

Technical Report CS-TR-2012-11: Monitoring Dense-Time, Continuous-Semantics, Metric Temporal Logic

Forward This document constitutes the ‘full version’ of the paper [1] that appeared in Runtime Verification 2012. It differs by the addition of appendices that provide proofs of the correctness and complexity.

There are also a number of corrections to the published version – most significantly a correction to the definition of the *until* and *since* operators. The published version of the paper implemented a definition of *since* that required that $\phi\mathcal{S}\psi \rightarrow \phi$ and neglected the requirement that $\psi \rightarrow \phi\mathcal{S}\psi$. Correspondingly, the transducer tables for *until* implemented $\phi\mathcal{U}\psi \rightarrow \phi$ and failed to implement $\psi \rightarrow \phi\mathcal{U}\psi$. The previously-implemented definitions may be valuable for the continuous semantics of metric temporal logic on dense time; The author’s preference for the continuous semantics rests on an assertion that they are more intuitive for dense time and a similar argument might be made for the previous definitions of *until* and *since*. However, operators implementing those definitions would be nothing more than a syntactic convenience since they are identical to $\phi \wedge \phi\mathcal{U}\psi$ and $\phi \wedge \phi\mathcal{S}\psi$, respectively, so their use would only serve to confuse the reader.

The treatment of the *eventually* operator from the original paper employed an initial-transition table that was only sensible in the case of the so-called instantaneous version (for which the maximum extent of its interval was zero). This is unneeded given that the ENQUEUE procedure ensures that only valid boolean signals result – regardless of the direction of the time shift of the output events – and so it has been removed. Transducer algorithms omitted from the published paper in the interest of brevity are reintroduced and the order of presentation of the transducer algorithms and tables has changed in the interest of clarity.

Monitoring Dense-Time, Continuous-Semantics, Metric Temporal Logic

Kevin Baldor^{1,2} and Jianwei Niu¹

¹ University of Texas at San Antonio, USA
{kbaldor,niu}@cs.utsa.edu

² Southwest Research Institute, San Antonio, USA

Abstract. The continuous semantics and dense time model most closely model the intuitive meaning of properties specified in metric temporal logic (MTL). To date, monitoring algorithms for MTL with dense time and continuous semantics lacked the simplicity the standard algorithms for discrete time and pointwise semantics. In this paper, we present a novel, transition-based, representation of dense-time boolean signals that lends itself to the construction of efficient monitors for safety properties defined in metric temporal logic with continuous semantics. Using this representation, we present a simple lookup-table-based algorithm for monitoring formulas consisting of arbitrarily nested MTL operators. We examine computational and space complexity of this monitoring algorithm for the past-only, restricted-future, and unrestricted-future temporal operators.

1 Introduction

Program monitoring has attracted interest as an alternative to model checking or theorem proving as these become impractical due to the size of the state space of a full program. A dynamic analysis, monitoring is limited to the detection of property violations that are actually observed in a program’s execution, and more fundamentally, to so called *safety properties*, those that can be falsified given only a finite number of program events.

Temporal logics such as linear temporal logic (LTL) [7] and computation tree logic (CTL) [3] provide an effective formal description for desired or undesired program behavior and are commonly used in monitoring applications. Each of these logics specifies constraints on the order of occurrence of events. For example, they can state that “after event p , event q must take place at some point in the future”. This is not an enforceable safety property but would become one if modified to state that q must occur within a certain period of time. To do so, these logics must be augmented with an explicit notion of time.

One such augmentation is metric temporal logic (MTL) [5]. It introduces limits on the periods of time over which a logical connective operates. For example, the property described in the preceding paragraph may be specified as $p \rightarrow \diamond_{[0,5]} q$. The subscript on the *eventually* operator (\diamond), is an interval – relative to the current time – in which q must hold in order for the statement to be true at the current time. Some notations support a number of subscript forms, but without loss of generality, we restrict our presentation to the use of intervals that may be closed or open on either end. Additionally, we admit intervals of the form $[a, a]$, though their use incurs a potential space penalty.

The runtime-verification community employs two time models for MTL: discrete and dense. Within the dense-time model, there are two semantics: point-based and continuous [8] [2]. We

concentrate on the latter, as in [2], Basin et al. assert that “Real-time logics based on a dense, interval-based time model are more natural and general than their counterparts based on a discrete or point-based model”. But in it, they present a monitoring algorithm that they describe as “conceptually simpler” for the point-based semantics than for the interval-based (continuous) semantics.

Our contribution with this paper is the introduction of a transition-based – rather than interval-based [2] – representation for the dense-time boolean signals that are a feature of the continuous semantics. With this representation, the output of all MTL connectives can be expressed as a simple lookup-table indexed by the input. Using this representation, we present a conceptually simple MTL monitoring algorithm modeled on the transducer-approach of [6] that reduces to something like the LTL-monitoring algorithm of [4] for past-only operators. We then observe the increase in space complexity of its extension to future MTL expressions.

2 Background

2.1 LTL and MTL

Logical expressions use the connectives for disjunction (\vee), *logical or*; conjunction (\wedge), *logical and*; and negation (\neg) to describe the relationship between logical statements at the current time. Linear Temporal Logic (LTL) augments them with a number of logical connectives that describe the relationship between logical expressions over time.

The past-only operator *historically* ($\Box\phi$) indicates that the expression ϕ has been true since time zero, *once* ($\Diamond\phi$) indicates that ϕ must have been true at some point in time since time zero, and *since* ($\phi \mathcal{S} \psi$) indicates that at some point in the past ψ must have been true and that ϕ must at least have been true at every point after that until the current time;

The future-only operator *henceforth* ($\Box\phi$) indicates that the expression ϕ is true now and will be true at all points in the future, *eventually* ($\Diamond\phi$) indicates that ϕ must be true at the current time or at some point in the future, and *until* ($\phi \mathcal{U} \psi$) indicates that at the current time or at some point in the future ψ must be true and that ϕ must at least have been true at every point between the current time and that point.

The semantics of LTL operate on a trace, a countably infinite sequence of truth values of atomic elements. LTL expressions are only interpreted as having a truth value at the instants of time corresponding to the elements of the trace. Beyond the order of events, the actual time at which the events plays no role in determining the truth of the LTL expressions.

The pointwise semantics of MTL are a natural extension of the semantics of LTL in that, while they augment the order constraints of LTL with true time constraints, the truth of an expression is only defined at discrete points in time. When used for monitoring, an MTL expression might be evaluated only when an input event arrives. This is more efficient than periodically re-evaluating expressions, but can lead to counter-intuitive results. In [2], Basin et al. discuss a number of such results. Perhaps most striking is that under pointwise semantics $\Diamond_{[0,1]}\Diamond_{[0,1]}\phi$ is not logically equivalent to $\Diamond_{[0,2]}\phi$. This is illustrated in the case that ϕ is true at time $\tau = 0$, and the next observation of the system takes place at time $\tau = 2$, $\Diamond_{[0,2]}\phi$ is *true* at time $\tau = 2$, but $\Diamond_{[0,1]}\Diamond_{[0,1]}\phi$ is *false* since the observations lack the ‘bridge’ at time $\tau = 1$ (for which $\Diamond_{[0,1]}\phi$ would evaluate to *true*) needed to declare $\Diamond_{[0,1]}\Diamond_{[0,1]}\phi$ to be *true* at time $\tau = 2$.

As observed by the authors of [2], adding additional sample points can restore the equivalence of these expressions at the cost of additional computation by the monitor. In a discrete-valued-time

system, this can be taken to the extreme of evaluating all expressions with each ‘clock tick’. Beyond the computational cost, this cannot be extended to the dense-time representation that best models external events for which there is no shared clock.

2.2 Continuous Semantics and Boolean Signals

Under continuous semantics, we avoid the ambiguity introduced by the selection of sample points. Loosely, the continuous semantics (MTL) assign a truth value to any expression for any point of time greater than zero. A more formal definition is presented in [2], but we will present enough here for the purpose of discussion. Essentially, the notion of a trace used in LTL and the pointwise semantics of MTL is replaced by a mapping from time $\tau \in \mathbb{R}_{\geq 0}$ to $\{true, false\}$ that the authors term a boolean signal. In their formulation, the boolean signal for the expression ϕ , denoted γ_ϕ , is the set of all points in time for which ϕ evaluates to true. In Figure 1 – expanded from the definitions given in [2] to include \mathcal{U}_I , the set of all signals in a model is written $\hat{\gamma}$, $\tau \in \mathbb{R}_{\geq 0}$ denotes the time for which the statement applies, and the subscript I is an interval on \mathbb{R} for which the operator applies.

$$\begin{aligned}
\hat{\gamma}, \tau \models p & \quad \text{iff } \tau \in \gamma_p \\
\hat{\gamma}, \tau \models \neg\phi & \quad \text{iff } \hat{\gamma}, \tau \not\models \phi \\
\hat{\gamma}, \tau \models \phi \wedge \psi & \quad \text{iff } \hat{\gamma}, \tau \models \phi \text{ and } \hat{\gamma}, \tau \models \psi \\
\hat{\gamma}, \tau \models \phi \mathcal{S}_I \psi & \quad \text{iff } \exists \tau' \in [0, \tau] \text{ such that } \tau - \tau' \in I, \hat{\gamma}, \tau' \models \psi, \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in (\tau', \tau] \text{ or}^3 \\
& \quad \exists \tau'' \in [0, \tau') \text{ such that } \tau - \tau'' \in I, \hat{\gamma}, \kappa \models \psi \forall \kappa \in [\tau'', \tau') \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau', \tau] \\
\hat{\gamma}, \tau \models \phi \mathcal{U}_I^4 \psi & \quad \text{iff } \exists \tau' \geq \tau \text{ such that } \tau' - \tau \in I, \hat{\gamma}, \tau' \models \psi, \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau, \tau') \text{ or} \\
& \quad \exists \tau'' > \tau' \text{ such that } \tau'' - \tau \in I, \hat{\gamma}, \kappa \models \psi \forall \kappa \in (\tau', \tau''] \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau, \tau']
\end{aligned}$$

Fig. 1. Continuous Semantics of MTL

3 Modeling Dense-Time Boolean Signals as Event Sequences

The dense model of time precludes a representation consisting of every value for which a boolean signal is true as there may be uncountably many such points. However, boolean signals are defined as satisfying the finite-variability condition that on any bounded interval there exists a finite number of non-overlapping intervals over which the signal is *true*. This lends itself to the representation used in [2] an at-most-countably-infinite set of non-overlapping intervals. Our representation differs in that we model boolean signals not as a series of intervals, but as a sequence of timed events denoting a transition from one truth value to another. For example, a signal described with the intervals $\{[1, 2], [3, 3], (4, 5), (5, 6)\}$ might be illustrated as



³ This clause was added to capture another boundary condition introduced by the dense time model

⁴ We introduce the *until* operator using the presentation of [2] to relate the definition of future operators in the standard definitions of the runtime-verification community.

where the higher line indicates *true* and the lower, *false*. The dots at the transition indicate whether the signal is considered *true* at the transition point. We represent this signal as the series of transitions

$$\{(\uparrow, 1), (\downarrow, 2), (\uparrow, 3), (\downarrow, 4), (\uparrow, 5), (\downarrow, 6)\}$$

This example exhausts all of the transition types required to describe boolean signals. In the following sections the additional events $(_, \tau)$ and $(\bar{_}, \tau)$ are used to indicate the lack of transition on one of the inputs of a binary operator. They are not strictly required to unambiguously describe a boolean signal, but are convenient for the implementation of the monitor.

