Moore-Machine Filtering for Timed and Untimed Pattern Matching

Masaki Waga and Ichiro Hasuo

Abstract—**Monitoring is an important body of techniques in runtime verification of real-time, embedded and cyber-physical systems. Mathematically, the monitoring problem can be formalized as a pattern matching problem against a pattern automaton. Motivated by the needs in embedded applications—especially the limited channel capacity between a sensor unit and a processor that monitors—we pursue the idea of *filtering* as preprocessing for monitoring. Technically, for a given pattern automaton, we present a construction of a *Moore machine* that works as a filter. The construction is automata-theoretic, and we find the use of Moore machines particularly suited for embedded applications, not only because their sequential operation is relatively cheap but also because they are amenable to hardware acceleration by dedicated circuits. We prove soundness (i.e. absence of lost matches), too. We work in two settings: in the *untimed* one, a pattern is an NFA; in the *timed* one, a pattern is a timed automaton. The extension of our untimed construction to the timed setting is technically involved, but our experiments demonstrate its practical benefits.**

Index Terms—**Monitoring, filtering, timed pattern matching, Moore machine, timed automaton, real-time system**

I. INTRODUCTION

A. Monitoring and (Timed) Pattern Matching

THE COMPLEXITY of *cyber-physical systems (CPS)* has been rapidly growing, due to increasingly advanced digital control that realizes not only enhanced efficiency (e.g. in cars' fuel consumption) but also totally new functionalities such as automatic driving. Getting those systems right is, therefore, a problem that is as important, and as challenging, as ever.

Due to such complexity of CPS, combined with other reasons such as black-box components provided by other suppliers, *formal verification* in the conventional sense is often hard to apply to real-world CPS. This has made researchers and practitioners turn to so-called *lightweight formal methods*. *Runtime verification* is a major branch therein, where execution traces of a given system is checked against a given specification [1]. Various algorithms for *monitoring* have been proposed for this purpose.

Mathematically speaking, one common formalization of the monitoring problem is as the problem of *pattern matching*.[1] In

[1]Another common formalization is as what we call the *pattern search* problem. Pattern search is easier than pattern matching, but is less informative. Their comparison is at the end of §I.

case an execution trace is given by a word $w = a_1 a_2 \ldots a_n$, the expected outcome is the set

$$
\mathrm{Match}(w, \mathsf{pat}) := \{(i, j) \mid w|_{[i,j]} \models \mathsf{pat}\} \\
(\text{where } w|_{[i,j]} = a_i a_{i+1} \ldots a_j) \tag{1}
$$

of pairs (i, j) of indices, the restriction of w to which satisfies the given pattern pat. The pattern pat can be given by a string, a set of strings, a regular expression, an automaton, etc.

Example I.1. Consider a word $w_1 = \mathsf{abb}\,\mathsf{bbb}\,\mathsf{aab}$ and a pattern given by a regular expression $\mathcal{A}_1 = \mathsf{aa}^* b$. There are three matches, and we have $\mathrm{Match}(w_1, \mathcal{A}_1) = \{(1, 2), (7, 9), (8, 9)\}$.

More interesting in the CPS context is the *timed* version of pattern matching. In one common formalization, an execution trace is given by a *timed word*—a sequence of time-stamped characters such as $w_2 = (\mathsf{a}, 0.1)(\mathsf{b}, 2.5)(\mathsf{a}, 3.5)(\mathsf{b}, 4.8)$. A pattern pat is then specified by a *timed automaton (TA)* $\mathcal{A}$ from [2], and we compute the set

$$
\mathrm{Match}(w, \mathcal{A}) := \left\{ (t, t') \in \mathbb{R}_{\geq 0}^2 \mid t < t' \text{ and } w|_{(t,t')} \in L(\mathcal{A}) \right\} \tag{2}
$$

of intervals (t, t'), the restriction of w to which is accepted by the TA $\mathcal{A}$. Unlike the untimed setting, a TA $\mathcal{A}$ allows one to express various real-time constraints, which leads to a much more refined analysis of execution traces of CPS.

Example I.2. Consider the timed word $w_2 = (\mathsf{a}, 0.1)(\mathsf{b}, 2.5)(\mathsf{a}, 3.5)(\mathsf{b}, 4.8)$, and the pattern "b occurs within two seconds after a does" (a TA for a pattern that is essentially the same is in Fig. 9). Any match contains the second succession of a and b—note that the first succession is too far apart. One such match is given by $w_2|_{(3,5)}$; but there are uncountably many such matches. The match set can be expressed symbolically as $\{(t, t') \mid 2.5 \leq t < 3.5, 4.8 < t'\}$.

Despite its obvious applications in various stages of CPS design and deployment, the study of timed pattern matching was started only recently [3]–[8]. Consequently, the use of timed pattern matching is quite limited in the current industry practice: for example, the existence of algorithms that can match against temporal specifications is not commonly known.

B. Remote Monitoring in Embedded Applications

In this paper, we propose *filtering* for (timed and untimed) pattern matching. It is preprocessing that is applied to an input word.

Our motivation comes from *embedded* applications. In embedded systems—an important aspect of CPS—it is common
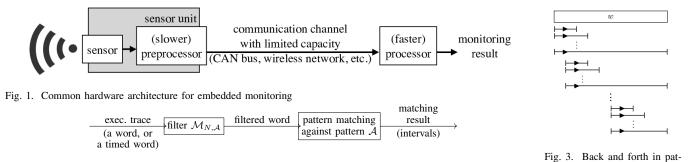
Fig. 1. Common hardware architecture for embedded monitoring

Fig. 2. Proposed monitoring workflow with filtering

Fig. 3. Back and forth in pattern matching

that a sensor unit and a processor (that conducts the computational task of monitoring) are placed physically apart. Moreover, the communication channel between them often has a limited capacity. See Fig. 1.

An example of such a situation is in a modern automobile, in which a sensor unit in the engine gathers data and sends them to a processor that is far apart (to avoid the engine's heat and vibration, for example). They are interconnected by a *controller area network (CAN)*, which is subject to severe performance limitation due to cost reduction. Another example is in an *IoT device* like a connected electronic appliance and a connected car. It continually sends its own status to a server, that is set up in the computer cloud and monitors the device. The wireless communication channel between them is limited e.g. by the battery capacity of the device.

C. Filtering for (Timed) Pattern Matching

A natural idea in such remote monitoring situations is to try to reduce the amount of data that is sent from the sensor unit to the processor, in such a way that does not affect the result of monitoring. Many sensor units come with processors within, and we can use them for such preprocessing. We assume that the preprocessor (within the sensor unit, Fig. 1) is much slower than the processor that conducts actual monitoring; therefore the preprocessing must be computationally cheap.

This leads to our proposed workflow shown in Fig. 2. There we apply computationally cheap *filtering* to the input word, in order to reduce the load on the communication channel to the processor, and to reduce the load on the processor as well.

D. Moore Machines as Filters

In this paper we focus on two settings of monitoring: (1) in the *untimed* setting, an execution trace is a word $w \in \Sigma^*$ and a pattern is a nondeterministic finite automaton (NFA) $\mathcal{A}$ over Σ; and (2) in the *timed* setting, an execution trace is a timed word and a pattern is a timed automaton (TA) $\mathcal{A}$. Our technical contribution is the construction of a filter $\mathcal{M}_{\mathcal{A},N}$ realized as a *Moore machine*, based on a pattern automaton $\mathcal{A}$ and a parameter $N \in \mathbb{N}$ called the *buffer size*. Moore machines are a well-known model of state-based computation: it is an automaton with an additional state-dependent *output* function. It operates in a nicely sequential and synchronous manner, reading one input character, moving to a next state and outputting one character. This feature is especially suited

to the *logic synthesis* of digital circuits [9], opening up the way to hardware acceleration by FPGA or ASIC.

Such sequential operation of Moore machines is in stark contrast with that of pattern matching. Since a pattern is given by an automaton $\mathcal{A}$ in our settings, the length of matches (i.e. $|w|_{[i,j]}| = j - i + 1$ such that $w|_{[i,j]} \in L(\mathcal{A})$) is not fixed. Therefore we have to try matching windows of different size at different positions, going back and forth over the input word w (see Fig. 3). This exhibits a qualitative difference between the (sequential) filtering task conducted by the (slower) preprocessor, and the (back-and-forth) pattern matching task conducted by the (faster) main processor. Our construction yields an (untimed) Moore machine as a filter $\mathcal{M}_{\mathcal{A},N}$, even in the timed setting.

The output of our filter $\mathcal{M}_{\mathcal{A},N}$ is the same as input (timed) word w, except that some characters are *masked* by a fresh character $\perp$. For example: $w = $ abb bbb aab is turned into ab$\perp$ $\perp\perp\perp$ aab under the pattern $\mathcal{A} = $ aa*b. By the binary representation of the length of successive $\perp$'s, the data size can be reduced exponentially. Moreover, in case we are only interested in the matched subwords $w|_{[i,j]}$ (but not in the indices i, j), we can further suppress successive $\perp$'s into one $\perp$. (Note that removing all $\perp$'s after the filtering stage can result in spurious matches in the pattern matching stage, see Fig. 2.)

Our Moore-machine filter $\mathcal{M}_{\mathcal{A},N}$ is constructed based on a pattern automaton $\mathcal{A}$ and a positive integer N for the *buffer size*. The parameter N allows a user to choose the balance between the computational cost of filtering (the greater N is, the more states $\mathcal{M}_{\mathcal{A},N}$ has) and the size of the filtered word (the greater N is, the more $\perp$, i.e. the smaller the filtered word is). This flexibility makes the algorithm suited for various hardware configurations.

We have implemented our construction; we present our experiment results for the *timed* setting (which is harder). Our examples come from the automotive domain. We observe that for realistic pattern TAs $\mathcal{A}$ and input timed words w, the filtered words can be 2–100 times shorter than the original word w. We also experimentally confirm that running a Moore machine $\mathcal{M}_{\mathcal{A},N}$ is cheap. Furthermore, we observe that having a filter (like in Fig. 2) accelerates the task of timed pattern matching itself, by 1.2–2 times.

In the theoretical analysis of our construction, we prove *soundness*: all the matches in the original input word are preserved by filtering. We show soundness for both of the

untimed and timed settings. We note, however, that soundness is satisfied by the trivial (identity) filter; so soundness itself does not speak much about the benefit of filtering. Besides our experiments, we present some theoretical results about the filtering performance in the untimed setting. They include the following: completeness (in the sense that *all* the unnecessary characters get masked) if $L(\mathcal{A})$ is finite (this is the setting of *multiple string matching* [10]); and *monotonicity* (bigger N leads to better filtering results). These results also suggest performance advantages of our *timed* construction since it shares the basic ideas with the untimed one.

Our construction of filters is automata-theoretic, with two principal steps of 1) equipping a buffer (of size N), and 2) determinization. For the second step in the timed setting, we employ *one-clock determinization* of timed automata (TA) [11, §5.3] that overapproximates a given TA by a deterministic and one-clock TA.

E. Contribution

Overall, our contribution is summarized as follows.

- A construction of a filter $\mathcal{M}_{\mathcal{A},N}$ for *untimed* pattern matching against $\mathcal{A}$. The filter is given by a Moore machine and thus operates in a simple, sequential and synchronous way. It is also amenable to hardware acceleration by logic circuits. The parameter N gives flexibility to a user in the trade-off between computational cost and the effect of filtering.
- A construction of a filter $\mathcal{M}_{\mathcal{A},N}$ for *timed* pattern matching. Given a timed automaton $\mathcal{A}$ as a pattern and the buffer size, we construct an (untimed) Moore machine $\mathcal{M}_{\mathcal{A},N}$. The construction extends the untimed version, and is more involved, employing zone-based abstraction of the pattern timed automaton. Given also its practical relevance, we consider the timed construction to be the main contribution of the paper.
- Soundness (preservation of all matches) is proved in the untimed and timed settings.
- In the timed setting we also prove some theoretical results about the performance of our filters.
- Implementation of the timed construction, and experiments that demonstrate the benefit of our filtering construction.
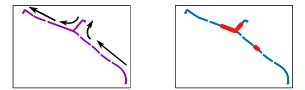
F. Pattern Matching vs. Pattern Search

Another mathematical formulation of monitoring—that is alternative to our choice of pattern matching—is what we call the *pattern search* problem. It asks if the match set (see (1–2)) is empty or not. Pattern search is appealing since it is easily reduced to the *membership problem*. Roughly speaking, given a pattern automaton $\mathcal{A}$, one first adds a self-loop to the initial state so that prefixes of an input word can be disregarded, and then monitors if any accepting state becomes active. Pattern search has been extensively studied in the monitoring context: see e.g. [12], [13].

Pattern *matching* is more expensive than pattern search, since it requires remembering indices (Fig. 3). It is nevertheless highly relevant to real-world monitoring applications—especially to *remote* monitoring such as in Example I.3 below.

Note that remote monitoring is often only *semi*-online: a log can arrive at a monitor in sizable chunks, in an intermittent manner. Then it is imperative to be able to single out which part of the received log is issuing an alert.

The relevance of pattern matching to monitoring applications is widely acknowledged in the community, as is witnessed by the recent proliferation of literature [3]–[6], [14]–[19].

Example I.3 (semi-online remote monitoring)**.** As a concrete example of remote monitoring (Fig. 1), let us imagine a semi-connected vehicle. It keeps driving logs in its memory, and sends them to a center over the Internet once it stops within the range of a known wireless network. Analysis of the logs is done in the center. A driving log is a timed word with information on the vehicle's position, velocity and throttle. One such timed word w, taken from ROSBAG STORE (rosbag.tier4.jp), looks like the left below after we plot the vehicle's position. (The plot is discontinuous: the absence of data can be attributed e.g. to loss of the GPS signals.)

