Metric Temporal Logic with Counting

Shankara Narayanan Krishna[1(✉)], Khushraj Madnani[1], and Paritosh K. Pandya[2]

[1] Department of Computer Science and Engineering, IIT Bombay, Mumbai 400 076, India
{krishnas,khushraj}@cse.iitb.ac.in

[2] School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai 400 005, India
pandya@tcs.tifr.res.in

Ability to count number of occurrences of events within a specified time interval is very useful in specification of resource bounded real time computation. In this paper, we study an extension of Metric Temporal Logic (MTL) with two different counting modalities called C and UT (until with threshold), which enhance the expressive power of MTL in orthogonal fashion. We confine ourselves only to the future fragment of MTL interpreted in a pointwise manner over finite timed words. We provide a comprehensive study of the expressive power of logic CTMTL and its fragments using the technique of EF games extended with suitable counting moves. Finally, as our main result, we establish the decidability of CTMTL by giving an equisatisfiable reduction from CTMTL to MTL. The reduction provides one more example of the use of temporal projections with oversampling introduced earlier for proving decidability. Our reduction also implies that MITL extended with C and UT modalities is elementarily decidable.

1 Introduction

Temporal logics provide constructs to specify qualitative ordering between events in time. But real time logics have the ability to specify quantitative timing constraints between events. Metric Temporal Logic MTL is amongst the best studied of real time logics. Its principal modality $a \, \mathsf{U}_I b$ states that an event b should occur in future within a time distance lying within interval I. Moreover, a should hold continuously till then.

In many situations, especially those dealing with resource bounded computation, the ability to count the number of occurrences of events becomes important. In this paper, we consider an extension of MTL with two counting modalities C and UT (until threshold) which provide differing abilities to specify constraints on counts on events in time intervals. The resulting logic is called CTMTL. Modality $\mathsf{C}_I^{\geq n}\, \phi$ states that the number of times formula ϕ holds in time interval I (measured relative to current time point) is at least n. This is a mild generalization of $\mathsf{C}_{(0,1)}^{\geq n}\, \phi$ modality studied by Rabinovich [2] in context of continuous time MTL. The UT modality $\phi \, \mathsf{U}_{I,\#\kappa \geq n}\, \psi$ is like MTL until but it additionally states that the number of time formula κ holds between now and time point

B. Jacobs and C. Löding (Eds.): FOSSACS 2016, LNCS 9634, pp. 335–352, 2016.
DOI: 10.1007/978-3-662-49630-5_20

where ψ holds is at least n. Thus it extends U to simultaneously specify constraint on time and count of subformula. Constraining U by count of subformula was already explored for untimed LTL by Laroussini et al. [7]. But the combination of timing and counting seems new. The following example illustrates the use of these modalities.

Example 1. We specify some constraints to be monitored by exercise bicycle electronics.

– Two minutes after the start of exercise, the heartbeat (number of pulses in next 60 seconds) should be between 90 and 120. This can be stated as

$$\Box(st \Rightarrow (\mathsf{C}^{\geq 90}_{[120,180]} pulse \wedge \mathsf{C}^{<120}_{[120,180]} pulse))$$

– Here is one exerise routine: After start of exercise, *slow_peddling* should be done for 1 kilometre (marked by odometer giving 1000 pulses) and this should be achieved in interval 1 to 2 min. After this *fast_peddling* should be done for 3 min. This can be specified as

$$\Box(st \Rightarrow slowpeddle \mathsf{U}_{[60,120],\#odo=1000}(\Box_{[0,180]} fastpeddle))$$

In specifying requirements over hybrid systems where count of events or mode changes are to be constrained by time intervals, our logic is quite relevant. In an HVAC case study, we have used these operators to specify properties such as no more than 3 switching on of motor are permitted in any one minute time interval. Similarly, in resource bounded computation, fairness constraints often need counting; e.g. "no more than 3 login attempted should be made in one minute". We believe that our operators are quite natural and useful in requirement modelling of real time systems.

The expressiveness and decidability properties of real time logics differ considerable based on nature of time. There has been considerable study of counting MTL in continuous time [3,12]. In this paper, we consider the case of pointwise time, i.e. CTMTL interpreted over finite timed words in a pointwise manner. We provide a comprehensive picture of expressiveness and decidability of CTMTL and its fragments in pointwise time and we find that this differs considerably when compared with continuous time.

As our first main result, we show that the C and the UT modalities both increase the expressive power of MTL but they are mutually incomparable. EF games are a classical technique used to study expressive power of logic. [10] have adapted EF games to MTL and shown a number of expressiveness results. In this paper, we further extend MTL EF games with counting moves corresponding to the C and UT modalities. We use the resulting EF theorem to characterise the expressive powers of several fragments of CTMTL.

One attraction of pointwise MTL over finite timed words is that its satisfiability is decidable [9] whereas continuous time MTL has undecidable satisfiability. As our second main result, we show that MTL extended with C and UT modalities also has decidable satisfiability. In order to prove this result, we give an equisatisfiable reduction from CTMTL to MTL. The reduction makes use of the notion of temporal projections modulo oversampling introduced earlier [4] where timed words satisfying original CTMTL formula have to be oversampled with additional time points to satisfy corresponding MTL formula. This result marks one more use of the technique of temporal projections. We note that our reduction can also be applied to MITL (with both U and S) extended with C and UT and it gives an equisatisfiable formula in MITL which is exponential in the size of original formula. Thus, we establish that CTMITL[U, S] has elementary satisfiability.

2 A Zoo of Timed Temporal Logics

In this section, we present the syntax and semantics of the various timed temporal logics we study in this paper. Let Σ be a finite set of propositions. A finite timed word over Σ is a tuple $\rho = (\sigma, \tau)$. σ and τ are sequences $\sigma_1\sigma_2\ldots\sigma_n$ and $t_1t_2\ldots t_n$ respectively, with $\sigma_i \in 2^{\Sigma} - \emptyset$, and $t_i \in \mathbb{R}_{\geq 0}$ for $1 \leq i \leq n$ and $\forall i \in dom(\rho)$, $t_i \leq t_{i+1}$, where $dom(\rho)$ is the set of positions $\{1, 2, \ldots, n\}$ in the timed word. An example of a timed word over $\Sigma = \{a, b\}$ is $\rho = (\{a, b\}, 0.3)(\{b\}, 0.7)(\{a\}, 1.1)$. ρ is strictly monotonic iff $t_i < t_{i+1}$ for all $i, i+1 \in dom(\rho)$. Otherwise, it is weakly monotonic. The set of finite timed words over Σ is denoted $T\Sigma^*$.

The logic MTL extends linear temporal logic (LTL) by adding timing constraints to the "until" modality of LTL. We parametrize this logic by a permitted set of time intervals denoted by $I\nu$. The intervals in $I\nu$ can be open, half-open or closed, with end points in $\mathbb{N} \cup \{0, \infty\}$. Such an interval is denoted $\langle a, b\rangle$. For example, $[3, 7), [5, \infty)$. Let $t + \langle a, b\rangle = \langle t + a, t + b\rangle$.

Metric Temporal Logic. Given Σ, the formulae of MTL are built from Σ using boolean connectives and time constrained version of the modality U as follows: $\varphi ::= a(\in \Sigma)|true|\varphi \wedge \varphi|\neg\varphi|\varphi\mathsf{U}_I\varphi$ where $I \in I\nu$. For a timed word $\rho = (\sigma, \tau) \in T\Sigma^*$, a position $i \in dom(\rho)$, and an MTL formula φ, the satisfaction of φ at a position i of ρ is denoted $(\rho, i) \models \varphi$, and is defined as follows:
$\rho, i \models a \leftrightarrow a \in \sigma_i$ and $\rho, i \models \neg\varphi \leftrightarrow \rho, i \nvDash \varphi$
$\rho, i \models \varphi_1 \wedge \varphi_2 \leftrightarrow \rho, i \models \varphi_1$ and $\rho, i \models \varphi_2$
$\rho, i \models \varphi_1\mathsf{U}_I\varphi_2 \leftrightarrow \exists j > i,\ \rho, j \models \varphi_2, t_j - t_i \in I$, and $\rho, k \models \varphi_1\ \forall\ i < k < j$
ρ satisfies φ denoted $\rho \models \varphi$ iff $\rho, 1 \models \varphi$. Let $L(\varphi) = \{\rho \mid \rho, 1 \models \varphi\}$ denote the language of a MTL formula φ. Two formulae φ and ϕ are said to be equivalent denoted as $\varphi \equiv \phi$ iff $L(\varphi) = L(\phi)$. Additional temporal connectives are defined in the standard way: we have the constrained future eventuality operator $\Diamond_I a \equiv true\mathsf{U}_I a$ and its dual $\Box_I a \equiv \neg\Diamond_I\neg a$. We also define the next operator as $\mathsf{O}_I\phi \equiv \bot\mathsf{U}_I\phi$. Weak versions of operators are defined as $\Diamond^{\mathsf{w}}_I a = a \vee \Diamond_I a, \Box^{\mathsf{w}}_I a \equiv a \wedge \Box_I a$, $a\mathsf{U}^{\mathsf{w}}_I b \equiv b \vee [a \wedge (a\mathsf{U}_I b)]$ if $0 \in I$, and $[a \wedge (a\mathsf{U}_I b)]$ if $0 \notin I$.

