

Towards a Digital Highway Code using Formal Modelling and Verification of Timed Automata

Gleifer Vaz Alves

Federal University of Technology
Parana, Brazil
gleifer@utfpr.edu.br

Maike Schwammburger*

University of Oldenburg
Oldenburg, Germany
schwammburger@informatik.uni-oldenburg.de

One of the challenges in designing safe, reliable and trustworthy Autonomous Vehicles (AVs) is to ensure that the AVs abide by traffic rules. For this, the AVs need to be able to understand and reason about traffic rules. In previous work, we introduce the spatial traffic logic USL-TR to allow for the unambiguous, machine-readable, formalisation of traffic rules. This is only the first step towards autonomous traffic agents that verifiably follow traffic rules. In this research preview, we focus on two further steps: a) retrieving behaviour diagrams directly from traffic rules and b) converting the behaviour diagrams into timed automata that are using formulae of USL-TR in guards and invariants. With this, we have a formal representation for traffic rules and can move towards the establishment of a Digital Highway Code. We briefly envision further steps which include adding environment and agent models to the timed automata to finally implement and verify these traffic rule models using a selection of formal verification tools.

1 Motivation

Our endeavor is to support the inclusion of traffic rules into the design process of Autonomous Vehicles (AVs). For this, we envision a Digital Highway Code, which is a machine-readable version of the natural language traffic rules as they exist for human drivers in various languages for various countries [13, 8]. Using such a Digital Highway Code, AVs would be able to understand and reason about traffic rules. With this, explainability of AV behaviour would be supported, e.g. if and why a traffic rule was violated.

We identify several necessary steps to get from natural language traffic rules to a Digital Highway Code. For this, we start with steps towards a formalisation of traffic rules, followed by some steps towards a verification of the formalised rules, as we visualise in Fig. 1. For a structural transition of natural language traffic rules to a formalised, digital, and machine-readable, format, we identify two characteristics in each traffic rule: *spatial* and *temporal aspects*. Examples for such spatial elements comprise that there is a traffic sign or a road junction ahead of an AV, or that a safe gap exists to move into. On the other hand, temporal aspects comprise timing behaviour (e.g. that an action does not happen immediately) and the sequentiality of a traffic rule, where, e.g., for changing a lane on a motorway, a car must first check that a sufficiently large free safe gap exists to change into, then set a turn signal, and only then change the lane.

For formalising spatial aspects of traffic rules, in previous work [20], we introduce the traffic logic *Urban Spatial Logic for Traffic Rules (USL-TR)*. In this research preview, we focus on the temporal aspects of traffic rules. We use selected rules from the UK Highway Code [13] as examples. Our formalisation of traffic rules follows two steps; first, we translate the traffic rules to behavioural diagrams,

*M. Schwammburger was supported by the German Research Council (DFG) in the PIRE Projects SD-SSCPS and ISCE-ACPS under grant no. FR 2715/4-1 and FR 2715/5-1.




Figure 1: The different steps from natural language traffic rules towards a Digital Highway Code.

i.e. UML activity diagrams [14]. With this, we capture the sequential aspects of traffic rules (e.g. that an action that is taken after another action or event). Note that these behaviour diagrams do not yet include formal aspects and are an intermediate step between an informal, often ambiguous, natural language traffic rule and a fully formalised traffic rule that is understandable by an agent steering an AV.

Next, we annotate the behaviour diagrams with formulae of USL-TR as guards and invariants to include the spatial traffic aspects into the behaviour diagrams. Further on, we add clock constraints and clock resets to specify timing behaviour of the traffic rule. The result is an extended timed automaton [1], representing both the spatial and the temporal aspects of the respective traffic rule. With this, we add a formalisation layer to traffic rules, allowing for extracting existent elements from them. Also, this formalised version of traffic rules allows for different verification approaches.