Let us say that we are interested in those road segments where the throttle is greater than a certain threshold for ten seconds or longer. We run timed pattern matching on w with a suitable pattern automaton $\mathcal{A}$. By mapping the identified time intervals to the position plot, we identify the road segments of our interests (right above, red, generated by the tool `MONAA` [20]).

G. Organization

We fix notations in §II. We present a construction of filter Moore machines for *untimed* pattern matching, in §III. We also prove some properties such as soundness. The same idea is used in §IV for a more complicated problem of filtering for *timed* pattern matching. Here we prove soundness. Implementation and experiment results for the timed case are presented in §V. We discuss related work in §VI. Most of the proofs are deferred to the appendix.

II. PRELIMINARIES

The set $\Sigma^* = \bigcup_{n \in \mathbb{N}} \Sigma^n$ is that of *words* over Σ. The length n of a word $w = a_1 a_2 \ldots a_n$ (where $a_i \in \Sigma$) is denoted by $|w|$.

For an NFA $\mathcal{A} = (\Sigma, S, s_0, S_F, E)$ and a string $w \in \Sigma^*$ over the common alphabet Σ, a *run* $\overline{s}$ of $\mathcal{A}$ over w is a sequence $\overline{s} = s_0, s_1, \ldots, s_{|w|}$ such that for each $i \in [1, |w|]$ we have $(s_{i-1}, w_i, s_i) \in E$. A run $\overline{s} = s_0, s_1, \ldots, s_{|w|}$ is *accepting* if we have $s_{|w|} \in S_F$.

The powerset of a set X is denoted by $\mathcal{P}(X)$. The disjoint union of X, Y is $X \amalg Y$. For an alphabet Σ, the set $\Sigma \amalg \{\bot\}$ augmented with a fresh symbol $\bot$ is denoted by $\Sigma_\bot$.

We will be using the set $\{1, 2, \ldots, N\}$ for the value domain of counters. This is denoted by $\mathbb{Z}/N\mathbb{Z}$, because we rely on its algebraic structure (such as addition modulo N).

A *Moore machine* is given by $\mathcal{M} = (\Sigma_{\text{in}}, \Sigma_{\text{out}}, Q, q_0, \Delta, \Lambda)$ where Σ_{in} and Σ_{out} are input and output alphabets, Q is a finite set of states, $q_0 \in Q$ is an initial state, $\Delta : Q \times \Sigma_{\text{in}} \to Q$ is a transition function, and $\Lambda : Q \to \Sigma_{\text{out}}$ is an output function. For $\mathcal{M}$ and an input word $w = a_1 a_2 \ldots a_n \in \Sigma_{\text{in}}^*$ (where $a_i \in \Sigma_{\text{in}}$), the *run* $\bar{q}$ of $\mathcal{M}$ over w is the sequence $\bar{q} = q_0 q_1 \ldots q_n \in Q^*$ satisfying $q_i = \Delta(q_{i-1}, a_i)$ for each $i \in [1, |w|]$. In this case, the *output word* $w' \in \Sigma_{\text{out}}^*$ of $\mathcal{M}$ over w is given by $w' = \Lambda(q_0)\Lambda(q_1) \ldots \Lambda(q_{n-1}) \in \Sigma_{\text{out}}^*$.

III. MOORE-MACHINE FILTERING FOR PATTERN MATCHING I: UNTIMED

A. Problem Formulation

Definition III.1 ((untimed) pattern matching). Given an NFA $\mathcal{A}$ over an alphabet Σ and a word $w = a_1 a_2 \ldots a_n \in \Sigma^*$, the *pattern matching* problem asks for the *match set* $\text{Match}(w, \mathcal{A}) = \{ (i, j) \in \mathbb{N}^2 \mid w|_{[i,j]} \in L(\mathcal{A}) \}$, where $w|_{[i,j]} = a_i a_{i+1} \ldots a_j$.

Our goal is the workflow shown in Fig. 2. We fix the input/output type of the filter in the following general definition.

Definition III.2 (Moore-machine filter for (untimed) pattern matching). Let $\mathcal{A}$ be an NFA over an alphabet Σ, and let N be a positive integer. A *filter for $\mathcal{A}$ with buffer size N* is a Moore machine $\mathcal{M} = (\Sigma_{\text{in}}, \Sigma_{\text{out}}, Q, q_0, \Delta, \Lambda)$ that satisfies the following.

- $\Sigma_{\text{in}} = \Sigma_{\text{out}} = \Sigma_{\perp}$.
- Let $w = a_1 \ldots a_n \in \Sigma^*$ be an arbitrary word, and consider the word $w \perp^N$ obtained by padding $\perp$'s in the end. We require that the output word of $\mathcal{M}$ over this word $w \perp^N$ be of the form $\perp^N w'$, where $w' = b_1 \ldots b_n$, and b_i is either $\perp$ or a_i for each $i \in [1, n]$. We say that the character a_i at the position i is *passed* if $b_i = a_i$; otherwise (i.e. if $b_i = \perp$) we say a_i is *masked*.

A filter $\mathcal{M}$ is *sound* if it preserves all matching intervals. That is, $b_k = a_k$ for each $k \in [1, n]$ such that $\exists i, j. (k \in [i, j] \wedge [i, j] \in \text{Match}(w, \mathcal{A}))$.

Explanations are in order, especially about the buffer size N and the padding by $\perp^N$ of the input and output words. The padding means filtering is done in the way depicted in Fig. 4, with a delay of $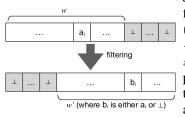N$ steps. This is because of how our Moore machine works (Fig. 5): it scans the input word from left to right; the machine stores N characters in the FIFO buffer in it (encoded in its state space Q); and the characters are output once they are dequeued from the FIFO buffer. Hence there is a delay of N steps. Initially, the buffer is filled with $\perp$ (Fig. 5, Step 0); this accounts for the prefix $\perp^N$ of the output word $\perp^N w'$. The padding $\perp^N$ at the end of the input word $w \perp^N$ is needed to dequeue the content of the buffer (from Step $n+1$ to $n+N$). In the course of the scanning in Fig. 5, some characters in $w = a_1 \ldots a_N$ get masked, although such masking is not explicit in the figure.

Fig. 4. Padding in filtering

B. Construction of Filter Moore Machine $\mathcal{M}_{\mathcal{A},N}$

Definition III.3 (the filter $\mathcal{M}_{\mathcal{A},N}$ for (untimed) pattern matching). Let Σ be an alphabet, N be a positive integer, and $\mathcal{A} = (\Sigma, S, s_0, S_F, E)$ be an NFA. We define a Moore machine

$$\mathcal{M}_{\mathcal{A},N} = (\Sigma_{\perp}, \Sigma_{\perp}, Q, q_0, \Delta, \Lambda)$$

as follows. Its input and output alphabets are both $\Sigma_{\perp}$.

The state space is $Q = \mathcal{P}\big(S \times (\mathbb{Z}/N\mathbb{Z})\big) \times \big((\Sigma_{\perp})^N \times \{\text{pass}, \text{mask}\}^N\big)$. Here $\mathbb{Z}/N\mathbb{Z}$ is the N-element set with addition modulo N (§II).

The initial state is $q_0 = \big(\{(s_0, 0)\}, (\perp, \ldots, \perp), (\text{mask}, \ldots, \text{mask})\big)$.

The transition $\Delta : Q \times \Sigma_{\perp} \to Q$ is as follows. For each $a \in \Sigma_{\perp}$,

$$\Delta\Big(\big(\mathcal{S}, (a_1, a_2, \ldots, a_N), (l_1, l_2, \ldots, l_N)\big), a \Big) \tag{3}$$
$$= \big(\mathcal{S}', (a_2, \ldots, a_N, a), \bar{l}\big), \quad \text{where}$$
$$\mathcal{S}' = \{ (s', (n \bmod N) + 1) \mid (s, n) \in \mathcal{S}, (s, a, s') \in E \} \tag{4}$$
$$\cup \{(s_0, 0)\}, \quad \text{and}$$

$$\bar{l} = \begin{cases} (\text{pass}, \ldots, \text{pass}) & \text{if } \exists s. (s, N) \in \mathcal{S}', \\ (l_2, l_3, \ldots, l_{N - \psi(\mathcal{S}') + 1}, \overbrace{\text{pass}, \ldots, \text{pass}}^{\psi(\mathcal{S}')}) & \text{else if } \exists n, s \in S_F. \\ & \quad (s, n) \in \mathcal{S}', \\ (l_2, l_3, \ldots, l_N, \text{mask}) & \text{otherwise.} \end{cases} \tag{5}$$

Here $\psi(\mathcal{S}') \in \mathbb{Z}_{>0}$ is $\psi(\mathcal{S}') = \max\{n \mid \exists s \in S_F. (s, n) \in \mathcal{S}'\}$.

Finally, the output function $\Lambda : Q \to \Sigma_{\perp}$ is defined as follows.

$$\Lambda\big(\mathcal{S}, (a_1, a_2, \ldots, a_N), (l_1, l_2, \ldots, l_N) \big) = \begin{cases} a_1 & \text{if } l_1 = \text{pass} \\ \perp & \text{if } l_1 = \text{mask} \end{cases} \tag{6}$$

We describe intuitions. The construction of $\mathcal{M}_{\mathcal{A},N}$ combines three building blocks: *determinization*, *counters* and a *buffer* of size N.

(Determinization) The pattern $\mathcal{A}$ is an NFA, but we want a *deterministic* Moore machine. This determinization accounts for the powerset construction $\mathcal{P}$ in the component $\mathcal{P}\big(S \times (\mathbb{Z}/N\mathbb{Z})\big)$ of the state space Q. For example, an element $\{(s_1, n_1), \ldots, (s_k, n_k)\}$ of this component means "the states $s_1, \ldots, s_k$ are active in the NFA $\mathcal{A}$." Observe that, in (4), the part regarding states $(s, s', \ldots)$ follows the usual determinization. An exception is the addition of $(s_0, 0)$ in (4): it is because a matching can start at any position of the input word.

(Counters) Moreover, each active state that traverses $\mathcal{A}$ keeps a *counter* for how many steps it has traveled since the initial state. This is the component $\mathbb{Z}/N\mathbb{Z}$ in the state space Q. The maximum value of those counters is the same as the buffer size N. Once this maximum is reached, a counter starts over from 1. See (4), where counters for active states are incremented modulo N.

(Buffer) A FIFO *buffer* of size N accounts for the second component $(\Sigma_{\perp})^N \times \{\text{pass}, \text{mask}\}^N$ of the state space Q. Each of the N cells stores a character from $\Sigma_{\perp}$ and a label (pass or mask). See (3) and (5), where the basic behavior of the buffer is to dequeue the leftmost element and to enqueue the read character on the right.
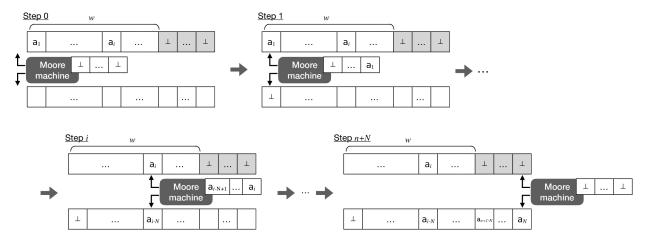
Fig. 5. Filtering by a Moore machine. Sometimes a character a_i is decided to be unnecessary and gets masked by $\perp$ (i.e. $b_i = \perp$), although such masking is not explicit in the figure.
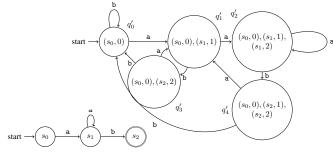

Fig. 6. Left: a pattern NFA $\mathcal{A}_0$ for $\texttt{aa*b}$. Right: the filter Moore machine $\mathcal{M}_{\mathcal{A}_0,2}$, its non-buffer part
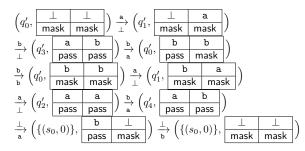

Fig. 7. The run of $\mathcal{M}_{\mathcal{A}_0,2}$ over the word $w = \texttt{abbbaab}$. The tables show the states of the buffer, enqueued from the right. In $\xrightarrow{a}{b}$, a is the input character and b is the output character.

It is the labels in the buffer (pass or mask) that determine whether to mask a character or not. The default label is mask (the third case in (5)), and if it remains unchanged for N steps then the corresponding character gets masked by $\perp$ in the output (the second case in (6)). The label can change from mask to pass for two different reasons (the first two cases in (5)).

- The second case in (5) is when it is detected that some characters towards the end of the buffer form a match against the pattern $\mathcal{A}$, leading to an accepting state $s \in S_F$ of $\mathcal{A}$. Then we mark those characters by pass, indicating that they need to be passed to pattern matching (Fig. 2). The number $\psi(\mathcal{S}')$ of characters to be passed is calculated using the counter n associated to the active state $s \in S_F$.
- On the first case in (5), its condition $\exists s. (s, N) \in \mathcal{S}'$ says that the counter for some active state s has reached the maximum N. In this case we are not sure whether this active state s of $\mathcal{A}$ will eventually reach an accepting state or not. To be on the safe side, we pass all the N characters to pattern matching without masking them. In the untimed setting, this is the only place where completeness of filtering is potentially lost.

In summary, in Def. III.3 we construct a Moore machine that works in the way depicted in Fig. 5. The Moore machine's state space combines the following: determinization of the pattern NFA $\mathcal{A}$; counters for the steps from the initial state; and a FIFO buffer that stores N characters labeled by pass or mask.