Theorem 1 [9]. *Satisfiability checking of* MTL *is decidable over finite timed words with non-primitive-recursive complexity.*

Metric Temporal Logic with Counting (CTMTL). We denote by CTMTL the logic obtained by extending MTL with the ability to count, by endowing it with two counting modalities C as well as UT.

Syntax of CTMTL: $\varphi ::= a(\in \Sigma)|true|\varphi \wedge \varphi|\neg\varphi|\varphi|\mathsf{C}_I^{\geq n}\varphi|\varphi\mathsf{U}_{I,\eta}\varphi$, where $I \in I\nu$, $n \in \mathbb{N} \cup \{0\}$ and η is a *threshold formula* of the form $\#\varphi \geq n$ or $\#\varphi < n$. The counting modality $\mathsf{C}_I^{\geq n}\varphi$ is called the C modality, while $\varphi\mathsf{U}_{I,\eta}\varphi$ is called the UT modality. Let $\rho = (\sigma, \tau) \in T\Sigma^*$, $i, j \in dom(\rho)$. Define

$$N^\rho[i, I](\varphi) = \{k \in dom(\rho) \mid t_k \in t_i + I \wedge \rho, k \models \varphi\}, \text{ and}$$
$$\rho[i, j](\varphi) = \{k \in dom(\rho) \mid i < k < j \wedge \rho, k \models \varphi\}.$$

Denote by $|N^\rho[i, I](\varphi)|$ and $|\rho[i, j](\varphi)|$ respectively, the cardinality of $N^\rho[i, I](\varphi)$ and $\rho[i, j](\varphi)$. $|N^\rho[i, I](\varphi)|$ is the number of points in ρ that lie in the interval t_i+I, and which satisfy φ, while $|\rho[i, j](\varphi)|$ is the number of points lying between i and j which satisfy φ. Define $\rho, i \models \mathsf{C}_I^{\geq n}\varphi$ iff $|N^\rho[i, I](\varphi)| \geq n$. Likewise, $\rho, i \models \varphi_1\mathsf{U}_{I,\#\varphi\geq n}\varphi_2$ iff $\exists j>i$, $\rho, j \models \varphi_2$, $t_j - t_i \in I$, and $\rho, k \models \varphi_1$, $\forall i < k < j$ and $|\rho[i, j](\varphi)| \geq n$.

Remark: The classical until operator of MTL is captured in CTMTL since $\varphi\mathsf{U}_I\psi \equiv \varphi\mathsf{U}_{I,\#true\geq 0}\psi$. We can express $\mathsf{C}_I^{\sim n}$ and $\#\varphi \sim n$ for $\sim \in \{\leq, <, >, =\}$ in CTMTL since $\mathsf{C}_I^{<n}\varphi \equiv \neg\mathsf{C}_I^{\geq n}\varphi$, $\mathsf{C}_I^{>n}\varphi \equiv \mathsf{C}_I^{\geq n+1}\varphi$, $\mathsf{C}_I^{\leq n}\varphi \equiv \neg\mathsf{C}_I^{\geq n+1}\varphi$ and $\#\varphi > n \equiv \#\varphi \geq n+1$, $\#\varphi \leq n \equiv \neg(\#\varphi > n+1)$. It can be shown that boolean combinations of threshold formulae are also expressible in CTMTL (see [6]). Thus, $a\mathsf{U}_{(1,2),(\#d=3\wedge\#\mathsf{C}_{(0,1)}^{<2}b)}\ c$ is expressible in CTMTL. The *nesting depth* of a CTMTL formula is the maximum nesting of C, UT operators. Formally,

- $depth(\varphi_1\mathsf{U}_{I,\#\varphi_3\sim n}\varphi_2) = max(depth(\varphi_1), depth(\varphi_2), depth(\varphi_3) + 1)$,
- $depth(C_I^{\geq n}\varphi) = depth(\varphi) + 1$, $depth(\varphi \wedge \psi) = max(depth(\varphi), depth(\psi))$,
- $depth(\neg\varphi) = depth(\varphi)$ and $depth(a) = 0$ for any $a \in \Sigma$.

For example, $depth(a\mathsf{U}_{[0,2],\eta}\mathsf{C}^{\geq 1}b)$ with $\eta = \#[a\mathsf{U}_{(0,1),\#[\mathsf{C}_{(0,1)}^{=2}a\wedge\Diamond_{(0,1),\#d=2}]\geq 1}c] < 7$ is 3. We obtain the following natural fragments of CTMTL as follows: We denote by CMTL, the fragment of CTMTL obtained by using the C modality and the U_I modality. Further, $\mathsf{C}_{(0,u)}$MTL denotes the subclass of CMTL where the interval I in $\mathsf{C}_I^{\sim n}\varphi$ is of the form $I = \langle 0, b\rangle$. When the interval is of the form $I = \langle 0, 1\rangle$, then we denote the class by $\mathsf{C}_{(0,1)}$MTL. Note that $\mathsf{C}_{(0,1)}$MTL is the class which allows counting in the next one unit of time. This kind of counting (unit counting in future and past) was introduced and studied in [2] in the continuous semantics. $\mathsf{C}_{(0,1)}$MTL is the pointwise counterpart of this logic, with only future operators. Clearly, $\mathsf{C}_{(0,1)}\mathsf{MTL} \subseteq \mathsf{C}_{(0,u)}\mathsf{MTL} \subseteq \mathsf{CMTL} \subseteq \mathsf{CTMTL}$. Restricting CTMTL to the UT modality, we obtain the fragment TMTL. Restricting the C modality to $\mathsf{C}_{(0,1)}$ or $\mathsf{C}_{(0,u)}$ and also allowing the UT modality, one gets the fragments $\mathsf{C}_{(0,1)}$TMTL and $\mathsf{C}_{(0,u)}$TMTL respectively. If we disallow the C modality, restrict the intervals

I appearing in the formulae to non-punctual intervals of the form $\langle a, b\rangle$ $(a \neq b)$, and restrict threshold formulae η to be of the form $\#true \geq 0$, then we obtain MITL.

3 Expressiveness Hierarchy in the Counting Zoo

In this section, we study the expressiveness and hierarchy of the logics introduced in Sect. 2. The main results of this section are the following:

Theorem 2. $\mathsf{MTL} \subset \mathsf{C}_{0,1)}\mathsf{MTL} \subset \mathsf{C}_{(0,u)}\mathsf{MTL} \subset \mathsf{TMTL} = \mathsf{C}_{(0,u)}\mathsf{TMTL} \subset \mathsf{CTMTL}$. *Moreover,* CMTL *and* TMTL *are incomparable, and* $\mathsf{C}_{(0,u)}\mathsf{MTL} \subset \mathsf{CMTL}$.

While Theorem 2 shows that there is an expressiveness gap between classical MTL and CTMTL, we show later that both these logics are equisatisfiable. Given $\varphi \in \mathsf{CTMTL}$, we can construct a formula $\psi \in \mathsf{MTL}$ such that φ is satisfiable iff ψ is. Note that our notion of equisatisfiability is a special one *modulo temporal projections*. If φ is over an alphabet Σ, ψ is constructed over a suitable alphabet $\Sigma' \supseteq \Sigma$ such that $L(\psi)$, when projected over to Σ gives $L(\varphi)$.

Theorem 3. *Satisfiability Checking of* CTMTL *is decidable over finite timed words.*

The rest of this paper is devoted to the proofs of Theorems 2 and 3. We establish Theorem 2 through Lemmas 1 to 4. To prove the separation between two logics, we define model-theoretic games.

3.1 CTMTL Games

Our games are inspired from the standard model-theoretic games [1,14]. The MTL games were introduced in [10]. We now extend these to CTMTL games.

Let (ρ_1, ρ_2) be a pair of timed words. We define a r-round k-counting pebble I_ν game on (ρ_1, ρ_2). The game is played on (ρ_1, ρ_2) by two players, the Spoiler and the Duplicator. The Spoiler will try to show that ρ_1 and ρ_2 are $\{r, k\}$-distinguishable by some formula in CTMTL[1] while the Duplicator will try to show that ρ_1, ρ_2 are $\{r, k\}$-indistinguishable in TMTL. Each player has r rounds and has access to a finite set of $\leq k$ pebbles from a box of pebbles $\mathcal{P}$ in each round of the game. Let I_ν be the set of permissible intervals allowed in the game.