Definition 1. A boolean signal is described by an infinite sequence of timed transitions (δ_i, τ_i) for which $\delta_0 \in \{_, \bar{_}, \uparrow, \downarrow\}$, $\delta_i \in \{\uparrow, \downarrow, \uparrow, \downarrow, \bar{_}, \bar{_}\} \forall i > 0$, and $\tau_i \in \mathbb{R}_{\geq 0}$, $\tau_0 = 0$, and $\tau_{i+1} > \tau_i \forall i \geq 0$. The timed transitions are subject to the further constraint types of adjacent transitions must agree in the sense that the incoming value of each transition must match the outgoing value of the previous transition. More formally,

$$\begin{aligned} \delta_i \in \{\bar{_}, \uparrow, \downarrow, \bar{_}\} &\rightarrow \delta_{i+1} \in \{\downarrow, \bar{_}, \bar{_}\} \text{ and} \\ \delta_i \in \{_, \downarrow, \bar{_}, \downarrow\} &\rightarrow \delta_{i+1} \in \{\uparrow, \bar{_}, \bar{_}\}. \end{aligned}$$

Definition 2. The truth of a signal γ at time τ is given by

$$\tau \in \gamma \doteq \begin{cases} \delta_k \in \{\bar{_}, \uparrow, \downarrow, \bar{_}\} & \exists k : \tau_k = \tau \\ \delta_k \in \{_, \downarrow, \bar{_}, \downarrow\} & \exists k : \tau_k < \tau \wedge (\tau_{k+1} > \tau \vee k = |\gamma|) \end{cases}$$

4 Monitor Construction

4.1 Supported Temporal Operators

Although the timed *until* and *since* are sufficient to capture MTL semantics, the treatment of their transitions is sufficiently complicated that we follow the approach of [6] and introduce timed *eventually* to enable the treatment of only the non-metric *until* and *since*. This is accomplished by exploiting the fact that timed *since* and *until* are redundant given timed *historically* (\boxminus_I), *once* (\diamond_I), *henceforth* (\square_I), and *future* (\diamond_I) since

$$\begin{aligned} \phi \mathcal{S}_{[a,b]} \psi &\leftrightarrow \boxminus_{[0,a]}(\phi \mathcal{S} \psi) \wedge \diamond_{[a,b]} \psi \text{ and} \\ \phi \mathcal{U}_{[a,b]} \psi &\leftrightarrow \square_{[0,a]}(\phi \mathcal{U} \psi) \wedge \diamond_{[a,b]} \psi \end{aligned}$$

and that \boxminus_I and \square_I are redundant since

$$\begin{aligned} \boxminus_I \phi &\leftrightarrow \neg \diamond_I \neg \phi \\ \square_I \phi &\leftrightarrow \neg \diamond_I \neg \phi \end{aligned}$$

Further, we observe that for the same input the output of $\diamond_{[a,a+\Delta]}$ and $\diamond_{[b,b+\Delta]}$ are both simply a time-shifted version of the output of $\diamond_{[0,\Delta]}$. By generalizing the intervals to support negative indices, we obtain

$$\diamond_{[a,b]}\phi \leftrightarrow \diamond_{[-b,-a]}\phi$$

As a result, we can monitor both future and past MTL using only the transducers for the operators \neg , \wedge , \mathcal{S} , \mathcal{U} , and \diamond_I , the formal definitions for which are given in Figure 2.

$$\begin{aligned}
\hat{\gamma}, \tau \models \neg\phi & \quad \text{iff} \quad \hat{\gamma}, \tau \not\models \phi \\
\hat{\gamma}, \tau \models \phi \wedge \psi & \quad \text{iff} \quad \hat{\gamma}, \tau \models \phi \text{ and } \hat{\gamma}, \tau \models \psi \\
\hat{\gamma}, \tau \models \diamond_I \phi & \quad \text{iff} \quad \exists \tau' \text{ such that } \tau' - \tau \in I, \hat{\gamma}, \tau' \models \phi \\
\hat{\gamma}, \tau \models \phi \mathcal{S} \psi & \quad \text{iff} \quad \exists \tau' \in [0, \tau] \text{ such that } \hat{\gamma}, \tau' \models \psi \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in (\tau', \tau] \text{ or} \\
& \quad \exists \tau'' \in [0, \tau') \text{ such that } \hat{\gamma}, \kappa \models \psi \forall \kappa \in [\tau'', \tau') \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau', \tau] \\
\hat{\gamma}, \tau \models \phi \mathcal{U} \psi & \quad \text{iff} \quad \exists \tau' \geq \tau \text{ such that } \hat{\gamma}, \tau' \models \psi \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in (\tau', \tau] \text{ or} \\
& \quad \exists \tau'' > \tau' \text{ such that } \hat{\gamma}, \kappa \models \psi \forall \kappa \in (\tau', \tau''] \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau, \tau']
\end{aligned}$$

Fig. 2. Semantics of Monitored MTL Connectives

4.2 Monitoring Algorithm

To monitor a formula ϕ , we begin by converting it into a parse tree. From this, we construct an array Φ consisting of one transducer for each node of the parse tree in reverse-topological-sort order. That is, for any node in the parse tree, its children appear before it in Φ . The transducers maintain some operation-specific fields, but each contains at least $\langle op, inputs, Q, I_{\text{valid}} \rangle$ where op identifies the operation, $inputs$ contains a pointer to the elements of Φ upon which it depends, Q is a queue containing the output of the transducer, and I_{valid} indicates the time interval over which the output of the transducer is valid.

The valid interval allows a transducer to ‘stop time’ while its state is undetermined. For example, the transducer for metric *eventually* uses this to apply a constant offset to all output transitions, whereas the *until* transducer may delay its output for an indeterminate period. Even the simple transducers such as negation and conjunction must be able to specify a valid interval since their input might do so.

Some transducers define additional state variables in addition to those mentioned above. The *since* transducer maintains a state indicating whether or not the latest transition left its output in the UP state; The *eventually* transducer maintains a timer that is used by the monitoring algorithm to call UPDATE again at some point in the future.

The monitoring procedure consists of gathering an ensemble of simultaneous transition events for the external inputs to the monitor and storing them in a container, Δ , that maps input variable names to transition types. If a timer expires, UPDATE may be called with no input transitions. The UPDATE function returns the next timer expiration time so that the monitor may call it when it

has expired. The following pseudocode describes the most general version of the update operation and the UPDATE procedure. Subsequent sections will describe the simplifications that are possible when supporting subsets of MTL.

```

1 function UPDATE( $\Phi$ ,  $\Delta$ ,  $\tau$ )
2   for  $\varphi \in \Phi$  do
3     if  $\varphi.op \in \text{input variables}$  then
4       if  $\varphi.op.id \in \Delta$  then
5         ENQUEUE( $\varphi_Q, \Delta[\varphi.op.id], \tau$ )
6          $\varphi.I_{\text{valid}} \leftarrow [0, \tau]$ 
7       else
8         while  $(\delta^*, \tau) \leftarrow \text{SYNC}(\varphi.inputs, \varphi.\tau_{\text{timer}})$  do
9           UPDATE $_{\varphi.op}$ ( $\varphi, \delta^*, \tau$ )
10      UPDATEVALIDINTERVAL $_{\varphi.op}$ ( $\varphi$ )
11  return  $\min(\{\varphi.\tau_{\text{timer}} \text{ for } \varphi \in \Phi\})$ 

```

The return value indicates the time at which the earliest timer will expire. The use of this is dependent upon the resources available to the monitor. For a real-time monitor, it will generally be used to schedule an operating-system or hardware timer. For an offline monitor, there may be no need to make use of the timer until the end of the input sequence at which point either each timer can be expired individually or a final UPDATE call can be made at a timestamp later than the latest timer.

The update procedure depends on five sub-procedures: The ENQUEUE, DEQUEUE, SYNC, and two operation-specific procedures, UPDATE $_{\varphi.op}$ and UPDATEVALIDINTERVAL $_{\varphi.op}(\varphi)$. The first and second are related to delivering input to the transducers that conforms to the rules of Definition 1 and that binary operators receive their input events in the correct order. The second two – actually a family of procedures, one for each transducer type – are used to produce the output of each transducer.

The ENQUEUE procedure supports the ‘time-shifting’ property of the *eventually* operator by ensuring that its output obeys the initial-transition properties of Definition 1. Lines 4 through 8 handle the case for which the time-shift produces a transition with timestamp less than zero, its potential initial transition is set to the constant value ($\bar{_}$ or $\bar{_}$) corresponding to the value of the signal immediately after the transition. In this case, the valid interval on the output will end at a timestamp less than zero, so this process can be repeated until the time-shift of the transducer is no longer able to produce transitions with timestamp less than zero. Similarly, if input arrives with a timestamp of zero, lines 10 through 18 ensure that the initial transition of the output is of a legal initial-transition type for which the truth value agrees with that of the input type for $\tau = 0$ and immediately afterward – that is, before the next transition event. Finally, if the time shift produces a first transition with timestamp greater than zero, lines 20 through 27 guarantee that there will be a constant-valued transition ($_$ or $\bar{_}$) matched to the current transition type being enqueued with timestamp zero added before adding the current transition. Note that if the first transition type is either $_$ or $\bar{_}$, only the zero-timestamped transition will be added because those transition types are only legal for the initial transition.

```

1 function ENQUEUE( $Q, \delta, \tau$ )
2   if  $\delta = \emptyset$  then return
3   case  $\tau < 0$ 
4     Clear( $Q$ )
5     if  $\delta \in \{\downarrow, \uparrow, \bar{\uparrow}\}$  then
6       APPEND( $Q, (\_, 0)$ )
7     else
8       APPEND( $Q, (\_, 0)$ )
9   case  $\tau = 0$ 
10    Clear( $Q$ )
11    case  $\delta \in \{\_, \bar{\_}\}$ 
12      APPEND( $Q, (\delta, 0)$ )
13    case  $\delta = \bar{\uparrow}$ 
14      APPEND( $Q, (\_, 0)$ )
15    case  $\delta \in \{\bar{\uparrow}, \uparrow\}$ 
16      APPEND( $Q, (\uparrow, 0)$ )
17    case  $\delta \in \{\downarrow, \bar{\downarrow}\}$ 
18      APPEND( $Q, (\bar{\downarrow}, 0)$ )
19    case  $\delta = \bar{\downarrow}$ 
20      APPEND( $Q, (\_, 0)$ )
21  case  $\tau > 0$ 
22    if EMPTY( $Q$ )  $\wedge Q.value = \emptyset$  then
23      if  $\delta \in \{\_, \bar{\_}\}$  then
24        APPEND( $Q, (\delta, 0)$ )
25      else if  $\delta \in \{\downarrow, \uparrow, \bar{\uparrow}\}$  then
26        APPEND( $Q, (\_, 0)$ )
27      else
28        APPEND( $Q, (\_, 0)$ )
29    APPEND( $Q, (\delta, \tau)$ )

```

Whereas the ENQUEUE procedure ensures that the initial conditions of the output of each stage of the transducer are correct, the DEQUEUE ensures that the output boolean signal will return the correct constant value if queried at a time for which there is no transition event even if all transitions have been consumed from its queue. This is simply a matter of storing the constant *true* or *false* transition depending on the truth value that will immediately follow the current transition. The simplicity of the DEQUEUE procedure depends on the SYNC procedure only calling it at the time of the next transition in its queue.

```

1 function DEQUEUE( $Q$ )
2    $(\delta, \tau) \leftarrow$  REMOVEFIRST( $Q$ )
3   case  $\delta \in \{\bar{\_}, \bar{\uparrow}, \uparrow, \bar{\downarrow}\}$ 
4      $Q.value \leftarrow \bar{\_}$ 
5   case  $\delta \in \{\_, \bar{\downarrow}, \downarrow, \bar{\uparrow}\}$ 
6      $Q.value \leftarrow \_$ 
7   return  $(\delta, \tau)$ 

```


The SYNC procedure handles the synchronization of multiple inputs of a single transducer stage. This allows the transducer UPDATE procedures to depend upon always receiving one transition on each of its inputs – even if those transitions are actually the constant values \top or \perp . As presented here, it allows the UPDATE procedure to treat unary and binary transducers identically. The *eventually* transducer is the only one for which the timer must be considered, so it is handled in the unary-transducer section. Timer expiration is effectively treated like a special case of a binary operator. If the time expires before the next input transition, the SYNC procedure returns the constant-value transition and the timestamp of the timer expiration. Otherwise, it returns the type and timestamp of the first input transition that is within the valid interval of the input signal.

```

1 function SYNC(inputs,  $\tau_{\text{timer}}$ )
2    $\varphi, \psi \leftarrow$  inputs
3   if  $\psi = \emptyset$  then ▷ one input
4     if  $\text{EMPTY}(\varphi.Q) \vee \text{HEAD}(\varphi.Q).\tau > \tau_{\text{timer}}$  then
5       if  $\tau_{\text{timer}} \in \varphi.I_{\text{valid}}$  then
6         return  $(\varphi.Q.value, \tau_{\text{timer}})$ 
7       return  $\emptyset$ 
8     return  $\text{DEQUEUE}(\varphi.Q)$ 
9   else ▷ two inputs
10     $I_{\text{valid}} \leftarrow \varphi.I_{\text{valid}} \cap \psi.I_{\text{valid}}$ 
11     $e_{\varphi} \leftarrow \text{HEAD}(\varphi.Q)$  if  $\text{HEAD}(\varphi.Q).\tau \in I_{\text{valid}}$ 
12     $e_{\psi} \leftarrow \text{HEAD}(\psi.Q)$  if  $\text{HEAD}(\psi.Q).\tau \in I_{\text{valid}}$ 
13    if  $e_{\varphi} = \emptyset \wedge e_{\psi} = \emptyset$  then
14      return  $\emptyset$ 
15    if  $e_{\varphi} \neq \emptyset \wedge e_{\psi} \neq \emptyset$  then
16      if  $e_{\varphi}.\tau = e_{\psi}.\tau$  then
17         $\text{DEQUEUE}(\varphi.Q)$ 
18         $\text{DEQUEUE}(\psi.Q)$ 
19        return  $((e_{\varphi}.\delta, e_{\psi}.\delta), e_{\varphi}.\tau)$ 
20      if  $e_{\varphi}.\tau < e_{\psi}.\tau$  then
21         $e_{\psi} \leftarrow \emptyset$ 
22      else
23         $e_{\varphi} \leftarrow \emptyset$ 
24    if  $e_{\varphi} \neq \emptyset$  then
25       $\text{DEQUEUE}(\varphi.Q)$ 
26      return  $((e_{\varphi}.\delta, \psi.Q.value), e_{\varphi}.\tau)$ 
27    else
28       $\text{DEQUEUE}(\psi.Q)$ 
29    return  $((\varphi.Q.value, e_{\psi}.\delta), e_{\varphi}.\tau)$ 

```

For binary transducers, the possible scenarios are: neither input has a transition; one input has a transition, but the other does not; both inputs have transitions, but they have different timestamps; and both inputs have transitions with the same timestamp. The SYNC procedure handles these cases in lines 10 through 12 by storing the contents of the head of each of the input queues if their respective timestamp are within the valid interval of both of the inputs.