Proposition III.4. *The Moore machine $\mathcal{M}_{\mathcal{A},N}$ is a filter for $\mathcal{A}$ with buffer size N, in the sense of Def. III.2.* □

In our implementation, we realize a filter *not* as a plain Moore machine with a state space $Q = \mathcal{P}(S \times (\mathbb{Z}/N\mathbb{Z})) \times ((\Sigma_\perp)^N \times \{\textsf{pass}, \textsf{mask}\}^N)$ as described in Def. III.2. Instead, we separate the state space Q into the "buffer part" $(\Sigma_\perp)^N \times \{\textsf{pass}, \textsf{mask}\}^N$ and the "non-buffer part" $\mathcal{P}(S \times (\mathbb{Z}/N\mathbb{Z}))$, and we generate the former buffer part on-the-fly. More specifically, the non-buffer part is constructed as a DFA all at once in the beginning, and this DFA dictates how to operate on the buffer part that is realized as an array of size N. See Example III.6 later for illustration.

Proposition III.5. *Let $\mathcal{A} = (\Sigma, S, s_0, S_F, E)$ be an NFA. For the induced filter Moore machine $\mathcal{M}_{\mathcal{A},N}$, the size of the non-buffer part $\mathcal{P}(S \times (\mathbb{Z}/N\mathbb{Z}))$ of its state space is bounded by $O(2^{N \cdot |S|})$.*

Therefore, the memory usage for the non-buffer part (including the transitions) is $O(2^{N \cdot |S|} \cdot |\Sigma|)$. The memory usage for the buffer part is $O(N \log |\Sigma|)$; overall, the space complexity of running our filter Moore machine $\mathcal{M}_{\mathcal{A},N}$ is $O(2^{N \cdot |S|} \cdot |\Sigma|)$. □

The space complexity is exponential in N, and this comes from the powerset construction for the non-buffer part $\mathcal{P}(S \times (\mathbb{Z}/N\mathbb{Z}))$. Experimentally, however, the memory consumption does not necessarily grow exponentially in N. This is because not all states in $\mathcal{P}(S \times (\mathbb{Z}/N\mathbb{Z}))$ are reachable. See RQ2 in §V.

Example III.6. Let us consider the pattern aa*b. It is expressed by the NFA $\mathcal{A}_0$ in Fig. 6. The filter Moore machine $\mathcal{M}_{\mathcal{A}_0,2}$ from Def. III.3 is depicted in Fig. 6, where buffer states are omitted. Its run over the word $w =$ abbbaab is shown in Fig. 7; its output word is $\bot\bot$ab$\bot\bot$aab, which means that the filtering result is ab$\bot\bot$aab.

C. Properties of the Moore Machine $\mathcal{M}_{\mathcal{A},N}$

Throughout the rest of this section, let $\mathcal{A}$ be a pattern NFA $\mathcal{A} = (\Sigma, S, S_0, E, S_F)$, N be a positive integer, and $\mathcal{M}_{\mathcal{A},N} = (\Sigma_\bot, \Sigma_\bot, Q, q_0, \Delta, \Lambda)$ be the filter Moore machine in Def. III.3. Let $w = a_1 a_2 \ldots a_n$ be a word over Σ, and $\bot^N w'$ be the output word of $\mathcal{M}_{\mathcal{A},N}$ over the input word $w\bot^N$. Let $w' = b_1 b_2 \ldots b_n$, where $b_i \in \Sigma_\bot$.

Theorem III.7 (soundness). *The Moore machine $\mathcal{M}_{\mathcal{A},N}$ is sound, in the sense of Def. III.2.* $\square$

Completeness holds if, the length of matches is bounded and the buffer size N is no shorter than the bound. This is essentially the setting of *multiple string matching* [10].

Theorem III.8 (completeness). *Assume we have $\max\{|w| \mid w \in L(\mathcal{A})\} \leq N < \infty$. Then we can construct an NFA $\mathcal{A}'$ with $L(\mathcal{A}) = L(\mathcal{A}')$, so that the filter Moore machine $\mathcal{M}_{\mathcal{A}',N}$ is complete. The latter means: if the index k satisfies $a_k = b_k$, then there is an interval $[i, j]$ such that $k \in [i, j]$ and $w|_{[i,j]} \in L(\mathcal{A})$.* $\square$

The intuition for monotonicity is that, the bigger the buffer size N' is, the filter $\mathcal{M}_{\mathcal{A},N'}$ masks more characters while its state space grows. A precise statement is more intricate, requiring the bigger buffer size N'' to be a multiple of the smaller one.

Theorem III.9 (monotonicity). *For any positive integer N', let $\mathcal{M}_{\mathcal{A},N'}$ be the filter Moore machine of Def. III.3, and $\bot^{N'} w'^{(N')}$ be the output word of $\mathcal{M}_{\mathcal{A},N'}$ over the input word $w\bot^{N'}$. Let $w'^{(N')} = b_1^{(N')} b_2^{(N')} \ldots b_n^{(N')}$ where $b_i^{(N')} \in \Sigma_\bot$. For any positive integers n, N' and any index k of w, $b_k^{(N')} = \bot$ implies $b_k^{(nN')} = \bot$.* $\square$

As mentioned in Prop. III.5, the state space of our filter Moore machine is exponentially bigger than that of $\mathcal{A}$. This is because of the powerset construction required by its deterministic branching. By sacrificing the execution time, one can determinize the NFA on the fly, which usually needs less memory space. See Appendix D of [21].

IV. MOORE-MACHINE FILTERING FOR PATTERN MATCHING II: TIMED

We present our main contribution, that is, the construction of filter Moore machines for *timed* pattern matching. While the basic ideas stay the same as in the untimed setting (§III), *determinization* poses a technical challenge, since timed automata (TA) cannot determinized in general [22]. Here we rely on *one-clock determinization* [11, §5.3]. The construction overapproximates reachability, hence we maintain soundness of filtering. Moreover, the local nature of the resulting TA—it has only one clock variable that is reset at every transition—allows us to devise a filter that is an *untimed, finite-state* Moore machine.

A. Problem Formulation

Definition IV.1 (timed words). Let Σ be an alphabet. A *timed word* over Σ is a sequence w of pairs $(a_i, \tau_i) \in \Sigma \times \mathbb{R}_{>0}$ satisfying $\tau_i < \tau_{i+1}$ for any $i \in [1, |w| - 1]$. Let $w = (\overline{a}, \overline{\tau})$ be a timed word. The set of timed words over Σ is denoted by $\mathcal{T}(\Sigma)$.

We denote the subsequence $(a_i, \tau_i), (a_{i+1}, \tau_{i+1}), \cdots, (a_j, \tau_j)$ by $w(i, j)$. For $t \in \mathbb{R}_{\geq 0}$, the t-*shift* of w is $(\overline{a}, \overline{\tau}) + t = (\overline{a}, \overline{\tau} + t)$ where $\overline{\tau} + t = (\tau_1 + t, \tau_2 + t, \cdots, \tau_{|\tau|} + t)$. For timed words $w = (\overline{a}, \overline{\tau})$ and $w' = (\overline{a'}, \overline{\tau'})$, their *absorbing concatenation* is $w \circ w' = (\overline{a} \circ \overline{a'}, \overline{\tau} \circ \overline{\tau'})$ where $\overline{a} \circ \overline{a'}$ and $\overline{\tau} \circ \overline{\tau'}$ are usual concatenations, and their *non-absorbing concatenation* is $w \cdot w' = w \circ (w' + \tau_{|w|})$. We note that the absorbing concatenation $w \circ w'$ is defined only when $\tau_{|w|} < \tau'_1$.

For a timed word $w = (\overline{a}, \overline{\tau})$ on Σ and $t, t' \in \mathbb{R}_{>0}$ satisfying $t < t'$, a *timed word segment* $w|_{(t,t')}$ is defined by the timed word $(w(i, j) - t) \circ (\$, t' - t)$ on the augmented alphabet $\Sigma \amalg \{\$\}$, where i, j are chosen so that $\tau_{i-1} \leq t < \tau_i$ and $\tau_j < t' \leq \tau_{j+1}$. Here the timed word $w(i, j) - t$ is the $(-t)$-shift of $w(i, j)$; and the fresh symbol \$ is called the *terminal character*.

Definition IV.2 (timed automaton). Let C be a finite set of *clock variables*, and $\Phi(C)$ denote the set of conjunctions of inequalities $x \bowtie c$ where $x \in C$, $c \in \mathbb{Z}_{\geq 0}$, and $\bowtie \in \{>, \geq, <, \leq\}$. A *timed automaton* $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$ is a tuple where Σ is an alphabet, S is a finite set of states, $s_0 \in S$ is an initial state, $S_F \subseteq S$ is a set of accepting states, and $E \subseteq S \times S \times \Sigma \times \mathcal{P}(C) \times \Phi(C)$ is a set of transitions. The components of a transition $(s, s', a, \lambda, \delta) \in E$ represent: the source, target, action, reset variables and guard of the transition, respectively.

We define a *clock valuation* ν as a function $\nu : C \to \mathbb{R}_{\geq 0}$. We define the t-*shift* $\nu + t$ of a clock valuation ν, where $t \in \mathbb{R}_{\geq 0}$, by $(\nu + t)(x) = \nu(x) + t$ for any $x \in C$. For a timed automaton $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$ and a timed word $w = (\overline{a}, \overline{\tau})$, a *run* of $\mathcal{A}$ over w is a sequence r of pairs $(s_i, \nu_i) \in S \times (\mathbb{R}_{\geq 0})^C$ satisfying the following: (initiation) s_0 is the initial state and $\nu_0(x) = 0$ for any $x \in C$; and (consecution) for any $i \in [1, |w|]$, there exists a transition $(s_{i-1}, s_i, a_i, \lambda, \delta) \in E$ such that $\nu_{i-1} + \tau_i - \tau_{i-1} \models \delta$ and $\nu_i(x) = 0$ (for $x \in \lambda$) and $\nu_i(x) = \nu_{i-1}(x) + \tau_i - \tau_{i-1}$ (for $x \notin \lambda$). A run only satisfying the consecution condition is a *path*. A run $r = (\overline{s}, \overline{\nu})$ is *accepting* if the last element $s_{|s|-1}$ of s belongs to S_F. The *language* $L(\mathcal{A})$ is defined to be the set $\{w \mid$ there is an accepting run of $\mathcal{A}$ over $w\}$ of timed words.

Here is our target problem. Its algorithms have been actively studied [3], [4], [6]; filtering Moore machines as preprocessors for those algorithms is this paper's contribution.

Definition IV.3 (timed pattern matching). Let $\mathcal{A}$ be a timed automaton, and w be a timed word, over a common alphabet Σ. The *timed pattern matching* problem requires all the intervals (t, t') for which the segment $w|_{(t,t')}$ is accepted by $\mathcal{A}$. That is,

it requires the *match set* $\mathrm{Match}(w, \mathcal{A}) = \{(t, t') \mid w|_{(t,t')} \in L(\mathcal{A})\}$.

B. One-Clock Determinization of TA

Among the three main building blocks for our untimed construction of filter (Def. III.3), *counters* and a *buffer* carry over smoothly to the current timed setting. For *determinization* we rely on the overapproximating notion in Def. IV.5; it is taken from [11, §5.3].

We start with some auxiliary notations.

Definition IV.4 (restriction $\nu|_C$, join $\nu \sqcup \nu'$). Let $\nu \colon C' \to \mathbb{R}_{\geq 0}$ be a clock valuation. The restriction of ν to $C \subseteq C'$ is denoted by $\nu|_C \colon C \to \mathbb{R}_{\geq 0}$. That is, $(\nu|_C)(x) = \nu(x)$ for each $x \in C$.

Let $\nu \colon C \to \mathbb{R}_{\geq 0}$ and $\nu' \colon C' \to \mathbb{R}_{\geq 0}$ be clock variations. We define their *join* $\nu \sqcup \nu' \colon C \amalg C' \to \mathbb{R}_{\geq 0}$ to be the following clock valuation over the disjoint union $C \amalg C'$.

$$(\nu \sqcup \nu')(x) = \begin{cases} \nu(x) & \text{if } x \in C \\ \nu'(x) & \text{if } x \in C'. \end{cases}$$

The function that maps x_i to r_i (for each $i \in \{1, \ldots, n\}$) is denoted by $[x_1 \mapsto r_1, \ldots, x_n \mapsto r_n]$.

Definition IV.5 (one-clock determinization). Let $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$ be a timed automaton (TA), and y be a fresh clock variable (i.e. $y \notin C$). A TA $\mathcal{A}' = (\Sigma, S', s_0', S_F', \{y\}, E')$ is a *one-clock determinization* of $\mathcal{A}$ if the following conditions hold.

1) Each element $\mathcal{S} \in S'$ of the (new, finite) state space S' is a finite set $\mathcal{S} = \{(s_1, Z_1), \ldots, (s_m, Z_m)\}$ of pairs (s_i, Z_i), where $s_i \in S$ is a state of $\mathcal{A}$, and Z_i is a subset of $(\mathbb{R}_{\geq 0})^{C \amalg \{y\}}$ given by a special polytope called a *zone* (see e.g. [23]).

2) For each transition $(\mathcal{S}, a, \delta, \lambda, \mathcal{S}') \in E'$ of $\mathcal{A}'$, the guard δ is a finite union of intervals of y. Moreover it respects enabledness of transitions E in $\mathcal{A}$. Precisely, for any $u, u' \in \mathbb{R}_{\geq 0}$ that satisfy δ, we have $E_a(\mathcal{S}, u) = E_a(\mathcal{S}, u')$, where the set $E_a(\mathcal{S}, u) \subseteq E$ is defined by

$$E_a(\mathcal{S}, u) = \{ (s, a, \delta', \lambda', s') \in E \mid \\ \exists (s, Z) \in \mathcal{S}. \exists \nu \in Z. \nu(y) = u \text{ and } \nu \text{ satisfies } \delta' \} .$$

3) Each transition of $\mathcal{A}'$ resets the unique clock variable y. That is, for each transition $(\mathcal{S}, a, \delta, \lambda, \mathcal{S}') \in E'$, we have $\lambda = \{y\}$.