We not only strive to formalise traffic rules, but also to analyse and verify them. With this, we can, for instance, identify conflicting traffic rules and show that autonomous agents can actually safely follow a Digital Highway Code. For the verification part, we suggest two directions: i) implement the timed automata representations directly into the model-checking tool UPPAAL [4] and ii), enhance the timed automata traffic rules with an agent and environment model. For the first direction, i), we suggest a parallel composition of different traffic rules to identify potentially conflicting traffic rules. For this direction, we intend to adapt an existing UPPAAL implementation of a timed automaton crossing controller that uses a predecessor logic of USL-TR [5]. For the second direction, ii), synchronisation channels are added to model the communication between the agent and the environment. Thus, we are not only considering the single agent endowed with Traffic Rules but also a necessary adaptation so that this agent works as an AV endowed with a Digital Highway Code, capable of communicating with the corresponding environment. Using this agent and environment model, we envision a formal verification of agents endowed with traffic rules that follows [2, 3] by using the MCAPL framework [12] and the agent programming language GWENDOLEN [11]. We aim to provide a formalised and verifiable Digital Highway Code with these two directions.

Related approaches include the so-called *Traffic Sequence Charts (TSCs)* [10]. TSCs are based on

Life Sequence Charts (LSCs) [9], thus also capable of describing dynamic system behaviour. An advantage of our approach, in comparison to TSCs, is that our timed automaton-based traffic rule formalisation allows the use of existing verification techniques like model-checking with UPPAAL. Next, while the industrial standardisation effort *OpenSCENARIO* [16] allows for the specification of dynamic traffic situations for simulation efforts, it lacks a formal semantics. In contrast, the timed automata specifications that we derive with our approach support the formal analysis of traffic rules.

In Sect. 2, we briefly introduce the logic USL-TR and in Sect. 3, we exemplarily show respective behaviour diagrams that we build for two of the traffic rules from the UK Highway Code. We add timing behaviour and USL-TR guards to the behaviour diagrams to derive our timed automata in Sect. 4. In Sect. 5, we give an outlook over the steps that are contained within the “verification” node of Fig. 1.

2 Spatial Aspects: USL-TR

In previous work [20], we introduced *Urban Spatial Logic for Traffic Rules (USL-TR)* for reasoning about traffic manoeuvres and rules at urban intersections. USL-TR is a spatial logic and is an extension of *Urban Multi-lane Spatial Logic (UMLSL)* that was introduced in [19].

Using USL-TR, one can specify spatial features at intersections, e.g. that the space in front of a car A is free in the traffic situation that is depicted in Fig. 2: $re(A) \frown free$. The atom $re(E)$ comprises the so-called reservation of space for car E . Moreover, the operator \frown is inspired by the chop operator from Interval Temporal Logic [17] and splits one larger horizontal interval of space into two smaller ones. Consider the following more complex formula $sg_I(A)$:

$$sg_I(A) \equiv \langle (re(A) \wedge \neg cs) \frown (free \wedge \neg cs) \frown (sg(A) \wedge cs) \rangle,$$

The formula $sg_I(A)$ is three-parted: we specify that first, there is a reservation of car A not yet on the intersection ($re(E) \wedge \neg cs$) followed by an interval of free space, also not yet on the intersection ($free \wedge \neg cs$), followed by a free safe gap on the intersection that is large enough for car A ($sg(A) \wedge cs$). The formula $sg_I(A)$ holds in the shaded part of Fig. 2. The abbreviation $sg(A)$ itself comprises the formula $sg(A) \equiv free \wedge \ell \geq size_A$, stating that there is a free space greater or equals the size of car A . For more details on USL-TR, we refer to [20].




Figure 2: Exemplary traffic situation.

Note that with the syntactical element $\ell \geq size_A$, we leave the actual size $size_A$ of the safe gap for car A abstract at this point of the specification. The actual semantics of $size_A$ depends on a variety of factors, e.g., the type of the road (asphalt, sand, ...), the weather (sunny, heavy rain, fog, ...) and the power of the brakes. As there exist specific traffic rules for safe gaps for such special cases, it will be

of interest to refine the semantics of a safe gap at a later point in our work. An approach to connect a predecessor logic of the spatial logic USL-TR with continuous models and closer-to-reality variable valuations is provided in [18].