A configuration of the game at the start of a round p is a pair of points (i_p, j_p) where $i_p \in dom(\rho_1)$ and $j_p \in dom(\rho_2)$. A configuration is called partially isomorphic, denoted $isop(i_p, j_p)$ iff $\sigma_{i_p} = \sigma_{j_p}$. Exactly one of the Spoiler or the Duplicator eventually wins the game. The initial configuration is (i_1, j_1), the

[1] ρ_1, ρ_2 are $\{r, k\}$-distinguishable iff there exists a CTMTL formula φ having $depth(\varphi) \leq r$ with max counting constant $\leq k$ in any threshold formula η or C modality in φ such that $\rho_1 \models \varphi$ and $\rho_2 \not\models \varphi$ or vice-versa.

starting positions of both the words, before the first round. A 0-round game is won by the Duplicator iff $isop(i_1, j_1)$. The r round game is played by first playing one round from the starting position. Either the Spoiler wins the round, and the game is terminated or the Duplicator wins the round, and now the second round is played from this new configuration and so on. The Duplicator wins the game only if he wins all the rounds. The following are the rules of the game in any round. Assume that the current configuration is (i_p, j_p).

– If $isop(i_p, j_p)$ is not true, then Spoiler wins the game, and the game is terminated. Otherwise, the game continues as follows:
– The Spoiler chooses one of the words by choosing ρ_x, for $x \in \{1, 2\}$. Duplicator has to play on the other word ρ_y, where $x \neq y$. Then Spoiler plays either a $\mathsf{U}_{I,\eta}$ round, by choosing an interval $I \in I_\nu$, and a number $c \leq k$ of counting pebbles to be used, or a $\mathsf{C}_I^{\sim c}$ round by choosing an interval $I \in I_\nu$ and a number $c \leq k$ of counting pebbles to be used. The number c is obtained from $\eta = \#\varphi \geq c$ or $\eta = \neg(\#\varphi \geq c)$.
  $\mathsf{U}_{I,\eta}$ *round*: Given the current configuration as (i_p, j_p) with $isop(i_p, j_p)$, then
  • Spoiler chooses a position $i'_p \in dom(\rho_x)$ such that $i_p < i'_p$ and $(t_{i'_p} - t_{i_p}) \in I$.
  • The Duplicator responds by choosing $j'_p \in dom(\rho_y)$ in the other word such that $j_p < j'_p$ and $(t_{j'_p} - t_{j_p}) \in I$. If the Duplicator cannot find such a position, the Spoiler wins the round and the game. Otherwise, the game continues and Spoiler chooses one of the following three options.
  • $\Diamond$ Part: The round ends with the configuration $(i_{p+1}, j_{p+1}) = (i'_p, j'_p)$.
  • U Part: Spoiler chooses a position j''_p in ρ_y such that $j_p < j''_p < j'_p$. The Duplicator responds by choosing a position i''_p in ρ_x such that $i_p < i''_p < i'_p$. The round ends with the configuration $(i_{p+1}, j_{p+1}) = (i''_p, j''_p)$. If Duplicator cannot choose an i''_p, the game ends with Spoiler's win.
  • Counting Part: First, Spoiler chooses one of the two words to play in the counting part. In his chosen word, Spoiler keeps $c \leq k$ pebbles from $\mathcal{P}$ at c distinct positions between the points j_p and j'_p (or i_p and i'_p depending on the choice of the word). In response, the Duplicator also keeps c pebbles from $\mathcal{P}$ at c distinct positions between the points i_p and i'_p (or j_p and j'_p) in his word. Spoiler then chooses a pebbled position say i''_p (note that $i_p < i''_p < i'_p$) in the Duplicator's word. In response, Duplicator chooses a pebbled position j''_p (note that $j_p < j''_p < j'_p$) in the Spoiler's word, and the game continues from the configuration $(i_{p+1}, j_{p+1}) = (i''_p, j''_p)$. At the end of the round, the pebbles are returned to the box of pebbles $\mathcal{P}$.

$\mathsf{C}_I^{\sim c}$*round*: Given the current configuration as (i_p, j_p) with $isop(i_p, j_p)$, Spoiler chooses an interval $I \in I_\nu$ as well as a number $c \leq k$. Spoiler then chooses one of the words to play (say ρ_1). From i_p, Spoiler places c pebbles from $\mathcal{P}$ in the points lying in the interval $t_{i_p} + I$. In response, Duplicator also places c pebbles from $\mathcal{P}$ in the points lying in $t_{j_p} + I$. Spoiler now picks a pebbled position j'_p in the word ρ_2, while Duplicator picks a pebbled position i'_p in the Spoiler's word. The round ends with the configuration (i'_p, j'_p). At the end of the round, the pebbles are returned to the box of pebbles $\mathcal{P}$.

Intuition on Pebbling: To give some intuition behind the pebbling, consider $\#\varphi \geq c$ or $\mathsf{C}_I^{\geq c}\varphi$. The idea behind $\mathsf{Spoiler}$ keeping c pebbles on his word in the chosen interval I is to say that these are the c points where φ evaluates to true. $\mathsf{Duplicator}$ is expected to find c such points in his word. If $\mathsf{Spoiler}$ suspects that in the $\mathsf{Duplicator}$'s word, there are $< c$ positions in I where φ holds good, he picks up the appropriate pebble at the position where φ fails. However, any pebbled position in $\mathsf{Spoiler}$'s word will satisfy φ. In this case, $\mathsf{Duplicator}$ loses. Similarly, if we have $\neg(\#\varphi \geq c)$, or $\mathsf{C}_I^{<c}\varphi$, then $\mathsf{Spoiler}$ chooses the word (say ρ_1) on which φ evaluates to true $\geq c$ times. Then $\mathsf{Duplicator}$ is on ρ_2. The idea is for $\mathsf{Spoiler}$ to find if there exist c or more positions in the interval I in ρ_1 where φ holds good, and if so, pebble those points. This is based on $\mathsf{Spoiler}$'s suspicion that there are at least c positions in I where φ evaluates to true, violating the formula. In response, $\mathsf{Duplicator}$ does the same on ρ_2. $\mathsf{Spoiler}$ will now pick any one of the c pebbles from ρ_2 and check for $\neg\varphi$. This is again based on $\mathsf{Spoiler}$'s belief that whichever c points $\mathsf{Duplicator}$ pebbles in ρ_2, $\neg\varphi$ will evaluate to true in at least one of them. If φ holds at all the c points in ρ_1, then $\mathsf{Duplicator}$ will lose on picking any pebble from ρ_1.

- We can restrict various moves according to the modalities provided by the logic. For example, in a $\mathsf{TMTL}[\Diamond_I]$ game, the possible rounds are $\Diamond_I$ and $\Diamond_{I,\eta}$. A CMITL game has only $\mathsf{U}_I, \mathsf{C}_I^{\geq n}$ rounds, with I_ν containing only non-punctual intervals.

Game equivalence: $(\rho_1, i_1) \approx_{r,k,I_\nu} (\rho_2, j_1)$ iff for every r-round, k-counting pebble CTMTL game over the words ρ_1, ρ_2 starting from the configuration (i_1, j_1), the $\mathsf{Duplicator}$ always has a winning strategy.

Formula equivalence: $(\rho_1, i_1) \equiv_{r,k,I_\nu}^{\mathsf{CTMTL}} (\rho_2, j_1)$ iff for every CTMTL formula φ of depth $\leq r$ having max counting constant $\leq k$ in the C, UT modalities, $\rho_1, i_1 \models \varphi \iff \rho_2, j_1 \models \varphi$. The proof of Theorem 4 can be found in [6].

Theorem 4. *$(\rho_1, i_1) \approx_{r,k,I_\nu} (\rho_2, j_1)$ iff $(\rho_1, i_1) \equiv_{r,k,I_\nu}^{\mathsf{CTMTL}} (\rho_2, j_1)$*

We now use these games to show the separation between various logics. For brevity, from here on, we omit I_ν from the notations $\equiv_{r,k,I_\nu}^{\mathsf{CTMTL}}$, $\equiv_{r,k,I_\nu}^{\mathsf{CMTL}}$, $\equiv_{r,k,I_\nu}^{\mathsf{TMTL}}$ and $\equiv_{r,I_\nu}^{\mathsf{MTL}}$.

Lemma 1. $\mathsf{CMTL} - \mathsf{TMTL} \neq \emptyset$

Proof. Consider the formula $\varphi = \mathsf{C}_{(1,2)}^{\geq 2} a \in \mathsf{CMTL}$. We show that for any choice of n rounds and k pebbles, we can find two words ρ_1, ρ_2 such that $\rho_1 \models \varphi, \rho_2 \not\models \varphi$, but $\rho_1 \equiv_{n,k}^{\mathsf{TMTL}} \rho_2$. Both ρ_1, ρ_2 are over $\Sigma = \{a\}$. Let $0 < \delta < \epsilon < \frac{1}{10^{10nk}}$ and $0 < \kappa < \frac{\epsilon - \delta}{2nk}$. Let l be the maximum constant in $\mathbb{N}$ appearing in the permissible intervals I_ν. Consider the word ρ_1 with $nl(k+1) = K$ unit intervals, with the following time stamps as depicted pictorially (Fig. 1) and in the table.