Lines 13 and 14 handle the case in which neither has a transition. Since no transducer will produce an output transition when the truth value of neither of its inputs changes, there is no need to process further. Lines 15 through 19 handle the case for which both inputs exhibit a simultaneous transition by consuming the first transition on each input and returning the pair of transitions types. Lines 20 through 23 convert the case for which both input have a valid transition, but at different timestamps, into effectively the case for which only the input with the earliest transition contained a transition. Lines 24 through 29 handle this and the case in which only one input exhibited a transition within the valid interval. It does this by consuming one whichever of the inputs has a transition and sending the appropriate constant-value transition type ($_$ or $\bar{_}$) for the other input.

5 Signal Transducer Tables

5.1 Negation and Conjunction

Figure 3 contains the transition tables for negation and conjunction. The INIT versions are used for the first transition and the regular version for all subsequent transitions. The non-transition states of the inputs are illustrated, but table entries for which there is no transition are omitted in the interest of readability. Note that for the INIT versions of the tables, the non-transition outputs (\neg and $_$) are shown because they are actually produced for the initial entry of a boolean signal.

```

1 function UPDATE $\neg$ ( $\varphi, \delta, \tau$ )
2   if  $\tau = 0$  then
3     ENQUEUE( $\varphi.Q, \text{NOT}_{\text{INIT}}[\delta], \tau$ )
4   else
5     ENQUEUE( $\varphi.Q, \text{NOT}[\delta], \tau$ )
6    $\varphi.I_{\text{valid}} \leftarrow [0, \tau]$ 
7 function UPDATEVALIDINTERVAL $\neg$ ( $\varphi$ )
8    $\phi.I_{\text{valid}} \leftarrow \varphi.\text{inputs}.I_{\text{valid}}$ 

1 function UPDATE $\wedge$ ( $\varphi, \delta_\phi, \delta_\psi, \tau$ )
2   if  $\tau = 0$  then
3     ENQUEUE( $\varphi.Q, \text{AND}_{\text{INIT}}[\delta_\phi, \delta_\psi], \tau$ )
4   else
5     ENQUEUE( $\varphi.Q, \text{AND}[\delta_\phi, \delta_\psi], \tau$ )
6    $\varphi.I_{\text{valid}} \leftarrow [0, \tau]$ 
7 function UPDATEVALIDINTERVAL $\wedge$ ( $\varphi$ )
8    $\phi, \psi \leftarrow \varphi.\text{inputs}$ 
9    $I_{\text{valid}} \leftarrow \phi.I_{\text{valid}} \cap \psi.I_{\text{valid}}$ 

```

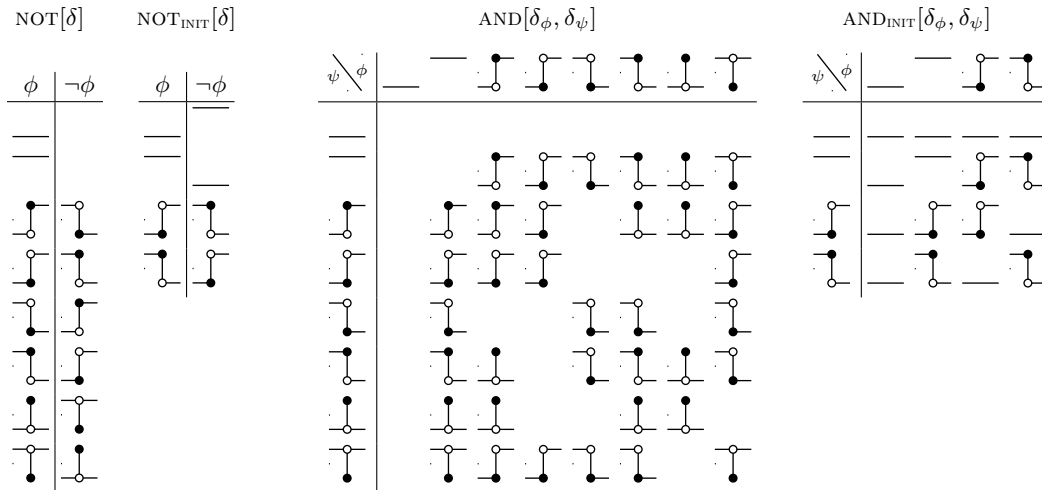


Fig. 3. Transition Tables for the Negation and Conjunction Operators

5.2 Since

The transducer for $\phi \mathcal{S} \psi$ is somewhat more complicated. The output transition for a given set of input transitions is influenced by whether or not the since operator is currently in the down or the up state. As in the previous section, omitted entries indicate states for which there is no transition. Some of the omitted entries indicate states that cannot be reached, but these are not specially indicated – nor is any special handling required to deal with unreachable states.

The transducer begins by applying the table $\text{SINCE}_{\text{INIT}}[\delta_\phi, \delta_\psi]$ for the initial transition, afterward the tables $\text{SINCE}_{\text{UP}}[\delta_\phi, \delta_\psi]$ and $\text{SINCE}_{\text{DOWN}}[\delta_\phi, \delta_\psi]$ are used. Which of the tables is to be used is determined by the type of transition last emitted. If it is in $\{\neg, \downarrow, \uparrow, \Uparrow\}$, then the UP table is used, otherwise, it is the DOWN table.

```

1 function UPDATE $\mathcal{S}$ ( $\varphi, \delta_\phi, \delta_\psi, \tau$ )
2   if  $\tau = 0$  then
3      $\delta' \leftarrow \text{SINCE}_{\text{INIT}}[\delta_\phi, \delta_\psi]$ 
4   else if  $\varphi.\text{state} = \text{UP}$  then
5      $\delta' \leftarrow \text{SINCE}_{\text{UP}}[\delta_\phi, \delta_\psi]$ 
6   else
7      $\delta' \leftarrow \text{SINCE}_{\text{DOWN}}[\delta_\phi, \delta_\psi]$ 
8   if  $\delta' \neq \emptyset$  then
9     ENQUEUE( $\varphi, Q, \delta', \tau$ )
10     $\varphi.\text{state} \leftarrow \begin{cases} \text{UP} & \delta' \in \{\neg, \downarrow, \uparrow, \Uparrow\} \\ \text{DOWN} & \text{else} \end{cases}$ 
11     $\varphi.I_{\text{valid}} \leftarrow [0, \tau]$ 
12 function UPDATEVALIDINTERVAL $\mathcal{S}$ ( $\varphi$ )
13    $\phi, \psi \leftarrow \varphi.\text{inputs}$ 
14    $I_{\text{valid}} \leftarrow \phi.I_{\text{valid}} \cap \psi.I_{\text{valid}}$ 

```




Fig. 4. Initialization Table for the *Since* Operator




Fig. 5. Transition Tables for the *Since* Operator

5.3 Eventually

The metric eventually operator uses the tables from Figure 6. Its transducer is distinguished from those introduced thus far by the addition of a timer used to generate a down event $b - a$ time units after the input transitions to the down state. All events emitted in response to an event at time τ are emitted at time $\tau - b$. Also, it can enter an indeterminate state in response to a down transition. When a down transition occurs, the transducer stores the potential down transition – the type of which is determined by the interval type – in the state variable δ_{\downarrow} . The actual transition emitted can be affected if an up transition occurs before or simultaneous with the timer expiration.

```

1 function UPDATE $_{\diamond_I}(\varphi, \delta, \tau)$ 
2    $\tau' \leftarrow \tau - b$  ▷ output time offset
3   if  $\tau = 0 \wedge \delta \in \{-, \_ \}$  then
4     ENQUEUE( $\varphi.Q, \delta, \tau'$ )
5   else if  $\varphi.\tau_{\text{timer}} = \emptyset$  then
6     if  $\delta \in \text{EVENTUALLY}_a$  then
7       ENQUEUE( $\varphi.Q, \text{EVENTUALLY}_a[\delta], \tau'$ )
8   else
9     if  $\tau = \varphi.\tau_{\text{timer}}$  then
10      if  $\delta = \emptyset$  then
11        ENQUEUE( $\varphi.Q, \varphi.\delta_{\downarrow}, \tau'$ )
12      else if  $\varphi.\delta_{\downarrow} = \_$  then
13        ENQUEUE( $\varphi.Q, \text{EVENTUALLY}_c[\delta], \tau'$ )
14       $\varphi.\delta_{\downarrow} \leftarrow \emptyset$ 
15       $\varphi.\tau_{\text{timer}} \leftarrow \emptyset$ 
16   if  $\delta \in \text{EVENTUALLY}_b$  then ▷ down transition
17      $\varphi.\delta_{\downarrow} \leftarrow \text{EVENTUALLY}_b[\delta]$ 
18      $\varphi.\tau_{\text{timer}} \leftarrow \tau + b - a$ 

```

```

1 function UPDATEVALIDINTERVAL $_{\diamond_I}(\varphi)$ 
2    $\phi \leftarrow \varphi.\text{inputs}$ 
3   switch  $\phi.I_{\text{valid}}$ 
4     case  $[0, i]$ 
5        $\varphi.I_{\text{valid}} \leftarrow [0, i - b]$ 
6     case  $[0, i)$ 
7        $\varphi.I_{\text{valid}} \leftarrow [0, i - b)$ 

```

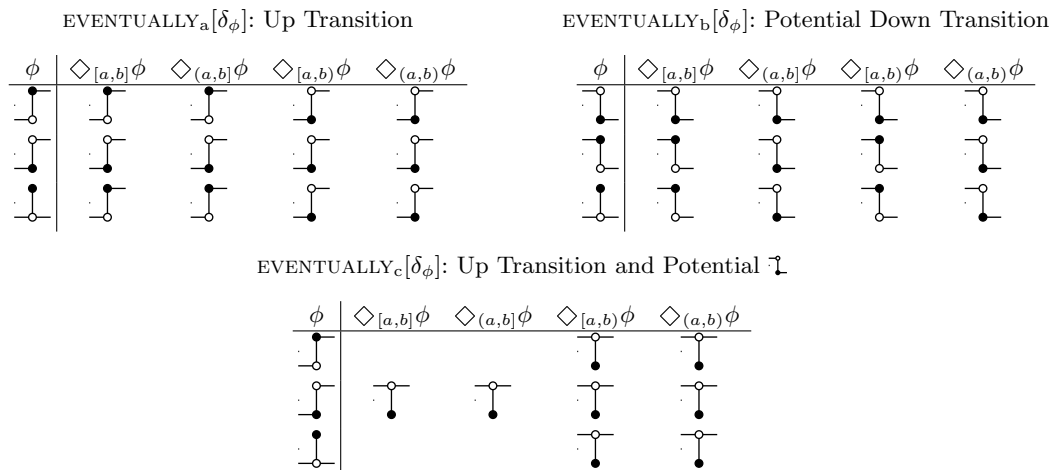


Fig. 6. Transition Tables for the *Eventually* Operator

5.4 Until

The transducer for the *until* operator can enter a more complicated indeterminate state than that of the non-metric *eventually* operator; It can maintain a different potential transition depending on whether it is ultimately found to be *true* or *false* at the point in time at which its output became uncertain. The simplest example occurs when monitoring $\phi \mathcal{U} \psi$ and ϕ becomes *true* with transition type \Downarrow at a time τ when ψ is *false*. If ψ becomes *true* before ϕ becomes *false*, then the transducer should emit the transition event (\Downarrow, τ) ; If ϕ becomes *true* before ψ becomes *false*, this potential transition is abandoned and the output remains *false* up to and including the current time. It is possible that the transducer will maintain two such potential transitions, δ_{\downarrow} or δ_{\uparrow} . We introduce the notations *non-deterministic up* \Downarrow and *non-deterministic down* \Uparrow in UNTIL_a to denote which of the potential transitions is to be emitted upon the arrival of input events.

```

1 function UPDATE $\mathcal{U}$ ( $\varphi, \delta_{\phi}, \delta_{\psi}, \tau$ )
2   if  $\tau = 0$  then
3     ENQUEUE( $\varphi.Q, \text{UNTIL}_{\text{INT}_a}[\delta_{\phi}, \delta_{\psi}], \tau$ )
4      $\varphi.\delta_{\uparrow} \leftarrow \text{UNTIL}_{\text{INT}_b}[\delta_{\phi}, \delta_{\psi}]$ 
5      $\varphi.\delta_{\downarrow} \leftarrow \text{UNTIL}_{\text{INT}_c}[\delta_{\phi}, \delta_{\psi}]$ 
6      $\varphi.\tau_{\text{pending}} \leftarrow \tau$ 
7   else
8     switch  $\text{UNTIL}_a[\delta_{\phi}, \delta_{\psi}]$ 
9        $\tau' \leftarrow \varphi.\tau_{\text{pending}}$ 
10      case  $\Downarrow$ 
11        ENQUEUE( $\varphi.Q, \varphi.\delta_{\uparrow}, \tau'$ )
12      case  $\Uparrow$ 
13        ENQUEUE( $\varphi.Q, \varphi.\delta_{\downarrow}, \tau'$ )
14      ENQUEUE( $\varphi.Q, \text{UNTIL}_b[\delta_{\phi}, \delta_{\psi}], \tau$ )
15       $\varphi.\delta_{\uparrow} \leftarrow \text{UNTIL}_c[\delta_{\phi}, \delta_{\psi}]$ 
16       $\varphi.\delta_{\downarrow} \leftarrow \text{UNTIL}_{\text{INT}_d}[\delta_{\phi}, \delta_{\psi}]$ 
17       $\varphi.\tau_{\text{pending}} \leftarrow \tau$ 
18   if  $\neg(\varphi.\delta_{\uparrow} = \varphi.\delta_{\downarrow} = \emptyset)$  then
19      $\varphi.I_{\text{valid}} \leftarrow [0, \tau)$ 