3 Sequential Aspects: Behaviour Diagrams

With USL-TR, we can formally specify a variety of spatial aspects of traffic rules. As shown in Fig. 1, our next step is to specify the sequential aspects that are generally contained within a traffic rule. For this, behaviour diagrams, e.g. UML activity diagrams [14] or simple flowcharts [15], are a favourable choice, as they allow to model dynamic and sequential system behaviour. In our case, such dynamic behaviour includes checking for safe gaps, while simultaneously watching out for other road users or reacting to dynamic traffic evolutions. Often, the behaviour that is described within a traffic rule also involves choices to be made, e.g. whether an intersection is entered or not. In this section, we showcase examples of how to extract behaviour diagrams from the UK Highway Code rules 170 and 171 [13]. For these two rules, we previously formalised their spatial aspects using USL-TR in [20]. In both rules, the overall goal of the AV is to safely enter a road junction.

Rule 170. We summarise this rule to the following three parts, thereby condensing cyclists, horse riders, motorcyclists and so forth under the general term of “road users” (cf. [13]): a) You should watch out for road users. b) Watch out for pedestrians crossing a road junction into which you are turning. If they have started to cross they have priority, so give way. and c) Do not cross or join a road until there is a (safe) gap large enough for you to do so safely.

In Fig. 3, we present the corresponding behaviour diagram for Rule 170, which describes the sequence of actions and choices that are included into rule 170. We identify a sequence of three states:

- “Away from Road Junction”: This is the initial state, where rule 170 does not yet apply, as no road junction is close.
- “At Road Junction”: When the AV arrives at an intersection, some actions are necessary.
- “On Road Junction”: After all parts of rule 170 have been applied, the AV can finally enter the intersection.




Figure 3: Rule 170: Behaviour Diagram

Note that in our meaning an AV is “at the road junction”, if it is closer to it than a distance constant d_c . This distance d_c depends on different factors, similarly to what we observed for the safe gap earlier on p. 80. While moving from “Away from Road Junction” to “At Road Junction” happens immediately (i.e. upon crossing that distance threshold), between the next states some choices need to be made; parts 1 and 2 of rule 170 specify to watch out for both pedestrians and road users (e.g. other AVs, cyclists, ...) and part 3 states that a road junction may only be entered, if there is a safe gap large enough to do so.

Rule 171. The rule reads as follows: “You must stop behind the line at a junction with a ‘Stop’ sign and a solid white line across the road. Wait for a safe gap in the traffic before you move off” [13]. Fig. 4 shows the behaviour diagram we construct for rule 171. The sequence of main states is the same as for the rule before, with the AV first being away from the road junction and then lastly being on the road junction. Between the first and the last state, a combined traffic sign needs to be noticed: A stop sign with a solid white line on the road. After that, the AV again needs to check for a safe gap before it may enter the road junction. We understand that it would be possible (in a further step) to combine the diagrams from Figures 3 and 4 into a single diagram since they share some states. However, we firstly wish to describe each rule separately.




Figure 4: Rule 171: Behaviour Diagram

4 Spatio-Temporal Aspects: Extended Timed Automata

We have captured the spatial elements of traffic rules using USL-TR (cf. [20]) and the sequential aspects using behaviour diagrams (cf. previous section). We now add time constraints to the behaviour diagrams to model all temporal aspects of traffic rules. We also add the spatial USL-TR aspects to the diagrams so that the resulting model is a timed automaton [1] that uses USL-TR formulae as guards and invariants. Thus, this timed automaton captures both spatial and temporal aspects of traffic rules. Such a type of extended timed automaton that uses USL-TR’s predecessor logic UMLSL as guards and invariants, as well as particular actions for traffic manoeuvres, has already been introduced in previous work [19]. This automaton type is named by *Automotive-Controlling Timed Automaton (ACTA)*.