Thus, ρ_1 and ρ_2 differ only in the interval (1,2): ρ_1 has two points in (1,2), while ρ_2 has only one. Thus, $\rho_1 \models \varphi, \rho_2 \not\models \varphi$.

Points in	ρ_1	ρ_2
(0,1)	$x_1 = 0.5, z_1 = 0.6, y_1 = 0.8$ and $2nk$ points between z_1, y_1 that are κ apart from each other	$x_1' = 0.5, z_1' = 0.6, y_1' = 0.8$ and $2nk$ points between z_1', y_1' that are κ apart from each other
(1,2)	$x_2 = 1.8 - \epsilon, z_2 = 1.8 + \epsilon$	$x_2' = 1.8 - \epsilon$
(2,3)	$e = 2.4 + n\epsilon, y_2 = 2.7 + n\epsilon$ and $2nk$ points between e and y_2 that are κ apart from each other	$z_2' = 2.4 + n\epsilon, y_2' = 2.7 + n\epsilon$ and $2nk$ points between z_2' and y_2' that are κ apart from each other
$(i, i+1)$ $3 \leq i \leq K-1$	$x_i = i + 0.4 + (n-i)\epsilon$ $z_i = i + 0.8 + (n+i)\epsilon + \delta$ $y_i = i + 0.8 + (n+i+1)\epsilon$ and $2nk$points between z_i, y_i that are κ apart from each other	$x_i' = i + 0.4 + (n-i)\epsilon$ $z_i' = i + 0.8 + (n+i)\epsilon + \delta$ $y_i' = i + 0.8 + (n+i+1)\epsilon$ and $2nk$ points between z_i, y_i that are κ apart from each other

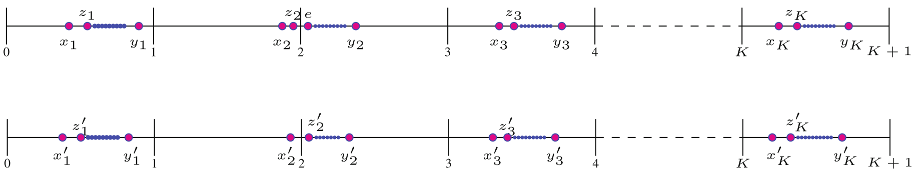

Fig. 1. Words showing $\mathsf{CMTL} - \mathsf{TMTL} \neq \emptyset$

Let $seg(i_p) \in \{0, 1, \ldots, K\}$ denote the left endpoint of the left closed, right open unit interval containing the point $i_p \in dom(\rho_1)$ or $dom(\rho_2)$. Our segments are [0,1), [1,2), ..., $[K, K+1)$. For instance, if the configuration at the start of the pth round is (i_p, j_p) with time stamps (1.2, 3), then $seg(i_p) = 1, seg(j_p) = 3$. The following lemma says that in any round of the game, Duplicator can either achieve the same segment in both the words, or ensure that the difference in the segments is at most 1. Moreover, by the choice of the words, there are sufficiently many segments on the right of any configuration so that Duplicator can always duplicate Spoiler's moves for the remaining rounds, preserving the lag of one segment.

Copy-cat strategy. Consider the pth round of the game with configuration (i_p, j_p). If Duplicator can ensure that $seg(i_{p+1}) - seg(i_p) = seg(j_{p+1}) - seg(j_p)$, then we say that Duplicator has adopted a *copy-cat* strategy in the pth round. We prove the following proposition to argue Duplicator's win.

Proposition 1. *For an n round* TMTL *game over the words ρ_1, ρ_2, the* Duplicator *always has a winning strategy such that for any $1 \leq p \leq n$, if (i_p, j_p) is the initial configuration of the p^{th} round, then $|seg(i_p) - seg(j_p)| \leq 1$. Moreover, when $|seg(i_p) - seg(j_p)| = 1$, then there are at least $(n-p)(l+1)$ segments to the right on each word after p rounds, for all $1 \leq p \leq n$.*

Proof. The initial configuration has time stamps (0,0). We will play a (n, k)-TMTL game on ρ_1, ρ_2. Assume that the Spoiler chooses ρ_1 while the Duplicator chooses ρ_2. Since the interval [1,2] is the only one different in both the words, it is interesting to look at the moves where the Spoiler chooses a point in interval (1,2). We consider the two situations possible for Spoiler to land up in a point in interval (1,2): he can enter interval (1,2) from some point in interval (0,1), or directly choose to enter interval (1,2) from the initial configuration with time stamps (0,0).

Situation 1: Consider the case when from the starting configuration (i_1, j_1) with time stamps (0,0), Spoiler chooses a $\mathsf{U}_{(1,2)\#a\sim c}$ move in ρ_1 and lands up at the point x_2 or z_2. In response, Duplicator has to come at the point x'_2 in ρ'_2. If (i'_1, j'_1) has time stamps (x_2, x'_2) and if Spoiler chooses to pebble between 0 and x_2, then Duplicator pebbles between 0 and x'_2; however, an identical configuration is obtained. Note that if Spoiler pebbles ρ_2, then Duplicator has it easy, since he will pebble the same positions in ρ_1. Let us hence consider obtaining the configuration (i'_1, j'_1) with time stamps (z_2, x'_2), and let Spoiler pebble ρ_1. Spoiler can keep a maximum of k pebbles in the points $x_1, \ldots, y_1, x_2$, while Duplicator keeps the same number of pebbles on the points $x'_1, \ldots, y'_1$. In this case, Spoiler has to a pick a pebbled position from among $x'_1, \ldots, y'_1$. In response, Duplicator will pick the same position from Spoiler's word and achieve an identical configuration. An interesting special case is when Spoiler keeps a single pebble at x_2 in ρ_1. In this case, Duplicator's best choice is to keep his pebble at x'_1, so that the next configuration (i_2, j_2) is one with time stamps (x_2, x'_1). x'_1 and x_2 are *topologically similar* in the sense that the distribution of points in subsequent segments have some nice properties as given below.

Topological Similarity of Words: Consider the $2nk + 3$ points $x_j < z_j < p_j^1 < \cdots < p_j^{2nk} < y_j$ in ρ_1, and $x'_{j-1} < z'_{j-1} < q_{j-1}^1 < \ldots < q_{j-1}^{2nk} < y'_{j-1}$ in ρ_2, for $j \in \{2, 3, 4, \ldots, K\}$. Define a function f that maps points in ρ_1 to topologically similar points in ρ_2. $f : \{x_j, z_j, p_j^1, \ldots, p_1^{2nk}, y_j\} \rightarrow \{x'_{j-1}, z'_{j-1}, q_{j-1}^1, \ldots, q_{j-1}^{2nk}, y'_{j-1}\}$ by $f(x_j) = x'_{j-1}, f(z_j) = z'_{j-1}, f(y_j) = y'_{j-1}$, $f(p_j^i) = q_{j-1}^i$. Let $g = f^{-1}$.

(a) The current configuration has timestamps $(x_2, x'_1) = (x_2, f(x_2))$. For $j \geq 2$, if Spoiler chooses to move to any $p \in \{z_j, y_j, x_{j+2}\}$ from x_2, then Duplicator can move to $f(p)$ from $f(x_2)$ since, for any time interval I, it can be seen that $p - x_2 \in I$ iff $f(p) - f(x_1) \in I$. Moreover, if Spoiler chooses to move to x_3 from x_2, then Duplicator can move to z'_2 from $f(x_2)$ since, $x_3 - x_2, z'_2 - f(x_2) \in (0, 1)$.

(b) We can extend (a) above as follows: Let the current configuration have timestamps $(p, f(p))$ or (x_3, z'_2). Then it can be seen that for any $q \in \{x_j, y_j, z_j\}$ and interval I, $q - p \in I$ iff $f(q) - f(p) \in I$, and $q - x_3 \in I$ iff $f(q) - z'_2 \in I$.

The facts claimed in (a) and (b) are evident from the construction of the timed words. They show that from a configuration (i_p, j_p), such that $seg(i_p) - seg(j_p) \leq 1$, Duplicator can always achieve an intermediate configuration (i'_p, j'_p) in any $\mathsf{U}_{I,\#a\sim c}$ such that $seg(i'_p) - seg(j'_p) \leq 1$. If Spoiler does not go for the until round or the counting round, then $(i_{p+1}, j_{p+1}) = (i'_p, j'_p)$. If Spoiler pebbles the points between i_p and i'_p (or j_p and j'_p), then Duplicator can always ensure that he pebbles points $f(P)$ in ρ_2 whenever Spoiler pebbles a set of points P in ρ_1. As a result, if Spoiler chooses a point $q = f(i) \in f(P)$ in ρ_2, then Duplicator can choose the point $g(q) = i \in P$ achieving the configuration $(i_{p+1}, j_{p+1}) = (g(q), q) = (i, f(i))$. By definition of f, g, we have $i_{p+1} - j_{p+1} \leq 1$. Note that Duplicator can also achieve an identical configuration if Spoiler moves ahead by several segments from i_p (thus, $i'_p >> i_p$), and pebbles a set of points that are also present between j_p and j'_p.