4) Each transition $(\mathcal{S}, a, \delta, \lambda, \mathcal{S}') \in E'$ of $\mathcal{A}'$ simulates transitions of $\mathcal{A}$. More precisely, let $(s, Z) \in \mathcal{S}$ and $(\nu \colon C \amalg \{y\} \to \mathbb{R}_{\geq 0}) \in Z$. Assume that $(s, \nu|_C) \xrightarrow{a, \tau} (s', \nu')$ is a (length-1) path of $\mathcal{A}$, for some $s' \in S$ and $\nu' \colon C \to \mathbb{R}_{\geq 0}$ (here τ is the dwell time). Then we require that there exist a zone $Z' \subseteq (\mathbb{R}_{\geq 0})^{C \amalg \{y\}}$ such that 1) $(s', Z') \in \mathcal{S}'$ and 2) the valuation $\nu' \sqcup [y \mapsto \tau]$, over the clock set $C \amalg \{y\}$, belongs to Z'.

5) $\mathcal{A}'$ is deterministic: for each state $\mathcal{S} \in S'$, each clock valuation $\nu \in (\mathbb{R}_{\geq 0})^{\{y\}}$, $a \in \Sigma$ and $\tau \in \mathbb{R}_{\geq 0}$ for the dwell time, a length-1 path from $(\mathcal{S}, \nu)$ labeled with a, τ is unique. That is, if $(\mathcal{S}, \nu) \xrightarrow{a, \tau} (\mathcal{S}', \nu')$ and $(\mathcal{S}, \nu) \xrightarrow{a, \tau}$

$(\mathcal{S}'', \nu'')$ are both paths in $\mathcal{A}'$, then $\mathcal{S}' = \mathcal{S}''$ and $\nu' = \nu''$. (Note that Cond. 3 forces $\nu' = \nu'' = [y \mapsto 0]$.)

6) The initial state s_0' of $\mathcal{A}'$ is given by $s_0' = \{(s_0, \{\mathbf{0}\})\}$. Here $\mathbf{0}$ is the valuation that maps every clock variable to 0.

7) A state $\mathcal{S}$ belongs to S_F' if and only if there exists $(s, Z) \in \mathcal{S}$ such that $s \in S_F$.

Proposition IV.6. *Let $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$ be a TA, and $\mathcal{A}' = (\Sigma, S', s_0', S_F', \{y\}, E')$ be a one-clock determinization of $\mathcal{A}$. Then $\mathcal{A}'$ satisfies the following properties.*

- *(Simulation) Let $w \in \mathcal{T}(\Sigma)$ be a timed word, and assume that there exists a run over w to a state $s \in S$ in $\mathcal{A}$. Then there exists $\mathcal{S} \in S'$ that satisfies the following: 1) $(s, Z) \in \mathcal{S}$ for some zone Z, and 2) there exists a run over w to $\mathcal{S}$ in $\mathcal{A}'$.*
- *(Language inclusion) In particular, $L(\mathcal{A}) \subseteq L(\mathcal{A}')$.* □

Note that Def. IV.5 gives a property and not a construction: there are multiple one-clock determinizations—of varying size and precision—of the same TA $\mathcal{A}$. In our implementation we use a specific construction presented in [11, §5.3.4]. A sketch is in Appendix D of [21].

C. Construction of Our Filter $\mathcal{M}_{\mathcal{A}, N}$

Definition IV.7 (the filter $\mathcal{M}_{\mathcal{A}, N}$ for timed pattern matching). Let $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$ be a TA, and $N \in \mathbb{Z}_{>0}$. The construction of the filter Moore machine $\mathcal{M}_{\mathcal{A}, N}$ is by the following steps.

In the first step we augment the original TA $\mathcal{A}$ with counters. Specifically, let $\mathcal{A}^{N\text{-ctr}} = (\Sigma_\perp, S \times [0, N], (s_0, 0), S_F^{N\text{-ctr}}, C, E^{N\text{-ctr}})$ be defined as follows. $S_F^{N\text{-ctr}} = \{(s_f, n) \mid s_f \in S_F, n \in [0, N]\}$, and

$$E^{N\text{-ctr}} = \{ ((s_0, 0), a, \mathbf{true}, C, (s_0, 0)) \mid a \in \Sigma_\perp \} \\ \cup \{ ((s, n), a, \delta, \lambda, (s', n+1)) \mid (s, a, \delta, \lambda, s') \in E, n \in [0, N-1] \} \\ \cup \{ ((s, N), a, \delta, \lambda, (s', 1)) \mid (s, a, \delta, \lambda, s') \in E \} .$$

In the second step we take a one-clock determinization (Def. IV.5) of $\mathcal{A}^{N\text{-ctr}}$. Let $\mathcal{A}^{N\text{-ctr-d}} = (\Sigma_\perp, S^{N\text{-ctr-d}}, s_0^{N\text{-ctr-d}}, S_F^{N\text{-ctr-d}}, \{y\}, E^{N\text{-ctr-d}})$ be the outcome.

Finally in the third step we define the Moore machine $\mathcal{M}_{\mathcal{A}, N}$:

$$\mathcal{M}_{\mathcal{A}, N} = \Big(\Sigma_\perp \times \mathbb{R}_{\geq 0}, \{\mathsf{pass}, \mathsf{mask}\}, S^{N\text{-ctr-d}} \times \{\mathsf{pass}, \mathsf{mask}\}^N, \\ \big(s_0^{N\text{-ctr-d}}, (\mathsf{mask}, \ldots, \mathsf{mask})\big), \Delta, \Lambda \Big) ,$$

where Δ and Λ are defined as follows. $\Delta\big((\mathcal{S}, \bar{l}), (a, \tau)\big) = (\mathcal{S}', \bar{l}')$, where $\mathcal{S}'$ is the unique successor of $\mathcal{S}$ in $\mathcal{A}^{N\text{-ctr-d}}$ under the character a and the dwell time τ (Def. IV.5), and $\bar{l}'$ is defined as follows.

$$\bar{l}' = \begin{cases} \mathsf{pass}^N & \text{if } \exists s, Z. ((s, N), Z) \in \mathcal{S} \\ l_2, l_3, \ldots, l_{N-\psi(\mathcal{S}')+1}, \overbrace{\mathsf{pass}, \ldots, \mathsf{pass}}^{\psi(\mathcal{S}')} & \text{else if } \exists((s, n), Z) \in \mathcal{S}. s \in S_F \\ l_2, l_3, \ldots, l_N, \mathsf{mask} & \text{otherwise} \end{cases}$$

Here $\psi(\mathcal{S}') = \max\{n \mid \exists s, Z. ((s, n), Z) \in \mathcal{S} \text{ and } s \in S_F\}$. We define $\Lambda\big((\mathcal{S}, (l_1, l_2, \ldots, l_N))\big) = l_1$.

Note that the resulting Moore machine takes timed words as input. This makes its input alphabet infinite (namely $\Sigma_\perp \times \mathbb{R}_{\geq 0}$). This is not a big issue in implementation: we note that the state space is still finite. Moreover, the output alphabet of $\mathcal{M}_{\mathcal{A},N}$ is a two-element set $\{\mathsf{pass}, \mathsf{mask}\}$, meaning that the filter Moore machine only outputs the masking information. A finite-state Moore machine is certainly not capable of buffering time-stamped characters, but the original timed word can be copied and suitable masking can be applied to it later. See the above figure.

Theorem IV.8 (soundness)**.** *Let $\mathcal{A}$ be a pattern TA $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$, N be a positive integer, and $\mathcal{M}_{\mathcal{A},N}$ be the filter Moore machine in Def. IV.7. Let $w = (a_1, \tau_1)(a_2, \tau_2) \ldots (a_n, \tau_n)$ be a timed word over Σ, and $\mathsf{mask}^N w'$ be the output word of $\mathcal{M}_{\mathcal{A},N}$ over the input word $w(\perp, \tau_n)^N$ (here the input and output words are padded by mask^N and $(\perp, \tau_n)^N$, much like in Fig. 4). Let $w' = b_1 b_2 \ldots b_n$, where $b_k \in \{\mathsf{pass}, \mathsf{mask}\}$.*

For any pair (i, j) of indices of w satisfying $w(i, j) - \tau_{i-1} \in L(\mathcal{A})$ and for any index $k \in [i, j]$ of w, we have $b_k = \mathsf{pass}$. $\square$
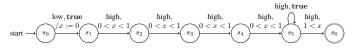
V. IMPLEMENTATION AND EXPERIMENTS

We implemented our filter construction for *timed* pattern matching. Our implementation suppresses successive $\perp$'s into two $\perp$'s, maintaining the timestamp of the first and the last $\perp$'s. We generate the buffer part of the state space Q on-the-fly (namely $\{\mathsf{pass}, \mathsf{mask}\}^N$ of Def. IV.7); this is as we discussed before Prop. III.5. We conducted experiments to answer the following research questions.

RQ1: Does our filter Moore machine mask many events?

RQ2: Is our filter Moore machine online capable? That is, does it work in linear time and constant space, with respect to the length of the input timed word?

RQ3: Does our filter Moore machine accelerate the whole task of timed pattern matching?

RQ4: Is our filter precise? That is, do many of the non-masked events contribute to actual matches?

RQ5: Is our filter responsive? That is, does it not cause large delays?

We implemented our filter construction in C++ and compiled them by clang-900.0.39.2. The tool's input consists of a pattern TA $\mathcal{A}$, the buffer size N and a timed word w; it outputs a filtered word. The experiments are done on Mac OS 10.13.4, MacBook Pro Early 2013 with 2.6 GHz Intel Core i5 processor and 8 GB 1600MHz DDR3 RAM. The benchmark problems we used are in Fig. 8—10. All of them are from automotive scenarios.

For the measurement of the execution time and the memory usage, we used GNU time and took the average of 20 executions. In each experiment, we measured the whole workflow: in RQ2, the time and memory usage include that for constructing a filter; and in RQ3, the time and memory usage account for filter construction, filtering, process intercommunication, and pattern matching. In the experiments for RQ3, we used MONAA [20], a

Fig. 8. TORQUE. The set W (length 242,808–4,873,207) of input words is generated by the automotive engine model `sldemo_enginewc.slx` [24] with random input. The pattern specifies four or more consecutive occurrences of high in one second. For $N = 10$, the size of our filter Moore machine $\mathcal{M}_{\mathcal{A},10}$, measured by the number of the reachable states in the non-buffer part $S^{N\text{-ctr-d}}$ (cf. Def. IV.7 and Prop. III.5), is 16.
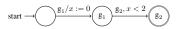

Fig. 9. GEAR. The set W (length 306–1,011,426) of input words is generated by the automatic transmission system model in [25]. The pattern, from ϕ_5^{AT} in [25], is for an event in which gear shift occurs too quickly (from the 1st to 2nd). For $N = 10$, the size of our filter Moore machine $\mathcal{M}_{\mathcal{A},10}$, measured in the same way of Fig. 8, is 3.

recent tool for timed pattern matching. See Appendix F of [21] for the detailed results.

A. RQ1: Filtering Rate

Fig. 11 shows the length of the filtered timed words by our filter Moore machine for each pattern timed automaton $\mathcal{A}$, buffer size N, and timed word $w \in W$.

We observe that the filtered word gets shorter as N is bigger. This concurs with our theoretical consideration in Theorem III.9 (although the result is for the untimed setting). It seems that peak performance is achieved with a relatively small N, such as $N = 10$. With $N = 10$ the length of the filtered timed word is about $1/3$, $1/2$, and $1/100$ of that of the original timed word in TORQUE, GEAR, and ACCEL respectively. For ACCEL our filter masks many characters, because the sizes of the alphabet and the pattern timed automaton are relatively large. This significant data reduction demonstrates the practical use of our filtering methodology in embedded usage scenarios (see Fig. 1).

B. RQ2: Speed and Memory Usage

Fig. 12 and 13 show the execution time and memory usage of our filter Moore machine for each pattern timed automaton $\mathcal{A}$, buffer size N, and timed word $w \in W$.

Fig. 10. ACCEL. The set W (length 708–1,739,535) of input words is generated by the same model as in GEAR. The pattern is from ϕ_8^{AT} in [25]: the gear shifts from 1st to 4th and RPM gets high in its course, but the velocity is low (i.e. the event $v \geq 100$ is absent). For $N = 10$, the size of our filter Moore machine $\mathcal{M}_{\mathcal{A},10}$, measured in the same way of Fig. 8, is 71.

Fig. 11. Length of the input and filtered timed words

Fig. 12. Execution time of our filter Moore machine (including time for construction of a filter)

Fig. 13. Memory usage of our filter Moore machine

Fig. 14. Execution time of the workflow in Fig. 2, with our filter and monitoring by MONAA

In Fig. 12, we observe that the execution time is linear to the length of the input word. In Fig. 13, the memory usage is more or less constant with respect to the length of the input word. These two results suggest that our filtering methodology is online capable.

The time for constructing a filter Moore machine is seen to be negligible: see the execution time for short input words in Fig. 12.

On the effect of varying the buffer size N, for smaller N the execution time was relatively large. This seems to be because fewer characters get masked, more characters are output, and this incurs I/O cost. On memory usage, despite the worst-case result in Prop. III.5 (exponential in N), the growth for bigger N was moderate. This is because not all states from the powerset construction are reachable.

