

Checking Metric Temporal Logic with TRACE

Martijn Hendriks^{*}, Marc Geilen[†], Amir R. B. Behrouzian[†], Twan Basten^{†*}, Hadi Alizadeh[†], Dip Goswami[†]

^{*}*Embedded Systems Innovation by TNO, Eindhoven, The Netherlands*

[†]*Eindhoven University of Technology, Eindhoven, The Netherlands*

Abstract—Execution traces, time-stamped sequences of events, provide a general, domain-independent, view on the behavior of systems. They enable analysis of metrics such as latency, pipeline depth and throughput. Often, however, it is not clear what such metrics exactly mean and ad hoc methods are used to compute them. Metric Temporal Logic (MTL) can be used to address this issue: it enables the formal specification of quantitative properties on execution traces. We thus have added an MTL checking capability to the TRACE tool, which is a tool for viewing and analyzing execution traces [1]. We use a recursive memoization algorithm that generates concise explanations of the truth value of the given MTL formula. These explanations can be visualized in the TRACE viewer to aid interpretation by the user.

Keywords—Metric temporal logic, execution trace, informative prefix, explanation, visualization

I. INTRODUCTION

The design of embedded systems requires careful analysis of trade-offs between properties such as cost, quality and performance. Many of these kind of system level metrics are hard to predict, which is why modeling and prototyping play an important role in the design process. Besides quantification of the system-level metrics of interest, models and prototypes can also be used to generate *execution traces*. These are time-stamped sequences of events and form a general and unified view on the behavior of systems and models of these systems. They can be used to investigate details of (unexpected) behavior, or to identify bottlenecks or optimization opportunities. The TRACE tool has been built to visualize execution traces as Gantt charts. Recently, several generic analysis methods have also been added [1], [2].

Although terms as “latency” and “throughput” are often used to indicate system qualities that can be extracted from execution traces, it is often not clear what is exactly meant. Formalisms for property specification with a well-defined syntax and semantics alleviate this problem. Metric temporal logic (MTL) [3] is such a formalism and it is suitable for execution traces, since it enables the expression of a wide variety of quantitative real-time properties for time-stamped event sequences. In this paper we explain how we have added the capability to interpret MTL formulas on execution traces to the TRACE tool.

Related work & contribution: Gantt charts as used by the TRACE tool for visualization are omnipresent and appear, e.g., in project planning ([4]), embedded-systems analysis

[5], and model checking tools [6]. There is also a plethora of domain specific visualization tools for execution traces through other means than Gantt charts, e.g., [7], [8].

Checking the validity of an MTL formula is a fundamental problem in model-checking and runtime verification. Tableau methods construct an automaton from the formula that accepts exactly those traces that are models of the formula. Such an automaton can then be used to verify properties of systems using model checking techniques, see, e.g., [9], [10], [11]. Runtime monitoring of formulas is investigated, e.g., in [12], [13], [14], [15], [16]. These techniques usually instrument source code with code that emits time-stamped events. The runtime stream of emitted time-stamped events is then checked, one event at a time, against the formula, which often takes the form of an automaton. These techniques are on-line and do not require the whole execution trace to be stored. Our situation is somewhat different from both the model-checking and the runtime verification settings in the sense that we have access to a single static finite execution trace in the TRACE tool. This execution trace does not grow anymore, but it can be (but does not need to be!) a prefix of a longer trace that just has not been recorded. In order to be able to deal with both situations (prefix or complete trace) and to generate a concise explanation of the truth value of the property, we have chosen to develop a straightforward recursive algorithm with memoization. It is directly based on the neutral, strong and weak MTL satisfaction relations defined in [16]. The algorithm is easy to implement and uses no complicated, error-prone data structures. The memoization allows us to construct an explanation of the truth value of the formula. This explanation can be used in the visualization to give feedback to the user. Our algorithm does not compute the truth values of all subformulas in all states in the trace as, e.g., the dynamic programming approaches that are suggested in [16], [17] do, but only the relevant ones. This makes the explanation more concise.

II. METRIC TEMPORAL LOGIC

In the remainder of this paper we assume the context of sets of states S , a set $\mathbf{AP}$ of atomic propositions, and a labeling function $l : S \rightarrow 2^{\mathbf{AP}}$ that assigns to a state $s \in S$ atomic propositions that are true in that state. Furthermore, we let $\mathbb{R}$ denote the set of real numbers.