```

```

1 function UPDATEVALIDINTERVAL $\mathcal{U}$ ( $\varphi$ )
2    $\phi, \psi \leftarrow \varphi.\text{inputs}$ 
3   if  $\varphi.\delta_{\uparrow} = \varphi.\delta_{\downarrow} = \emptyset$  then
4      $\varphi.I_{\text{valid}} \leftarrow \phi.I_{\text{valid}} \cap \psi.I_{\text{valid}}$ 

```

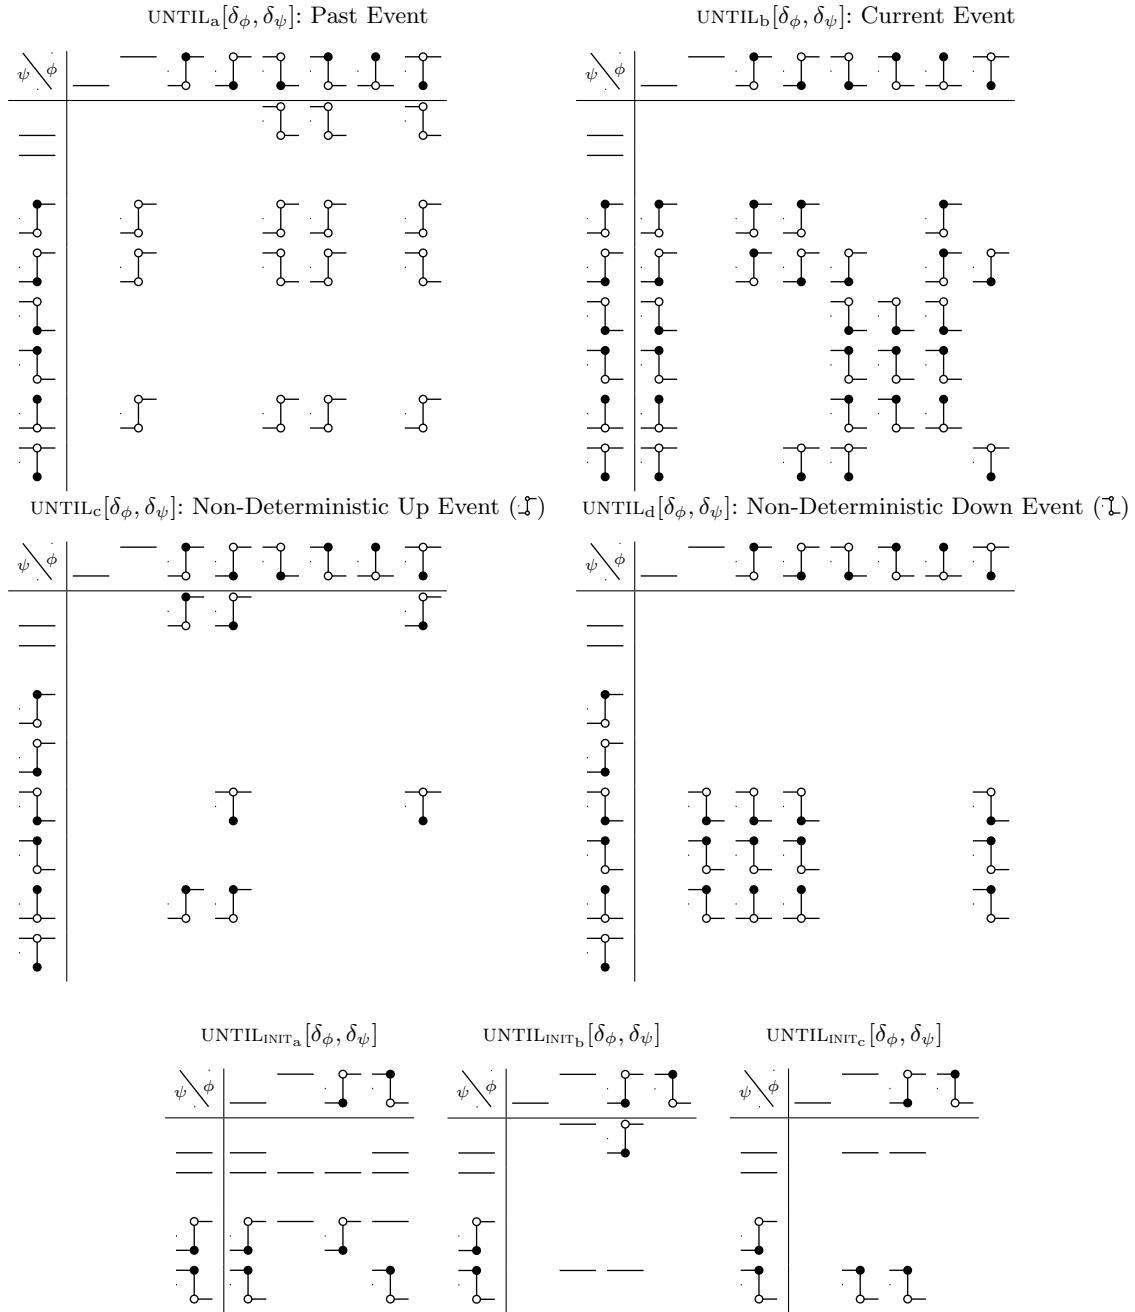


Fig. 7. Transition Tables for the *Until* Operator, $\phi \mathcal{U} \psi$

6 Correctness

Theorem 1. *The UPDATE procedure applied to the above transducer tables correctly models the semantics of Figure 1.*

Proof. See appendix A.

7 Monitoring-Algorithm Complexity

7.1 Instantaneous Transducers

The instantaneous transducers are defined as those for which all emitted transitions take place at the current time, that is, with the same τ as that of the event that caused them. They comprise \neg , \wedge , \mathcal{S} , and $\diamond_{[-a,0]}$.

The UPDATE procedure can be simplified in that there is no need to keep track of the *valid* intervals. Without the potential for delayed output, the queues will never grow larger than one element and can be replaced with single values.

Theorem 2. *The runtime to monitor expression ϕ that consists of only instantaneous operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs.*

Proof. The proof is provided in appendix B, but is reasonably clear from the pseudocode if it is given that ENQUEUE, DEQUEUE, and SYNC run in constant time.

Theorem 3. *The space required to monitor expression ϕ that consists of only instantaneous operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|)$.*

Proof. The proof is provided in appendix B and centers on a proof that the queues for the instantaneous transducers do not grow larger than one element.

7.2 Strictly Past Transducer

The transducer for the operator $\diamond_{[a,b]}$ for $a \leq b < 0$ operates identically to that of $\diamond_{[b-a,0]}$ except that the timestamp of the output that results from an event at time τ is $\tau - b$. The corresponding statement is true for other intervals with open bounds as well. This introduces the need to maintain valid ranges and queues to store the output of the intermediate stages to support operators with multiple inputs.

Theorem 4. *The runtime to monitor expression ϕ that consists of only past-time operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Theorem 5. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,b]}$ where $a < b < 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi| \lceil \frac{|b|}{b-a} \rceil)$.*

Theorem 6. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,a]}$ where $a < 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

7.3 Restricted Future

The next increment in monitor complexity introduces the metric eventually operator $\diamond_{[a,b]}$ with $b > 0$. From the UPDATE procedure, we see that it introduces a delay in its output events relative to the input events that produces them. This adds no computational complexity, but reduces the guarantees that can be made about space complexity even when $a \neq b$.

Theorem 7. *The runtime to monitor expression ϕ that consists of past-time and restricted-future operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Theorem 8. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,b]}$ where $a \leq b$ and $b > 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

7.4 Unrestricted Future

The introduction of the *until* operator presents two challenges related to the fact that it may remain in an indeterminate state for an arbitrary length of time. The unchanged space complexity belies the fact that whereas a bound may be placed on the growth of the size of the monitor for restricted-future operators if a limit can be placed on the number of transitions within any time interval, no such limit can be placed on the size of the monitor for unrestricted-future operators. Also, it is possible to construct liveness properties that can not be falsified by a monitoring procedure.

Theorem 9. *The runtime to monitor expression ϕ that consists of past-time and restricted-future operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Theorem 10. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,b]}$ where $a \leq b$ and $b > 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

8 Conclusion

We have presented straightforward procedures for monitoring dense-time continuous-semantics MTL formulae as well as the tradeoffs in runtime and space complexity incurred as the expressiveness of the supported formulae increases.

We have included the unrestricted-future operators to demonstrate support for full MTL but also because policy writers may find them to be the most natural way of representing the policy that they wish to enforce. That said, they must be used with care as they introduce the ability to describe pure liveness properties for which no truth value will ever be determined, such as $\neg\diamond\neg\diamond\phi$.

Future work may include mechanisms for trimming the boolean signals of subexpressions that cannot affect the truth of the full monitored expression as well as the augmentation to the unrestricted-future monitoring algorithm to support the extension of the valid-interval for binary operators in cases for which its output value can be determined based on only the one of its inputs for which the valid-interval extends further in time. For example, a conjunction for which one of its input is unknown beyond τ , but the other is known to be *false* over the entire interval $[\tau, \text{current time})$.

More immediately, we intend to pursue a VHDL implementation of the subset of MTL for which size restrictions can be guaranteed.

9 Acknowledgements

Jianwei Niu is supported in part by NSF award CNS-0964710 and the UTSA research award TRAC-2008.

References

1. Kevin Baldor and Jianwei Niu. Monitoring dense-time, continuous-semantics, metric temporal logic. In Shaz Qadeer and Serdar Tasiran, editors, *Runtime Verification*, volume 7687 of *Lecture Notes in Computer Science*, pages 245–259. Springer Berlin Heidelberg, 2013.
2. David Basin, Felix Klaedtke, and Eugen Zălinescu. Algorithms for monitoring real-time properties. In *Proceedings of the 2nd International Conference on Runtime Verification (RV)*, volume 7186 of *Lecture Notes in Computer Science*, pages 260–275. Springer-Verlag, 2012.
3. E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Trans. Program. Lang. Syst.*, 8(2):244–263, April 1986.
4. Klaus Havelund and Grigore Rosu. Synthesizing monitors for safety properties. In *In Tools and Algorithms for Construction and Analysis of Systems (TACAS02)*, pages 342–356. Springer, 2002.
5. Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2:255–299, 1990. 10.1007/BF01995674.
6. Oded Maler, Dejan Nickovic, and Amir Pnueli. From mtl to timed automata. In Eugene Asarin and Patricia Bouyer, editors, *Formal Modeling and Analysis of Timed Systems*, volume 4202 of *Lecture Notes in Computer Science*, pages 274–289. Springer Berlin / Heidelberg, 2006. 10.1007/11867340_20.
7. Amir Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57, 1977.
8. Pavithra Prabhakar and Deepak D’Souza. On the expressiveness of mtl with past operators. In *Proceedings of the 4th international conference on Formal Modeling and Analysis of Timed Systems, FORMATS’06*, pages 322–336, Berlin, Heidelberg, 2006. Springer-Verlag.