Rule 170. In Fig. 5, we show the timed automaton that corresponds to the previously introduced behaviour diagram (cf. Fig. 3). This timed automaton contains a set of locations (L_0, L_1, \dots) that mainly conform to the states from the behaviour diagram; only the state L_2 represents a decision process: “Is there a safe gap?” (previously seen in Fig. 3).




Figure 5: Rule 170: Timed Automaton

In timed automata, taking a transition is a process that happens immediately, without any passing of time. However, in reality, most actions and decisions of an AV will take time, even if it may only be some milliseconds. Thus, we add a clock variable x to our timed automaton for rule 170 to specify that checking for a safe gap for pedestrians and road users takes some time (from the time interval $[t_1, t_2]$) and does not happen immediately.

The USL-TR guard $\neg\exists c : pc(c) \wedge \neg pa(E)$ on the transition from location L_1 to L_2 checks that there does not exist a potential collision with any road user ($\neg\exists c : pc(c)$) and that no pedestrian is ahead ($pa(E)$). Only if this guard holds it is possible to move from location L_1 to L_2 . In the guard on the transition to location L_3 , we add the check for a safe gap on an intersection $sg_I(E)$ to the guards that we already described for the previous transition. This assures that the location L_3 is only entered if the conditions above are satisfied and additionally there exists a safe gap large enough for the car E on the road junction. If any of these conditions do not hold, the transition back to location L_1 is taken, and all the checks will be done again before the AV may safely enter the road junction.

Rule 171. Fig.6 depicts the timed automaton that we construct for Rule 171. Notice that to move from location L_1 to L_2 the guard $ob(Stop) \wedge ob(SWL)$ should be satisfied, i.e. the two traffic signs “Stop” and “Solid With Line” need to be acknowledged by the AV. We again add a clock variable x and clock constraints (t_1, t_2) to model that stopping at the traffic signs and checking for a safe gap does not happen immediately. Again, the AV may only move onto the intersection if a safe gap in the intersection has been identified ($sg_I(E)$).




Figure 6: Rule 171: Timed Automaton

In this section, we introduced extended timed automata representing the UK Highway Code traffic rules 170 and 171. In future work, we intend to formalise further traffic rules from the UK Highway Code, but also driving regulations from other countries (e.g. from the German “Straßenverkehrsordnung” [8]) in the same manner. For now, we leave the actions of the AVs abstract (e.g. stopping at a traffic sign or setting a turn signal). However, the ACTA formalism from [19] allows adding such controller actions. Equally, for the agent and environment models that we intend to build for the verification, we will include these actions. In the next section, we describe these future steps toward our agent and traffic rule verification endeavour.

5 On the Verification of Traffic Rules

Following our diagram from Fig. 1, starting from the timed automata that we obtained in the previous section, we intend to combine two different methods for the verification of traffic rules:

- Checking for conflicts and inconsistencies between traffic rules via UPPAAL model-checking [4]
- Adding agent and environment models for the verification of traffic-rule-following agents.

Identifying conflicts and inconsistencies with UPPAAL. Traffic rules from the known driving regulations are prone to having inconsistencies and conflicts with each other. Consider the following example: worldwide, most driving regulations contain a rule stating that cars must stop at a red traffic light. However, many driving regulations also include the traffic rule of the “green arrow for turning right” (“left” in case of the UK Highway Code). This rule allows turning right at a red traffic light if there is enough space. Generally, these two rules might be in conflict with one another, as one allows something that the other rule forbids. To altogether find such conflicts, we intend to analyse parallel compositions of traffic rule timed automata (including those that we introduce in Sect.4). We will start from a UPPAAL implementation of a protocol for performing turn manoeuvres at intersections from previous work [5], which is modelled as an ACTA and is close to the timed automaton for rule 170 (cf. Fig. 5), but it does not yet precisely follow traffic rules. After finding conflicting rules with the help of UPPAAL, we need to consider how AVs may solve such conflicts. While humans use common sense for this task, we envision a situation-dependent prioritisation of AV traffic rules.