C. RQ3: Acceleration of Timed Pattern Matching

Fig. 14 shows the execution time of the workflow in Fig. 2, where the filter is given by our algorithm and pattern matching is done by a recent tool MONAA [20]. More sp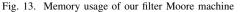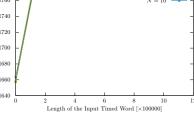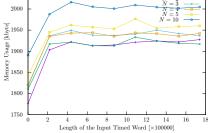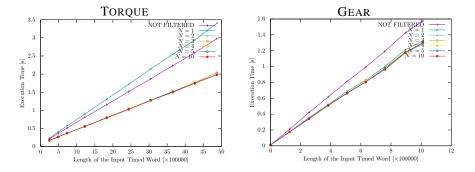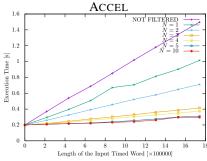ecifically, we connect the standard output of the filter Moore machine to the standard input of MONAA by Unix pipeline. This way the filter and MONAA run parallelly on different cores.

We observe that, via filtering, the overall performance of timed pattern matching is accelerated when N is large enough (say $N = 10$). For TORQUE and GEAR, the speedup is 1.2 times; for ACCEL, it is roughly twice. This acceleration suggests that our filtering methodology is potentially beneficial independent of the architecture assumption in Fig. 1: when a log is enormous monitoring can take hours or days; by running a filter in parallel the time can be shortened.

D. RQ4: Precision

For the three examples TORQUE, GEAR, and ACCEL, and $N = 10$, the ratios of non-masked events contributing to actual matches were 0.34%, 99%, and 92% respectively. Therefore the precision differs greatly depending on patterns. It should also be noted that, even in the low-precision example (TORQUE), our filter successfully reduced the log size by approximately three times (Fig. 11).

Most of the imprecision in the current timed setting is attributed to the laxness of one-clock determinization (Def. IV.5). For example, the TA for TORQUE (Fig. 8) requires four consecutive occurrences of high within one second, using the same clock x. The best overapproximation by one-clock determinization—in which all clock variables get reset after each transition—is to require an occurrence of high in each of four consecutive time segments of length ≤ 1. This is a much looser requirement than the original, and explains the comparatively poorer precision in the example TORQUE.

E. RQ5: Responsiveness

For the three examples TORQUE, GEAR, ACCEL, and the buffer size $N = 10$, we calculated the average delay caused by our filter, which is (the execution time)/$|w| \times N$. The results were $2.2\,\mu s$, $3.1\,\mu s$, and $0.91\,\mu s$ respectively. Although these delays highly depend on the computation power of the processor, the delays were tiny. We conclude our filter is responsive enough.

VI. RELATED WORK

Efficiency of pattern matching has been actively pursued in the fields of database [17], [18] and networking [12], [19], [26]–[29]. In these fields, issues in hardware architecture—such as speed gap between L1/L2 caches and main memory—constitute situations similar to that of embedded monitors that we discussed in the above.

Many works in these application domains deal with (potentially multiple) strings as patterns. There the main source of inspiration is classic algorithms such as Boyer–Moore [30], Commentz-Walter [31], and Aho–Corasick [10]. Examples are e.g. in [32]. Many algorithms for patterns given by regular expressions or automata (instead of strings) also rely on those string-matching techniques. See e.g. [17].

In database and networking, pattern matching against regular expressions is mainly approached by application-specific heuristics that often take machine architecture into account. For example, in [19], [26] pattern regular expressions undergo application of some application-specific rewrite rules. The hierarchical structure of L1/L2 caches and main memory is exploited in [17], [18], while the features of *content-addressable memory* (CAM)—an alternative to RAM often used in network devices—are exploited in [12]. The work [29] extends the formalism of DFA by auxiliary variables for the purpose of moderating state space inflation when multiple patterns are combined.

Pre-filtering before actual pattern matching has been considered in the above lines of works [17], [26], [32]. The principal difference between those works and ours is that their filters produce *match candidates*, data that include explicit indices for potential matches. For this reason, the second step of the workflow (what we call pattern matching) is called *verification* in those papers. In contrast, our filter only *masks* the input word. This is because our goal—motivated by embedded applications—is not only in the matching speed but also in reducing the amount of data passed from the sensor to the pattern matching unit. This choice enables us to use Moore machines, too, which can be readily implemented on FPGA and ASIC.

Monitoring over a real-time temporal logic property [33], [34] and timed pattern matching [3], [14] are relatively new topics. They have been mainly pursued in the context of cyber-physical systems, although its applications in database and networking are very likely. In timed pattern matching, specifications can be given by timed automata (as we do in the current work), *timed regular expressions* [35] and *metric temporal logic* formulas [36]. Algorithms for timed pattern matching against specifications in these formalisms have been actively pursued in [3]–[5], [7], [8]. Besides, the line of work [6], [14], [20] has pursued acceleration of timed pattern matching, by combining *shift table* techniques (like in Boyer–Moore [30] and Franek–Jennings–Smyth [37]) and timed automata [6], [14]. This idea of automata-theoretic shift tables is pioneered in [38].

VII. Conclusions and Future Work

Motivated by the recent rise of monitoring needs in embedded applications, we presented the construction of filtering Moore machines for (untimed and timed) pattern matching. The construction is automata-theoretic, realizing filters as Moore machines.

We will pursue embedded implementation of the proposed technique. In particular, we will investigate hardware acceleration by FPGA or ASIC, exploiting that our filters are Moore machines [9].

On the theoretical side, we plan to further investigate the relationship among: automata theory, lightweight formal methods for cyber-physical and embedded systems, and other fields like database and networking. The papers in §VI seem to suggest that there are many techniques waiting for being exported from one field to another.

Besides, a generalization to a distributed setting was suggested by anonymous reviewers. It is often the case with the actual embedded systems, and we believe it will be interesting future work.

References

[1] G. Reger and K. Havelund, Eds., *RV-CuBES 2017. An International Workshop on Competitions, Usability, Benchmarks, Evaluation, and Standardisation for Runtime Verification Tools, September 15, 2017, Seattle, WA, USA*, ser. Kalpa Publications in Computing, vol. 3. EasyChair, 2017. [Online]. Available: http://www.easychair.org/publications/volume/RV-CuBES_2017

[2] R. Alur and D. L. Dill, "A theory of timed automata," *Theor. Comput. Sci.*, vol. 126, no. 2, pp. 183–235, 1994. [Online]. Available: http://dx.doi.org/10.1016/0304-3975(94)90010-8

[3] D. Ulus, T. Ferrère, E. Asarin, and O. Maler, "Timed pattern matching," in *Formal Modeling and Analysis of Timed Systems - 12th International Conference, FORMATS 2014, Florence, Italy, September 8-10, 2014. Proceedings*, ser. Lecture Notes in Computer Science, A. Legay and M. Bozga, Eds., vol. 8711. Springer, 2014, pp. 222–236. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-10512-3_16

[4] ——, "Online timed pattern matching using derivatives," in *Tools and Algorithms for the Construction and Analysis of Systems - 22nd International Conference, TACAS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, ser. Lecture Notes in Computer Science, M. Chechik and J. Raskin, Eds., vol. 9636. Springer, 2016, pp. 736–751. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-49674-9_47

[5] D. Ulus, "Montre: A tool for monitoring timed regular expressions," in *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I*, ser. Lecture Notes in Computer Science, R. Majumdar and V. Kuncak, Eds., vol. 10426. Springer, 2017, pp. 329–335. [Online]. Available: https://doi.org/10.1007/978-3-319-63387-9_16

[6] M. Waga, I. Hasuo, and K. Suenaga, "Efficient online timed pattern matching by automata-based skipping," in *Formal Modeling and Analysis of Timed Systems - 15th International Conference, FORMATS 2017, Berlin, Germany, September 5-7, 2017, Proceedings*, ser. Lecture Notes in Computer Science, A. Abate and G. Geeraerts, Eds., vol. 10419. Springer, 2017, pp. 224–243. [Online]. Available: https://doi.org/10.1007/978-3-319-65765-3_13

[7] A. Bakhirkin, T. Ferrère, O. Maler, and D. Ulus, "On the quantitative semantics of regular expressions over real-valued signals," in *Formal Modeling and Analysis of Timed Systems - 15th International Conference, FORMATS 2017, Berlin, Germany, September 5-7, 2017, Proceedings*, ser. Lecture Notes in Computer Science, A. Abate and G. Geeraerts, Eds., vol. 10419. Springer, 2017, pp. 189–206. [Online]. Available: https://doi.org/10.1007/978-3-319-65765-3_11

[8] E. Asarin, O. Maler, D. Nickovic, and D. Ulus, "Combining the temporal and epistemic dimensions for MTL monitoring," in *Formal Modeling and Analysis of Timed Systems - 15th International Conference, FORMATS 2017, Berlin, Germany, September 5-7, 2017, Proceedings*, ser. Lecture Notes in Computer Science, A. Abate and G. Geeraerts, Eds., vol. 10419. Springer, 2017, pp. 207–223. [Online]. Available: https://doi.org/10.1007/978-3-319-65765-3_12

[9] R. B. Reese, *Introduction to Logic Synthesis Using Verilog HDL*. Morgan and Claypool Publishers, 2006.

[10] A. V. Aho and M. J. Corasick, "Efficient string matching: An aid to bibliographic search," *Commun. ACM*, vol. 18, no. 6, pp. 333–340, 1975.

[11] M. Krichen and S. Tripakis, "Conformance testing for real-time systems," *Formal Methods in System Design*, vol. 34, no. 3, pp. 238–304, 2009. [Online]. Available: https://doi.org/10.1007/s10703-009-0065-1

[12] C. R. Meiners, J. Patel, E. Norige, E. Torng, and A. X. Liu, "Fast regular expression matching using small tcams for network intrusion detection and prevention systems," in *19th USENIX Security Symposium, Washington, DC, USA, August 11-13, 2010, Proceedings*. USENIX Association, 2010, pp. 111–126. [Online]. Available: http://www.usenix.org/events/sec10/tech/full_papers/Meiners.pdf

[13] A. Bauer, M. Leucker, and C. Schallhart, "Runtime verification for LTL and TLTL," *ACM Trans. Softw. Eng. Methodol.*, vol. 20, no. 4, pp. 14:1–14:64, 2011. [Online]. Available: http://doi.acm.org/10.1145/2000799.2000800

[14] M. Waga, T. Akazaki, and I. Hasuo, "A boyer-moore type algorithm for timed pattern matching," in *Formal Modeling and Analysis of Timed Systems - 14th International Conference, FORMATS 2016, Quebec, QC, Canada, August 24-26, 2016, Proceedings*, ser. Lecture Notes in Computer Science, M. Fränzle and N. Markey, Eds., vol. 9884. Springer, 2016, pp. 121–139. [Online]. Available: https://doi.org/10.1007/978-3-319-44878-7_8

[15] D. Ulus and O. Maler, "Specifying timed patterns using temporal logic," in *HSCC*. ACM, 2018, pp. 167–176.

[16] T. Ferrère, O. Maler, D. Nickovic, and D. Ulus, "Measuring with timed patterns," in *CAV (2)*, ser. Lecture Notes in Computer Science, vol. 9207. Springer, 2015, pp. 322–337.

[17] R. Kandhan, N. Teletia, and J. M. Patel, "Sigmatch: Fast and scalable multi-pattern matching," *PVLDB*, vol. 3, no. 1, pp. 1173–1184, 2010. [Online]. Available: http://www.comp.nus.edu.sg/~vldb2010/proceedings/files/papers/R104.pdf

[18] A. Majumder, R. Rastogi, and S. Vanama, "Scalable regular expression matching on data streams," in *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, J. T. Wang, Ed. ACM, 2008, pp. 161–172. [Online]. Available: http://doi.acm.org/10.1145/1376616.1376635

[19] F. Yu, Z. Chen, Y. Diao, T. V. Lakshman, and R. H. Katz, "Fast and memory-efficient regular expression matching for deep packet inspection," in *Proceedings of the 2006 ACM/IEEE Symposium on Architecture for Networking and Communications Systems, ANCS 2006, San Jose, California, USA, December 3-5, 2006*, L. N. Bhuyan, M. Dubois, and W. Eatherton, Eds. ACM, 2006, pp. 93–102. [Online]. Available: http://doi.acm.org/10.1145/1185347.1185360

[20] M. Waga, I. Hasuo, and K. Suenaga, "monaa," https://github.com/MasWag/monaa, 2017.

[21] M. Waga and I. Hasuo, "Moore-machine filtering for timed and untimed pattern matching (extended ver.)," *Available at request*, 2018.

[22] S. Tripakis, "Folk theorems on the determinization and minimization of timed automata," *Inf. Process. Lett.*, vol. 99, no. 6, pp. 222–226, 2006. [Online]. Available: https://doi.org/10.1016/j.ipl.2006.04.015

[23] G. Behrmann, P. Bouyer, K. G. Larsen, and R. Pelánek, "Lower and upper bounds in zone-based abstractions of timed automata," *STTT*, vol. 8, no. 3, pp. 204–215, 2006. [Online]. Available: http://dx.doi.org/10.1007/s10009-005-0190-0

[24] *Simulink User's Guide*, The MathWorks, Inc., Natick, MA, USA, 2015.

[25] B. Hoxha, H. Abbas, and G. E. Fainekos, "Benchmarks for temporal logic requirements for automotive systems," in *1st and 2nd International Workshop on Applied veRification for Continuous and Hybrid Systems, ARCH@CPSWeek 2014, Berlin, Germany, April 14, 2014 / ARCH@CPSWeek 2015, Seattle, WA, USA, April 13, 2015.*, ser. EPiC Series in Computing, G. Frehse and M. Althoff, Eds., vol. 34. EasyChair, 2014, pp. 25–30. [Online]. Available: http://www.easychair.org/publications/paper/250954

[26] T. Liu, Y. Sun, A. X. Liu, L. Guo, and B. Fang, "A prefiltering approach to regular expression matching for network security systems," in *Applied Cryptography and Network Security - 10th International Conference, ACNS 2012, Singapore, June 26-29, 2012. Proceedings*, ser.