A Correctness

This section demonstrates that the update procedure applied to the transducer tables models the continuous semantics of MTL. The transducer approach allows this proof to be broken into a proof of correctness for each transducer alone. That is, once we have established that the update procedure presents the output of each stage to the input of subsequent stages in a timely manner, the correctness of the output follows from the correctness of each stage.

$$\hat{\gamma}, \tau \models \neg\phi \quad \text{iff} \quad \hat{\gamma}, \tau \not\models \phi \quad (1)$$

$$\hat{\gamma}, \tau \models \phi \wedge \psi \quad \text{iff} \quad \hat{\gamma}, \tau \models \phi \text{ and } \hat{\gamma}, \tau \models \psi \quad (2)$$

$$\hat{\gamma}, \tau \models \diamond_I \phi \quad \text{iff} \quad \exists \tau' \text{ such that } \tau' - \tau \in I, \hat{\gamma}, \tau' \models \phi \quad (3)$$

$$\hat{\gamma}, \tau \models \phi \mathcal{S} \psi \quad \text{iff} \quad \exists \tau' \in [0, \tau] \text{ such that } \hat{\gamma}, \tau' \models \psi \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in (\tau', \tau] \text{ or} \quad (4)$$

$$\exists \tau'' \in [0, \tau') \text{ such that } \hat{\gamma}, \kappa \models \psi \forall \kappa \in [\tau'', \tau') \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau', \tau] \quad (5)$$

$$\hat{\gamma}, \tau \models \phi \mathcal{U} \psi \quad \text{iff} \quad \exists \tau' \geq \tau \text{ such that } \hat{\gamma}, \tau' \models \psi \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in (\tau', \tau] \text{ or} \quad (6)$$

$$\exists \tau'' > \tau' \text{ such that } \hat{\gamma}, \kappa \models \psi \forall \kappa \in (\tau', \tau''] \text{ and } \hat{\gamma}, \kappa \models \phi \forall \kappa \in [\tau, \tau'] \quad (7)$$

In the following proofs, we will use the notation γ_ϕ to represent the transition-sequence model of the signal that describes the truth evaluation of the signal ϕ . To paraphrase definition 2, we have the truth evaluation of ϕ at time τ represented as $\tau \in \gamma_\phi$ when γ_ϕ contains a transition at time τ that is in $\{-, \uparrow, \downarrow\}$ or γ_ϕ does not contain a transition at time τ , but the transition immediately preceding time τ is in $\{-, \uparrow, \downarrow, \uparrow\}$.

We also will speak of transitions taking place at a specific time τ and will discuss the truth value immediately preceding, τ^- , and immediately following, τ^+ , this time. By immediately preceding, we mean any time in (τ', τ) where τ' is the time of the transition with the maximum timestamp that is less than τ . Similarly, immediately following indicates any timestamp in (τ, τ'') where τ'' is the time of the transition with the minimum timestamp that is greater than τ if such a transition exists or in (τ, ∞) if it does not. When discussing initial transitions for which $\tau = 0$, we will use 0^+ to represent the time immediately after $\tau = 0$.

Lemma 1. *The application of the tables NOT_{INIT} and NOT in the update procedure satisfies the negation operation (1).*

Proof. Taking γ_ϕ as the input, the output, $\gamma_{\neg\phi}$, of the transducer must be *true* only at those time point for which γ_ϕ is *false*.

We begin with $\tau = 0$, when γ_ϕ must contain a transition in $\{-, _, \uparrow, \downarrow\}$. Because this transition takes place at time $\tau = 0$, the table NOT_{INIT} will be applied to produce the first output transition. From Figure 3, we see that the corresponding output for each of the inputs is $\{-, _, \downarrow, \uparrow\}$. Definition 1 gives that $\gamma_{\neg\phi}, 0 \not\models \phi$. That is, at time $\tau = 0$, if the input is *true*, the output is *false*, and vice versa. We also see that any input transition that leaves the input *true* at time 0^+ leaves the output *false* until the next transition and vice versa.. Thus, we know that the transducer produces correct output for time $\tau \in [0, \tau')$, where τ' is the time of the second transition of the input.

Next, we consider subsequent transitions beginning with the first transition after $\tau = 0$. The previous analysis established that immediately prior to the next transition either γ_ϕ or $\gamma_{\neg\phi}$ is *true*,

but not both. Let us first consider the case where γ_ϕ is *false* at τ^- . From Definition 1, we have that the next transition must be in $\{\mathcal{F}, \mathcal{I}, \mathcal{L}\}$. The table NOT in Figure 3 gives the corresponding outputs as $\{\mathcal{L}, \mathcal{L}, \mathcal{F}\}$. Observe that in each case, the value at τ^- is *false* for γ_ϕ and *true* for $\gamma_{-\phi}$ and that the truth evaluation of each is the opposite of the other at the time of the transition and at times τ and τ^+ . Similarly, when γ_ϕ is *true* at τ^- , its next transition must be in $\{\mathcal{L}, \mathcal{L}, \mathcal{F}\}$ and the table NOT gives the resulting output transitions as $\{\mathcal{F}, \mathcal{I}, \mathcal{L}\}$. As before, these satisfy the negation property. In either case, we are left at τ^+ in a state such that the value of $\gamma_{-\phi}$ is the opposite of that of γ_ϕ , just as we had for time 0^+ .

Having shown that an initial transition in $\{_, _, \mathcal{F}, \mathcal{L}\}$ will be produced at time $\tau = 0$ that satisfies the negation property w.r.t. the input initial transition and that from any state for which the negation property holds a legal input transition will result in a legal output transition that satisfies the negation property, we can state that the update procedure applied to the NOT_{INIT} and NOT tables models the negation operation.

Lemma 2. *The application of the tables AND_{INIT} and AND in the update procedure satisfies the conjunction operation (2).*

Proof. Taking γ_ϕ and γ_ψ as the inputs, the output, $\gamma_{\phi \wedge \psi}$, of the transducer must be *true* only at those time point for which both γ_ϕ and γ_ψ are *true*.

We first consider the initial state. We begin by noting that the first row and first column of AND_{INIT} – those that correspond to either γ_ϕ or γ_ψ having the initial state $_$ – consist entirely of output transitions of type $_$. This is correct because if either input is *false* at all three τ^- , τ , and τ^+ then the output must also be *false* at τ^- , τ , and τ^+ . Next, we consider the second row and second column – those that correspond to either γ_ϕ or γ_ψ having the initial state $\bar{_}$. In this row or column, note that the output transition type is identical to that of the input that is not constant on that row. This is correct because if either input is *true* at all three τ^- , τ , and τ^+ then the truth – and thus transition type – is completely determined by the truth of the other input. Next note the diagonal for which γ_ϕ and γ_ψ have the same initial transition type. The output transition type in these cases is identical to that of each of the inputs. Finally, consider the case where the initial transition types of γ_ϕ and γ_ψ are $(\mathcal{F}, \mathcal{L})$ or vice versa. The output transition type for each is $_$, which is correct because in both of these cases, the value of each input is the opposite of that of the other at each time τ^- , τ , and τ^+ .

Changes in the state of either input are handled by selecting the output transition type from AND. First note that it is symmetrical about the diagonal defined by each input having the same transition type. In the following analysis, we will consider only the upper triangle, but all of the arguments will hold for the lower triangle. Also, as this is evaluated only at time $\tau > 0$, we needn't concern ourselves with the states for which both of the inputs are in $\{_, \bar{_}\}$ as the update procedure will not be called if neither input experiences a transition that results in a change to its truth value at τ or τ^+ .

To avoid tedious analysis, we assert that all entries in the table for which an output transition is provided can be seen to be correct by observing that the output transition type is determined by the truth of the input transitions at times τ^- , τ , and τ^+ . That is, the output transition will only be *true* at τ^- if both input transitions are *true* at τ^- and so on. The entries for which there is no output transition given are those for which the transition type would have been $_$, but in either case no output transition will be produced because the transition model for boolean signals forbids the redundant production of $_$ or $\bar{_}$ for any $\tau > 0$.

Lemma 3. *The application of the tables $\text{EVENTUALLY}_a[\phi]$, $\text{EVENTUALLY}_b[\phi]$, and $\text{EVENTUALLY}_c[\phi]$ in the update procedure satisfies the eventually operation (3).*

Proof. Since the metric *eventually* operator can introduce a time shift, there is less of a clear distinction between the initial transition and the subsequent transitions. The legality of the output of the transducer is guaranteed by the ENQUEUE operation on its output. However, since the initial transition can exhibit the constant-valued transition types $_$ and $\bar{_}$, lines 3 and 4 are required to handle it. Quite simply, if the input is known to be true at time zero, then the output must also be true at an offset of b time units, and vice versa. This is correct because for an output at timestamp $-b$ the interval to check for input is $[b - a, 0]$ and the input will either be *true* at all point in that interval or *false* at all points in that interval. Thus, it doesn't matter whether the interval on the metric operator is closed at the minimum or maximum extent.

Subsequent transitions are handled as follows.

The algorithm uses the fact that the timer is set or not set to determine the state of the transducer. The timer is only set when a 'down' transition, one for which the input is *false* at time τ or τ^+ , occurs. Thus, if it is set, the output must be in the *true* state, but will become *false* when the timer expires unless an 'up' transition arrives on its input before or simultaneous with the timer expiration.

We will begin with the handling of an 'up' transition on the input, one for which the input is *true* at time τ or τ^+ , when the input has been *false* for longer than the last $b - a$ time units. As the output of the transducer must be *false* at such a time, this must yield an output transition that is *true* at time $(\tau - b)^+$. The output at time $(\tau - b)$ is determined by whether the interval on the operator is closed at the maximal value and whether the input is *true* at time τ . Inspection of table EVENTUALLY_a bears this out and lines 5 and 6 ensure that it is applied only in this case.

If the input exhibits a 'down' transition it becomes possible that the output will also exhibit a 'down' transition, but if an up transition occurs within $b - a$ time units this 'down' transition will not take place. This is handled by storing the type of down transition that will take place $b - a$ (the width of the interval) time units later if no such 'up' events arrive and saving the timestamp $\tau + b - a$. Similar to the case for the 'up' transitions, the potential down transitions are characterized by being *true* at time $(\tau + b - a)^-$ and *false* at time $(\tau + b - a)^+$, but the truth at time $(\tau + b - a)$ is determined by the truth of the input at time τ and by whether the interval is closed at its minimal value. Inspection of EVENTUALLY_b bears this out.

If there is a timer set that it has not expired, the arrival of an 'up' transition will simply cancel the timer as the output will remain true for all future time unless a 'down' event arrives. This is accomplished by line 9, which guarantees that no output transition is produced, and lines 14 and 15, which cancel the timer.

Should the timer expire before the arrival of any input transitions, the stored 'down' transition is emitted with the timer-expiration timestamp shifted – as always – by b time units. This is handled by lines 9 through 11 with lines 14 and 15 canceling the timer as above.

If an 'up' input transition arrives precisely at the time of a timer expiration then we know that the output must be *true* at time $(\tau - b)^+$. If the potential down transition was of type $\bar{\downarrow}$, then this would result in the non-transition $\bar{_}$, so line 14 ensures that the update procedure emits no transition in this case. If, however, the potential down transition was of type \downarrow it is possible that the output will be *false* at time $(\tau - b)$ producing an output transition of type \uparrow . This is determined just as it was for simple 'up' transitions and table EVENTUALLY_c illustrates the input transitions and interval types for which this is the case.

Note that \downarrow is both an ‘up’ transition and a ‘down’ transition as we have defined them for this proof. Thus conditional on line 16 is checked even for algorithmic states for which only an ‘up’ transition is possible.

Lemma 4. *The application of the tables $\text{SINCE}_{\text{INIT}}$, SINCE_{UP} , and $\text{SINCE}_{\text{DOWN}}$ in the update procedure satisfies the since operation (4).*

Proof. From the initial state, we begin by observing the degenerate behavior of $\phi\mathcal{S}\psi$ – it is true at any time τ for which ψ is *true* regardless of the value of ϕ at time τ or at any point $\tau' > \tau$. Otherwise, $\phi\mathcal{S}\psi$ can only be *true* at time τ if ψ has been *true* at some time $\tau' < \tau$. As a result, the initial transition type of $\gamma_{\phi\mathcal{S}\psi}$ is equal to the initial transition of γ_{ψ} for all but two cases, those for which ψ is \downarrow and $\phi \in \{-, \uparrow\}$. In each of these cases, the output transition is of type $-$. Its truth at time τ^- and τ is given by the truth of the initial transition of ψ at those times, but the truth of the output at time τ^+ is given by the truth of the initial transition of ϕ at τ^+ .

At time $\tau = 0^+$, it is possible to be in any of the reachable states of the *since* operator: $\gamma_{\phi\mathcal{S}\psi} = _$, where $(\gamma_{\phi}, \gamma_{\psi})$ may be either $(_, _)$ or $(-, _)$, and $\gamma_{\phi\mathcal{S}\psi} = -$, where $(\gamma_{\phi}, \gamma_{\psi})$ will be in $\{(_, -), (-, -), (-, _)\}$.