Figure 1: A TRACE Gantt-chart view of an execution in which 10 objects are processed.

A *trace* is a possibly infinite sequence $s_0, s_1, \dots$ of states in S . A *timed trace* is a possibly infinite sequence of tuples of $S \times \mathbb{R}$, say $\rho = (s_0, t_0), (s_1, t_1), \dots$, such that $t_i \leq t_{i+1}$. In the context of a timed trace, we call the state-time tuples also “states”, and ρ_i refers to (s_i, t_i) .

The syntax of MTL formulas ranged over by ϕ is inductively defined as follows:

$$\phi := \text{true} \mid p \mid \phi \wedge \phi \mid \neg \phi \mid \phi \mathbf{U}_I \phi \quad p \in \mathbf{AP}$$

where $I \subseteq [0, \infty)$ is an interval (open, closed, or half open) on $\mathbb{R}$ where the end points are defined by real numbers or ∞ . We let Φ denote the set of all MTL formulas. We assume the existence of comparison relations that check whether a real number is larger or smaller than all elements in the interval.

LTL is a fragment of MTL in which all intervals are of the form $[0, \infty)$. In this case, the intervals are omitted in the notation. Below we define the satisfaction relation $\models$ for MTL.

Definition 1 (MTL semantics). Let $\rho = (s_0, t_0), (s_1, t_1), \dots$ be an infinite timed trace. The satisfaction relation on the elements $\rho_i = (s_i, t_i)$ is defined inductively as follows:

- $\rho_i \models \text{true}$
- $\rho_i \models p$ iff $p \in l(s_i)$
- $\rho_i \models \phi_1 \wedge \phi_2$ iff $\rho_i \models \phi_1$ and $\rho_i \models \phi_2$
- $\rho_i \models \neg \phi$ iff $\rho_i \not\models \phi$
- $\rho_i \models \phi_1 \mathbf{U}_I \phi_2$ iff $\exists j$ s.t. $j \geq i$, $\rho_j \models \phi_2$, $t_j - t_i \in I$, and $\rho_k \models \phi_1$ for all $i \leq k < j$

Note that formulas are interpreted relative to timed states and not to arbitrary moments in time. We say that a timed trace ρ satisfies an MTL formula ϕ , denoted by $\rho \models \phi$, if and only if $\rho_0 \models \phi$. Some useful and well-known abbreviations are the following: $\mathbf{F}_I \phi = \text{true} \mathbf{U}_I \phi$ (finally ϕ), and $\mathbf{G}_I \phi = \neg \mathbf{F}_I \neg \phi$ (globally ϕ). The satisfaction relation $\models$ can also be defined for finite traces by restricting the scope of the existential quantifier for the until operator to the trace length. We denote the finite satisfaction relation by $\models_f$.

Example 1. We consider a pipelined processing system consisting of seven tasks, A – G, that work on a stream

of input objects. Through the use of the OCTOSIM discrete-event simulator [18] we have access to finite timed traces of this system. For instance, Fig. 1 shows a Gantt-chart representation, using the TRACE tool, for processing of 10 input objects. Atomic propositions in this setting are of the form N , N^v or N_i^v , where N is the name of the task, $v \in \{s, e\}$ indicates whether it is the start event of the task or the end event, and i is the object number. For instance, G_0^e specifies the end of task G for the first object in the stream, and G^e refers to end events of activities with name G. A number of useful MTL properties that can be used to analyze timed traces of the system for, e.g., 1000 input objects, are shown below (object ids range from 0 to 999).

The first property formalizes that the first object (with id 0) has been completely processed within 25 time units.

$$\mathbf{F}_{[0,25]} G_0^e \quad (1)$$

The second property formalizes that the total execution time is at most 6500 time units.

$$\mathbf{F}_{[0,6500]} G_{999}^e \quad (2)$$

The third property is a variation of the second property that, however, does not refer to the object identifier 999. Instead, a temporal operator is used to ensure that there will be no more end events of activity G after 6500 time units.

$$\mathbf{F}_{[0,6500]} \mathbf{G}_{(0,\infty)} \neg G^e \quad (3)$$

The fourth property formalizes that the per object processing time is at most 70 time units.