Situation 2: Starting from (i_1, j_1) with time stamps (0,0), if the Spoiler chooses a $\mathsf{U}_{(0,1),\#a\sim c}$ move and lands up at some point between x_1 and y_1, Duplicator will play copy-cat and achieve an identical configuration. Consider the case when Spoiler lands up at y_1[2]. In response, Duplicator moves to y'_1. From configuration (i_2, j_2) with time stamps (y_1, y'_1), consider the case when Spoiler initiates a $\mathsf{U}_{(1,2),\#a\sim c}$ and moves to $z_2 = 1.8 + \epsilon < 2$. In response, Duplicator moves to the point $z'_2 = 2.1 > 2$. A pebble is kept at the inbetween positions x_2, x'_2 respectively in ρ_1, ρ_2. When Spoiler picks the pebble in Duplicator's word, then we obtain the configuration (i_3, j_3) with time stamps (x_2, x'_2). If Spoiler does not get into the counting part/until part, the configuration obtained has time stamps (z_2, z'_2), with the lag of one segment ($seg(i_3) = 1$, $seg(j_3) = 2$, $seg(j_3)$-$seg(i_3)$=1). We show in [6] that from (i_3, j_3) with time stamps either (x_2, x'_2) or (z_2, z'_2), Duplicator can either achieve an identical configuration, or achieve a configuration with a lag of one segment.

From situations (1), (2) in Proposition 1, we know that either Duplicator achieves an identical configuration, in which case there is no segment lag, or there is a lag of at most one segment. The length of the words are $lnk + nl = K$. If Spoiler always chooses bounded intervals (of length $\leq l$), then Duplicator respects his segment lag of 1, and the maximum number of segments that can be explored in either word is at most $nl < K$. In this case, after p rounds, there are at least $K - pl \geq nlk + nl - pl \geq (n - p)(l + 1)$ segments to the right of ρ_1 and $K - pl + 1$ segments to the right of ρ_2. If Spoiler chooses an unbounded interval in any round, then Duplicator can either enforce an identical configuration in both situations 1 and 2, or obtain one of the configurations with time stamps $(p, f(p))$, $f(p) \neq x'_2$, or (z_2, x'_2) or (x_2, x'_2), from where it is known that Duplicator wins.

Lemma 2. $\mathsf{MTL} \subset \mathsf{C}_{(0,1)}\mathsf{MTL} \subset \mathsf{C}_{(0,u)}\mathsf{MTL}$

Proof. We show that the formula $\varphi = \mathsf{C}^{=2}_{(0,1)} a \in \mathsf{C}_{(0,1)}\mathsf{MTL}$ cannot be expressed in MTL. Likewise, the formula $\varphi = \mathsf{C}^{\geq 2}_{(0,2)} a \in \mathsf{C}_{(0,u)}\mathsf{MTL}$ cannot be expressed in $\mathsf{C}_{(0,1)}\mathsf{MTL}$. A detailed proof of these are given in [6].

Lemma 3. *(i)* $\mathsf{C}_{(0,u)}\mathsf{MTL} \subset \mathsf{TMTL} = \mathsf{C}_{(0,u)}\mathsf{TMTL} = \mathsf{C}_{(0,1)}\mathsf{TMTL}$ *and*
(ii) $\mathsf{C}_{(0,u)}\mathsf{MTL} \subset \mathsf{CMTL}$.

[2] The argument when Spoiler lands up at x_1 or a point in between x_1, y_1 is exactly the same.

Proof. (i) The first containment as well as the last two equalities follows from the fact that the counting modality $\mathsf{C}^{\geq n}_{\langle 0,j\rangle}\varphi$ of $\mathsf{C}_{(0,\mathsf{u})}\mathsf{MTL}$ can be written in TMTL as $\Diamond_{\langle 0,j\rangle,\#\varphi\geq n} true$. The strict containment of $\mathsf{C}_{(0,\mathsf{u})}\mathsf{MTL}$ then follows from Lemma 4. (ii) We know that $\mathsf{C}_{(0,\mathsf{u})}\mathsf{MTL} \subseteq \mathsf{CMTL}$. This along with (i) and Lemma 1 gives the strict containment.

Lemma 4. $\mathsf{TMTL} - \mathsf{CMTL} \neq \emptyset$

Proof. Consider the formula $\varphi = \Diamond_{(0,1),\#a\geq 3} b \in \mathsf{TMTL}$. We show that for any choice of n rounds and k pebbles, we can find two words ρ_1, ρ_2 such that $\rho_2 \models \varphi, \rho_1 \not\models \varphi$, but $\rho_1 \equiv^{\mathsf{CMTL}}_{n,k} \rho_2$. The words can be seen in Fig. 2 and the details in [6].

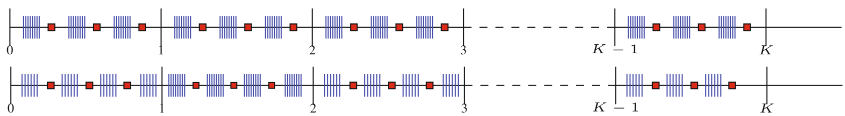

Fig. 2. The red square represents a, the block of blue lines represents a block of b's. There are 3 a's in each unit interval of both ρ_1 and ρ_2. The difference is that ρ_1 has 3 blocks of b's in each unit interval, while ρ_2 has 4 blocks of b's in each unit interval except the last. Clearly, $\rho_2 \models \varphi, \rho_1 \not\models \varphi$ (Color figure online).

4 Satisfiability Checking of Counting Logics

In this section, we show that CTMTL has a decidable satisfiability checking. For this, given a formula in CTMTL we synthesize an equisatisfiable formula in MTL, and use the decidability of MTL. We start discussing some preliminaries. Let Σ, X be finite sets of propositions such that $\Sigma \cap X = \emptyset$.

1. (Σ, X)-*simple extensions.* A (Σ, X)-simple extension is a timed word $\rho' = (\sigma', \tau')$ over $X \cup \Sigma$ such that at any point $i \in dom(\rho')$, $\sigma'_i \cap \Sigma \neq \emptyset$. For $\Sigma = \{a, b\}, X = \{c, d\}$, $(\{a\}, 0.2)(\{b, c, d\}, 0.3)(\{b, d\}, 1.1)$ is a (Σ, X)-simple extension. However, $(\{a\}, 0.2)(\{c, d\}, 0.3)(\{b, d\}, 1.1)$ is not.
2. *Simple Projections.* Consider a (Σ, X)-simple extension ρ. We define the *simple projection* of ρ with respect to X, denoted $\rho \setminus X$ as the word obtained by erasing the symbols of X from each σ_i. Note that $dom(\rho) = dom(\rho \setminus X)$. For example, if $\Sigma = \{a, c\}$, $X = \{b\}$, and $\rho = (\{a, b, c\}, 0.2)(\{b, c\}, 1)(\{c\}, 1.3)$, then $\rho \setminus X = (\{a, c\}, 0.2)(\{c\}, 1)(\{c\}, 1.3)$. $\rho \setminus X$ is thus, a timed word over Σ. If the underlying word ρ is *not* a (Σ, X)-simple extension, then $\rho \setminus X$ is *undefined.*
3. (Σ, X)-*oversampled behaviours.* A (Σ, X)-oversampled behaviour is a timed word $\rho' = (\sigma', \tau')$ over $X \cup \Sigma$, such that $\sigma'_1 \cap \Sigma \neq \emptyset$ and $\sigma'_{|dom(\rho')|} \cap \Sigma \neq \emptyset$. Oversampled behaviours are more general than simple extensions since they allow occurrences of new points in between the first and the last position. These new points are called *oversampled points.*

All other points are called *action points*. For $\Sigma = \{a, b\}, X = \{c, d\}$, $(\{a\}, 0.2)(\{c, d\}, 0.3)(\{a, b\}, 0.7)(\{b, d\}, 1.1)$ is a (Σ, X)-oversampled behaviour, while $(\{a\}, 0.2)(\{c, d\}, 0.3)(\{c\}, 1.1)$ is not.

4. *Oversampled Projections.* Given a (Σ, X)-oversampled behaviour $\rho' = (\sigma', \tau')$, the oversampled projection of ρ' with respect to Σ, denoted $\rho' \downarrow X$ is defined as the timed word obtained by deleting the oversampled points, and then erasing the symbols of X from the action points. $\rho = \rho' \downarrow X$ is a timed word over Σ.