Agent and Environment Models. We intend to obtain the corresponding Agent and Environment models from the extended timed automata. When getting the Agent model, it is possible to translate it into a sketch of a *Belief-Desire-Intention (BDI)* agent [7]. A BDI agent is defined by a set of plans, where each plan usually comprises the following: **i.** a trigger event; **ii.** guards **iii.** and a set of actions. The first two elements of the BDI agent can be obtained from the guards represented in the timed automata using USL-TR. E.g., the guard $\neg\exists c : pc(c) \wedge \neg pa(E)$ (cf. Fig. 5) would generate a trigger event `~potential-collision` and a guard `~pedestrian-ahead` for the BDI agent plan. The third element can be obtained from the locations in the timed automata. E.g., the location L_2 is responsible for checking for a safe gap; as a result, it would generate an action for the BDI agent plan named `checkForSafeGap`. Wrapping it all together, we have the following BDI agent plan:

```
~potential-collision : ~pedestrian-ahead <- checkForSafeGap;
```

We could then endow the agent with traffic rules and implement the corresponding agent’s environment. Therefore, extending the agent-based approach for formally verifying the agent’s behaviour against traffic rules previously developed in [3].

A Digital Highway Code. In this research preview, we sketch a methodology (cf. Fig. 1) in which we start with natural language traffic rules and, step by step, develop a formalised and verifiable digital version of traffic regulations. By developing techniques for finding and resolving conflicts between traffic rules and introducing verifiable, traffic rule-following, autonomous agents, we pave the way toward a Digital Highway Code. Such a Code will allow AVs to understand, follow and explain traffic rules, which will be necessary for a successful, safe and desirable joint future for humans and AVs.

However, the research is in an early stage and there are several interesting aspects that we will need to keep in mind for a further development of our approach:

- Some traffic rules for humans are obsolete for AVs, because they simply do not apply to AVs or because AVs follow them implicitly, “by design”. For instance, a traffic rule that forbids checking your smartphone while driving, which invariantly holds for an AV.
- A mechanism is needed to handle “broken” traffic rules. E.g., if a traffic light is defect or a traffic sign is not recognisable due to graffiti or dirt, another traffic rule or plan of action must exist to prevent the AV from stopping forever.
- As we sketched before, if some traffic rules are in conflict, the AV needs to be able to prioritise traffic rules to assess which rule must be obeyed and which rule may be temporarily ignored (so called “trolley problems” or “moral dilemmas” cf. [6]). Such a prioritisation of traffic rules cannot be defined globally, but will be situation-dependent.