$$\bigwedge_{i=0}^{999} \mathbf{G}(A_i^s \Rightarrow \mathbf{F}_{[0,70]} G_i^e) \quad (4)$$

The fifth property formalizes that the throughput is at least 10/65 in every window of 10 consecutive end events of task G.

$$\bigwedge_{i=0}^{989} \mathbf{G}(G_i^e \Rightarrow \mathbf{F}_{[0,65]} G_{i+10}^e) \quad (5)$$

The sixth property formalizes that the pipeline is always at least four objects deep, except for the start and end areas in the trace (note that this is an LTL property).

$$\bigwedge_{i=9}^{990} \mathbf{G}(A_{i+3}^s \Rightarrow \mathbf{F}_I G_i^e) \quad (6)$$

The seventh property formalizes that after any end event of task G, another end event of task G happens within 3 and 15 time units.

$$\mathbf{G}(G^e \Rightarrow \mathbf{F}_{[3,15]} G^e) \quad (7)$$

These examples illustrate the flexibility and expressive power of MTL. The formalism allows us to define what we exactly mean with, e.g., pipeline depth, latency and throughput.

In the present paper we assume that we have a finite execution trace of some system. This trace can be obtained from a real system, but also, for instance, from a discrete-event simulation model. We distinguish two situations: (i) the trace represents the full execution of some process, or (ii) the trace is a prefix of some ongoing, possibly infinite, process. An example of the first situation is the execution trace of an image processing pipeline that processes 10 images and then is done. In this case, we can apply the MTL semantics for finite traces. An example of the second situation is a part of an execution obtained from a running web server. In this case, however, application of the finite MTL semantics is not appropriate, because there is an unknown extension of the trace that can affect the truth value of the property. For this situation, we use the notion of *informative prefixes* [19] which is explained in more detail in the next section.

III. INFORMATIVE PREFIXES

Consider a finite prefix ρ of some timed trace and an MTL formula ϕ . Then we say that ρ is a *bad* prefix if and only if every extension of ρ dissatisfies ϕ . Dually, ρ is a *good* prefix if and only if every extension of ρ satisfies ϕ . A *neutral* prefix is neither good nor bad. Intuitively, an *informative* prefix tells the whole story about the (dis)satisfaction of an MTL formula [19]. For instance, the prefix $(p, 0), (p, 1), (p, 2), (q, 3)$ is bad for $\mathbf{G}p$ and it is also informative. The prefix is also bad for $\mathbf{F}(p \wedge \neg p)$, but not informative because the dissatisfaction for any extension depends on the unsatisfiability of $p \wedge \neg p$. This information is not to be found in the prefix itself.

Below we follow the approach of [16] (which uses the truncated semantics of [20]) to define strong and weak satisfaction relations for MTL formulas and timed traces, $\models^+$ and $\models^-$ respectively. If $\rho \models^+ \phi$ then for any finite extension of ρ , say ρ' , holds that $\rho' \models_f \phi$. Thus, the finite prefix proves the satisfaction of the formula and it is stable. We say that ρ is an *informative good prefix* for ϕ . On the other hand, $\rho \not\models^- \phi$ implies that no finite extension of ρ exists that models ϕ . In that case, ρ is an *informative bad prefix* for ϕ . Otherwise, we say that ρ is non-informative.

Definition 2. Let $\rho = (s_0, t_0), \dots, (s_n, t_n)$ be a finite prefix of a timed trace. The strong and weak satisfaction relations $\models^+$ and $\models^-$ on the elements $\rho_i = (s_i, t_i)$ are defined inductively as follows:

- $\rho_i \models^+ \text{true}$
- $\rho_i \models^+ p$ iff $p \in l(s_i)$
- $\rho_i \models^+ \phi_1 \wedge \phi_2$ iff $\rho_i \models^+ \phi_1$ and $\rho_i \models^+ \phi_2$
- $\rho_i \models^+ \neg \phi$ iff $\rho_i \not\models^- \phi$
- $\rho_i \models^+ \phi_1 \mathbf{U}_I \phi_2$ iff $\exists i \leq j \leq n$ $\rho_j \models^+ \phi_2$, $t_j - t_i \in I$, and $\rho_k \models^+ \phi_1$ for all $i \leq k < j$