A *temporal projection* is either a simple projection or an oversampled projection. We now define *equisatisfiability modulo temporal projections*. Given MTL formulae ψ and ϕ, we say that ϕ is equisatisfiable to ψ *modulo temporal projections* iff there exist disjoint sets X, Σ such that (1) ϕ is over Σ, and ψ over $\Sigma \cup X$, (2) For any timed word ρ over Σ such that $\rho \models \phi$, there exists a timed word ρ' such that $\rho' \models \psi$, and ρ is a temporal projection of ρ' with respect to X, (3) For any behaviour ρ' over $\Sigma \cup X$, if $\rho' \models \psi$ then the temporal projection ρ of ρ' with respect to X is well defined and $\rho \models \phi$.

If the temporal projection used above is a simple projection, we call it *equisatisfiability modulo simple projections* and denote it by $\phi = \exists X.\psi$. If the projection in the above definition is an oversampled projection, then it is called *equisatisfiability modulo oversampled projections* and is denoted $\phi \equiv \exists \downarrow X.\psi$. Equisatisfiability modulo simple projections are studied extensively [5,11,13]. It can be seen that if $\varphi_1 = \exists X_1.\psi_1$ and $\varphi_2 = \exists X_2.\psi_2$, with X_1, X_2 disjoint, then $\varphi_1 \wedge \varphi_2 = \exists (X_1 \cup X_2).(\psi_1 \wedge \psi_2)$ [8].

As in the case of simple projections, equisatisfiability modulo oversampled projections are also closed under conjunctions when one considers the relativized formulae. For example, consider a formula $\varphi = \Box_{(0,1)} a$ over $\Sigma = \{a, d\}$. Let $\psi_1 = \Box_{(0,1)}(a \vee b) \wedge \Diamond_{(0,1)}(b \wedge \neg a)$ be a formula over the extended alphabet $\{a, b, d\}$ and $\psi_2 = \Box(c \leftrightarrow \Box_{(0,1)} a) \wedge c$ over the extended alphabet $\{a, c, d\}$. Note that $\varphi = \exists \downarrow \{b\}.\psi_1$ and $\varphi = \exists \downarrow \{c\}.\psi_2$ but $\varphi \wedge \varphi \neq \exists \downarrow \{b, c\}.(\psi_1 \wedge \psi_2)$ as the left hand side evaluates to φ which is satisfiable while the right hand side is unsatisfiable. This is due to the presence of a *non-action* point where only b holds. But this can easily be fixed by relativizing all the formulae over their respective action points. ψ_1 is relativized as $\lambda_1 = \Box_{(0,1)}(act_1 \rightarrow (a \vee b)) \wedge \Diamond_{(0,1)}(act_1 \wedge b \wedge \neg a)$ and ψ_2 is relativized as $\lambda_2 = \Box(act_2 \rightarrow (c \leftrightarrow \Box_{(0,1)}(act_2 \rightarrow a))) \wedge act_2 \wedge c$ where $act_1 = b \vee d \vee a$ and $act_2 = a \vee c \vee d$. Now, $\varphi \wedge \varphi = \exists \downarrow \{b, c\}.(\lambda_1 \wedge \lambda_2)$. The relativized forms of ψ_1, ψ_2 are called their *Oversampled Normal Forms* with respect to Σ and denoted $ONF_\Sigma(\psi_1)$ and $ONF_\Sigma(\psi_2)$. Then it can be seen that $\varphi_1 \wedge \varphi_2 = \exists \downarrow \{b, d\}.[ONF_\Sigma(\psi_1) \wedge ONF_\Sigma(\psi_2)]$, and $\varphi_1 = \exists \downarrow \{b\}.ONF_\Sigma(\psi_1)$, $\varphi_2 = \exists \downarrow \{d\}.ONF_\Sigma(\psi_2)$. The formal definition of $ONF_\Sigma(\varphi)$ for a formula φ over $\Sigma \cup X$ can be found in [6]. Equisatisfiability modulo oversampled projections were first studied in [4] to eliminate non-punctual past from MTL over timed words. We use equisatifiability modulo simple projections to eliminate the C modality and oversampled projections to eliminate the UT modality from CTMTL.

Elimination of Counting Modalities from CTMTL. In this section, we show how to eliminate the counting constraints from CTMTL over strictly monotonic timed words. This can be extended to weakly monotonic timed words.

Given any CTMTL formula φ over Σ, we "flatten" the C, UT modalities of φ and obtain a flattened formula. As an example, consider the formula $\varphi = a\mathsf{U}_{[0,3]}(c \wedge \mathsf{C}^{=1}_{(2,3)} d\mathsf{U}_{(0,1),\#(d\wedge\mathsf{C}^{=1}_{(0,1)}e)\geq 1}\mathsf{C}^{\geq 2}_{(0,1)}e])$. Replacing the counting modalities with fresh witness propositions w_1, w_2, we obtain $\varphi_{flat} = [a\mathsf{U}_{[0,3]}(c \wedge w_1)] \wedge T$ where $T = T_1 \wedge T_2 \wedge T_3 \wedge T_4$, with $T_1 = \Box^{\mathsf{w}}[w_1 \leftrightarrow \mathsf{C}^{=1}_{(2,3)}w_2]$, $T_2 = \Box^{\mathsf{w}}[w_2 \leftrightarrow d\mathsf{U}_{(0,1),\#w_4\geq 1}w_3]]$, $T_3 = \Box^{\mathsf{w}}[w_4 \leftrightarrow (d\wedge\mathsf{C}^{=1}_{(0,1)}e)]$, and $T_4 = \Box^{\mathsf{w}}[w_3 \leftrightarrow \mathsf{C}^{\geq 2}_{(0,1)}e]$. Each temporal projection T_i obtained after flattening contains either a C modality or a UT modality. In the following, we now show how to obtain equisatisfiable MTL formulae corresponding to each temporal projection.

Lemma 5. *The formula $\mathsf{C}^{\geq n}_{\langle l,\infty)}b$ is equivalent to MTL formula $F_{\langle l,\infty)}(b \wedge F(b \wedge \ldots Fb)))$.*

We now outline the steps followed to obtain an equisatisfiable formula in MTL, assuming $\mathsf{C}^{\geq n}_{\langle l,\infty)}b$ modalities have been eliminated using Lemma 5.

1. *Flattening*: Flatten χ obtaining χ_{flat} over $\Sigma \cup W$, where W is the set of witness propositions used, $\Sigma \cap W = \emptyset$.
2. *Eliminate Counting*: Consider, one by one, each temporal definition T_i of χ_{flat}. Let $\Sigma_i = \Sigma \cup W \cup X_i$, where X_i is a set of fresh propositions, $X_i \cap X_j = \emptyset$ for $i \neq j$.
   - If T_i is a temporal projection containing a C modality of the form $\mathsf{C}^{\sim n}_{\langle l,u\rangle}$, or a UT modality of the form $x\mathsf{U}_{I,\#b\leq n}y$, then Lemma 6 synthesizes a formula $\zeta_i \in \mathsf{MTL}$ over Σ_i such that $T_i \equiv \exists X_i.\zeta_i$.
   - If T_i is a temporal projection containing a UT modality of the form $x\mathsf{U}_{I,\#b\geq n}y$, Lemma 7 gives $\zeta_i \in \mathsf{MTL}$ over Σ_i such that $ONF_\Sigma(T_i) \equiv \exists \downarrow X_i.\zeta_i$.
3. *Putting it all together*: The formula $\zeta = \bigwedge_{i=1}^{k} \zeta_i \in \mathsf{MTL}$ is such that

$$\bigwedge_{i=1}^{k} ONF_\Sigma(T_i) \equiv \exists \downarrow X. \bigwedge_{i=1}^{k} \zeta_i \text{ where } X = \bigcup_{i=1}^{k} X_i.$$

Lemma 6. *1. Consider a temporal definition $T = \Box^{\mathsf{w}}[a \leftrightarrow \mathsf{C}^{\geq n}_{[l,u)}b]$, built from $\Sigma \cup W$. Then we synthesize a formula $\zeta \in \mathsf{MTL}$ over $\Sigma \cup W \cup X$ such that $T \equiv \exists X.\zeta$.*
2. Consider a temporal definition $T = \Box^{\mathsf{w}}[a \leftrightarrow x\mathsf{U}_{I,\#b\leq n}y]$, built from $\Sigma \cup W$. Then we synthesize a formula $\zeta \in \mathsf{MTL}$ over $\Sigma \cup W \cup X$ such that $T \equiv \exists X.\zeta$.

Proof. 1. Lets consider intervals of the form $[l, u)$. Our proof extends to all intervals $\langle l, u\rangle$. Consider $T = \Box^{\mathsf{w}}[a \leftrightarrow \mathsf{C}^{\geq n}_{[l,u)}b]$. Let $\oplus$ denote addition modulo $n + 1$.

Intuitively, we run a global modulo $n+1$ counter B (encoded using propositional variables $b_0, \ldots, b_n \in X$) which is initialized to 0 and incremented modulo $n+1$ at every position in timed word where b occurs, else it remains same. This is enforced by (a) and (b) below. In any interval I, there are at least n bs iff counter takes all the values of the set $\{0, \ldots, n\}$ in interval I. This is checked in (c) below.