For any time τ for which one or both of the inputs exhibits a transition, if $\gamma_{\phi\mathcal{S}\psi} = \textit{false}$ at τ^- , the table $\text{SINCE}_{\text{DOWN}}$ is used to determine the output transition. Conversely, if $\gamma_{\phi\mathcal{S}\psi} = \textit{true}$ at τ^- , the table $\text{SINCE}_{\text{DOWN}}$ is used to determine the output transition.

For the *false* case, since $\gamma_{\phi\mathcal{S}\psi}$ will be *true* whenever γ_{ψ} is *true*, it is impossible for γ_{ψ} to exhibit any of the transitions in $\{-, \downarrow, \uparrow, \Uparrow\}$. As such, the table $\text{SINCE}_{\text{DOWN}}$ has no entries for these transitions on ψ . Also, if $\gamma_{\psi} = _$, $\gamma_{\phi\mathcal{S}\psi}$ must also be $_$ regardless of the transition type of γ_{ϕ} because if γ_{ψ} is *false* the only way for $\gamma_{\phi\mathcal{S}\psi}$ to be *true* is for γ_{ϕ} to be *true* and have been *true* since some time in the past at which γ_{ψ} was *true*. No such time can exist at time τ or τ^+ when γ_{ψ} is $_$ at time τ . The transition type $_$ is not legal at time $\tau > 0$, so $\text{SINCE}_{\text{DOWN}}$ has no entries for the row for which the γ_{ψ} is $_$. Since $\gamma_{\phi\mathcal{S}\psi}$ must be *true* whenever γ_{ψ} is *true* but can not become *true* at any point that γ_{ψ} is *false*, if γ_{ψ} is \uparrow or \Uparrow $\gamma_{\phi\mathcal{S}\psi}$ will exhibit the same transition type. Similarly, when γ_{ψ} is \downarrow at time τ , $\gamma_{\phi\mathcal{S}\psi}$ is guaranteed to be *false* at time τ^- and *true* at time τ . However, its truth at time τ^+ will depend upon the truth of γ_{ϕ} at τ^+ because this satisfies the condition that ϕ is *true* and has been *true* since a point in time when ψ was *true*.

The case for which $\gamma_{\phi\mathcal{S}\psi}$ is *true* at τ^- uses SINCE_{UP} and is complicated enough that we will consider each transition type on γ_{ψ} separately:

For $\gamma_{\psi} = _$ we first have that γ_{ϕ} mustn't be in $\{_, -\}$ as one of the inputs must exhibit a transition. We also know that γ_{ϕ} must be *true* at τ^- , so the transitions in $\{\uparrow, \Uparrow, \downarrow\}$ are also impossible. Of those that remain, the output remains *true* until the first instant that γ_{ϕ} becomes *false*, and then remains *false* thereafter; that is, \downarrow yields \downarrow , \uparrow yields \downarrow , and \Uparrow yields \downarrow .

For $\gamma_{\psi} = -$ no transition on γ_{ϕ} can make $\gamma_{\phi\mathcal{S}\psi}$ false, so SINCE_{UP} lacks transitions for all transitions of γ_{ϕ} in this case.

For $\gamma_{\psi} = \uparrow$ the truth of γ_{ψ} at τ or τ^+ guarantees that $\gamma_{\phi\mathcal{S}\psi}$ will be *true* at those times. As we are considering cases for which $\gamma_{\phi\mathcal{S}\psi}$ is *true* at τ^- , SINCE_{UP} lacks transitions for all transitions of γ_{ϕ} in this case as well.

For $\gamma_{\psi} = \Uparrow$ The truth of γ_{ψ} at τ^+ guarantees that $\gamma_{\phi\mathcal{S}\psi}$ will be *true* at τ^+ and the fact that $\gamma_{\phi\mathcal{S}\psi}$ is *true* at τ^- means that γ_{ϕ} must be *true* at τ^- . Of the transitions on γ_{ϕ} for which this is the case, any for which γ_{ϕ} is *true* at time τ , will yield no transition on the output. The only remaining transitions on γ_{ϕ} are \downarrow and \Uparrow , both of which yield \Uparrow because of the guarantee that $\gamma_{\phi\mathcal{S}\psi}$ be *true* at τ^+ .

For $\gamma_\psi = \bar{\perp}$, $\gamma_{\phi\mathcal{S}\psi}$ will also be $\bar{\perp}$ unless the transition of γ_ϕ is *true* at time τ . If γ_ϕ is *true* at both times τ and τ^+ , as is the case of $\bar{\perp}$ and $\bar{\uparrow}$, then the output will remain *true* at both times as well and thus will produce no output transition. If γ_ϕ is *true* at time τ but not at time τ^+ , the output transition will be $\bar{\perp}$.

By a similar argument, $\gamma_\psi = \bar{\perp}$ will result in $\gamma_{\phi\mathcal{S}\psi}$ exhibiting the transition $\bar{\perp}$, except for those transitions on γ_ϕ for which it is *true* for τ^+ that result in no transition.

For $\gamma_\psi = \perp$, γ_ϕ must be *true* at a time τ^- so we can disregard all transitions for which this is not the case. The fact that γ_ψ is *true* at time τ guarantees that the output will also be *true* at time τ . Whenever γ_ϕ is *true* at time τ^+ , $\gamma_{\phi\mathcal{S}\psi}$ will be satisfied at time τ^+ . For those transitions on γ_ϕ for which this is not the case, $\bar{\perp}$ and $\bar{\perp}$, the output will be *false* at time τ^+ , resulting in an output transition of $\bar{\perp}$ in both cases.

For $\gamma_\psi = \bar{\uparrow}$, the output must be *true* at both time τ^- and τ^+ . For any transition on γ_ϕ for which ϕ is *true* at time τ , the output will also be *true* at time τ resulting in no transition on the output. For all other transitions on γ_ϕ the output will be *false* at time τ , yielding an output transition of type $\bar{\uparrow}$.

Lemma 5. *The application of the tables $\text{UNTIL}_{\text{INIT}_{a-c}}$ and UNTIL_{a-d} in the update procedure satisfies the until operation (5).*

Proof. The *until* transducer will enter an indeterminate state whenever γ_ψ is *true* but γ_ϕ is *false*. The transition type, if any, emitted at the time at which this first becomes true is determined by whether or not γ_ψ becomes *true* before γ_ϕ becomes *false*. Thus, the monitoring algorithm must wait before emitting any transition until it is clear which is the case. The update procedure for the transducer uses two state variables to keep track of the potential outcomes of entering this state: δ_\uparrow , for when γ_ψ becomes *true* before γ_ϕ becomes *false*, and δ_\downarrow , for when it does not. Either of these transitions could be \emptyset even when in the indeterminate state, but at least one of them will hold a potential transition – otherwise there would be no uncertainty about the transition type.

The rules for determining the transition types are as follows:

1. If γ_ψ is *true* at time τ^- , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ^- .
2. If γ_ψ is *true* at time τ , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ .
3. If γ_ψ is *true* at time τ^+ , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ^+ .
4. If γ_ϕ is *true* at time τ^- , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ^- if it is found to be true at time τ .
5. If γ_ϕ is *true* at time τ , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ if it is found to be true at time τ^+ .
6. If γ_ψ is *false* at time τ^+ and γ_ϕ is *true* at time τ^+ , $\gamma_{\phi\mathcal{U}\psi}$ is *true* at time τ^+ only if it is found to be *true* at time τ_{next}^- (i.e. immediately before the time of the next transition).
7. In all other cases, $\gamma_{\phi\mathcal{U}\psi}$ is *false*.

Base Case (Initial Transition)

For the initial transition, there is no need to consider the possibility of being in an indeterminate state before the first transition, so there are only three tables used. $\text{UNTIL}_{\text{INIT}_a}$ describes the transition type to be emitted immediately, but since it is possible that the initial transitions can result in an indeterminate state, $\text{UNTIL}_{\text{INIT}_b}$ and $\text{UNTIL}_{\text{INIT}_c}$ are used to produce δ_\uparrow and δ_\downarrow , respectively.

By rule 6, when δ_ψ is $_$, $\delta_{\phi\mathcal{U}\psi}$ will be indeterminate for $\delta_\phi \in \{\bar{\perp}, \bar{\uparrow}\}$

By rule 7, when δ_ψ is $_$, $\delta_{\phi\mathcal{U}\psi}$ will be $_$ for $\delta_\phi \in \{_, \bar{\perp}\}$

By rules 1-3, when δ_ψ is $\bar{\perp}$, $\delta_{\phi\mathcal{U}\psi}$ will be $\bar{\perp}$.

By rules 3 and 7, when δ_ψ is $\bar{\uparrow}$, $\delta_{\phi\mathcal{U}\psi}$ will be $\bar{\uparrow}$ for $\delta_\phi \in \{_, \bar{\uparrow}\}$.

By rules 3 and 5, when δ_ψ is \Downarrow , $\delta_{\phi\mathcal{U}\psi}$ will be \neg for $\delta_\phi \in \{\neg, \Downarrow\}$.

By rule 6, when δ_ψ is \Downarrow , $\delta_{\phi\mathcal{U}\psi}$ will be indeterminate for $\delta_\phi \in \{\neg, \Downarrow\}$

By rules 1,2, and 7, when δ_ψ is \Downarrow , $\delta_{\phi\mathcal{U}\psi}$ will be \Downarrow for $\delta_\phi \in \{\neg, \Downarrow\}$

Initial indeterminate states

By rules 4 through 6, when δ_ψ is $_$ and δ_ϕ is \neg , $\delta_{\phi\mathcal{U}\psi}$ will be \neg if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *true* at time τ_{next}^- .

By rule 7, when δ_ψ is $_$ and δ_ϕ is \neg , $\delta_{\phi\mathcal{U}\psi}$ will be $_$ if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *false* at time τ_{next}^- .

By rules 7 and 6, when δ_ψ is $_$ and δ_ϕ is \Downarrow , $\delta_{\phi\mathcal{U}\psi}$ will be \Downarrow if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *true* at time τ_{next}^- .

By rule 7, when δ_ψ is $_$ and δ_ϕ is \Downarrow , $\delta_{\phi\mathcal{U}\psi}$ will be \Downarrow if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *false* at time τ_{next}^- .

By rules 1,2, and 6, when δ_ψ is \Downarrow and δ_ϕ is in $\{\neg, \Downarrow\}$, $\delta_{\phi\mathcal{U}\psi}$ will be \neg if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *true* at time τ_{next}^- .

By rules 1,2, and 6, when δ_ψ is \Downarrow and δ_ϕ is in $\{\neg, \Downarrow\}$, $\delta_{\phi\mathcal{U}\psi}$ will be \Downarrow if $\gamma_{\phi\mathcal{U}\psi}$ is found to be *false* at time τ_{next}^- .

Inductive Step (Next Transition)

Resolution of indeterminate transition

Table UNTIL_a is used to determine which, if either, of the nondeterministic transitions was actually correct. By rule 6, the state of the prior transition can only be indeterminate if, for previous transition time τ_{prev}^+ , γ_ϕ is *true* at τ_{prev}^+ and γ_ψ is *false* at τ_{prev}^+ . As such, the only possible transitions that could resolve the indeterminate state are those for which γ_ϕ is *true* at τ_{prev}^+ and γ_ψ is *false* at τ^- .

When δ_ψ is $_$, none of the conditions in rules 1 or 4 apply for any of the three of the applicable transitions on γ_ψ so all result in the nondeterministic-down transition being selected.

When δ_ψ is \Downarrow , rule 4 is satisfied for all applicable transitions on γ_ψ so all result in the nondeterministic-up transition being selected.

When δ_ψ is \Downarrow , rule 5 – and therefore 4 – is satisfied for the transitions \neg and \Downarrow on γ_ψ so both result in the nondeterministic-up transition being selected, but for transitions \Downarrow and \Downarrow , neither of the rules 1, 4 applies so both result in the nondeterministic-down transition being selected.

When δ_ψ is \Downarrow , rule 4 is satisfied for all applicable transitions on γ_ψ so all result in the nondeterministic-up transition being selected.

Production of determinate current transition

When δ_ψ is $_$, no rule provides a deterministic transition.

When δ_ψ is \neg , rules 1 through 3 guarantee that $\gamma_{\phi\mathcal{U}\psi}$ is true at times τ^- , τ , and τ^+ . As \neg is only legal for an initial transition, no transition is produced.

When δ_ψ is \Downarrow , rules 2 and 3 guarantee that $\gamma_{\phi\mathcal{U}\psi}$ will be *true* at time τ and τ^+ . For transitions in which δ_ϕ is *false* at time τ^- , rule 10 applies for time τ^- , yielding a transition of type \Downarrow . For the other cases, the non-deterministic up will have been produced at time τ_{prev} and the current time will not produce any transition as the \neg transition is illegal for any but the initial transition.