References

- [1] Rajeev Alur & David L. Dill (1994): *A Theory of Timed Automata*. *Theoretical Computer Science* 126(2), pp. 183–235, doi:10.1016/0304-3975(94)90010-8.
- [2] Gleifer Vaz Alves, Louise Dennis & Michael Fisher (2020): *Formalisation and Implementation of Road Junction Rules on an Autonomous Vehicle Modelled as an Agent*. In: *Formal Methods. FM 2019 International Workshops*, Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 217–232, doi:10.1007/978-3-030-54994-7_16.
- [3] Gleifer Vaz Alves, Louise Dennis & Michael Fisher (2021): *A Double-Level Model Checking Approach for an Agent-Based Autonomous Vehicle and Road Junction Regulations*. *Journal of Sensor and Actuator Networks* 10(3), doi:10.3390/jsan10030041. Available at <https://www.mdpi.com/2224-2708/10/3/41>.
- [4] Gerd Behrmann, Alexandre David & Kim G. Larsen (2004): *A Tutorial on Uppaal*. In: *Formal Methods for the Design of Real-Time Systems*, Springer, pp. 200–236, doi:10.1007/978-3-540-30080-9_7.
- [5] Christopher Bischofink & Maike Schwammberger (2019): *Verification of Fair Controllers for Urban Traffic Manoeuvres at Intersections*. In: *Formal Methods. FM 2019 International Workshops - Porto, Portugal, October 7-11, 2019, Revised Selected Papers, Part I, Lecture Notes in Computer Science* 12232, Springer, pp. 249–264, doi:10.1007/978-3-030-54994-7_18.
- [6] Jean-François Bonnefon, Azim Shariff & Iyad Rahwan (2016): *The social dilemma of autonomous vehicles*. *Science* 352(6293), pp. 1573–1576, doi:10.1126/science.aaf2654. arXiv:<http://science.sciencemag.org/content/352/6293/1573.full.pdf>.
- [7] Michael E. Bratman (1987): *Intentions, Plans, and Practical Reason*. Harvard University Press, doi:10.2307/2185304.
- [8] Bundesrepublik Deutschland (2013): *Straßenverkehrsordnung (StVO)*. Available at https://www.gesetze-im-internet.de/stvo_2013/index.html.
- [9] Werner Damm & David Harel (2001): *LSCs: Breathing Life into Message Sequence Charts*. *Formal Methods in System Design* 19(1), pp. 45–80, doi:10.1023/A:1011227529550.
- [10] Werner Damm, Eike Möhlmann, Thomas Peikenkamp & Astrid Rakow (2018): *A Formal Semantics for Traffic Sequence Charts*. In: *Principles of Modeling - Essays Dedicated to Edward A. Lee on the Occasion of His 60th Birthday*, Springer, pp. 182–205, doi:10.1007/978-3-319-95246-8_11.
- [11] Louise A. Dennis (2017): *Gwendolen Semantics: 2017*. Technical Report ULCS-17-001, University of Liverpool, Department of Computer Science.
- [12] Louise A. Dennis, Michael Fisher, Matthew P. Webster & Rafael H. Bordini (2012): *Model Checking Agent Programming Languages*. *Automated Software Engineering* 19(1), pp. 5–63, doi:10.1007/s10515-011-0088-x.
- [13] Department for Transport (2017): *Using the road (159 to 203) - The Highway Code - Guidance - GOV.UK*. Available at <https://www.gov.uk/guidance/the-highway-code/using-the-road-159-to-203>.
- [14] Martin Fowler (2003): *UML Distilled: A Brief Guide to the Standard Object Modeling Language*, 3 edition. Addison-Wesley Longman Publishing Co., Inc., USA.
- [15] Frank B. Gilbreth & Lillian M. Gilbreth (1921): *Process Charts*. The American Society of Mechanical Engineers, New York, USA. Available at <https://www.thegilbreths.com/resources/Gilbreth-Process-Charts-1921.pdf>.
- [16] IRES Simulationstechnologie GmbH. (2022): *OpenSCENARIO*. Available at <http://www.openscenario.org>. Accessed: 2022-08-05.
- [17] Ben Moszkowski (1985): *A Temporal Logic for Multilevel Reasoning About Hardware*. *Computer* 18(2), pp. 10–19, doi:10.1109/MC.1985.1662795.
- [18] Ernst-Rüdiger Olderog, Anders P. Ravn & Rafael Wisniewski (2017): *Linking Discrete and Continuous Models, Applied to Traffic Manoeuvres*. In Michael G. Hinchey, Jonathan P. Bowen & Ernst-Rüdiger Olderog,

- editors: *Provably Correct Systems*, NASA Monographs in Systems and Software Engineering, Springer, pp. 95–120, doi:10.1007/978-3-319-48628-4_5.
- [19] Maïke Schwammberger (2018): *An abstract model for proving safety of autonomous urban traffic*. *Theoretical Computer Science* 744, pp. 143–169, doi:10.1016/j.tcs.2018.05.028.
- [20] Maïke Schwammberger & Gleifer Vaz Alves (2021): *Extending Urban Multi-Lane Spatial Logic to Formalise Road Junction Rules*. In: *Proceedings Third Workshop on Formal Methods for Autonomous Systems, FMAS 2021, Virtual, 21st-22nd of October 2021, EPTCS* 348, pp. 1–19, doi:10.4204/EPTCS.348.1.