(a) *Construction of a* $(\Sigma \cup W, X)$- *simple extension.* We introduce a fresh set of propositions $X = \{b_0, b_1, \ldots, b_n\}$ and construct a simple extension $\rho' = (\sigma', \tau')$ from $\rho = (\sigma, \tau)$ as follows:
   - $C1$: $\sigma'_1 = \sigma_1 \cup \{b_0\}$. If $b_k \in \sigma'_i$ and if $b \in \sigma_{i+1}$, $\sigma'_{i+1} = \sigma_{i+1} \cup \{b_{k\oplus 1}\}$.
   - $C2$: If $b_k \in \sigma'_i$ and $b \notin \sigma_{i+1}$, then $\sigma'_{i+1} = \sigma_{i+1} \cup \{b_k\}$.
   - $C3$: σ'_i has exactly one symbol from X for all $1 \leq i \leq |dom(\rho)|$.

(b) *Formula specifying the above behaviour.* The variables in X help in counting the number of b's in ρ. $C1$ and $C2$ are written in MTL as follows:
   - $\delta_1 = \bigwedge_{k=0}^{n} \Box^{\mathsf{w}}[(\mathsf{O}b \wedge b_k) \rightarrow \mathsf{O}b_{k\oplus 1}]$ and $\delta_2 = \bigwedge_{k=0}^{n} \Box^{\mathsf{w}}[(\mathsf{O}\neg b \wedge b_k) \rightarrow \mathsf{O}b_k]$

(c) *Marking the witness* 'a' *correctly at points satisfying* $\mathsf{C}^{\geq n}_{[l,u)}b$. The index i of b_i at a chosen point gives the number of b's seen so far since the previous occurrence of b_0. From a point i, if the interval $[t_i+l, t_i+u)$ has k elements of X, then there must be k b's in $[t_i+l, t_i+u)$. To mark the witness a appropriately, we need to check the number of times b occurs in $[t_i+l, t_i+u]$ from the current point i. A point $i \in dom(\rho')$ is marked with witness a iff all variables of X are present in $[t_i+l, t_i+u)$, as explained in MTL by $\kappa = \Box^{\mathsf{w}}[a \leftrightarrow (\bigwedge_{k=1}^{n} \Diamond_{[l,u)} b_k)]$.

   $\zeta = \delta_1 \wedge \delta_2 \wedge \kappa$ in MTL is equisatisfiable to T modulo simple projections.

2. The proof is similar to the above, details are in [6].

Lemma 7. *Consider a temporal definition* $T = \Box^{\mathsf{w}}[a \leftrightarrow x\mathsf{U}_{I,\#b\geq n}y]$*, built from* $\Sigma \cup W$*. Then we synthesize a formula* $\psi \in \mathsf{MTL}$ *over* $\Sigma \cup W \cup X$ *such that* $ONF_\Sigma(T) \equiv \exists \downarrow X.\psi$ *where* $ONF_\Sigma(T)$ *is* T *relativized with respect to* Σ*.*

Proof. If I is of the form $\langle l, \infty)$, then $x\mathsf{U}_{\langle l,\infty),\#b\geq n}y \equiv x\mathsf{U}_{\langle l,\infty)}y \wedge x\mathsf{U}_{\#b\geq n}y$. The untimed threshold formula $x\mathsf{U}_{\#b\geq n}y$ can be straightforwardly rewritten in LTL (see [7]).

The case of bounded intervals is the most complex and requires oversampling. Timed word ρ is oversampled at every integer valued time point upto the maximum time stamp in ρ to give ρ'. All integer points are marked by a unique proposition from $C = \{c_0, \ldots c_u\}$ in a circular fashion with first point being marked as c_0. Each c_i is associated with a bounded counter B^i , which saturates at maximum value n (encoded using propositions $b^i_0, \ldots, b^i_n \in X$). This counter is reset to 0 at each occurrence of c_i and it is incremented each time a b occurs. This encoding is explained in (O1) and (O2) below and enforced by formula η given below. Let the resultant word after the markings be ρ'.

Now, by semantics, $\rho', j \models x\mathsf{U}_{I,\#b\geq n}y$ iff for some $p > j$, $\rho', p \models y$, and x holds invariantly between j and p, and $\#b(\rho'[j,p]) \geq n$, i.e. number of times b holds between j and p is $\geq n$. This happens iff

- either (case i) there is a nearest previous integer point to p called α (marked c_i for some i) and there exist integers $g, h : 0 \leq g, h \leq n$ such that $\#b(\rho'[\alpha, p] \geq g$, and $\#b(\rho'[j, \alpha,] \geq h$ and $g + h \geq n$. This happens iff for some i and $0 \leq g, h \leq n$, we have $\rho', j \models x\mathsf{U}_I(y \wedge B^i = g)$ and $\rho', j \models x \wedge \neg c_i \mathsf{U}_{\#b \geq max(0,h)} c_i$ and $g + h \geq n$. This is encoded by the formula δ below.
- or, (case ii) j and p both lie inside a unit interval bounded by some $[c_{i-1}, c_i]$. In this case, untimed LTL formula $\rho', j \models (x \wedge (\neg \vee c_i))\mathsf{U}_{\#b=n} y)$ holds. Note that the formulae are not yet relativized, the relativized ones with details are given below.

(1) We consider the case when the interval I is bounded and left closed right open of the form $[l, u)$. Our reduction below can be adapted to other kinds of bounded intervals. Let $L = u-l$. Define $s \boxplus t = min(s+t, n)$, and $s \oplus t = (s+t) \bmod (u+1)$.

- $O1$**:** $C = \{c_0, c_1, \ldots, c_u\}$. A point i of ρ is marked c_g iff $t_i \bmod u = g$. In the absence of such a point i (such that t_i is an integer value $k < t_{|dom(\rho)|}$), we add a new point i to $dom(\rho)$ with time stamp t'_i and mark it with c_g iff $t'_i \bmod u = g$. Let $\rho_c = (\sigma^c, \tau^c)$ denote the word obtained from ρ after this marking. $O1$ can be easily coded in MTL (say η_1).
- $O2$**:** $B = \cup_{i=0}^{u} B^i$, where $B^i = \{b_0^i, b_1^i, \ldots b_n^i\}$. All the points of ρ_c marked c_i are marked as b_0^i. Let p, q be two integer points such that p is marked c_i, q is marked $c_{i \oplus L}$, and no point between p, q is marked $c_{i \oplus L}$. p, q are L apart from each other. Let $p < r < q$ be such that $b_g^i \in \sigma_r^c$ for some g. If $c_{i \oplus L} \notin \sigma_{r+1}^c$ and $b \in \sigma_{r+1}^c$, then the point $r + 1$ is marked $b_{g \boxplus 1}^i$. If $c_{i \oplus L}, b \notin \sigma_{r+1}^c$, then the point $r + 1$ is marked b_g^i. Each B^i is a set of counters which are reset at c_i and counts the number of occurrences of b up to the threshold n between a c_i and the next occurrence of $c_{i \oplus L}$. Starting at a point marked c_i with counter b_0^i, the counter increments up to n on encountering a b, until the next $c_{i \oplus L}$. Further, we ensure that the counter B^i does not appear anywhere from $c_{i \oplus L}$ to the next c_i. Let the resultant word be ρ_b. Let ρ' be the word obtained after all the markings.

(2) *Formula for specifying above behaviour*. We give following MTL formulae to specify $O2$. $\eta_2 = \bigwedge_{i=0}^{u} (\kappa_{2i}(1) \wedge \kappa_{2i}(2) \wedge \kappa_{2i}(3))$ encodes $O2$ where

$\eta_{2i}(1) = \Box^{\mathsf{w}}(c_i \rightarrow b_0^i) \wedge \bigwedge_{k=0}^{n} \Box^{\mathsf{w}}[(\mathsf{O}(b \wedge \neg c_{i \oplus L}) \wedge b_k^i) \rightarrow \mathsf{O} b_{k \boxplus 1}^i]$,

$\eta_{2i}(2) = \bigwedge_{k=0}^{n} \Box^{\mathsf{w}}[(\mathsf{O}(\neg b \wedge \neg c_{i \oplus L}) \wedge b_k^i) \rightarrow \mathsf{O} b_k^i]$ and

$\eta_{2i}(3) = \bigwedge_{i=0}^{u} \Box^{\mathsf{w}}[c_{i \oplus L} \rightarrow (\neg c_i \wedge \neg b^i)\mathsf{U} c_i]$, where $b^i = \bigvee_{k=0}^{u} b_k^i$. Let $\eta = \eta_1 \wedge \eta_2$.

(3) *Marking the witness 'a' correctly at points satisfying* $x\mathsf{U}_{I,\#b \geq n} y$. As explained above there are two possible cases: Let $act = \bigvee(\Sigma \cup W)$.