and

- $\rho_i \models^- \text{true}$

- $\rho_i \models^- p$ iff $p \in l(s_i)$
- $\rho_i \models^- \phi_1 \wedge \phi_2$ iff $\rho_i \models^- \phi_1$ and $\rho_i \models^- \phi_2$
- $\rho_i \models^- \neg \phi$ iff $\rho_i \models^+ \phi$
- $\rho_i \models^- \phi_1 \mathbf{U}_I \phi_2$ iff either (1) $\exists i \leq j \leq n$ $\rho_j \models^- \phi_2$, $t_j - t_i \in I$, and $\rho_k \models^- \phi_1$ for all $i \leq k < j$ or (2) $t_n - t_i < \sup(I)$ and $\rho_k \models^- \phi_1$ for all $i \leq k \leq n$

Note that the definitions for $\neg$ constructs use the dual satisfaction relation, and that the definition of $\models^-$ for the until operator consists of a disjunction that includes the case in which the end of the trace is reached.

Proposition 1 ([16]). Let ρ be a finite timed trace and let i be a position in that trace. Then $\rho_i \models^+ \phi \Rightarrow \rho_i \models_f \phi$ and $\rho_i \models_f \phi \Rightarrow \rho_i \models^- \phi$.

Proposition 2 ([16]). Let ρ be a prefix of some finite timed trace ρ' . Then $\rho_i \models^+ \phi \Rightarrow \rho'_i \models^+ \phi$ and $\rho_i \not\models^- \phi \Rightarrow \rho'_i \not\models^- \phi$.

IV. AN ALGORITHM FOR INFORMATIVE PREFIXES

In this section we present an algorithm that computes the informative value of a finite prefix of some timed trace for a given formula by straightforward application of the semantics in Def. 2. (A detailed description of the algorithm and proofs of theorems can be found in [21].) Below we use standard tree-valued logic over the values $\{T, F, U\}$, where T stands for true, F for false and U for unknown.

The algorithm consists of two mutually recursive procedures, shown in Alg. 1 and Alg. 2. Both have four arguments: A finite timed trace $\rho = (s_1, t_1), \dots, (s_n, t_n)$, an MTL formula ϕ , a state index $1 \leq i \leq n$, and a table $v : \Phi \times \mathbb{N} \mapsto \{T, F, U, ?\}$ ($v(\phi, i) = ?$ means that *compute* has not yet been called for formula ϕ and index i).

The user calls *getOrCompute*($\rho, \phi, 0, v_0$) in order to compute the truth value of formula ϕ on trace ρ , where v_0 initially maps all subformula-state combinations to ?. The *getOrCompute* function is a wrapper around the *compute* function, and it performs memoization of computed values using the lookup table. The main *compute* algorithm follows the semantics of the satisfaction relation in a straightforward way for all but the case of the temporal operator. Alg. 2 shows how this case is handled. When the call to *getOrCompute*($\rho, \phi, 0, v_0$) returns, the table v_0 contains the values that have been computed in order to reach the conclusion. This serves as a basis for the visualization of the explanation of the formula.

Theorem 1. Let ρ be a finite timed trace of length n , let ϕ be an MTL formula, and let $0 \leq i \leq n$. Then we have that *getOrCompute*(ρ, ϕ, i, v_0) = $T \Leftrightarrow \rho_i \models^+ \phi$, and also that *getOrCompute*(ρ, ϕ, i, v_0) = $F \Leftrightarrow \rho_i \not\models^- \phi$.

Theorem 2. Let ρ be a finite timed trace and let ϕ be an MTL formula. The time complexity of the call to *getOrCompute*($\rho, \phi, 0, v_0$) is $\mathcal{O}(|\phi| \cdot |\rho|^2)$, where $|\phi|$ denotes the size of the formula, and $|\rho|$ is the length of the trace.