Lecture Notes in Computer Science, F. Bao, P. Samarati, and J. Zhou, Eds., vol. 7341. Springer, 2012, pp. 363–380. [Online]. Available: https://doi.org/10.1007/978-3-642-31284-7_22

[27] X. Zhou, B. Xu, Y. Qi, and J. Li, "MRSI: A fast pattern matching algorithm for anti-virus applications," in *Seventh International Conference on Networking (ICN 2008), 13-18 April 2008, Cancun, Mexico.* IEEE Computer Society, 2008, pp. 256–261. [Online]. Available: https://doi.org/10.1109/ICN.2008.119

[28] O. Erdogan and P. Cao, "Hash-av: fast virus signature scanning by cache-resident filters," *IJSN*, vol. 2, no. 1/2, pp. 50–59, 2007. [Online]. Available: https://doi.org/10.1504/IJSN.2007.012824

[29] R. Smith, C. Estan, S. Jha, and S. Kong, "Deflating the big bang: fast and scalable deep packet inspection with extended finite automata," in *Proceedings of the ACM SIGCOMM 2008 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Seattle, WA, USA, August 17-22, 2008*, V. Bahl, D. Wetherall, S. Savage, and I. Stoica, Eds. ACM, 2008, pp. 207–218. [Online]. Available: http://doi.acm.org/10.1145/1402958.1402983

[30] R. S. Boyer and J. S. Moore, "A fast string searching algorithm," *Commun. ACM*, vol. 20, no. 10, pp. 762–772, 1977. [Online]. Available: http://doi.acm.org/10.1145/359842.359859

[31] B. Commentz-Walter, "A string matching algorithm fast on the average," in *Automata, Languages and Programming, 6th Colloquium, Graz, Austria, July 16-20, 1979, Proceedings*, ser. Lecture Notes in Computer Science, H. A. Maurer, Ed., vol. 71. Springer, 1979, pp. 118–132. [Online]. Available: https://doi.org/10.1007/3-540-09510-1_10

[32] L. Salmela, J. Tarhio, and J. Kytöjoki, "Multipattern string matching with q-grams," *ACM Journal of Experimental Algorithmics*, vol. 11, 2006. [Online]. Available: http://doi.acm.org/10.1145/1187436.1187438

[33] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems, Joint International Conferences on Formal Modelling and Analysis of Timed Systems, FORMATS 2004 and Formal Techniques in Real-Time and Fault-Tolerant Systems, FTRTFT 2004, Grenoble, France, September 22-24, 2004, Proceedings*, ser. Lecture Notes in Computer Science, Y. Lakhnech and S. Yovine, Eds., vol. 3253. Springer, 2004, pp. 152–166. [Online]. Available: https://doi.org/10.1007/978-3-540-30206-3_12

[34] D. A. Basin, F. Klaedtke, S. Müller, and E. Zalinescu, "Monitoring metric first-order temporal properties," *J. ACM*, vol. 62, no. 2, pp. 15:1–15:45, 2015. [Online]. Available: http://doi.acm.org/10.1145/2699444

[35] E. Asarin, P. Caspi, and O. Maler, "Timed regular expressions," *J. ACM*, vol. 49, no. 2, pp. 172–206, 2002. [Online]. Available: http://doi.acm.org/10.1145/506147.506151

[36] R. Alur and T. A. Henzinger, "Back to the future: Towards a theory of timed regular languages," in *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992.* IEEE Computer Society, 1992, pp. 177–186. [Online]. Available: https://doi.org/10.1109/SFCS.1992.267774

[37] F. Franek, C. G. Jennings, and W. F. Smyth, "A simple fast hybrid pattern-matching algorithm," *J. Discrete Algorithms*, vol. 5, no. 4, pp. 682–695, 2007. [Online]. Available: https://doi.org/10.1016/j.jda.2006.11.004

[38] B. W. Watson and R. E. Watson, "A boyer-moore-style algorithm for regular expression pattern matching," *Sci. Comput. Program.*, vol. 48, no. 2-3, pp. 99–117, 2003. [Online]. Available: https://doi.org/10.1016/S0167-6423(03)00013-3

[39] C. Dima, "Real-time automata and the kleene algebra of sets of real numbers," in *STACS 2000, 17th Annual Symposium on Theoretical Aspects of Computer Science, Lille, France, February 2000, Proceedings*, ser. Lecture Notes in Computer Science, H. Reichel and S. Tison, Eds., vol. 1770. Springer, 2000, pp. 279–289. [Online]. Available: https://doi.org/10.1007/3-540-46541-3_23

[40] A. Abate and G. Geeraerts, Eds., *Formal Modeling and Analysis of Timed Systems - 15th International Conference, FORMATS 2017, Berlin, Germany, September 5-7, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10419. Springer, 2017. [Online]. Available: https://doi.org/10.1007/978-3-319-65765-3

Masaki Waga received the B.S and M.S degrees in computer science from the University of Tokyo, Tokyo, Japan, in 2016 and 2018, respectively. He is currently pursuing the Ph.D degree in informatics with SOKENDAI, Kanagawa, Japan.

His current research interests include runtime verification of cyber-physical systems and theory of timed automata.

Ichiro Hasuo is an Associate Professor at National Institute of Informatics (NII), Tokyo, Japan, where he also leads JST ERATO Metamathematics for Systems Design Project. He received the BSc and MSc degrees from the University of Tokyo (2002) and Tokyo Institute of Technology (2004), respectively, and received the PhD degree (cum laude) from Radboud University Nijmegen, the Netherlands (2008). Prior to joining NII, he was an Assistant Professor at RIMS, Kyoto University, and a Lecturer and an Associate Professor at the University of Tokyo. His research interests are in mathematical (logical, algebraic and categorical) structures in formal methods; abstraction and generalization of deductive and automata-theoretic verification techniques; and their application to cyber-physical systems.

APPENDIX

Notation A.1 ($s_0 \not\xrightarrow{w} \cdot$). *Let $\mathcal{A} = (\Sigma, S, s_0, E, S_F)$ be an NFA, and $w = a_1 \ldots a_n \in \Sigma^*$ be a word. The fact that reading w in $\mathcal{A}$ leads to no active states, that is, $\{ s \in S \mid \exists s_1, \ldots, s_{n-1} \in S. (s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \ldots \xrightarrow{a_{n-1}} s_{n-1} \xrightarrow{a_n} s \text{ in } \mathcal{A}) \} = \emptyset$, is denoted by $s_0 \not\xrightarrow{w} \cdot$.*

Lemma A.2. *Let $\mathcal{A}$ be an NFA $\mathcal{A} = (\Sigma, S, s_0, E, S_F)$, and N be a positive integer. Let $\mathcal{M}_{\mathcal{A},N} = (\Sigma_\perp, \Sigma_\perp, Q, q_0, \Delta, \Lambda)$ be the Moore machine in Def. III.3. of $\mathcal{A}$, $w = a_1 a_2 \ldots a_n$ be a word over Σ, and $\perp^N w'$ be the output word of $\mathcal{M}_{\mathcal{A},N}$ over the input word $w \perp^N$ where $w' = b_1 b_2 \ldots b_n$. For any index k of w, we have $b_k = \perp$ if and only if for any $i \in [1, k]$ and $k' \in [k, k+N-1]$, we have either $(k'-i) \bmod N < k'-k$ or $w|_{[i,k']} \notin L(\mathcal{A})$, and for any $k' \in [k, k+N-1]$ and $m \in \mathbb{Z}_{>0}$, we have $s_0 \xrightarrow{w|_{[k'-mN+1,k']}} \cdot$.*

Proof. Let $(\mathcal{S}_0, \bar{a}_0, \bar{l}_0), (\mathcal{S}_1, \bar{a}_1, \bar{l}_1), \ldots, (\mathcal{S}_{|w|+N}, \bar{a}_{|w|+N}, \bar{l}_{|w|+N}) \in Q^*$ be the run of $\mathcal{M}$ over $w \perp^N$. By the definition of Λ, we have $b_k = \perp$ if and only if we have $l_{k+N,1} = \mathsf{mask}$, which holds if and only if we have

$$\forall k' \in [k, k+N-1], (s, n) \in \mathcal{S}_{k'}. (s \notin S_F \lor n \leq k'-k) \land n \neq N \tag{7}$$

by the definition of Δ. The first part $\forall k' \in [k, k+N-1], (s, n) \in \mathcal{S}_{k'}. s \notin S_F \lor n < k'-k$ of (7) holds if and only if for any $i \in [1, k]$, we have either $(k'-i) \bmod N < k'-k$ or $w|_{[i,k']} \notin L(\mathcal{A})$. The second part $\forall k' \in [k, k+N-1], (s, n) \in \mathcal{S}_{k'}. n \neq N$ holds if and only if for any $k' \in [k, k+N-1]$ and $m \in \mathbb{Z}_{>0}$, we have $s_0 \xrightarrow{w|_{[k'-mN+1,k']}} \cdot$. $\square$

A. Proof of Thm. III.7

Proof. Let (i, j) be a pair of indices of w satisfying $w|_{[i,j]} \in L(\mathcal{A})$ and $s_i, s_{i+1}, \ldots, s_j \in S^*$ be a sequence of states such that we have $(s_0, a_i, s_i) \in E$, $s_j \in S_F$, and for each $k \in [i+1, j]$, $(s_{k-1}, a_k, s_k) \in E$ holds. Let $(\mathcal{S}_0, \bar{a}_0, \bar{l}_0), (\mathcal{S}_1, \bar{a}_1, \bar{l}_1), \ldots, (\mathcal{S}_{n+N}, \bar{a}_{n+N}, \bar{l}_{n+N}) \in Q^*$ be the run of $\mathcal{M}_{\mathcal{A},N}$ over $w \perp^N$. By the definition of the transition function Δ of $\mathcal{M}_{\mathcal{A},N}$, for each $k \in [i, j]$, we have $(s_k, n_k) \in \mathcal{S}_k$ where $n_k \in \mathbb{Z}/N\mathbb{Z}$. Thus for each $m \in \mathbb{Z}_{\geq 0}$ satisfying $i + Nm \leq j$, we have $\bar{a}_{i+Nm+N-1} = (a_{i+Nm}, a_{i+Nm+1}, \cdots, a_{i+Nm+N-1})$ and $\bar{l}_{i+Nm+N-1} = (\mathsf{pass}, \mathsf{pass}, \ldots, \mathsf{pass})$. Also, because of $s_j \in S_F$, we have $\bar{a}_{j-N+1} = (a_{j-N+1}, a_{j-N+1}, \cdots, a_j)$ and for each $m \in [N - ((j-i) \bmod N), N]$ we have $l_{j-N+1,m} = \mathsf{pass}$. Thus, for any index $k \in [i, j]$ of w, we have $a_k = b_k$. $\square$

B. Proof of Thm. III.8

Proof. Since $\max\{|w| \mid w \in L(\mathcal{A})\} < \infty$, there is an NFA $\mathcal{A}'$ such that $L(\mathcal{A}) = L(\mathcal{A}')$ holds and there is no transition from any accepting states. Let $\mathcal{A}'$ be such an NFA and $(\mathcal{S}_0, \bar{a}_0, \bar{l}_0), (\mathcal{S}_1, \bar{a}_1, \bar{l}_1), \ldots, (\mathcal{S}_{|w|+N}, \bar{a}_{|w|+N}, \bar{l}_{|w|+N}) \in Q^*$ be the run of $\mathcal{M}_{\mathcal{A},N}$ over $w \perp^N$. Since the length of a run of $\mathcal{A}'$ is not more than $N+1$ and the run is accepting if its length is $N+1$, $a_k = b_k$ holds if and only if there exist $i \in [1, k]$ and $j \in [k, k+N-1]$, we have $w|_{[i,j]} \in L(\mathcal{A})$, by Lemma A.2. $\square$

C. Proof of Thm. III.9

Proof. By Lemma A.2, if we have $b_k^{(N)} = \perp$, we have

$$\forall i \in [1, k], j \in [k, k+N-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor (j-i) \bmod N < j-k \tag{8}$$

and

$$\forall j \in [k, k+N-1], m \in \mathbb{Z}_{>0}. s_0 \xrightarrow{w|_{[j-mN+1,j]}} \cdot . \tag{9}$$

Because of (8), we have either

$$\forall i \in [1, k], j \in [k, k+N-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor (j-i) \bmod nN < j-k \tag{10}$$

or

$$\forall i \in [1, k], j \in [k, k+N-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor \exists p \in [1, n-1]. pN \leq j-i-qn \tag{11}$$

where q is the positive integer satisfying $0 \leq j-i-qnN < nN$. Because of $pN \leq j-i-qnN < Np+j-k \iff k \leq j-i-qnN-pN+k < j$ and $pN \leq j-i-qnN < Np+j-k \leq pN + N$, (11) implies $\forall i \in [1, k], j \in [k, k+N-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor \exists q \in \mathbb{Z}_{>0}. k \leq j-i-qN+k < j$, which is equivalent to $\forall i \in [1, k], j \in [k, k+N-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor \exists q \in \mathbb{Z}_{>0}. k < i+qN \leq j$. Because of (9), if there exists $q \in \mathbb{Z}_{>0}$ satisfying $k < i+qN \leq j$, we have $s_0 \xrightarrow{w|_{[i,i+qN]} \cdot w|_{[i+qN+1,j]}} \cdot$. Thus, (11) implies $w|_{[i,j]} \notin L(\mathcal{A})$ and we have (10). Besides, by (9), for any $j' \in [k, k+nN-1]$ and $m' \in \mathbb{Z}_{>0}$, we have $s_0 \xrightarrow{w|_{[j'-mN+1,j'']}} \cdot$ where $j'' = k + ((j'-k) \bmod N)$, which leads $s_0 \xrightarrow{w|_{[j'-mN+1,j'']} \cdot w|_{[j''+1,j']}} \cdot$.