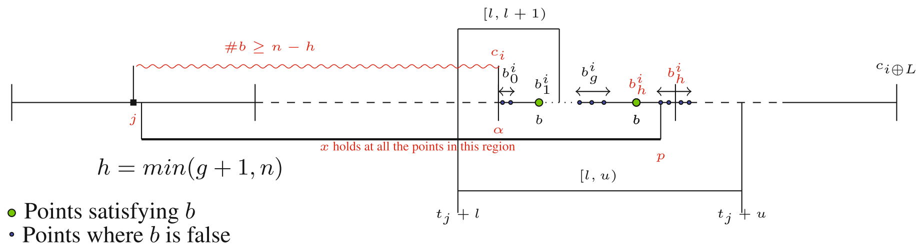

Fig. 3. Illustration of point j, p and the point α such that $t_\alpha = \lceil t_j + l \rceil$. α is marked with some c_i since it is an integer time point. The counting of b's is reset at c_i, starting with b_0^i, and continues till $c_{i \oplus L}$. h is the count of b's between α and p. To satisfy $\Box(a \leftrightarrow x\mathsf{U}_{[l,u),\#b \geq n} y)$ at j, we check that the number of b's between j and α is $\geq n-h$ when b is not true at α, and is $\geq n-h-1$ when b is true at α.

– Case(i): The formula is $\delta = \bigvee_{0 \leq g,h \leq n \wedge g+h \geq n} \delta_1(g) \wedge \delta_2(h)$ where
  - $\delta_1(g) = (x \wedge \neg act)\mathsf{U}_I(y \wedge b_g)$ and
  - $\delta_2(h) = [\{(\neg act \vee x) \wedge \neg c_i\}\mathsf{U}_{\#b \geq h-1}(c_i \wedge b)] \vee [\{(\neg act \vee x) \wedge \neg c_i\}\mathsf{U}_{\#b \geq h}(c_i \wedge \neg b)]$ for $h > 0$ and $\delta_2(0) = true$

– Case(ii): This formula takes care of the case when timestamps of points j and p have the same integral part. Thus $\lambda = (x \vee \neg act) \wedge \neg c)\mathsf{U}_{\#_b \geq n} y$, where $c = \bigvee c_k$ captures the required condition. Note that this case is only applicable when the intervals are of the $\langle 0, u \rangle$. For other type of intervals we omit this case.

The formula $\zeta = \Box^{\mathsf{w}}(a \leftrightarrow \delta \vee \lambda)$ specifies the marking of witness correctly in ρ'.

Thus we obtain the MTL formula $\psi = \eta \wedge \zeta$. Thus, by above construction, we get a formula ψ such that $ONF_\Sigma(T) \equiv \exists \downarrow X.\psi$, where $X = C \cup B$.

5 Discussion and Related Work

Within temporal and real time logics, the notion of counting has attracted considerable interest. Laroussini *et al.* extended untimed LTL with threshold counting constrained until operator. They showed that the expressiveness of LTL is not increased by adding threshold counting but the logic becomes exponentially more succinct. Hirshfeld and Rabinovich introduced $\mathsf{C}_{(0,1)}$ operator in continuous timed QTL and showed that it added expressive power. They also showed that in continuous time, more general $\mathsf{C}_{\langle l,u \rangle}$ operator can be expressed with just $\mathsf{C}_{(0,1)}$. Building upon this, Hunter showed that MTL with $\mathsf{C}_{(0,1)}$ operator is expressively complete w.r.t. $\mathsf{FO}[<,+1]$. Thus it can also express UT operator which is straightforwardly modelled in $\mathsf{FO}[<,+1]$.

In this paper, we have explored the case of MTL with counting operators over timed words interpreted in pointwise manner. We have shown that both C_I and UT operators add expressive power to MTL. Moreover, the two operators are independent in the sense that neither can be expressed in terms of the other and MTL. (We use prefixes C and T to denote a logic extended with C and UT operators respectively). We have also shown that $CMTL_{0,1} \subset CMTL_0 \subset CMTL_I \subset CTMTL$. Moreover, it is easy to show (see [6]) that $\mathsf{CTMTL} \subset \mathsf{TPTL}^1$. All these expressiveness results straightforwardly carry over to MTL over infinite timed words. Thus, pointwise semantics exhibits considerable complexity in expressiveness of operators as compared to continuous time semantics where all these logics are equally expressive. While this may arguably be considered a shortcoming of the pointwise models of timed behaviours, the pointwise models have superior decidability properties making them more amenable to algorithmic analysis. MTL already has undecidable satisfiability in continuous time whereas it has decidable satisfiability over finite timed words in pointwise semantics.

In this paper, we have shown that MTL extended with C and UT operators also has decidable satisfiability. The result is proved by giving an equisatisfiable reduction from CTMTL to MTL using the technique of oversampling projections. This technique was introduced earlier [4] and used to show that $\mathsf{MTL}[\mathsf{U}_I, \mathsf{S}_{\mathsf{np}}]$ with non-punctual past operator is also decidable in pointwise semantics. The current paper marks one more use of the technique of oversampling projections. A closer examination of our reduction from CTMTL to MTL shows that it can be used in presence of any other operator. Also, it does not introduce any punctual use of U_I in the reduced formula. The reduced formula is exponentially larger than the original formula (assuming binary encoding of integer constants). All this implies that $\mathsf{CTMTL}[\mathsf{U}_I, \mathsf{S}_{\mathsf{np}}]$ is also decidable over finite timed words. Thus, we significantly extend the frontier of decidable real time logics. Moreover, $\mathsf{CTMITL}[\mathsf{U}_{\mathsf{NS}}, \mathsf{S}_{\mathsf{NS}}]$ can be equisatisfiably reduced to $\mathsf{MITL}[\mathsf{U}_{\mathsf{np}}, \mathsf{S}_{\mathsf{np}}]$ and it is decidable with at most 2-$\mathsf{EXPSPACE}$ complexity. The exact complexity of satisfiability checking of CTMITL is open although $\mathsf{EXPSPACE}$ lowerbound trivially follows from MITL and counting LTL which are syntactic subsets.

In another line of work involving counting and projection, Raskin [13] extended MITL and event clock logic with ability to count by extending these logics with automaton operators and adding second order quantification. The expressiveness was shown to be that of recursive event clock automaton. These logics were able to count over the whole model rather than a particular timed interval. The resultant logic cannot specify constraints such as "within a time unit $(0, 1)$ the number of occurrences of a particular formula is k" but it can also incorporate modulo counting. Thus Raskin's logics and the CTMTL are expressively incomparable.

References

1. Etessami, K., Wilke, T.: An until hierarchy for temporal logic. In: LICS, pp. 108–117 (1996)
2. Hirshfeld, Y., Rabinovich, A.: An expressive temporal logic for real time. In: Královič, R., Urzyczyn, P. (eds.) MFCS 2006. LNCS, vol. 4162, pp. 492–504. Springer, Heidelberg (2006)

3. Hunter, P.: When is metric temporal logic expressively complete? In: CSL, pp. 380–394 (2013)
4. Krishna, S.N., Madnani, K., Pandya, P.K.: Partially punctual metric temporal logic is decidable. In: TIME, pp. 174–183 (2014)
5. Kini, D.R., Krishna, S.N., Pandya, P.K.: On construction of safety signal automata for $MITL[\mathcal{U}, \mathcal{S}]$ using temporal projections. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 225–239. Springer, Heidelberg (2011)
6. Krishna, S.N., Madnani, K., Pandya, P.K.: Metric temporal logic with counting (2015). CoRR, abs/1512.09032
7. Laroussinie, F., Meyer, A., Petonnet, E.: Counting LTL. In: TIME, pp. 51–58 (2010)
8. Madnani, K., Krishna, S.N., Pandya, P.K.: Partially punctual metric temporal logic is decidable (2014). http://arxiv.org/abs/1404.6965
9. Ouaknine, J., Worrell, J.: On the decidability of metric temporal logic. In: LICS, pp. 188–197 (2005)
10. Pandya, P.K., Shah, S.S.: On expressive powers of timed logics: comparing boundedness, non-punctuality, and deterministic freezing. In: Katoen, J.-P., König, B. (eds.) CONCUR 2011. LNCS, vol. 6901, pp. 60–75. Springer, Heidelberg (2011)
11. Prabhakar, P., D'Souza, D.: On the expressiveness of MTL with past operators. In: Asarin, E., Bouyer, P. (eds.) FORMATS 2006. LNCS, vol. 4202, pp. 322–336. Springer, Heidelberg (2006)
12. Rabinovich, A.: Complexity of metric temporal logics with counting and the Pnueli modalities. Theoret. Comput. Sci. **411**(22–24), 2331–2342 (2010)
13. Raskin, J.F.: Logics, automata and classical theories for deciding real time. Ph.D. thesis, Universite de Namur (1999)
14. Straubing, H.: Finite Automata, Formal Logic and Circuit Complexity. Birkhauser, Boston (1994)