Algorithm 1: $getOrCompute(\rho, \phi, i, v)$

```
1  $result \leftarrow v(\phi, i)$ 
2 if  $result = ?$  then
3    $result \leftarrow compute(\rho, \phi, i, v)$ 
4    $v(\phi, i) \leftarrow result$ 
5 return  $result$ 
```

Algorithm 2: $compute(\rho, \phi, i, v)$

```
1 // We have  $\phi = \phi_1 \mathbf{U}_I \phi_2$ 
2  $c_1 \leftarrow false, c_{11} \leftarrow true$ 
3  $c_2 \leftarrow true, c_{22} \leftarrow false, c_3 \leftarrow t_n - t_i \geq sup(I)$ 
4 for  $j = i$  to  $n$  do
5    $r_1 \leftarrow getOrCompute(\rho, \phi_1, j, v)$ 
6    $r_2 \leftarrow getOrCompute(\rho, \phi_2, j, v)$ 
7    $c_1 \leftarrow c_1 \vee (r_2 = T \wedge t_j - t_i \in I \wedge c_{11})$ 
8    $c_2 \leftarrow c_2 \wedge (r_2 = F \vee t_j - t_i \notin I \vee c_{22})$ 
9    $c_3 \leftarrow c_3 \vee (r_1 = F)$ 
10   $c_{11} \leftarrow c_{11} \wedge (r_1 = T)$ 
11   $c_{22} \leftarrow c_{22} \vee (r_1 = F)$ 
12 if  $c_1$  then
13   return  $T$ 
14 else if  $c_2 \wedge c_3$  then
15   return  $F$ 
16 else
17   return  $U$ 
```

We have applied two optimizations to Alg. 2 to (i) improve the performance, and (ii) to reduce the number of computed entries in the lookup table v to make the explanation for the user more concise. The first optimization consists of lazy evaluation of conjunction arguments. When one of the conjuncts is F , then the other one does not need to be computed. When both conjuncts are unknown, we first compute the one with the smallest formula depth (based on the number of nested temporal operators).

The second optimization consists of bounding rules for the temporal operator. The evaluation of the temporal operator in Alg. 2 is written for clarity. Clearly, the fact that the loop always iterates to the end of the trace is not very efficient. We can use the following observations to short-circuit the loop in several cases:

- Once c_1 becomes *true*, it remains *true*. Therefore, the return statement guarded by c_1 can be pulled into the loop. Furthermore, the evaluation of r_2 can wait until after this if-then construct.
- Similarly, once c_{22} or c_3 becomes *true* it remains *true*. This also holds for $t_j - t_i > sup(I)$, which implies $t_j - t_i \notin I$. This gives that if at the end of the loop both c_2 and c_3 hold, and this is the last iteration or c_{22} holds or $t_j - t_i > sup(I)$ holds, we have that in the

next iterations (if any), c_2 (and c_3) also hold, and we can return F .

- If c_1 is *false* now, and c_{11} is *false* or $t_j - t_i > sup(I)$ then c_1 will remain *false*. Furthermore, if c_2 is *false* then it will remain *false*. This implies that the cases to return T or F will not be satisfied, and we can return U .

(The resulting algorithm is shown in [21].)

V. IMPLEMENTATION IN TRACE

A. The TRACE tool

The TRACE tool [1] is a generic visualization tool for activities and their resource usage. It is centered around the concept of a *claim*, which represents an activity on a single resource. A claim has a start and end time stamp, a resource, and a number of configurable attributes (e.g., the name of the activity), which are represented by a key-value mapping of attribute names to attribute values. A set of claims is visualized by TRACE as a Gantt chart. The tool offers many features, e.g., filtering, zooming, grouping, coloring based on attribute values, and to do analysis such as critical-path analysis and difference analysis [2].

B. MTL in TRACE

To specify MTL formulas in the TRACE context, we first define both states and atomic propositions as key-value mappings. We say that a state defined by mapping M_1 satisfies atomic propositions defined by mapping M_2 if and only if $M_2 \subseteq M_1$. Second, we convert a claim c with start time-stamp t_s and end time-stamp t_e to two state-time tuples:

- $(m \cup \{mtl \mapsto s\}, t_s)$ and
- $(m \cup \{mtl \mapsto e\}, t_e)$,

where m is the key-value mapping given by the attributes of the claim and its values. Note that the additional *mtl* element in the mapping is used to indicate the start or the end of the claim. Third, we order the set of all state-time tuples associated with a set of claims according to their time-stamps, which gives a timed trace. Note, however, that we arbitrarily order events with identical time-stamps even though different orderings are generally distinguishable by MTL formulas. This is a consequence of the fact that the TRACE input format does not support information on the order of events with an identical time-stamp.