Because of (9), for any $i \in [1, k]$ there exists $k' \in [k, k+N-1]$ such that we have $s_0 \xrightarrow{w|_{[i,k']}} \cdot$. Thus, for any $i \in [1, k]$ and $j \in [k+N, k+nN-1]$, we have $w|_{[i,j]} \notin L(\mathcal{A})$. Thus, we have both

$$\forall i \in [1, k], j \in [k, k+nN-1]. w|_{[i,j]} \notin L(\mathcal{A}) \lor (j-i) \bmod nN < j-k$$

and

$$\forall j \in [k, k+nN-1], m \in \mathbb{Z}_{>0}. s_0 \xrightarrow{w|_{[j-mnN+1,j]}} \cdot ,$$

and we have $b_k^{(nN)} = \perp$. $\square$

D. Proof of Thm. IV.8

Proof. If we have $w(i, j) - \tau_{i-1} \in L(\mathcal{A})$, there is an accepting run $((s_0, s_i, s_{i+1}, \ldots, s_j), (\nu_0, \nu_i, \nu_{i+1}, \ldots, \nu_j))$ of $\mathcal{A}$ over $w(i, j) - \tau_{i-1}$. Let $\bar{\nu} = (\nu_0, \nu_i, \nu_{i+1}, \ldots, \nu_j)$. By the definition of $\mathcal{A}^{N\text{-ctr}}$, there is an accepting run $(((s_0, 0), (s_i, n_i), (s_{i+1}, n_{i+1}), \ldots, (s_j, n_j)), \bar{\nu})$ of $\mathcal{A}^{N\text{-ctr}}$ over $w(i, j) - \tau_{i-1}$ where for any $k \in [i, j]$, $n_k = (k - i \bmod N) + 1$. Let $((\mathcal{S}_0, \mathcal{S}_1, \ldots, \mathcal{S}_j), (\nu'_0, \nu'_1, \ldots, \nu'_j))$ be the run of $\mathcal{A}^{N\text{-ctr-d}}$ over $w(1, j)$. By Prop. IV.6, for any $k \in [i, j]$ we have $((s_k, n_k), Z_k) \in \mathcal{S}_k$ where $s_k = (k - i \bmod N) + 1$ and $s_j \in S_F$. Let $((\mathcal{S}_0, \bar{l}_0), (\mathcal{S}_1, \bar{l}_1), \ldots, (\mathcal{S}_{n+N}, \bar{l}_{n+N}))$ be the run of $\mathcal{M}_{\mathcal{A},N}$ over $w(\perp, \tau_n)^N$. By the definition of the filter Moore machine $\mathcal{M}_{\mathcal{A},N}$, for any $m \in \mathbb{Z}_{\geq 0}$ satisfying $i + Nm - 1 \leq j$, we have $b_{i+Nm-1} = N$ and $\bar{l}_{i+Nm-1} = \mathsf{pass}^N$. Also, because of $\mathcal{S}_j \in S'_F$, we have $\bar{l}_j = l_2, l_3, \ldots, l_{N-\psi(\mathcal{S}_j)+1}, \overbrace{\mathsf{pass}, \ldots, \mathsf{pass}}^{\psi(\mathcal{S}_j)}$ where for any $k \in$

$[2, N - \psi(\mathcal{S}_j) + 1]$, l_k is either pass or mask. Thus, for any $k \in [i, j]$, we have $b_k = $ pass. $\qquad\square$

See Alg. 1.

We shall now sketch the one-clock determinization. We can split the one-clock determinization into two steps: 1) overapproximation of a timed automaton by a *real-time automaton (RTA)*; and 2) determinization of the resulting real-time automaton.

E. Overapproximation of TA by RTA

Real-time automata (RTA) [39] form a subclass of TA. Intuitively, in an RTA the enabledness of a transition is determined solely by the dwell time at the source state of the transition. This amounts, in technical terms, to the condition that an RTA has only one clock variable, and the clock is reset after each transition.

Algorithm 2, that gives an RTA $\mathcal{A}^{\mathrm{rt}}$ such that $L(\mathcal{A}) \subseteq L(\mathcal{A}^{\mathrm{rt}})$, is essentially the construction described in [11, §5.3.4]. The difference is as follows. In Line 8 we use zones and their representation by *difference bound matrices (DBM)* for constructing new guards, while in [11, §5.3.4] a region-like construction is suggested (the latter should be more expensive) . We use DBM also in Lines 6, 7 and 9; this is the same as in [11, §5.3.4].

In Line 9, we take a clock valuation $\nu \in Z'$ and first reset all the clocks in λ and additionally y, and then take t-shift for arbitrary t. We note that each endpoint of the interval I is an integer or $+\infty$, and hence the condition "$y \in I$" (e.g. in Line 10) can be represented by a constraint in $\Phi(\{y\})$.

Termination of Algorithm 2 can be established much like the termination of usual zone automata constructions.

F. Determinization of RTA

Determinization of RTA is possible unlike general TA. This is because each transition resets the unique clock y, and therefore the enabledness of a transition is determined locally.

For an RTA $\mathcal{A} = (\Sigma, S, s_0, S_F, \{y\}, E)$, its determinization $\mathcal{A}^{\mathrm{d}} = (\Sigma, \mathcal{P}(S), \{s_0\}, S_F^{\mathrm{d}}, \{y\}, E^{\mathrm{d}})$ can be defined as follows. The accepting states are defined by $S_F^{\mathrm{d}} = \{\mathcal{S} \in \mathcal{P}(S) \mid \mathcal{S} \cap S_F \neq \emptyset\}$, as usual. On transitions, for any $\mathcal{S} \in \mathcal{P}(S)$ and $a \in \Sigma$, let $\mathbb{I}_{\mathcal{S},a} = \{I_1, I_2, \ldots, I_n\}$ be the coarsest partition of $\mathbb{R}_{\geq 0}$ such that for any $i \in \{1, 2, \ldots, n\}$ and $(s, a, y \in I, \{y\}, s') \in E$, we have either $I_i \subseteq I$ or $I_i \cap I = \emptyset$. This is used in the definition of the transition set E^{d}:

$$E^{\mathrm{d}} = \big\{\, (\mathcal{S}, a, y \in I', \mathcal{S}') \mid I' \in \mathbb{I}_{\mathcal{S},a},$$
$$\mathcal{S}' = \{s' \in S \mid \exists s \in \mathcal{S}.\, \exists (s, a, y \in I'', \{y\}, s') \in E.\ I' \subseteq I''\} \,\big\} \ .$$

See Table I–VI.

TABLE I
FILTER (TORQUE)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]	$|w'|$
242,808	1	0.12	1,796.80	242,808
487,021	1	0.25	1,791.20	487,021
732,281	1	0.36	1,796.40	732,281
1,221,149	1	0.60	1,787.40	1,221,149
1,830,434	1	0.93	1,791.20	1,830,434
2,436,963	1	1.21	1,801.80	2,436,963
3,045,927	1	1.54	1,796.20	3,045,927
3,654,904	1	1.87	1,787.00	3,654,904
4,265,044	1	2.16	1,790.20	4,265,044
4,873,207	1	2.44	1,788.80	4,873,207
242,808	2	0.05	1,795.20	74,333
487,021	2	0.11	1,798.20	149,805
732,281	2	0.16	1,793.60	226,107
1,221,149	2	0.26	1,796.00	372,876
1,830,434	2	0.40	1,795.60	557,182
2,436,963	2	0.53	1,793.40	740,611
3,045,927	2	0.67	1,802.60	926,892
3,654,904	2	0.80	1,797.00	1,110,315
4,265,044	2	0.93	1,791.40	1,294,699
4,873,207	2	1.06	1,797.40	1,479,893
242,808	3	0.06	1,809.60	74,274
487,021	3	0.12	1,806.80	149,698
732,281	3	0.18	1,797.80	225,932
1,221,149	3	0.29	1,803.40	372,592
1,830,434	3	0.44	1,814.40	556,749
2,436,963	3	0.58	1,810.80	740,047
3,045,927	3	0.74	1,802.40	926,197
3,654,904	3	0.88	1,799.80	1,109,452
4,265,044	3	1.04	1,809.20	1,293,686
4,873,207	3	1.16	1,803.60	1,478,727
242,808	4	0.06	1,808.00	74,305
487,021	4	0.11	1,801.40	149,743
732,281	4	0.16	1,796.60	225,993
1,221,149	4	0.27	1,801.60	372,702
1,830,434	4	0.41	1,801.00	556,904
2,436,963	4	0.54	1,806.00	740,241
3,045,927	4	0.68	1,804.20	926,438
3,654,904	4	0.82	1,807.60	1,109,733
4,265,044	4	0.95	1,813.40	1,294,033
4,873,207	4	1.08	1,811.00	1,479,095
242,808	5	0.06	1,806.00	74,133
487,021	5	0.11	1,805.40	149,439
732,281	5	0.16	1,809.80	225,548
1,221,149	5	0.27	1,812.20	371,902
1,830,434	5	0.41	1,813.80	555,670
2,436,963	5	0.54	1,809.60	738,656
3,045,927	5	0.68	1,809.80	924,412
3,654,904	5	0.81	1,815.40	1,107,300
4,265,044	5	0.96	1,821.80	1,291,174
4,873,207	5	1.08	1,820.60	1,475,874
242,808	10	0.06	1,812.60	74,119
487,021	10	0.11	1,819.60	149,425
732,281	10	0.16	1,823.80	225,513
1,221,149	10	0.27	1,815.60	371,832
1,830,434	10	0.41	1,817.20	555,572
2,436,963	10	0.53	1,816.60	738,537
3,045,927	10	0.68	1,824.20	924,272
3,654,904	10	0.82	1,812.20	1,107,153
4,265,044	10	0.96	1,816.60	1,290,985
4,873,207	10	1.08	1,818.80	1,475,657

Algorithm 1 A Filtering Algorithm for Pattern Matching with On-The-Fly Determinization

---

Require: Word $w \in \Sigma^*$, NFA $\mathcal{A} = (\Sigma, S, s_0, E, S_F)$, $N \in \mathbb{Z}_{>0}$
Ensure: $\bot^{N-1} \cdot w' \in \Sigma_\bot$ is the output word of $\mathcal{M}_{\mathcal{A},N}$ given the input word $w \cdot \bot^{N-1}$
1: $currSN \leftarrow \{(s_0, 0)\}; \bar{a} \leftarrow \bot^N; \bar{l} \leftarrow \mathsf{mask}^N$
2: **for** $i = 1; i \leq |w|; i \leftarrow i+1$ **do**
3:      $nextSN \leftarrow \{(s_0, 0)\} \cup \{(s', n+1 \bmod N) \mid \exists (s,n) \in currSN, s' \in S. (s, w(i), s') \in E\}$
4:      $\bar{a} \leftarrow a_2, a_3, \ldots, a_N, w_i$
5:      **if** $\exists (s, N) \in nextSN$ **then**
6:          $\bar{l} \leftarrow \mathsf{pass}^N$
7:      **else if** $\exists (s, n) \in nextSN. s \in S_F$ **then**
8:          $maxN \leftarrow \max\{n \mid (s,n) \in nextSN, s \in S_F\}$
9:          $\bar{l} \leftarrow l_2, l_3, \ldots, l_{N-maxN+1}, \mathsf{pass}^{maxN}$
10:      **else**
11:          $\bar{l} \leftarrow l_2, l_3, \ldots, l_N, \mathsf{mask}$
12:      **if** $l_1 = \mathsf{pass}$ **then**
13:          $w' \leftarrow w' \cdot a_1$
14:      **else**
15:          $w' \leftarrow w' \cdot \bot$
16:      $currSN \leftarrow nextSN$
17: **for** $i = 1; i \leq N-1; i \leftarrow i+1$ **do**
18:      **if** $l_i = \mathsf{pass}$ **then**
19:          $w' \leftarrow w' \cdot a_i$
20:      **else**
21:          $w' \leftarrow w' \cdot \bot$

Algorithm 2 An over-approximation of a timed automaton by a real-time automaton.