When δ_ψ is \Downarrow , rule 3 guarantees that $\gamma_{\phi\mathcal{U}\psi}$ will be *true* at time τ^+ . If δ_ϕ is *false* at time τ , rule 10 applies at times τ^- and τ , yielding a transition of type \Downarrow . For those cases in which δ_ϕ is *false* at τ^- , but *true* at τ , rule 5 results in a transition of type \Downarrow . For the remaining cases, δ_ϕ is *true* at

both times τ^- and τ , these will both have selected a non-deterministic up transition at time τ_{prev} , resulting in the illegal transition $\bar{\text{---}}$.

When δ_ψ is $\bar{\text{---}}$, a determinate current transition can only be produced if δ_ϕ is *false* at τ^+ . In these cases, rules 4 through 6 can not apply, so rules 1 and 10 give that all transitions are of type $\bar{\text{---}}$.

When δ_ψ is $\bar{\text{---}}$, a determinate current transition can only be produced if δ_ϕ is *false* at τ^+ . In these cases, rules 1, 2, and 10 give that all transitions are of type $\bar{\text{---}}$.

When δ_ψ is $\bar{\text{---}}$, a determinate current transition can only be produced if δ_ϕ is *false* at τ^+ . If δ_ϕ is *true* at τ^- , rules 2, 4, and 10 yield a transition of type $\bar{\text{---}}$. Otherwise, rules 2 and 10 yield a transition of type $\bar{\text{---}}$.

Production of indeterminate transitions

Tables UNTIL_c and UNTIL_d yield the nondeterministic up and down transitions, respectively. By rule 6, the state of the current transition can only be indeterminate if γ_ϕ is *true* at τ^+ and γ_ψ is *false* at τ^+ .

When δ_ψ is $\bar{\text{---}}$, δ_ϕ must be a real transition, that is not $\bar{\text{---}}$. Of those transition types, previous arguments show that $\gamma_{\phi\mathcal{U}\psi}$ will have been found to have been *false* at τ_{prev}^+ and will thus $\delta_{\phi\mathcal{U}\psi}$ be *false* at τ^- , so only nondeterministic up transitions can be produced. By definition these will be *true* at time τ^+ . Rule 5 gives the truth of $\delta_{\phi\mathcal{U}\psi}$ as equal to that of γ_ϕ at time τ .

When δ_ψ is $\bar{\text{---}}$, rule 1 gives that $\delta_{\phi\mathcal{U}\psi}$ must be *true* at time τ^- . Since a transition of $\bar{\text{---}}$ is illegal, the only possible nondeterministic up transition is $\bar{\text{---}}$. It will be produced whenever δ_ϕ is *false* at time τ . All nondeterministic down transitions will be of type $\bar{\text{---}}$ because rule 5 can not be satisfied for a down transition.

When δ_ψ is $\bar{\text{---}}$, rules 1 and 2 give that $\delta_{\phi\mathcal{U}\psi}$ must be *true* at times τ^- and τ . As such, no legal up transition can be produced and the only possible down transition is of type $\bar{\text{---}}$.

When δ_ψ is $\bar{\text{---}}$, rule 2 gives that $\delta_{\phi\mathcal{U}\psi}$ must be *true* at time τ . If δ_ϕ were *true* at τ^- , rule 4 would give that $\gamma_{\phi\mathcal{U}\psi}$ is *true* at τ^- and yield the illegal $\bar{\text{---}}$ transition. For the other indeterminate cases, a nondeterministic up transition of type $\bar{\text{---}}$ will be produced. The nondeterministic down transitions are similarly constrained. If δ_ϕ is *true* at time τ^- , the transition type will be $\bar{\text{---}}$, otherwise it will be $\bar{\text{---}}$.

B Complexity

This section proves the runtime and space complexity of the monitoring algorithm as applied to significant subsets of MTL: instantaneous transducers, limited-future transducers, and unlimited-future transducers (full MTL). The following proofs make use of the following notations. When monitoring a property for which the formula is ϕ , $|\phi|$ indicates the number of operators in the formula. $|\gamma|$ is the number of transitions on all input signals.

We begin with a few lemmata that are valid for all subsets and are used in the proofs for the specific subsets of MTL.

Lemma 6. *The ENQUEUE procedure for runs in constant time for each input transition.*

Proof. The ENQUEUE procedure contains no loops and simply performs a number of comparison and queue operations. Given that appending to, emptying, and checking the emptiness of a queue can be implemented to execute in constant time, the ENQUEUE operation itself executes in constant time.

Lemma 7. *The DEQUEUE procedure for runs in constant time for each output transition.*

Proof. The DEQUEUE procedure contains no loops and simply extracts the first transition from the queue and stores the constant value associated with the extracted transitions truth at time τ^+ . Given that extracting the first element of a queue can be implemented to execute in constant time, the DEQUEUE operation itself executes in constant time.

Lemma 8. *The SYNC procedure for runs in constant time for each output transition.*

Proof. The SYNC procedure contains no loops. Lines 3 through 8 handle the single-input case. As the *eventually* operator has only one input, it ensures that the next transition event is before or simultaneous with the current timer. To do so, it uses a new operation HEAD which simply returns the content of the head of the queue without actually removing it. It is taken as given that this can be implemented to run in constant time. Lines 10 through 29 handle the two input case. It introduces no new operators and simply compares valid intervals to determine which, if any, input transitions should be dequeued. In each case, all operations execute in constant time, so the SYNC operation does so as well.

Lemma 9. *The update procedure for negation runs in constant time for each input transition.*

Proof. For each input transition, exactly one output transition is produced either on line 3 or 5 depending on the value of the timestamp τ . The type of each output transition is determined by a table lookup that is accomplished in constant time and supplied as an argument to an ENQUEUE operation that Lemma 6 gives as executing in constant time.

The UPDATEVALIDINTERVAL for negation simply copies the valid interval – guaranteed to be contiguous – of its input and thus executes in constant time.

Lemma 10. *The length, measured in number of transitions, of the negation of a signal is equal to that of the signal itself. $|\gamma_{\neg p}| = |\gamma_p|$*

Proof. As described in the previous proof, for each input transition, exactly one output transition is produced either on line 3 or 5 depending on the value of the timestamp τ .

Lemma 11. *The update procedure for conjunction runs in constant time for each pair of input transitions.*

Proof. For each pair of input transitions, exactly one output transition is produced either on line 3 or 5 depending on the value of the timestamp τ . The type of each output transition is determined by a table lookup that is accomplished in constant time and supplied as an argument to an ENQUEUE operation that Lemma 6 gives as executing in constant time.

The UPDATEVALIDINTERVAL for conjunction simply computes the intersection of the valid intervals of its inputs. As they are guaranteed to be contiguous, this requires only the comparison of the lower and upper limits of each and thus executes in constant time.

Lemma 12. *The length of the conjunction of two signals is less than or equal to the sum of the lengths of the input signals. $|\gamma_{p \wedge q}| \leq |\gamma_p| + |\gamma_q|$*

Proof. The longest series of transitions occurs when no two input transitions are simultaneous. In this case, each input transition results in one output transition. If any two input transitions are simultaneous, the pair will result in only one output transition.

Lemma 13. *The update procedure for the since operator runs in constant time for each pair of input transitions.*

Proof. For each pair of input transitions, exactly one output transition – or non-transition – is produced either on line 3, 5, or 7 depending on the value of the timestamp τ and the current state of the *since* transducer. The type of each output transition is determined by a table lookup that is accomplished in constant time and – if it is not null – supplied as an argument to an ENQUEUE operation that Lemma 6 gives as executing in constant time and used to determine the next state of the transducer based on the transition type.

The UPDATEVALIDINTERVAL for *since* simply computes the intersection of the valid intervals of its inputs. As they are guaranteed to be contiguous, this requires only the comparison of the lower and upper limits of each and thus executes in constant time.

Lemma 14. *the length of the output of the since operator is less than or equal to the length of the input signal for the second argument. $|\gamma_{p \mathcal{S} q}| \leq |\gamma_q|$*

Proof. This is the result of the dominance of the second argument on the *since* operation. The formula $p \mathcal{S} q$ is *true* whenever q is *true*. Thus, $|\gamma_{p \mathcal{S} q}|$ can be equal to $|\gamma_q|$. If p is *true* at all times, $\gamma_{p \mathcal{S} q}$ will also be *true* at all times and then $|\gamma_{p \mathcal{S} q}| = 1$ regardless of $|\gamma_q|$. $\gamma_{p \mathcal{S} q}$ can only become *true* at a point in time that γ_q is *true*. The truth of γ_p can only keep $\gamma_{p \mathcal{S} q}$ *true* after γ_q has become *false*, and it will only do so until γ_p becomes *false* for the first time after a down transition on γ_q . Thus, γ_p can only delay a down transition that would have resulted from a down transition on γ_q or remove a pair of transitions (down then up) on γ_q if γ_p remains *true* between them.

Lemma 15. *The update procedure for the eventually operator runs in constant time for input or timer expiration.*

Proof. For each input transition or timer expiration, At most one call to the ENQUEUE operation is made on a transition type obtained from a table lookup. In some cases, the δ_\downarrow and τ_{timer} may be updated as many as two times. As all of these operations execute in constant time, the entire update procedure does so as well.

The UPDATEVALIDINTERVAL for *eventually* simply updates its interval based on that of its input with an offset to the maximum extent. The switch simply ensures that the inclusion of the maximum extent matches that of the input interval.

Lemma 16. *the length of the output of the eventually operator is no more than twice the length of its input. $|\gamma_{\diamond_{IP}}| \leq 2|\gamma_p|$*

Proof. The introduction of the timer allows that an event can occur at a time for which there is no input. However, this can only take place after a \uparrow , \downarrow , or \downarrow transition for which no \uparrow , \downarrow , or \downarrow occurs within $|I|$ time units. If such an input transition occurs within $|I|$ time units, neither transition will produce an output transition. No \uparrow input transition ever produces an output transition. Thus, the maximum number of output events produced by an up event (\uparrow or \downarrow) followed by a down event (\downarrow or \uparrow) is two – no net gain. Only the \downarrow transition can produce more than one output transition. When it is not followed by an up transition within $|I|$ time units, it produces exactly one up transition and one down transition. As such, the worst case behavior is achieved when all input events are \downarrow transitions separated by greater than $|I|$ time units and it will be precisely twice the length of the input.

Lemma 17. *The length of the output of cascaded eventually operators is also not more than twice the length of the input.*

Proof. As described above, only the \downarrow transition can produce more than one output transition. The types of the two output transitions that it produces are exactly one transition in $\{\uparrow, \downarrow\}$ and one transition in $\{\uparrow, \downarrow\}$. Neither of these transitions will result in more than one transition in subsequent *eventually* transducers. As such the length of the output of a sequence of *eventually* operators is less than or equal to twice the length of the input to the sequence of *eventually* operators.

Lemma 18. *The runtime cost of the introduction of clocks is asymptotically linear with product of the size of the formula, $|\phi|$, and the size of the input, $|\gamma|$.*

Proof. The number of clocks is limited by the number of *eventually* operators in the formula, $|\phi_{\diamond_I}|$, as each can only maintain one clock at any given time. Each *eventually* operator can produce at most one clock per input event. Determining the next clock event is done as part of the update procedure pass over all transducers. Its runtime cost is constant with each transducer, so the cost is in $O(|\phi|)$. Thus the additional overhead associated with clocks is limited to $O(|\gamma| \cdot |\phi| + |\gamma| \cdot |\phi_{\diamond_I}|)$ which is in $O(|\gamma| \cdot |\phi|)$.

Lemma 19. *The update procedure for the until operator runs in constant time for each pair of input transitions.*

Proof. For each input transition or timer expiration, at most two calls to the ENQUEUE operation are made on a transition type obtained from a table lookup, both δ_{\uparrow} and δ_{\downarrow} are set as is the τ_{pending} . As all of these operations execute in constant time, the entire update procedure does so as well.

In the worst case, the UPDATEVALIDINTERVAL for *until* simply computes the intersection of the valid intervals of its inputs. As they are guaranteed to be contiguous, this requires only the comparison of the lower and upper limits of each and thus executes in constant time. If either δ_{\uparrow} or δ_{\downarrow} is set, even this action is not taken.

Lemma 20. *the length of the output of the until operator is less than or equal to the length of the input signal for the second argument. $|\gamma_{pUq}| \leq |\gamma_q|$*

Proof. This follows the same logic as the proof of Lemma 14. The output must be *true* at all points for which γ_q is *true* but the truth of γ_p can only serve to reduce the number of transitions on the output if it remains *true* between two intervals during which γ_q is *true*.

B.1 Instantaneous Transducers

The instantaneous operators are: negation, conjunction, since, and the metric eventually, \diamond_I for which the upper bound of the interval I is zero. They all have the property that every output transition that happens in response to an input transition at time τ will take place at time τ . This simplifies the ENQUEUE procedure because it is never necessary to store more than one transition on the output of any transducers, the one for the current time. Similarly, DEQUEUE will operate on a single value simply storing the appropriate $\bar{}$ or $_$ transition in the single storage cell for the transducer's output after removing the stored transition. Each of these operations takes constant time.