Besides the core constructs of MTL given in Sec. II, we have added the logical operators *or* and *implies* and the temporal operators *globally* (represented by $\mathbf{G}$) and *finally* (represented by $\mathbf{F}$) for convenience to the expression language. Furthermore, we have added a top-level conjunction construct in combination with integer expressions to be able to conveniently specify a set of formulas. An example that contains many of these ingredients is Prop. 5 in Ex. 1, which looks as follows in the syntax of the TRACE MTL expression language:



Figure 2: A TRACE view in the ECLIPSE Integrated Development Environment.

```

/\ (i=0...989)
  G (
    {name=G, id=i, mtl=e}
    =>
    F_[0, 65] {name=G, id=(i+10), mtl=e}
  )

```

The top-level conjunction instantiates the following formula (which starts with element **G**) for every value of $0 \leq i \leq 989$. It thus creates 990 formulas which express that the distance between the end of activity **G** with id i and the end of activity **G** with id $i + 10$ is at most 65 for all $0 \leq i \leq 989$. This construct is very convenient when claims are associated, e.g., with an object identifier, image number, etc.

C. User interface

Figure 2 shows the ECLIPSE IDE with the TRACE plugin installed. The window has (1) a project explorer view of the files in the workspace, (2) a number of TRACE toolbar items, (3) the main Gantt-chart view, (4) the MTL explanation view, and (5) a concrete explanation of the property being analysed overlayed on the Gantt-chart view. In this case, the Gantt chart visualizes a (part of a) run of the system from Ex. 1 for 1000 objects while Prop. 7 is being analysed.

The project explorer associates files with an *mtl* extension with the MTL dialog. Double clicking an *mtl* file when a trace is open, opens the MTL dialog with the contents of the *mtl* file. The MTL dialog has several configuration options: (i) whether to apply it to the set of filtered claims or to the whole set of claims, (ii) whether to interpret the trace as a prefix or not, (iii) whether to generate explanations of

computed values. If the OK button of the MTL dialog is pressed, the MTL specification is checked against the current trace.

We generate explanations in two forms. First, the claims that are relevant for the truth value of the formula can be highlighted. This is a rather straightforward visualization based on a marking of states and their claims during the run of the algorithm. Nevertheless, it is often very useful and allows us to zoom into relevant parts of the trace quickly for diagnosis. The second form consists of an annotation of (part of) the time axis with the truth values of all subformulas of the formula that is checked. Also this annotation is constructed on-the-fly during the run of the algorithm. This annotation allows the user to trace the result according to the semantics. Figure 2 shows the user interface after checking Prop. 7 and after visualizing the second type of explanation. Below the time axis are the three subformulas of the implication. A red bar means that the property is not satisfied in any state in that time interval, and a green bar means that it is satisfied. A blue bar indicates that the property may or may not be satisfied by an arbitrary extension of this prefix. For this property, the key is that the implication $\{name=G, mtl=e\} \Rightarrow F_{[3, 15]} \{name=G, mtl=e\}$ holds for every end event of **G** but the last one. However, an extension of the trace could have more end events of task **G** within the indicated interval. Therefore, $F_{[3, 15]} \{name=G, mtl=e\}$ may or may not be satisfied by the last state in the prefix, hence the blue marking.

D. Performance and scalability

Despite the time complexity of $\mathcal{O}(|\phi| \cdot |\rho|^2)$ we believe that the algorithm is practically feasible. For instance, Props. 1–3 and 7 can be checked for a trace of 25,000 objects (approximately 10^6 states) within a second on a regular laptop computer. (Note that the trace is an informative good prefix for Props. 1 and 2, and non-informative for Props. 3 and 7.) Properties 4–6 can be checked in approximately 15 seconds for a trace of 2500 objects (approximately 10^5 states). The 2500 separate formulas per property are checked in parallel by our implementation using all available CPU cores.

Finally, we note that the worst-case quadratic complexity of our algorithm is not observed for the properties above, but it can be observed when checking formulas such as the following pathological example: $\mathbf{GGGG}^{true}$.

VI. CONCLUSIONS

Metric Temporal Logic provides a flexible mechanism to specify quantitative properties of execution traces. We have shown how we have added an MTL checking capability to the TRACE tool using a recursive memoization algorithm that can generate concise explanations of the truth value of the given MTL formula. These explanations can be visualized in the TRACE viewer to aid the user in understanding. The combination of a generic and flexible visualization layer for execution traces and non-trivial, domain independent, analysis capabilities such as critical-path analysis, trace difference analysis and now MTL checking, that the TRACE tool provides, is something we have not seen elsewhere. The TRACE tool is available through <http://trace.esi.nl/>.