Require: a timed automaton $\mathcal{A} = (\Sigma, S, s_0, S_F, C, E)$
Ensure: a real-time automaton $\mathcal{A}^{\mathrm{rt}} = (\Sigma, S^{\mathrm{rt}}, s_0^{\mathrm{rt}}, S_F^{\mathrm{rt}}, \{y\}, E^{\mathrm{rt}})$
1: $s_0^{\mathrm{rt}} \leftarrow (s_0, \{\mathbf{0} + t \mid t \in \mathbb{R}_{\geq 0}\}); \quad S^{\mathrm{rt}} \leftarrow \{s_0^{\mathrm{rt}}\}; \quad currS \leftarrow \{s_0^{\mathrm{rt}}\}$
2: **while** $currS \neq \emptyset$ **do**
3:      $nextS \leftarrow \emptyset$
4:      **for** $(s, Z) \in currS$ **do** $\triangleright$ Z is a zone in $(\mathbb{R}_{\geq 0})^{C \amalg \{y\}}$
5:          **for** $(s, a, \delta, \lambda, s') \in E$ **do**
6:                  $Z' \leftarrow \{\nu \mid \nu \in Z, \nu \models \delta\}$
7:                  **if** $Z' \neq \emptyset$ **then**
8:                        $I \leftarrow \{\nu|_{\{y\}} \mid \nu \in Z'\}$
9:                        $Z'' \leftarrow \{\nu[x \mapsto 0]_{x \in (\lambda \amalg \{y\})} + t \mid \nu \in Z', t \in \mathbb{R}_{\geq 0}\}$
10:                        $E^{\mathrm{rt}} \leftarrow E^{\mathrm{rt}} \cup \{((s, Z), a, y \in I, \{y\}, (s', Z''))\}$
11:                        **if** $(s', Z'') \notin S^{\mathrm{rt}}$ **then**
12:                            $nextS \leftarrow nextS \cup \{(s', Z'')\}$
13:      $currS \leftarrow nextS; \quad S^{\mathrm{rt}} \leftarrow S^{\mathrm{rt}} \cup nextS$
14: $S_F^{\mathrm{rt}} \leftarrow \{(s, Z) \in S^{\mathrm{rt}} \mid s \in S_F\}$

TABLE II
FILTER (GEAR)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]	$|w'|$
306	1	< 0.01	1,658.80	201
127,552	1	0.05	1,779.40	75,666
255,750	1	0.09	1,779.00	152,584
383,168	1	0.13	1,785.00	229,136
508,756	1	0.17	1,779.40	301,952
632,484	1	0.22	1,785.80	369,940
758,500	1	0.26	1,779.60	444,020
894,692	1	0.31	1,784.00	534,112
1,011,426	1	0.35	1,781.60	591,858
306	2	< 0.01	1,657.20	59
127,552	2	0.04	1,781.20	64,245
255,750	2	0.08	1,781.80	129,929
383,168	2	0.12	1,781.00	195,143
508,756	2	0.15	1,781.60	256,385
632,484	2	0.19	1,781.40	312,581
758,500	2	0.22	1,779.00	375,159
894,692	2	0.27	1,776.40	454,755
1,011,426	2	0.30	1,778.40	500,001
306	3	< 0.01	1,664.80	59
127,552	3	0.04	1,785.40	64,245
255,750	3	0.08	1,793.60	129,929
383,168	3	0.12	1,785.40	195,143
508,756	3	0.16	1,791.00	256,385
632,484	3	0.20	1,786.60	312,581
758,500	3	0.24	1,791.00	375,159
894,692	3	0.29	1,790.00	454,755
1,011,426	3	0.32	1,786.80	500,001
306	4	< 0.01	1,657.40	59
127,552	4	0.04	1,778.40	64,245
255,750	4	0.08	1,776.80	129,929
383,168	4	0.12	1,783.20	195,143
508,756	4	0.16	1,780.60	256,385
632,484	4	0.19	1,777.20	312,581
758,500	4	0.23	1,782.40	375,159
894,692	4	0.27	1,778.80	454,755
1,011,426	4	0.31	1,784.60	500,001
306	5	< 0.01	1,665.60	59
127,552	5	0.04	1,789.60	64,245
255,750	5	0.08	1,791.80	129,929
383,168	5	0.12	1,798.00	195,143
508,756	5	0.16	1,790.60	256,385
632,484	5	0.19	1,792.60	312,581
758,500	5	0.23	1,790.20	375,159
894,692	5	0.28	1,784.80	454,755
1,011,426	5	0.31	1,789.80	500,001
306	10	< 0.01	1,663.20	59
127,552	10	0.04	1,783.40	64,245
255,750	10	0.08	1,781.20	129,929
383,168	10	0.12	1,788.20	195,143
508,756	10	0.16	1,782.20	256,385
632,484	10	0.19	1,783.00	312,581
758,500	10	0.23	1,786.20	375,159
894,692	10	0.28	1,785.60	454,755
1,011,426	10	0.31	1,783.80	500,001

TABLE III
FILTER (ACCEL)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]	$|w'|$
708	1	0.00	1,776.60	677
218,247	1	0.10	1,903.40	190,726
436,611	1	0.19	1,921.80	380,611
655,237	1	0.29	1,913.20	572,777
870,967	1	0.38	1,914.80	761,240
1,087,411	1	0.48	1,921.40	950,127
1,304,404	1	0.57	1,924.40	1,140,025
1,527,632	1	0.67	1,922.00	1,333,230
1,739,525	1	0.76	1,927.60	1,520,334
708	2	0.01	1,810.40	287
218,247	2	0.06	1,917.00	106,107
436,611	2	0.13	1,921.80	215,710
655,237	2	0.18	1,913.00	323,056
870,967	2	0.24	1,912.60	423,590
1,087,411	2	0.30	1,933.40	517,707
1,304,404	2	0.36	1,924.40	621,700
1,527,632	2	0.43	1,918.80	752,312
1,739,525	2	0.47	1,917.20	828,584
708	3	0.01	1,817.40	154
218,247	3	0.03	1,936.40	16,226
436,611	3	0.06	1,949.40	33,137
655,237	3	0.09	1,938.00	48,475
870,967	3	0.11	1,937.60	63,868
1,087,411	3	0.14	1,940.40	81,909
1,304,404	3	0.17	1,949.80	98,221
1,527,632	3	0.20	1,942.00	113,203
1,739,525	3	0.22	1,936.80	131,265
708	4	0.01	1,820.40	143
218,247	4	0.03	1,935.40	13,678
436,611	4	0.05	1,943.00	27,300
655,237	4	0.07	1,943.80	40,133
870,967	4	0.09	1,935.60	54,500
1,087,411	4	0.12	1,943.60	66,614
1,304,404	4	0.14	1,941.60	80,141
1,527,632	4	0.16	1,935.60	93,248
1,739,525	4	0.18	1,942.20	106,825
708	5	0.01	1,842.80	1
218,247	5	0.02	1,945.00	2,844
436,611	5	0.04	1,962.20	6,190
655,237	5	0.06	1,957.00	8,531
870,967	5	0.08	1,952.80	10,509
1,087,411	5	0.10	1,976.60	13,096
1,304,404	5	0.12	1,954.60	15,790
1,527,632	5	0.14	1,958.40	19,545
1,739,525	5	0.15	1,960.20	21,170
708	10	0.01	1,890.60	1
218,247	10	0.03	1,987.80	2,340
436,611	10	0.04	2,016.20	5,287
655,237	10	0.06	2,005.20	7,068
870,967	10	0.08	2,000.80	8,514
1,087,411	10	0.10	2,009.40	10,058
1,304,404	10	0.11	2,004.40	12,129
1,527,632	10	0.13	2,001.80	16,199
1,739,525	10	0.15	2,004.60	16,277

TABLE IV
FILTERED MONAA (TORQUE)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]
242,808	1	0.24	3,725.00
487,021	1	0.40	3,709.20
732,281	1	0.57	3,695.00
1,221,149	1	0.90	3,715.20
1,830,434	1	1.31	3,724.80
2,436,963	1	1.72	3,696.20
3,045,927	1	2.14	3,692.60
3,654,904	1	2.54	3,710.40
4,265,044	1	2.94	3,721.20
4,873,207	1	3.40	3,721.00
242,808	2	0.17	3,686.00
487,021	2	0.27	3,698.60
732,281	2	0.37	3,693.00
1,221,149	2	0.56	3,711.40
1,830,434	2	0.79	3,709.60
2,436,963	2	1.04	3,712.00
3,045,927	2	1.27	3,706.60
3,654,904	2	1.50	3,722.20
4,265,044	2	1.75	3,703.80
4,873,207	2	2.04	3,716.80
242,808	3	0.17	3,715.80
487,021	3	0.27	3,711.20
732,281	3	0.37	3,695.40
1,221,149	3	0.56	3,696.00
1,830,434	3	0.81	3,700.80
2,436,963	3	1.04	3,706.20
3,045,927	3	1.28	3,701.00
3,654,904	3	1.52	3,716.80
4,265,044	3	1.77	3,698.40
4,873,207	3	2.03	3,719.00
242,808	4	0.17	3,698.80
487,021	4	0.27	3,704.40
732,281	4	0.37	3,692.20
1,221,149	4	0.56	3,697.40
1,830,434	4	0.80	3,704.20
2,436,963	4	1.03	3,705.80
3,045,927	4	1.27	3,701.20
3,654,904	4	1.51	3,705.60
4,265,044	4	1.74	3,722.80
4,873,207	4	1.99	3,707.80
242,808	5	0.17	3,688.40
487,021	5	0.27	3,691.00
732,281	5	0.37	3,691.60
1,221,149	5	0.55	3,701.80
1,830,434	5	0.80	3,722.80
2,436,963	5	1.03	3,707.20
3,045,927	5	1.28	3,713.00
3,654,904	5	1.51	3,697.80
4,265,044	5	1.75	3,708.00
4,873,207	5	1.99	3,720.00
242,808	10	0.17	3,703.80
487,021	10	0.27	3,701.40
732,281	10	0.37	3,697.80
1,221,149	10	0.55	3,708.40
1,830,434	10	0.79	3,702.80
2,436,963	10	1.02	3,710.60
3,045,927	10	1.28	3,703.20
3,654,904	10	1.50	3,710.00
4,265,044	10	1.75	3,731.60
4,873,207	10	1.98	3,714.80

TABLE V
FILTERED MONAA (GEAR)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]
306	1	0.01	2,163.60
127,552	1	0.18	2,250.00
255,750	1	0.36	2,258.60
383,168	1	0.52	2,273.20
508,756	1	0.69	2,278.20
632,484	1	0.83	2,260.00
758,500	1	1.00	2,266.00
894,692	1	1.21	2,261.20
1,011,426	1	1.32	2,255.20
306	2	0.01	2,172.60
127,552	2	0.18	2,273.60
255,750	2	0.34	2,266.20
383,168	2	0.51	2,265.80
508,756	2	0.67	2,265.40
632,484	2	0.81	2,253.20
758,500	2	0.96	2,267.40
894,692	2	1.17	2,257.00
1,011,426	2	1.29	2,250.20
306	3	0.01	2,167.20
127,552	3	0.18	2,289.80
255,750	3	0.34	2,279.00
383,168	3	0.51	2,265.20
508,756	3	0.67	2,280.60
632,484	3	0.80	2,266.80
758,500	3	0.98	2,247.20
894,692	3	1.19	2,258.20
1,011,426	3	1.29	2,243.80
306	4	0.01	2,155.80
127,552	4	0.18	2,273.40
255,750	4	0.35	2,265.20
383,168	4	0.51	2,261.20
508,756	4	0.67	2,273.40
632,484	4	0.81	2,269.00
758,500	4	0.97	2,259.80
894,692	4	1.18	2,247.40
1,011,426	4	1.29	2,265.60
306	5	0.01	2,162.00
127,552	5	0.18	2,250.00
255,750	5	0.35	2,258.00
383,168	5	0.51	2,266.80
508,756	5	0.66	2,262.20
632,484	5	0.81	2,266.40
758,500	5	0.98	2,274.80
894,692	5	1.18	2,258.40
1,011,426	5	1.29	2,262.40
306	10	0.01	2,160.20
127,552	10	0.17	2,286.00
255,750	10	0.34	2,265.20
383,168	10	0.51	2,271.60
508,756	10	0.67	2,261.60
632,484	10	0.81	2,282.40
758,500	10	0.96	2,275.20
894,692	10	1.17	2,249.60
1,011,426	10	1.30	2,265.40

TABLE VI
FILTERED MONAA (ACCEL)

$|w|$	N	Execution Time [s]	Memory Usage [kbyte]
708	1	0.20	4,953.80
218,247	1	0.30	4,976.20
436,611	1	0.40	4,968.00
655,237	1	0.51	4,964.40
870,967	1	0.67	4,995.00
1,087,411	1	0.71	4,971.00
1,304,404	1	0.81	4,995.80
1,527,632	1	0.90	5,004.00
1,739,525	1	1.01	4,957.40
708	2	0.20	4,960.40
218,247	2	0.26	4,987.40
436,611	2	0.33	4,973.40
655,237	2	0.39	4,987.20
870,967	2	0.46	4,982.00
1,087,411	2	0.52	4,964.60
1,304,404	2	0.58	4,972.80
1,527,632	2	0.65	4,974.60
1,739,525	2	0.71	4,987.40
708	3	0.20	4,959.20
218,247	3	0.22	4,999.80
436,611	3	0.25	4,966.80
655,237	3	0.28	4,970.40
870,967	3	0.30	4,982.20
1,087,411	3	0.33	4,987.20
1,304,404	3	0.36	4,967.20
1,527,632	3	0.39	4,967.40
1,739,525	3	0.42	4,982.20
708	4	0.20	4,949.80
218,247	4	0.21	4,976.80
436,611	4	0.24	4,975.40
655,237	4	0.26	4,973.20
870,967	4	0.28	4,989.00
1,087,411	4	0.30	4,975.00
1,304,404	4	0.33	4,957.00
1,527,632	4	0.35	4,990.40
1,739,525	4	0.37	4,988.80
708	5	0.20	4,955.20
218,247	5	0.21	4,985.00
436,611	5	0.22	4,975.80
655,237	5	0.23	4,978.00
870,967	5	0.24	4,994.00
1,087,411	5	0.26	4,972.60
1,304,404	5	0.28	4,996.80
1,527,632	5	0.30	4,983.80
1,739,525	5	0.31	4,983.20
708	10	0.20	4,952.60
218,247	10	0.21	4,979.40
436,611	10	0.22	5,001.00
655,237	10	0.22	4,979.00
870,967	10	0.23	4,989.80
1,087,411	10	0.24	4,980.00
1,304,404	10	0.26	4,969.40
1,527,632	10	0.30	4,998.40
1,739,525	10	0.30	4,988.40