_{\geq 0}}(v, V)$ is no more difficult than validity, because one can encode the bounded variability requirement as in Lemma 12. Therefore, $BV_{\mathbb{R}_{\geq 0}}(v, V)$ is decidable (nonprimitive recursive) over finite words [23], and is **RE** over infinite words [22]. One can prove matching lower bounds along the lines of the proof of Lemma 11: the abstraction of clock valuations into clock regions used in the construction of [23] is such that the time difference between any pair of consecutive timestamps in any word that satisfies a formula ϕ is bounded above by a finite constant δ_ϕ that depends only on ϕ . Therefore, $\overline{BV}_{\mathbb{R}_{\geq 0}}(0, \delta_\phi)$ for ϕ reduces to non-satisfiability of ϕ , giving the matching lower bounds through complement. Finally, $BV_{\mathbb{R}_{\geq 0}}$ is **RE** under the pointwise semantics: by definition of **RE** as existential quantification over a recursive relation (finite words); and by dovetailing through the possible values v, V and enumerating $BV_{\mathbb{R}_{\geq 0}}(v, V)$ for them (infinite words). Finding matching lower bounds belongs to future work, which may exploit the connection between MTL over infinite words under the pointwise semantics and channel machines with insertion errors [22].

VIII. CONCLUSIONS

The strong negative results of the paper need not be the deathblow to leveraging bounded variability to simplify temporal reasoning. From a broader perspective, we can still look at the glass as half-full: while deciding bounded variability is intractable in general, there are situations where the physical requirements of a system include a notion of *finite speed*, which bounded variability naturally embodies. For such systems, there is still hope of using the expressiveness of MTL without succumbing to the dark side of intractability.

Acknowledgements. Thanks to the anonymous reviewers of TIME 2014 for useful comments; in particular, reviewer 4's insightful observations helped simplify some of the proofs.

REFERENCES

- [1] R. Alur, T. Feder, and T. A. Henzinger, "The benefits of relaxing punctuality," *JACM*, vol. 43, no. 1, pp. 116–146, 1996.

- [2] R. Alur and T. A. Henzinger, “Real-time logics: Complexity and expressiveness,” *Inf. Comp.*, vol. 104, no. 1, pp. 35–77, 1993.
- [3] —, “A really temporal logic,” *JACM*, vol. 41, no. 1, pp. 181–204, 1994.
- [4] M. Bauland, T. Schneider, H. Schnoor, I. Schnoor, and H. Vollmer, “The complexity of generalized satisfiability for linear temporal logic,” *Logical Methods in Computer Science*, vol. 5, no. 1, 2009.
- [5] P. Bouyer, N. Markey, J. Ouaknine, and J. Worrell, “The cost of punctuality,” in *LICS*, 2007, pp. 109–120.
- [6] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, and G. Sciacicco, “The dark side of interval temporal logic: Sharpening the undecidability border,” in *TIME*. IEEE, 2011, pp. 131–138.
- [7] S. Demri and P. Schnoebelen, “The complexity of propositional linear temporal logics in simple cases,” *Inf. Comput.*, vol. 174, no. 1, pp. 84–103, 2002.
- [8] D. D’Souza and P. Prabhakar, “On the expressiveness of MTL in the pointwise and continuous semantics,” *STTT*, vol. 9, no. 1, pp. 1–4, 2007.
- [9] E. A. Emerson, “Temporal and modal logic,” in *Handbook of Theoretical Computer Science*. Elsevier, 1990, pp. 996–1072.
- [10] M. Fränzle, “Model-checking dense-time duration calculus,” *Formal Asp. Comput.*, vol. 16, no. 2, pp. 121–139, 2004.
- [11] C. A. Furia, D. Mandrioli, A. Morzenti, and M. Rossi, *Modeling Time in Computing*, ser. Monographs in Theoretical Computer Science. An EATCS series. Springer, 2012.
- [12] C. A. Furia and M. Rossi, “MTL with bounded variability: Decidability and complexity,” in *FORMATS*, ser. LNCS, vol. 5215. Springer, 2008, pp. 109–123.
- [13] —, “A theory of sampling for continuous-time metric temporal logic,” *ACM Transactions on Computational Logic*, vol. 12, no. 1, pp. 1–40, 2010, article 8.
- [14] C. A. Furia and P. Spoletini, “On relaxing metric information in linear temporal logic,” in *TIME*. IEEE, 2011, pp. 72–79.
- [15] —, “Automata-based verification of linear temporal logic models with bounded variability,” in *TIME*. IEEE, 2012, pp. 89–96.
- [16] Y. Hirshfeld and A. M. Rabinovich, “Logics for real time: Decidability and complexity,” *Fundam. Inform.*, vol. 62, no. 1, pp. 1–28, 2004.
- [17] P. Hunter, J. Ouaknine, and J. Worrell, “Expressive completeness for metric temporal logic,” in *LICS*. IEEE, 2013, pp. 349–357.
- [18] R. Koymans, “Specifying real-time properties with metric temporal logic,” *Real-Time Systems*, vol. 2, no. 4, pp. 255–299, 1990.
- [19] C. Lutz, D. Walther, and F. Wolter, “Quantitative temporal logics over the reals: PSPACE and below,” *Inf. Comput.*, vol. 205, no. 1, pp. 99–123, 2007.
- [20] O. Maler, D. Nickovic, and A. Pnueli, “Real time temporal logic: Past, present, future,” in *FORMATS*, ser. LNCS, vol. 3829. Springer, 2005, pp. 2–16.
- [21] M. L. Minsky, *Computation: Finite and Infinite Machines*. Prentice Hall, 1967.
- [22] J. Ouaknine and J. Worrell, “On metric temporal logic and faulty Turing machines,” in *FoSSaCS*, ser. LNCS, vol. 3921. Springer, 2006, pp. 217–230.
- [23] —, “On the decidability and complexity of metric temporal logic over finite words,” *Logical Methods in Computer Science*, vol. 3, no. 1, 2007.
- [24] —, “Some recent results in metric temporal logic,” in *FORMATS*, ser. LNCS, vol. 5215. Springer, 2008, pp. 1–13.
- [25] C. Papadimitriou, *Computational Complexity*. Addison-Wesley, 1994.
- [26] H. Rogers, Jr., *Theory of Recursive Functions and Effective Computability*. MIT Press, 1987.
- [27] A. P. Sistla and E. M. Clarke, “The complexity of propositional linear temporal logics,” *JACM*, vol. 32, no. 3, pp. 733–749, 1985.
- [28] T. Wilke, “Specifying timed state sequences in powerful decidable logics and timed automata,” in *FTRTFT*, ser. LNCS, vol. 863. Springer, 1994, pp. 694–715.