ACKNOWLEDGEMENT

This research is supported by the ARTEMIS joint undertaking under grant agreement no 621439 (ALMARVI).

REFERENCES

- [1] Embedded Systems Innovation by TNO, “TRACE website,” <http://trace.esi.nl>.
- [2] M. Hendriks, J. Verriet, T. Basten, B. Theelen, M. Brassé, and L. Somers, “Analyzing execution traces – critical-path analysis and distance analysis,” *Submitted to STTT*, 2015.
- [3] R. Alur and T. Henzinger, “Real-time logics: complexity and expressiveness,” *Information and Computation*, vol. 104, pp. 390–401, 1993.
- [4] Atlassian, “Gantt plugins for JIRA,” <https://marketplace.atlassian.com/search?application=jira&q=gantt>.
- [5] National Instruments, “Profiling VI Execution With the LabVIEW Desktop Execution Trace Toolkit,” <http://www.ni.com/white-paper/8083/en/>.
- [6] G. Behrmann et al., “Uppaal 4.0,” in *Proceedings of the 3rd international conference on the Quantitative Evaluation of Systems*, ser. QEST ’06. IEEE Computer Society, 2006.
- [7] Intel, “Intel VTune Amplifier 2016,” <https://software.intel.com/en-us/intel-vtune-amplifier-xe>.
- [8] Microsoft, “Windows Performance Analyzer,” <https://msdn.microsoft.com/en-us/library/hh448170.aspx>.
- [9] R. Alur, T. Feder, and T. Henzinger, “The benefits of relaxing punctuality,” *Journal of the ACM*, vol. 43, no. 1, Jan. 1996.
- [10] M. Geilen, “An improved on-the-fly tableau construction for a real-time temporal logic,” in *Computer Aided Verification*, ser. Lecture Notes in Computer Science, vol. 2725. Springer Berlin Heidelberg, 2003.
- [11] D. Ničković and N. Piterman, “From mtl to deterministic timed automata,” in *Proceedings of the 8th International Conference on Formal Modeling and Analysis of Timed Systems*, vol. 6246. Springer-Verlag, 2010.
- [12] I. Lee, S. Kannan, M. Kim, O. Sokolsky, and M. Viswanathan, “Runtime assurance based on formal specifications,” in *In Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, 1999.
- [13] D. Drusinsky, “The temporal rover and the atg rover,” in *Proceedings of the 7th International SPIN Workshop on SPIN Model Checking and Software Verification*, 2000, pp. 323–330.
- [14] K. Havelund and G. Roşu, “Testing linear temporal logic formulae on finite execution traces,” Tech. Rep., 2001.
- [15] P. Thati and G. Roşu, “Monitoring algorithms for metric temporal logic specifications,” *Electronic Notes in Theoretical Computer Science*, vol. 113, pp. 145–162, 2005.
- [16] H. Ho, J. Ouaknine, and J. Worrell, “Online monitoring of metric temporal logic,” in *Runtime Verification*, ser. Lecture Notes in Computer Science. Springer International Publishing, 2014, vol. 8734.
- [17] N. Markey and J.-F. Raskin, “Model checking restricted sets of timed paths,” *Theoretical Computer Science*, vol. 358, no. 2, 2006.
- [18] Hendriks, M. et al., “A blueprint for system-level performance modeling of software-intensive embedded systems,” *International Journal on Software Tools for Technology Transfer*, 2014.
- [19] O. Kupferman and M. Vardi, “Model checking of safety properties,” *Formal Methods in System Design*, vol. 19, no. 3, 2001.
- [20] C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, and D. Van Campenhout, “Reasoning with temporal logic on truncated paths,” in *Computer Aided Verification*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2003, vol. 2725.
- [21] M. Hendriks, M. Geilen, A. Behrouzian, T. Basten, H. Alizadeh, and D. Goswami, “Checking metric temporal logic with TRACE,” Tech. Rep. ESR-2016-01, 2016, <http://www.es.ele.tue.nl/esreports.php>.