Theorem 2. *The runtime to monitor expression ϕ that consists of only instantaneous operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. The naïve update procedure presented in this paper consists of visiting each transducer and running its update procedure to determine whether it can produce output based on the input that has arrived for each input transition. Assuming the worst case input – that for which no input transitions are simultaneous – this establishes a floor of $|\phi|n$ update operations. The timers of the *eventually* operator can cause the update procedure to be called at points for which there are no input transitions. In the worse case, all of the operators in ϕ are *eventually* operators and the input and timer expirations are such that no two events are simultaneous, this gives a maximum number of transducer update procedure calls of $2|\phi|n$. As lemmas 9, 11, 13, and 15 give that the update for each transducer runs in constant time, this yields an asymptotic runtime of $O(|\phi|n)$.

Theorem 3. *The space required to monitor expression ϕ that consists of only instantaneous operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|)$.*

Proof. The monitor for any formula consists of a tree formed by transducers. The leaves of this tree represent the input signals. Assuming that we maintain a separate queue for each transducer input that comes from an external signal we have a maximum of $|\phi| + 1$ input queues – produced when all operators are binary. Each transducer output also requires a queue for a total of maximum of $2|\phi| + 1$. At any point in time τ the valid interval for all input signals will be $[0, \tau]$. For instantaneous transducers, this results in valid intervals of the transducer outputs that are also always of the form $[0, \tau]$. Inspection of the SYNC procedure shows that if all inputs and transducers have valid intervals that include the current time, then no transition on any queue with timestamp less than or equal to τ will be left on the queue pending some future event on the other input. That is, any external input transitions, with timestamp $\tau = \text{now}$, will be immediately consumed and the output of each transducer will also have a timestamp of τ and will thus be immediately consumed by subsequent transducers. Consequently, no queue will ever contain more than one transition. Thus the size of all queues must be in $O(|\phi|)$. Some transducers also maintain state information, but it is of a constant size; The *eventually* operator maintains a timer and *since* maintains a boolean state. The size of all such storage will be in $O(|\phi|)$, so the entire update procedure requires storage that is in $O(|\phi|)$.

B.2 Strictly Past Transducer

The transducer for the operator $\diamond_{[a,b]}$ for $a \leq b < 0$ operates identically to that of $\diamond_{[b-a,0]}$ except that the timestamp of the output that results from an event at time τ is $\tau - b$. The corresponding statement is true for other intervals with open bounds as well. Since b is defined as being less than zero, this output will take place at a time that is in the future. This introduces the need to maintain valid ranges and queues to store the output of the intermediate stages to support operators with multiple inputs as we may need to wait $-b$ time units.

Theorem 4. *The runtime to monitor expression ϕ that consists of only past-time operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. This follows from the proof of the runtime for instantaneous transducers.

Lemma 21. *For the output of the metric eventually operator of the form $\diamond_{[a,b]}$, the duration of any interval – other than those beginning at time zero – during which it is true must be $b - a$ time units. More formally, $\forall i : (\delta_i, \tau_i) \in \gamma_{\diamond_{[a,b]}}$, $\delta_i \in \{\downarrow, \uparrow, \Uparrow\} \rightarrow \tau_{i+1} - \tau_i \geq b - a$.*

Proof. First note that all up transitions, produced on lines 4, 7, or 13 of `UPDATEMTLFI`, are the result of an input up transition at time τ and are produced at time $\tau - b$. The only way to produce a down transition is to have set a timer in response to an input transition on lines 17 and 18. The shortest-duration interval between an up transition and its subsequent down transition will occur when the input both causes an up transition and results in the setting of a timer to produce a down transition. For such an input transition that takes place at time τ , its resulting up transition will have a timestamp of $\tau' = \tau - b$ and a timer will be set for time $\tau + b - a$. When that timer expires, and the update procedure produces a down transition, it will do so with a timestamp $\tau'' = (\tau + b - a) - b = \tau - a$. The difference between the timestamps of these events will be $\tau'' - \tau' = (\tau - a) - (\tau - b) = b - a$.

The reason that an interval beginning at time zero might be of shorter duration than $b - a$ is that for if $b > 0$, the timestamps produced by the *eventually* operator might be negative and lines 3 through 8 of the `ENQUEUE` procedure will suppress their production – effectively only emitting the portion of the *true* interval for which time timestamps are non-negative.

Lemma 22. *The output of cascaded eventually operators of the form $\diamond_{[a_1,b_1]}$ followed by $\diamond_{[a_2,b_2]}$ is equivalent to that of $\diamond_{[a_1+a_2,b_1+b_2]}$.*

Proof. From Figure 2, we have

$$\hat{\gamma}, \tau \models \diamond_I \phi \quad \text{iff} \quad \exists \tau' \text{ such that } \tau' - \tau \in I, \hat{\gamma}, \tau' \models \phi$$

which, for cascaded *eventually* operators gives

$$\hat{\gamma}, \tau \models \diamond_{[a_1,b_1]} \diamond_{[a_2,b_2]} \phi \quad \text{iff} \quad \exists \tau', \tau'' \text{ such that } \tau' - \tau \in [a_2, b_2] \text{ and } \tau'' - \tau' \in [a_1, b_1], \hat{\gamma}, \tau'' \models \phi$$

The constraints can be rewritten as

$$\begin{aligned} a_1 &\leq \tau'' - \tau' && \leq b_1 \\ a_2 &\leq \tau' - \tau && \leq b_2. \end{aligned}$$

Summing these lines yields

$$a_1 + a_2 \leq \tau'' - \tau \leq b_1 + b_2,$$

which gives

$$\begin{aligned} \hat{\gamma}, \tau \models \diamond_{[a_1, b_1]} \diamond_{[a_2, b_2]} \phi & \quad \text{iff} \quad \exists \tau'' \text{ such that } \tau'' - \tau \in [a_1 + a_2, b_1 + b_2], \hat{\gamma}, \tau'' \models \phi \\ \hat{\gamma}, \tau \models \diamond_{[a_1 + a_2, b_1 + b_2]} \phi & \quad \text{iff} \quad \exists \tau'' \text{ such that } \tau'' - \tau \in [a_1 + a_2, b_1 + b_2], \hat{\gamma}, \tau'' \models \phi \end{aligned}$$

Theorem 5. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a, b]}$ where $a < b < 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi| \lfloor \frac{|b|}{b-a} \rfloor)$.*

Proof. The increased space requirement comes from the fact that the timestamp of the output produced by the *eventually* transducer in response to an event (transition or timer expiration) at time τ will have timestamp $\tau - b$. Since b is known to be negative in this case, these events take place at a future time compared to other inputs to the monitor. Subsequent binary transducers, such as conjunction, must wait until the valid intervals of both of its inputs contain the timestamp of any input transition before it can be acted upon. For a concrete example, consider

$$p \wedge \diamond_{[a, b]} q.$$

In this case, the output queue of the *eventually* operator will have to store all transitions until the SYNC operation finds that the valid interval of the p input contains their timestamp. As the p input's valid interval will always be $[0, \tau_{\text{current}}]$. That is, it will potentially store as many transitions as the *eventually* operator can produce within $|b|$ time units. Lemma 21 gives a minimum width of $b - a$ for any interval for which the output of the *eventually* operator is *true*. The caveat about intervals beginning at time zero does not apply to the past-time version, so we know that there can be only $\lfloor \frac{|b|}{b-a} \rfloor$ such intervals. The maximum number of transitions that this could entail would be $2 \lfloor \frac{|b|}{b-a} \rfloor$, but this is not important to the asymptotic performance. The worst case storage is achieved when the number of mismatched input valid intervals is maximized which arises in the formula

$$\bigwedge_{i=0}^k p_i \wedge \diamond_{[a, b]} q_i.$$

This produces $2k$ conjunction operators and k eventually operators, thus the number of eventually operators is in $\Theta(|\phi|)$ giving a worst-case space complexity in $\Theta(|\phi| \lfloor \frac{|b|}{b-a} \rfloor)$.

Cascading *eventually* operators of the form $\diamond_{[a, b]}$ does not increase this asymptotic behavior because Lemma 22 gives that this is equivalent to a single *eventually* operator $\diamond_{[2a, 2b]}$. The space requirement for that case would be in $O(|\phi| \lfloor \frac{2|b|}{2(b-a)} \rfloor) = O(|\phi| \lfloor \frac{|b|}{b-a} \rfloor)$.

For formulas containing *eventually* operators with different intervals, it suffices to consider only the *eventually* operator for which the ratio of $|b|$ to $b - a$ is maximized. This is because the ratio for the cascaded *eventually* operators must be less than that of one of its inputs. This can be shown by contradiction. To simplify the presentation, we will introduce $n_1 = |b_1|$ and $d_1 = b_1 - a_1$ as the numerator and denominator of the ratio for the first of the *eventually* operators and n_2 and d_2 for the second. To show the contradiction, we assume that the cascaded *eventually* operators do entail a higher space complexity than either of the individual *eventually* operators alone.

$$\begin{array}{ll}
\frac{n_1}{d_1} < \frac{n_1 + n_2}{d_1 + d_2} & \frac{n_2}{d_2} < \frac{n_1 + n_2}{d_1 + d_2} \\
n_1(d_1 + d_2) < d_1(n_1 + n_2) & n_2(d_1 + d_2) < d_2(n_1 + n_2) \\
n_1d_1 + n_1d_2 < d_1n_1 + d_1n_2 & n_2d_1 + n_2d_2 < d_2n_1 + d_2n_2 \\
n_1d_2 < d_1n_2 & n_2d_1 < d_2n_1 \\
\frac{n_1}{d_1} < \frac{n_2}{d_2} & \frac{n_2}{d_2} < \frac{n_1}{d_1}
\end{array}$$

which establishes the contradiction.

Theorem 6. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,a]}$ where $a < 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. The transducer for this case is not implemented according to the algorithm described for the general case. For each input transition (δ, τ) , the corresponding output is $(\delta, \tau - a)$. A time-shifted version of the input. Again, as $a < 0$, this new input will be in the future. However, unlike the previous case, there is nothing to limit the number of transitions that can occur within those $|a|$ time units. In the worst case, a formula such as

$$\bigwedge_{i=0}^k p \wedge \diamond_{[a,a]} p$$

for which all n input takes place within $|a|$ time units. The queues will have to contain $k \cdot n$ unprocessed transitions at some point. As k grows, $|\phi|$ grows linearly, so the worst case has a space complexity of $O(|\phi|n)$.

B.3 Restricted Future

The next increment in monitor complexity introduces the metric eventually operator $\diamond_{[a,b]}$ with $b > 0$. From the UPDATE procedure, we see that it introduces a delay in its output events relative to the input events that produces them. This adds no computational complexity, but reduces the guarantees that can be made about space complexity even when $a \neq b$.

Theorem 7. *The runtime to monitor expression ϕ that consists of past-time and restricted-future operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. This follows from the proof of the runtime for instantaneous transducers.

Theorem 8. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,b]}$ where $a \leq b$ and $b > 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. With b greater than zero, the valid interval for the $\diamond_{[a,b]}$ operator will be b time units earlier than that of its input. In the expression $p \wedge \diamond_{[a,b]}q$, the queue containing the transitions of the p input will have to store all transitions that have taken place within the last b time units because of the behavior of the SYNC procedure on the inputs to the conjunction operator. There is no restriction on the number of transitions that can take place on the input p within b time units, so the queue can potentially contain all n transitions. As with the previous proof, formulas of the form

$\bigwedge_{i=0}^k p \wedge \diamond_{[a,a]}p$ can produce a worse-case space complexity of $O(|\phi|n)$.

B.4 Unrestricted Future

Theorem 9. *The runtime to monitor expression ϕ that consists of past-time and restricted-future operators on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. The introduction of the unrestricted-future expands the monitoring capability to full MTL with the introduction of the *until* operator. Lemma 20 gives that the addition of the *until* operator doesn't entail any increase in the number of output transitions over those of the input. Lemma 19 gives that each call to the update procedure for the *until* operator executes in constant time. With the addition of these, the proof follows the form of that for the instantaneous transducers.

Theorem 10. *The space required to monitor expression ϕ that consists of only past-time operators with $\diamond_{[a,b]}$ where $a \leq b$ and $b > 0$ on input signals $\hat{\gamma}$ from time zero until time τ is in $O(|\phi|n)$, for $n =$ the sum of the number of transitions on all inputs with timestamp less than or equal to τ .*

Proof. This proof follows from that of Theorem 8, but using an expression of the form $p \wedge r\mathcal{U} q$. The worst-case runtime is the same, but the inability to limit the space complexity is somewhat stronger as the valid interval of the *until* operator can have an upper limit that is an arbitrary amount of time difference from the current time. Even if there were some limit that could be placed on the number of input transitions within a fixed period, there would still be no limit on the worst-case space complexity for full MTL.